2025 edition

NetMob

BOOK OF ABSTRACTS

Conservatoire National des Arts et Métiers - CNAM
Paris, France

The primary conference on the analysis of mobile phone datasets in social, urban, societal and industrial problems.

October 8-10 2025

Inria | le cnam | ifp Energies nouvelles | orange

Université Gustave Eiffel | FRANCE 2030 PROGRAMME DE RECHERCHE DIGITALISATION ET DECARBONATION DES MOBILITES | berlin | Qualcomm

about NetMob

The proliferation of mobile phone datasets has facilitated an enhanced comprehension of human communication, social interaction, urban movement, and mobilization. NetMob serves as the preeminent conference dedicated to the analysis of such datasets within the context of social, urban, societal, and industrial challenges. This conference convenes a diverse array of researchers, policymakers, and industry specialists to engage in discourse regarding the most recent advancements in mobility data and its application across a wide spectrum of contexts.

The NetMob 2025 Data Challenge provided researchers with a unique GNSS-based mobility dataset from the Île-de-France region. Over 3,300 volunteers carried BT-Q1000XT Bluetooth® GNSS receivers, which capture position and velocity with high precision through DGPS and Doppler shift measurements. Data was collected between October 2022 and May 2023, with full consent, anonymization, and validation via follow-up interviews. The result is a rich, 7-day, high-resolution view of mobility behavior across Greater Paris. Learn more: https://arxiv.org/abs/2506.05903.

program committee

Abdelmounaim Bouroudi, INRIA France
Albert Ali Salah, Utrecht University, The Netherlands
Aline Carneiro Viana, Inria, France
Alexandre Chasse, IFP Energies Nouvelles, France
Andrea Araldo, IMT, France
Anne Josiane Kouam, TU Berlin, Germany
Antonio Loureiro, UFMG, Brazil
Chen Zhong, UCL, United Kingdom
Daphné Tuncer, ENPC, France
Fernando Molano Ortiz, INRIA, France
Giovanna Miritello, Vodafone, United Kingdom
Giulia Pullano, Georgetown University, USA
Jussara Almeida, UFMG, Brazil
Laetitia Gauvin, IRD, France
Lei Dong, Peking University, China
Luca Pappalardo, ISTI-CNR, Italy
Luca Vassio, Politecnico di Torino, Italy
Marco Fiore, IMDEA Networks Institute, Spain
Mariano Beiro, UBA, Argentina
Nadjib Achir, INRIA, France
Pedro Velloso, CNAM, France
Razvan Stanica, INSA, France
Riccardo Di Clemente, Northeastern University London, United Kingdom
Takahiro Yabe, NYU, United States
Veena Mendiratta, Northwestern University, USA
Wellington Lobato, INRIA, France
Yoshihide Sekimoto, University of Tokyo, Japan

steering committee

Esteban Moro Egido - Northeastern University
Vanessa Frias-Martinez - University of Maryland
Leo Ferres - IDSUDD, ISI Foundation & Telefónica
Vincent Blondel - University of Louvain
Marco Fiore - IMDEA Networks Institute & Net AI
Renaud Lambiotte - University of Oxford
Alex 'Sandy' Pentland - MIT

organizing committee

Aline Carneiro Viana - INRIA
Nadjib Achir - INRIA
Razvan Stanica - INSA Lyon
Alexandre Chasse - IFP Energies Nouvelles
Pedro Velloso - CNAM
Daphné Tuncer - ENPC
Cedric Adjih - INRIA
Anne Josiane Kouam - TU Berlin
Abdelmounaim Bouroudi - INRIA
Fernando Molano Ortiz - INRIA
Wellington Lobato - INRIA

keynotes

Mobility data for pandemic response

Dr. Vittoria Colizza, INSERM

Abstract: The COVID-19 pandemic harshly tested decades of preparedness and scientific progress. Yet, its unprecedented nature catalyzed a major shift in the use of mobility data from mobile phones—from a traditional approach for mapping spatial disease transmission to an operational proxy for inferring social distancing and behavioral adaptation. Drawing on our experience during the pandemic response, this talk will explore how mobility data became essential for real-time assessment of intervention impact, monitoring of behavioral response, and informing public health decision-making—alongside its constraints and limitations.

Bio: Vittoria Colizza is director of Research specializing in the mathematical and computational modeling of infectious diseases. She uses mobility data to anticipate the spread of epidemics and pandemics and to guide public health policies in their prevention and control strategies. Her contributions have been recognized with numerous distinctions: Knighthood of the Order of Merit of the Italian Republic in 2020; the Irène Joliot-Curie Prize awarded by the French Academy of Sciences and the French Ministry of Research in 2021; and Knighthood of the Order of the Légion d'honneur in 2023.

Collecting and sharing human data with meaningful privacy protection

Dr. Luc Rocher, University of Oxford

Abstract: Mobile phone data has revolutionised our understanding of human mobility, social networks, and urban dynamics. However, the privacy protections we rely on when working with call detail records, location traces, and behavioural data may be far weaker than commonly assumed. In this talk, Luc Rocher will examine vulnerabilities in anonymisation techniques applied to large-scale datasets and discuss how to best use sensitive data. The talk will address critical questions for the mobile phone data research community: What makes some individuals more vulnerable than others to privacy attacks? How can we balance usability and reproducibility with genuine data protection?

Bio: Luc Rochier is a senior researcher and expert on issues related to the use of personal data. A specialist in privacy and algorithms, Luc's research focuses on the risks associated with the massive use of personal data and on the means of making digital technologies more responsible, secure, and in service of the public interest. Luc's research has been cited by the European Commission, the OECD, the World Bank, the World Economic Forum, the FTC, by European data protection authorities, in legal cases in the United States, and has led to amendments to the UK's data protection bill.

Contents

1 Main Conference: Talks **8**

1.1 Use of mobile phone data to measure behavioral response to SMS evacuation alerts 9
1.2 Evaluating the Predictability of Disaster Evacuation Behavior using LLM Agents . . . 11
1.3 Disruption and Adaptation of Social Connectedness at Third Places after Disaster Displacement . . . 13
1.4 Targeting mobile money cash transfers for emergency relief using Call Detail Records in the Democratic Republic of the Congo (DRC) . . . 15
1.5 On the utility of Digital Contact Tracing on empirical contact network . . . 17
1.6 Changes in social contact time following digital SARS-CoV-2 post-exposure notification: A retrospective case-control study at a US university. . . . 19
1.7 PRIMO-Scrap: towards automating the cartography of smartphone apps for mobility . 21
1.8 Economic spillover effects of electric vehicle charging station placement on local businesses: a staggered adoption study . . . 23
1.9 Understanding the Impacts of EV Charging Stations on Surrounding Businesses . . . 25
1.10 Causal inference in the city: Improving Urban Policy Evaluation Through Mobility-Aware Methods . . . 27
1.11 The Urban Impact of AI: Modeling Feedback Loops in Next-Venue Recommendation . 29
1.12 Mobile Phone Data Sheds Light on the Localized Activity Patterns of Informal Workers 31
1.13 Revealing Hidden Parental Mobility Behaviors Through Large-scale GPS Data . . . 33
1.14 Measuring Bias in CDR Mobility Indicators Arising from Multiple SIM Ownership . . . 35
1.15 Maintaining Connectivity in Subterranean Emergency Response: A 5G and Wi-Fi Mesh Mobility Dataset . . . 37
1.16 The Anatomy of Olympic Games: a Mobile Traffic Demand Perspective . . . 39
1.17 Exploring urban event dynamics using mobile application measurements . . . 41
1.18 Unequal Journeys to Food Markets: Continental-Scale Evidence from Open Data in Africa . . . 43
1.19 An Urban Geography of Mobile Traffic Demands . . . 45
1.20 Exploring the Digital Life of Urban Parks Through Mobile Traffic . . . 47
1.21 Spatial Heterogeneity in Human Mobility Responses to London's Ultra-Low Emission Zone Expansion . . . 49
1.22 The temporal dimension of experienced segregation in cities . . . 51
1.23 Mobility behavior moderates residential partisan segregation in US metropolitan areas 53
1.24 Tracing Behavioural Changes of High Street Visitors in London Using Mobile Location Data . . . 55
1.25 When Proximity Falls Short: Inequalities in Commuting and Accessibility in Santiago, Chile . . . 57
1.26 Hybrid work has increased income segregation at third places in US cities . . . 59
1.27 Beyond Distance: Mobility Neural Embeddings Reveal Visible and Invisible Barriers in Urban Space . . . 61
1.28 Correcting temporal bias in mobility data using time-use surveys . . . 63
1.29 Towards Assessing Accessibility Resulting from Integrating Demand-Responsive and Conventional Public Transport with Travel Time Uncertainty . . . 65
1.30 Beyond Proximity: Rethinking Urban Accessibility Through Human Mobility Patterns in United States . . . 67
1.31 Understanding Urban-Rural Disparities in Mobility Inefficiency for Colombia, Mexico, and India . . . 69
1.32 Potential route diversification in road networks . . . 71
1.33 Urban semantic: dynamic land use through human mobility patterns . . . 73

1.34 Working behaviour and the spatial complexity of economic activity 75
1.35 Bias estimates for MPD-derived mobility statistics in low-and middle-income countries 77
1.36 Predicting Vehicle Trajectories based on Latitude and Longitude Positioning Historical Data 79
1.37 Temporal Convolutional Network Forecasting of Tourism Demand Indicators Based on GPS Data: A Case Study in Japan 81
1.38 Parameter-free robust stop-detection with HDBSCAN 83
1.39 Tackling the Ambiguity of Stop Detection in Human Mobility Data: A Model-Agnostic Evaluation Framework 85
1.40 TransGM: Transferable Gravity Models for Adaptive Urban Policy 87
1.41 Characterizing User Exposure in Mobility Data: Insights from Varying User Behavioral Profiles 89
1.42 Understanding collective and individual phenomena in urban mobility using mobile phone datasets 91

2 Data Challenge: Talks **93**
2.1 Protecting participants or population? Comparison of k-anonymous Origin-Destination matrices 94
2.2 On commuting duration and the daily travel time budget 96
2.3 On the Relationship between Space-Time Accessibility and Leisure Activity Participation 98
2.4 A Simulation Study on Equitable Mobility During City Emergencies, Focusing on Vulnerable Groups 100
2.5 Summarizing Trajectories to Uncover Mobility Behaviors and Urban Contexts 102
2.6 Persons with reduced mobility facing invisible obstacles: A contrastive analysis on mobility patterns and behaviours based on GNSS data 104
2.7 Integrating Temporal Context into Mobility Embeddings: A Doc2Vec Approach to User Profiling 106
2.8 Who Breathes More Smog? Modeling Individual PM2.5 Exposure and Inequality with GPS Mobility Data 108
2.9 Day and Night in Paris: Vibrancy and Perceived Security Analysis from Mobility Data 110
2.10 Machine Learning-Based NetMob25 Mobility Model for Network Simulators 112
2.11 Mobility Under Pressure: Socioeconomic Inequities in Urban Travel Behavior During Urban Strikes 114
2.12 Mobility Patterns, Segregation, and Pollution Exposure in the Context of the Parisian Low-Emission Zone 116
2.13 Impact of income-stratified mobility mixing and vaccination uptake on epidemic progression: a modeling approach on emerging pandemic flu 118
2.14 Commuting Patterns Across Household Structures: An Empirical Analysis of Time and Mode Choices in Île-de-France 120
2.15 Post-Disaster Route Reliability Assessment Using Drone-Aided Mobility Data from NetMob25 122

3 Main Conference: Posters **124**
3.1 Modeling Base Station Metadata Geolocation 125
3.2 Cross-City Travel Mode Identification from Sparse Mobile Phone Data with Transformer Adapters 127
3.3 Estimating Traffic Flows from Vehicle Trajectories Based on Sparse Mobile Phone Geolocation Data 129
3.4 The Promise and Pitfalls of Mobile Phone Data for Recreational Fishing: Evaluating Potential, Biases, and Representativeness 131

3.5 Modeling Mobility with Graph Neural Networks to Predict Neighbourhood Culture . . . 133
3.6 Mapping Urban Mobility with Interest Networks: Cross-Platform and User Behavior Insights . . . 135
3.7 Evaluating Deep Learning for Transport Mode Detection on Crowdsourced Smartphone GPS Data . . . 137
3.8 Modeling Urban Air Quality Using Taxis as Sensors . . . 139
3.9 Measuring and Mapping Patterns of Space Utilization for Characterizing Human Mobility in Epidemiological Contexts Using a Novel Digital Data Set . . . 141
3.10 Decomposing International Migration Flows with Mobile Trace Data: Trends, Seasonality, and Shocks . . . 143
3.11 How People Navigate in Crowds Indoors: Evidence from Trajectories in UCL PEARL . 145
3.12 Enhancing mobile network data with GPS and satellite imagery for accurate population detection . . . 147
3.13 The (un)predictability of the earthquake-induced mobility . . . 149
3.14 Mixture models uncover polycentric urban structure from GPS mobility data . . . 151
3.15 Activity Space Segregation Through the Lens of Income Outliers in London . . . 153
3.16 Spatial Modeling of NOx Emissions from Vessel Trajectories Using AIS Data . . . 155
3.17 Robustness of Stop Detection in Sparse Location Datasets . . . 157

4 Data Challenge: posters **159**
4.1 Persona-Aware Mobility Management: Detecting Sustained Demand, Speed, and Coverage Shifts . . . 160
4.2 Identifying First/Last Mile Deserts in the Paris Region Using Mobile GPS Data . . . 162
4.3 Segmenting Human Mobility Behavior Under Internet Quality Pressure in Paris . . . 164
4.4 Behavioral Segmentation of Older Adult Mobility in Metropolitan Areas . . . 166
4.5 Beyond Averages: Inequalities in Mobility-Based PM2.5 Exposure . . . 168
4.6 Decoding Urban Interest: The Roles of Purpose, Transportation, and Scale . . . 170
4.7 Urban Real Accessibility Evaluation Based on GPS Data . . . 172
4.8 Optimizing Ad-Placement for Socio-Demographic Segments Using Mobility Data . . . 174
4.9 Measuring Individual Exploration in Human Mobility . . . 176
4.10 The City as Experienced: Measuring Contextual Exposure . . . 178
4.11 Zero-Shot Transportation Mode Classification in Île-de-France . . . 180
4.12 Gendered Mobility Differences in the Era of Flexible Work . . . 182
4.13 Generation of Realistic Geolocated Agendas in a Digital Twin . . . 184
4.14 Divided by Design: Transport, Amenities, and Inequality . . . 186
4.15 Understanding Urban Trips: Mode and Purpose Prediction . . . 188
4.16 Revisiting Algorithms for Activity Type Estimation from GPS Data . . . 190
4.17 Improving Activity Classification Using Detailed Demographics . . . 192
4.18 Multi-Modal Mobility for Generating Electricity . . . 194
4.19 Beyond Willpower: Constraints and Mobility-Related CO_2 . . . 196
4.20 Higher-Order Network Analysis of Human Mobility Data . . . 198
4.21 How Daily Mobility Practices Reshape Greenspace Exposure . . . 201
4.22 Graphing the Jam: GNSS-Powered Traffic Congestion Modeling . . . 203
4.23 Optimizing Store Location Based on Urban Accessibility . . . 205
4.24 When Movement Reveals Poverty . . . 207
4.25 Evaluating Socio-Spatial Inequalities in Paris's 15-Minute City . . . 209
4.26 Paris as a 15-Minute City: An Explainable AI Perspective . . . 211
4.27 The Privacy Bounds of Fine-Grained Human Mobility . . . 213
4.28 Advanced Crime–Mobility Relationship Modeling . . . 215
4.29 Discovering Functional Urban Zones via Heterogeneous Graph Convolution . . . 217

4.30 Gendered Mobility Patterns in Île-de-France from GNSS Traces 219

Main Conference

talks

Use of mobile phone data to measure behavioral response to SMS evacuation alerts

ERICK ELEJALDE[1,*], TIMUR NAUSHIRVANOV[2], KYRIAKI KALIMERI[3], ELISA OMODEI[2], MÁRTON KARSAI[2,4], LORETO BRAVO[5], AND LEO FERRES[3,5,*]

[1] L3S Research Center, Leibniz University Hannover, Hannover, Germany
[2] Department of Network and Data Science, Central European University, Vienna, Austria
[3] ISI Foundation, Torino, Italy
[4] National Laboratory for Health Security, HUN-REN Rényi Institute of Mathematics, Budapest, Hungary
[5] Institute of Data Science, Universidad del Desarrollo, Santiago, Chile

1. INTRODUCTION

The pervasive adoption of mobile technology has fundamentally transformed emergency alert systems. With 112 mobile-cellular subscriptions per 100 inhabitants globally and 133 in Chile as of June 2024 [1], Short Message Service (SMS) notifications have emerged as a widely adopted and viable method for mass communication during emergencies [2]. A key advantage of SMS-based alert systems is their functionality on any mobile phone (e.g., does not require internet connectivity), making them particularly effective in reaching diverse populations, including those in remote or low-resource settings. Chile's Sistema de Alerta de Emergencia (SAE), implemented after the 2010 earthquake, ensures broad accessibility of these critical warnings. Despite the proven utility of SMS and alert systems [3], their effectiveness in prompting protective action remains an area that requires further investigation. Prior evaluations of messaging effectiveness have primarily relied on surveys [4]. While valuable, this approach is prone to self-reporting biases regarding actual human behavior during disasters, and it often lacks the high-resolution data necessary to accurately assess evacuation patterns. Consequently, significant gaps persist in understanding how individuals respond to emergency messages in real-world wildfire events.

We address these critical gaps by leveraging anonymized mobile network data from the devastating Valparaíso wildfires in Chile on February 2, 2024, which resulted in 137 deaths and affected over 16,000 people [5]. The research examines population-wide behavioral responses to evacuation alerts, focusing on three key questions: (1) How does the timing and sequence of emergency alerts influence evacuation behavior? (2) To what extent do socioeconomic factors shape evacuation patterns? (3) What are the unintended consequences of non-targeted alerts on population movement? By integrating high-resolution behavioral data with policy-relevant questions, this work aims to optimize emergency communication strategies, thereby enhancing public safety and resilience in the face of escalating climate-driven hazards [6]. The use of mobile network data represents a methodological advancement, moving beyond self-reported intentions to capture objective, real-time population movement [7].

2. DATA AND METHODS

The study uses anonymized eXtended Detail Records (XDRs) from Telefónica Movistar, a major South American mobile operator with a 28% market share as of June 2024. We aggregate network activity into 15-minute intervals by counting distinct phone IDs per tower. A decrease in connectivity at specific towers indicates population displacement, serving as a proxy for mobility. Two periods were analyzed: a baseline (January 25-28, 2024), representing normal usage, and the fire week (February 1-4, 2024), during a wildfire event. Evacuation alerts from the National Emergency Office of Chile (SENAPRED) provided data on dates and warned locations, which were geocoded to determine their latitude and longitude. Towers within a 5 km radius of these points were labeled as "warned towers," while others served as controls. This method acknowledges the imprecision of alerts, addressing the challenge of public interpretation that could lead to unintended over-evacuations. The socioeconomic status of each tower was estimated based on their geolocation, using 2017 census data from the National Institute of Statistics of Chile, categorized into Low, Medium, and High Socioeconomic Groups (SEGs) based on higher education levels [8].

We introduce the Relative Evacuation Index (REX), a normalized metric to quantify population displacement. Defined as $REX = (n_t - n_b)/(n_t + n_b)$ (where n_t is the number of unique mobile devices at time t, and n_b is the corresponding count for the same tower, on the same day of the week, and exact time of the day during the baseline period), REX maps relative population change onto the interval $[-1, 1]$. This enhances the reliability and comparability of evacuation rates across diverse geographical and demographic units, making the findings more robust and generalizable.

A quasi-experimental design, specifically a Controlled Interrupted Time Series (CITS) analysis [9], was employed to evaluate the causal effect of the warning alert intervention on the Relative Evacuation Index (REX). This methodology allowed for a comparison between affected (treatment) and non-affected (control) cellular towers. To account for response latency, a delayed intervention effect of 1 hour and 15 minutes was incorporated into the model, identified as optimal through model calibration [10, 11]. The CITS model includes terms for time, intervention, group (affected/control), and their interactions. The Newey-West variance estimator was used to correct for heteroskedasticity and autocorrelation, ensuring robust inference.

3. RESULTS

The analysis of mobile network connectivity during the Valparaíso wildfires revealed clear behavioral changes in response to evacuation alerts. The first alert, issued on February 2, 2024, at 16:45, caused an immediate decline in connectivity, dropping to an average of 132 connections and indicating rapid population displacement from the area. Some regions experienced over an 80% decrease in connectivity within 1.5 hours of the alert (Figure 2). While subsequent alerts were issued, connectivity levels gradually returned to pre-alert figures within 8 hours, aligning with typical circadian rhythms by 00:00-01:00 on February 3. Interestingly, the first alert on February 3 at 10:45 had no noticeable impact on connectivity, highlighting diminishing effectiveness and suggesting alert fatigue. The alert fatigue observed here likely stems from a complex mix of perceived risk and message redundancy. Future alert systems should adopt adaptive mes-

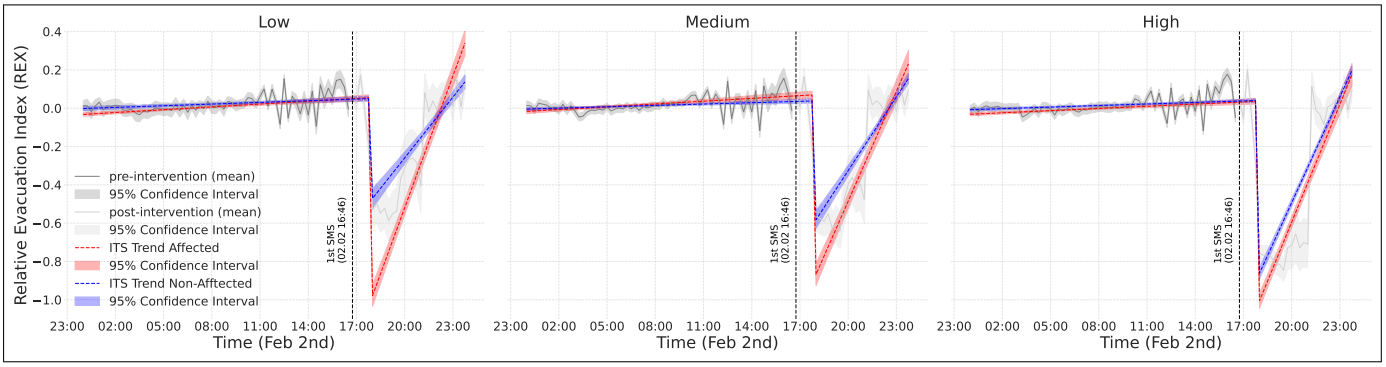

Fig. 1. Controlled Interrupted Time Series (CITS) stratified by three SEGs (Low (left), Medium (center), High (right) for February 2nd, 2024.

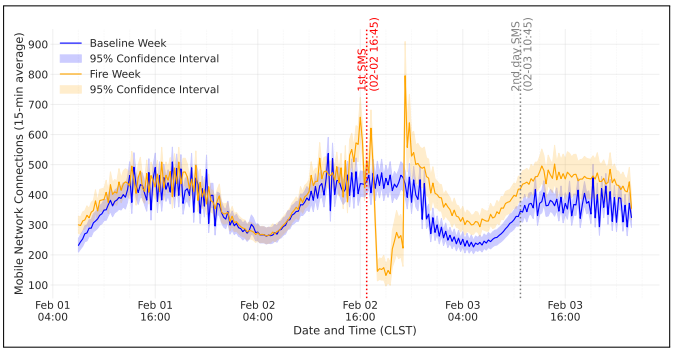

Fig. 2. Mobile phone tower connectivity patterns during the Valparaiso wildfire evacuation event.

saging strategies that vary in content, urgency, or target areas based on the evolving disaster context and past alert history.

Significant socioeconomic disparities were evident in evacuation patterns and recovery following an alert. High-income regions responded most vigorously to the alert, achieving higher evacuation levels faster. By 18:00 (1 hour 15 minutes post-alert), over 95% of towers in high socioeconomic groups (SEGs) had reached at least 50% evacuation rate, compared to 92% in low and 82% in medium areas. The starkest difference was observed at the 85% evacuation rate, with 86% of high SEG towers, 74% of low SEG towers, and 69% of medium SEG towers achieving this level. Using the Controlled Interrupted Time Series (CITS) model confirmed these disparities, revealing that higher socioeconomic status correlated with quicker evacuation responses to warnings (Figure 1). Post-intervention, recovery trends showed lower SEGs returning more quickly; however, this did not translate into a swifter recovery to pre-intervention levels. Lower SEGs displayed a sustained pattern of hyperactivity during the second day, indicating that socioeconomic status affects evacuation and recovery. Thus, equitable emergency planning should address both material and informational disparities to enhance trust and improve access to resources.

A notable finding highlighted evacuation behavior in non-warned areas. Both warned and non-warned towers reacted strongly to the evacuation alert, with the REX showing significant movement: -0.47 for low SEG, -0.56 for medium SEG, and -0.81 for high SEG. This indicates many chose to evacuate regardless of direct warnings. Evacuation was particularly pronounced in high SEG towers, which showed less difference between warned and non-warned areas—only about a 15% variation—suggesting a tendency for precautionary evacuations. In contrast, lower-income populations relied more on government alerts, exhibiting a stronger differential response (Figure 1). High SEG populations may even "overreact" to alerts, evacuating without direct threats. These behaviors pose challenges for disaster management, potentially leading to congestion and resource strain on roads meant for areas in immediate need. This highlights a critical need for more granular, precise geo-targeting of alerts and potentially public education campaigns that balance individual safety with collective efficiency in disaster response.

With climate change increasing wildfire severity, optimizing emergency communication systems is essential for population safety. Our findings contribute to evidence supporting the benefits of integrating mobile connectivity patterns into disaster management—a critical adaptation pathway for societies facing increasing climate-driven hazards.

REFERENCES

1. Subsecretaría de Telecomunicaciones de Chile (Subtel), "Abonados móviles en chile - junio 2024," Subtel - Gobierno de Chile (2024). Accessed: 2025-03-11.
2. G. Haddow and K. S. Haddow, *Disaster communications in a changing media world* (Butterworth-Heinemann, 2013).
3. J. Cinnamon, S. K. Jones, and W. N. Adger, "Evidence and future potential of mobile phone data for disease disaster management," Geoforum **75**, 253 (2016).
4. E. J. Carlson, H. Bean, C. Ratcliff, *et al.*, "Do 360-character wireless emergency alert messages work better than 90-character messages? testing the risk communication consensus," J. contingencies crisis management **32**, e12587 (2024).
5. M. Cursino, "Chile forest fires: More than 100 dead in valparaíso region," BBC News (2024).
6. M. Burke, S. Heft-Neal, J. Li, *et al.*, "Exposures and behavioural responses to wildfire smoke," Nat. Hum. Behav. **6**, 1351 (2022). Publisher: Nature Publishing Group.
7. T. Yabe, N. K. Jones, P. S. C. Rao, *et al.*, "Mobile phone location data for disasters: A review from natural hazards and epidemics," Comput. Environ. Urban Syst. **94**, 101777 (2022).
8. T. Naushirvanov, E. Elejalde, K. Kalimeri, *et al.*, "Evacuation patterns and socioeconomic stratification in the context of wildfires," EPJ Data Sci. **14**, 23 (2025).
9. J. Lopez Bernal, S. Cummins, and A. Gasparrini, "The use of controls in interrupted time series studies of public health interventions," Int. J. Epidemiol. **47**, 2082 (2018).
10. R. F. Fahy, G. Proulx, and others, "Toward creating a database on delay times to start evacuation and walking speeds for use in evacuation modeling," in *2nd international symposium on human behaviour in fire,* (Boston, MA, USA, 2001), pp. 175–183.
11. Y. Liu, Z. Zhang, and Z. Mao, "Analysis of influencing factors in pre-evacuation time using Interpretive Structural Modeling," Saf. Sci. **128**, 104785 (2020).

Evaluating the Predictability of Disaster Evacuation Behavior using LLM Agents

DongHak Lee[1,2], Jiayi Weng[1], Vaidehi Raipat[1,2], and Takahiro Yabe[1,2]

[1] Center for Urban Science and Progress, Tandon School of Engineering, New York University, 370 Jay St. Brooklyn, NY 11201
[2] Department of Technology Management and Innovation, Tandon School of Engineering, New York University, 2 Metrotech, Brooklyn, NY 11201
[*] takahiroyabe@nyu.edu

1. INTRODUCTION

Large language models (LLMs) have been shown to mimic social behavior in different contexts (see [1]) given its rich contextual information embedded into the model. Park et al. even showed that endogenous interactions among LLM agents result in emerging social dynamics often observed in human communities, such as information propagation and collective decision making. Motivated by this breakthrough, several works have leveraged LLM agents for human mobility prediction [2–4]. However, very few works have attempted to apply LLM agents for predicting human mobility patterns during disaster events. To what extent these generative AI models are able to replicate and predict mobility and behavior during crisis situations is unknown and deserves investigation [5], for many downstream and practical applications like evacuation shelter prepositioning, emergency resource allocation, and recovery planning. To address this gap, we simulate evacuation dynamics during disaster scenarios using an LLM-based agent simulation framework, and systematically compare their predicted evacuation behaviors against real-world ground-truth evacuation patterns derived from mobility data. This approach enables a comprehensive evaluation of the capabilities and limitations of LLM agents in replicating human decision-making and movement patterns during crisis. To formally guide our experiments, we pose the following research questions. To guide our experiments, we ask whether LLM agents with synthetic personas can make realistic evacuation decisions based on pre-disaster mobility and demographic data, and how their predictability can be improved. We also ask whether these agents can meaningfully incorporate sociodemographic context when making decisions, and whether this context improves prediction quality. Finally, we consider whether fine-tuning the LLM produces outputs that align with real-world evacuation behavior.

2. DATA AND METHODOLOGY

A. Synthetic Persona and Ground Truth Data

To simulate individual decision-making, we generate synthetic agents using SafeGraph POI visit data and Census-based demographics. Mobility flows inform lifestyle routines, while ACS variables (e.g., population, income, education, race) enrich agent heterogeneity.

To evaluate model predictability, we use anonymized GPS data from Cuebiq, capturing device-level movement during the 2021 Marshall Fire in Boulder County, CO (Dec 30, 2021 – Jan 1, 2022) and 2022 Mill Fire in Siskiyou County, CA (Sep 2, 2022 - Sep 13, 2022). Here ground-truth evacuation is defined as spending the night (8PM–7AM) away from home, a criterion also embedded in the LLM prompt for consistency.

B. LLM Agent-Based Simulation Framework

We simulate 6,355 synthetic agents (2% of the local population) sampled from SafeGraph POI visit data, which provides aggregated mobility flows across Census Block Groups (CBGs) in Boulder County, Colorado. Each agent is assigned a persona prompt that combines mobility behavior, sociodemographic traits from the American Community Survey (ACS), and proximity to the fire.

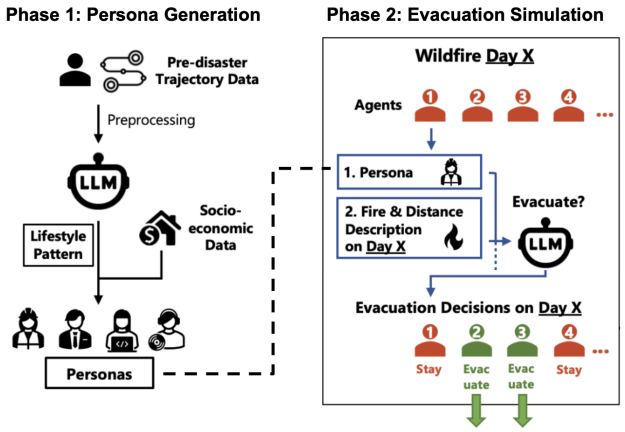

Fig. 1. LLM Agent-Based Simulation Framework. The model simulates the movement of a synthetic population. It consists of agents constructed using large language models which are fed with specific personas, socioeconomic characteristics, and disaster information.

These prompts are passed to a Meta-Llama-3.1-8B-Instruct model to generate binary evacuation decisions. Figure 1 illustrates this simulation framework, which supports scalable and interpretable modeling, as well as experiments to assess the impact of contextual cues on prediction quality.

C. Fine-tuning and Scalability

While zero-shot evaluation shows that pretrained LLMs can partially capture evacuation dynamics, we explore ways to improve predictive performance using synthetic data. Specifically, we fine-tune the Meta-Llama-3.1-8B-Instruct model using Low-Rank Adaptation (LoRA), which updates a small set of rank-decomposed weight matrices rather than all model parameters.

To construct training data, we generate synthetic agents based on 2021 Marshall Fire mobility and sociodemographic data, assigning each agent a binary evacuation outcome so that aggregated rates match ground-truth estimates from Cuebiq GPS data. Each training example consists of a persona prompt describing pre-disaster routines, ACS-based demographics, and fire proximity paired with the evacuation label. The LoRA-tuned model is trained to map these cues to evacuation decisions.

We first evaluate the adapted model on a held-out set of synthetic agents from the Marshall Fire and compare predictions with observed evacuation rates. We then conduct cross-event evaluation by generating synthetic personas for the 2022 Mill

Fire in Siskiyou County, CA, and assessing whether the model reproduces evacuation patterns. These experiments demonstrate the scalability of the fine-tuned model, yielding robust predictions in previously unseen disaster scenarios.

3. SIMULATION RESULTS

A. Zero-Shot LLM Agents

We begin by evaluating the zero-shot performance of the pre-trained LLM, using prompts that include only lifestyle-based mobility patterns and distance to the fire. As shown in Figure 2, the model tends to overestimate evacuation rates, relying heavily on proximity. The correlation with observed behavior is moderate ($r = 0.485$), indicating limited alignment without tuning or additional context.

Next, we test whether adding sociodemographic context improves prediction quality. Prompts are enriched with Census-derived sociodemographic variables alongside fire proximity. This results in a stronger correlation with ground-truth evacuation rates ($r = 0.577$), suggesting that the LLM can incorporate demographic signals to reason beyond simple distance-based heuristics.

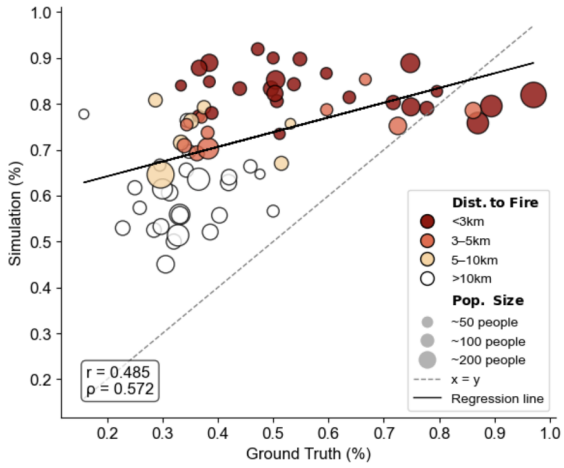

Fig. 2. Zero-shot prediction of evacuation without sociodemographic context. Each point represents a Census Tract, colored by distance to fire.

B. LoRA Fine-Tuned LLM Agents

To test whether synthetic personas can help LLM agents better reproduce real-world evacuation behavior, we fine-tune the model using the Low-Rank Adaptation (LoRA) method. We generate synthetic personas for 30% of the population and assign evacuation labels to each agent, matching the observed evacuation rate within each CBG.

After fine-tuning the LLM on synthetic personas labeled with ground-truth evacuation outcomes—and enriching prompts with sociodemographic context such as median income, total population, and white population percentage—the model achieves a strong correlation of $r = 0.790$ with observed evacuation rates.

These results suggest that LLM agents can not only respond to contextual cues but also internalize empirical behavioral patterns through lightweight tuning. The fine-tuned model captures key real-world trends such as the distance-decay effect in evacuation behavior, reinforcing the potential of LLMs as interpretable, data-driven simulators in crisis scenarios. As a next step, we plan to evaluate the fine-tuned model on synthetic personas for the 2022 Mill Fire to test its ability to generalize to a new disaster event.

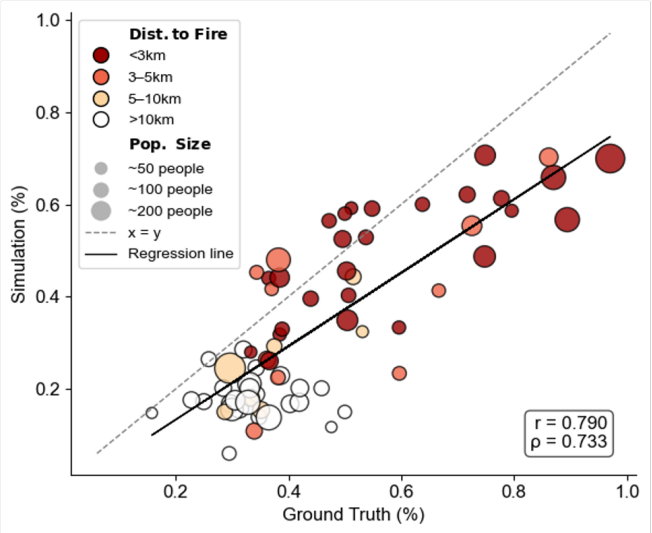

Fig. 3. LLM prediction after LoRA fine-tuning with extended sociodemographic context.

4. CONCLUSION AND DISCUSSION

This study demonstrates that large language models (LLMs) can serve as realistic and interpretable agents for simulating evacuation behavior during disasters. Our study highlights the potential of LLMs not just as text generators, but as cognitive agents capable of learning and replicating human decision-making behavior anchored in real-world data.

We find that incorporating contextual sociodemographic information such as population size, racial composition, and income improves prediction quality compared to using proximity alone. Moreover, fine-tuning LLMs with synthetic persona data—sampled to reflect empirical evacuation rates—further enhances predictive accuracy while preserving interpretability.

Building on this foundation, we will next look into enhancing LLM reasoning capabilities such as incorporating chain-of-thought prompting and multi-step decision processes to better capture complex human behavior. We also plan to test the approach across multiple cities, scaling to larger populations and more diverse scenarios to systematically evaluate its robustness and transferability.

Funding. T.Y. and D.L. acknowledge support by the National Science Foundation under grant number CMMI 2343646.

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. J. S. Park, J. O'Brien, C. J. Cai, *et al.*, "Generative agents: Interactive simulacra of human behavior," in *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST)*, (ACM, 2023).
2. C. Gao, X. Lan, N. Li, *et al.*, Humanit. Soc. Sci. Commun. **1**, 1 (2024).
3. J. Wang, R. Jiang, C. Yang, *et al.*, "Large language models as urban residents: An llm agent framework for personal mobility generation," in *Advances in Neural Information Processing Systems*, , vol. 37 (NeurIPS, 2024).
4. J. Piao, Y. Ye, J. Li, *et al.*, arXiv preprint arXiv:2502.08691 (2025).
5. Z. Wu, R. Peng, X. Han, *et al.*, arXiv preprint arXiv:2311.06330 (2023).

Disruption and Adaptation of Social Connectedness at Third Places after Disaster Displacement

VAIDEHI RAIPAT[1,2] AND TAKAHIRO YABE[1,2,*]

[1] Department of Technology Management and Innovation, Tandon School of Engineering, New York University
[2] Center for Urban Science and Progress, Tandon School of Engineering, New York University
[*] ty2520@nyu.edu

1. INTRODUCTION

Change in population dynamics such as displacement and return trigger changes in social and spatial patterns in communities. Disaster-induced population displacement could significantly disrupt community-level social network structures, causing fragmentation and weakening of existing social ties [1, 2]. Some studies suggest that communities characterized by dense social networks enable rapid recovery of their network functionality following perturbations. They also suggest that perturbations tend to decrease density and increase the heterogeneity in complex networks. [3–5]. However, fewer studies have quantified the changes in social network structures resulting from post-disaster population dynamics [2, 6].

In this study, we aim to quantify the impact of disaster-induced displacement on the social connectedness of affected communities in the aftermath of a disaster. We analyze the dynamics after the Marshall Fire, a fast-moving, three-day-long wildfire that occurred between December 30, 2021 and January 1, 2022, and led to widespread evacuations and displacement in various cities of Boulder county. Specifically, we ask: 1) Do disasters disrupt or reinforce socio-spatial network structures?; 2) Is the post-disaster connectivity in the social network random or structured?; and 3) How do these changes impact the levels of segregation in affected communities?

Our findings reveal that disaster-induced displacement results in significant fragmentation of social networks in affected communities. However, these changes are not random, they exhibit patterns of structural adaptability and reorganization as the actual post-disaster network maintains connectivity more than the random baseline model. We also find that as compared to a random baseline, socially similar ties in the post-disaster network were more likely to increase. This indicates increased social segregation after disaster.

2. RESULTS

A. Identifying existing and potential social connections

Third spaces are essential for creating strong bonding ties by building a sense of belonging, identity and diversity in a community [7, 8]. We define social connectedness as potential and existing social ties. Such ties are captured from repeated spatio-temporal co-location in socially meaningful places that facilitate voluntary, non-obligatory social behavior. Social points-of-interest (POIs) have been defined as cafes, bars, parks, gyms, or recreational venues that individuals visit during evening hours and weekends. For this study, we focus on individuals residing in Boulder County, Colorado, before the occurrence of the Marshall Fire.

We leverage anonymized mobile phone GPS trace data provided by Cuebiq, which captures individual user trajectories in the form of 'stops'. Each stop includes spatial coordinates, a timestamp, and dwell time, allowing us to reconstruct detailed patterns of mobility and place-based interactions both before and after the disaster. We construct a weighted social network for the pre-disaster period by linking individuals who co-visit these POIs and spend at least 15 minutes in the same space, over a span of one month. The edge weights are proportional to the frequency of these interactions, capturing both existing and potential social ties. Then, we aggregate the co-location network into a Census Block Group (CBG) level contact network to avoid any noise in the individual level data. In the contact-network the connectivity between CBGs depends on the number of individual level ties between them. The weights of the edges are the sum of the weights of individual level ties. We then build a similar post-disaster network, focusing on the same set of individuals who were residents of the disaster-affected area before disaster. This enables us to isolate the structural transformation of the local social network caused specifically by the disruptions associated with the disaster. The post-disaster network clearly shows visible fragmentation and reduction in connectivity compared to the pre-disaster network, as shown in Figure 1. This highlights the structural breakdown of community-level social ties following the disaster.

B. Post-disaster fragmentation of social connectedness

To measure fragmentation, we compare structural metrics of pre-disaster network with post-disaster contact network, and find that there are significant changes in network structure in the aftermath of the disaster 1c). We see that 329 of 618 pre-disaster nodes (individuals) were lost from the dataset and the number of edges dropped from 13757 to 2225 leading to a drop in mean degree as well (8.53 to 2.68). Then we compute the Degree Centrality, Closeness Centrality and Clustering Coefficient for both pre and post-disaster networks and find a steep drop in their values. Degree Centrality, drops from 44.52 to 15.39 (Figure 2a). Closeness centrality also experiences a drop in the mean value from 0.44 to 0.36 (Figure 2b). This decline reflects a significant increase in path lengths, indicating that CBGs in the post-disaster network require more steps, on average, to reach one another. Clustering coefficient also decreases from 0.49 to 0.39, reflecting a reduction in triadic closures, suggesting fewer "friends of friends" connections in the post-disaster period (Figure 2c). Altogether, these changes suggest that disaster-induced displacement severely fragments the existing social connectedness.

C. Adaptability within the network in response to disruption

To evaluate if these network changes are the result of an underlying social processes or simply the result of random population loss, we construct a randomized post-disaster baseline by removing the same number of nodes from the pre-disaster network as observed in the actual post-disaster scenario. Here, rather than using the real evacuated nodes, we randomly select nodes for removal controlling on total node loss per CBG. Then we aggregate the individual level co-location network to a CBG level contact network. This process is bootstrapped 20 times, allowing us to generate a distribution of centrality measures

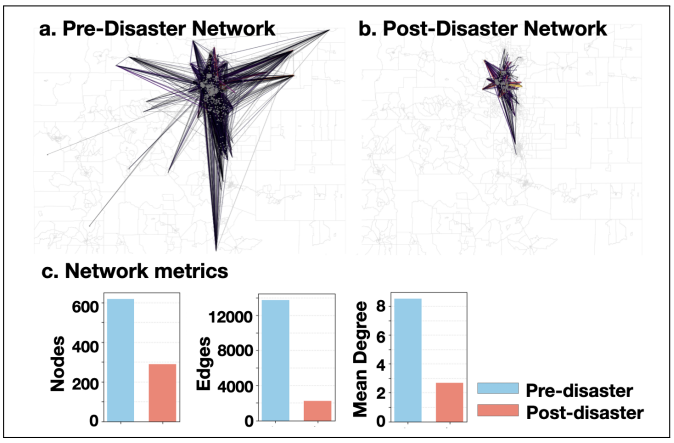

Fig. 1. Significant post-disaster fragmentation of social connectedness. Graphs show the pre- and post-disaster networks to visually illustrate the impact of disaster on social networks. An edge between two CBGs exists if at least one user from each CBG is co-located and edge weight is sum of the co-location frequencies between all user–user pairs. **a.** Pre-disaster contact network. **b.** Post-Disaster contact network. **c.** Comparing change in number of nodes, edges and mean degree in pre and post periods. The post-disaster network reveals a visibly fragmented structure with fewer edges indicating disruption in habitual encounters and a fragmentation of social connectivity.

and network properties with a robust confidence interval for comparison against the actual contact network. We find that, the raw degree centrality in random node loss is much lower than in actual post-disaster scenario (44.5210 in pre-disaster network to 15.39 post-disaster, and 8.72 in randomized baseline) (Figure 2a), suggesting that there may be underlying structure in who evacuates and which ties are preserved. Closeness centrality experiences a substantial drop post-disaster (pre-disaster: 0.44, post-disaster: 0.36, random baseline: 0.28, $p < 0.001$), indicating increased average path lengths and reduced global efficiency (Figure 2b). We find that closeness in the post-disaster network is higher than in the randomized scenario highlighting that the remaining nodes maintain better global reachability than a random removal. Clustering coefficient also declines significantly from 0.49 pre-disaster to 0.39 post-disaster and a lowest of 0.31 in the random baseline model($p = 0.05$) (Figure 2c), reflecting a loss in triadic closures. The steep drop in the randomized secnario suggests that the actual post-disaster network preserves more local clustering than expected under random loss. This therefore reinforces that there exists a non-random self-organization in preservation of social connectedness post-disaster.

D. Post-disaster induced fragmentation increases segregation

To evaluate whether network changes led to greater connectivity among socio-demographically similar communities than expected under the random baseline, we computed similarity between every unique CBG pair using cosine similarity of sociodemographic variables (race, income, and age). We then fit a linear regression model of the form ($W_{i,j} = \beta_1 S_{i,j} + \beta_2 d_{i,j}$), where $W_{i,j}$ is weight of edges between CBGs i and j, $S_{i,j}$ is the Sociodemographic similarity between CBG's i and j and, $d_{i,j}$ is the euclidean distance between them. We find that Socio-demographic similarity (controlling for distance) is associated with the strength of ties ($p - value > 0.001$). Coefficient for socio-demographic similarity for the post-disaster scenario ($\beta_1 = 0.012202$) is much

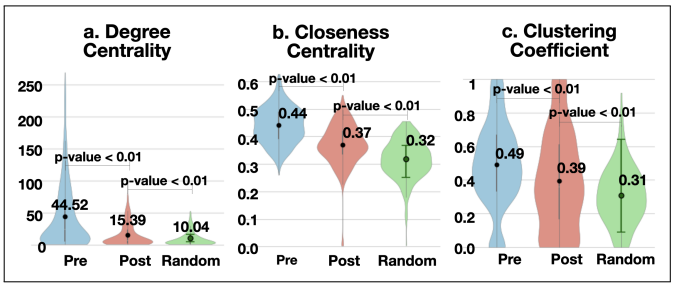

Fig. 2. Structural adaptability in preserving neighborhood-level social ties. Violin plots compare the distribution of Degree centrality, closeness centrality and clustering coefficient values across the pre-disaster, post-disaster, and random baseline networks.

higher than for the random baseline model ($\beta_1 = 0.005399$) indicating that social connectivity among socio-demographically similar CBGs increased disproportionately in the post-disaster scenario relative to random expectations.

3. DISCUSSION

Our findings demonstrate that disaster-induced population displacement leads to significant disruption of social connectedness within affected communities. This aligns with prior research suggesting that disasters fragment the social fabric by displacing residents from everyday social setting. [2, 8, 9]. However, our analysis goes beyond previous work by showing that this disruption is not purely random. While displacement results in sharp decreases in network size and connectivity, the structure of the remaining network exhibits characteristics of adaptability. We observe that closeness centrality and clustering coefficients in the actual post-disaster network are consistently higher than in randomized counterfactual scenarios. These patterns suggest that some social ties are more likely to persist even under external stressors like evacuation. This structural adaptability indicates that communities reorganize in ways that preserve connectivity and local support structures despite perturbations. As next steps, we plan to analyze whether these effects differ for ties across different socio-demographic groups.

REFERENCES

1. D. P. Aldrich and M. A. Meyer, Am. Behav. Sci. **59**, 254 (2015).
2. T. Yabe, K. Tsubouchi, N. Fujiwara, *et al.*, J. Royal Soc. Interface **17**, 20190532 (2020).
3. J. Gao, B. Barzel, and A.-L. Barabási, Nature **530**, 307 (2016).
4. M. E. Newman and M. Girvan, Phys. review E **69**, 026113 (2004).
5. A. M. Sadri, S. V. Ukkusuri, and H. Gladwin, Nat. Hazards Rev. **18**, 04017005 (2017).
6. H. Deng, D. P. Aldrich, M. M. Danziger, *et al.*, Humanit. Soc. Sci. Commun. **8**, 1 (2021).
7. R. Oldenberg, *The Great Good Place: Cafés, coffee shops, bookstores, bars, hair salons, and other hangouts at the heart of a community* (Berkshire, 2023).
8. R. Chetty, M. O. Jackson, T. Kuchler, *et al.*, Nature **608**, 122 (2022).
9. D. P. Aldrich, *Building resilience: Social capital in post-disaster recovery* (University of Chicago Press, 2012).

Funding. T.Y. and V.R. acknowledge support by the National Science Foundation under grant number CMMI 2343646.

Acknowledgment. The authors would like to thank Spectus who kindly provided us with the mobility dataset for this research through their Data for Good program.

Disclosures. The authors declare no conflicts of interest.

Targeting mobile money cash transfers for emergency relief using Call Detail Records in the Democratic Republic of the Congo (DRC)

Galina Veres[1,*], James Harrison[1], Sarah Allen[1], Marie Sophie Pocha[2], Véronique Lefebvre[1]

[1]*Flowminder Foundation, Magnus Ladulåsgatan 1, 118 65 Stockholm, Sweden*
[2]*GiveDirectly, Inc.,P.O. Box 3221, New York, New York 10008,*
*galina.veres@flowminder.org

Received XX Month XXXX; revised XX Month, XXXX; accepted XX Month XXXX; posted XX Month XXXX (Doc. ID XXXXX); published XX Month XXXX

1. INTRODUCTION

Armed conflicts and instability in the DRC often cause large-scale internal displacement, including in remote areas where identifying and targeting internally displaced persons (IDPs) by field operations is slow, costly, and geographically limited. To address these challenges, Flowminder has developed a novel targeting method using daily anonymised statistics derived from Call Detail Records (CDR) to detect displacements in near real-time. Together with GiveDirectly and a mobile network operator (MNO), we created and operated a live CDR-driven cash transfer pipeline based on this novel targeting method, as a pilot project in the provinces of North and South Kivu (DRC). The targeting pipeline was repeated for 8 armed attacks within 5 months, reaching over 5,600 displaced subscribers in total. Crucially, the method and pipeline dramatically improved response times: after the first attack, 490 displaced people received $125 each within 7 days, compared to the typical 100–130 day delays of traditional aid. Our CDR-driven remote cash transfer programming pipeline is currently operational to respond to ongoing displacement events in the North-East of the DRC, enabling rapid and targeted humanitarian cash transfers when and where needed.

2. A CDR-DRIVEN CASH TRANSFER PIPELINE

GiveDirectly operates unconditional cash transfer, including in emergencies, to people displaced by crisis and conflicts. Identification and targeting displaced people often take a long time with difficulties to reach remote and unsafe areas. CDR data contain information which can help to detect displacements remotely if the area of interest has mobile phone coverage before, during and after an incident. We developed a CDR-driven cash transfer pipeline which allows rapid response in conflicts and crisis situations. The pipeline consists of five steps:

1. **Incident Identification**: GiveDirectly identifies displacement-triggering events of interest through local situational awareness and intelligence gathering.
2. **Calculation of mobility statistics from CDRs**: The MNO uses SQL queries provided by Flowminder to aggregate CDR traces across subscribers and compute mobility statistics, updated every day for a time period spanning the event and for the region of interest. Any statistics lower than 15 subscribers is redacted for privacy purposes so that the dataset can be shared with Flowminder for analysis.
3. **Displacement Inference**: Using a novel methodology described below, Flowminder detects the locations, timings and number of probable displacements related to the events of interest and potentially newly detected events. Displacement estimates are subsequently approved by GiveDirectly, and shared with the MNO.
4. **Data Sharing**: From the estimates and locations computed by Flowminder the MNO selects the subscribers most likely to be displaced and shares the list with GiveDirectly.
5. **Verification of the subscriber selection**: GiveDirectly verifies identities and eligibility of subscribers by contacting them.
6. **Cash transfer**: GiveDirectly enrolls individuals and initiates remote cash transfers.

3. METHODOLOGY FOR DETECTING DISPLACEMENTS

Displacement patterns are inferred from changes in mobile subscribers' "stay locations" and network activity. Four key metrics are used: **Relocations In** (a number of subscribers whose stay location (main location over the past X days) changes to a given cluster of cell towers on a specific day); **Relocations Out** (a number of subscribers whose stay location changes away from a given cluster on a specific day); from these we can compute the net relocation flow to a given cluster (in - out). Then **Pairwise Consecutive Relocations** (the count of subscriber movements from one cluster to another) provide the origin and destination of relocation flows. Finally **Active Cell Counts** (a number of cells with at least one active subscriber for a given day in a cluster) enable us to check for network outages and potential call redirections, which risk being confused for movements.

Our **Displacement Detection method** from CDRs is based on the hypothesis that forced displacements will manifest as statistically significant outliers in mobility flows

(changes in 'stay' location) following specific incidents (e.g., conflict events). Clusters of cell towers exhibiting outliers in flows (relocations) are automatically flagged as 'affected' by the conflict, in particular those with a deficit in net relocations and excess relocations out, and from these the methodology estimates: the number of potentially displaced subscribers and the outlier magnitude, to prioritise clusters with the most significant outlier flows.

However, mobility outliers can be produced by network outages, which themselves can stem from the conflict events. To rule out effects of outages and call redirections we only select locations where the number of active cells remains stable over the time period of interest. Additional validation involves identifying host clusters from the pairwise relocations, as statistically significant excesses of arrivals from affected clusters [1]. Detection of multiple host clusters, if supported by local knowledge, strengthens confidence in inferred displacements.

All inferred displacements are then reviewed by GiveDirectly, incorporating field intelligence. Validated events and estimates trigger secure and anonymised subscriber selection by the MNO, and ID transfers from the MNO to GiveDirectly for identity verification and potential cash transfers.

4. RESULTS: DISPLACEMENT INFERENCE

In order to prevent fraud through exposure of the detailed targeting methodology, as previously experienced by GiveDirectly [2], we are unable to disclose further details and results from the methodology in use in our operational CDR-driven cash transfer pipeline.

However, we have applied the same displacement inference methodology to CDRs from another MNO in the DRC (Vodacom), and can show example results (Fig. 1). Fig. 1 shows inferred displacement flows between health zones of South Kivu during the week of 8–14 April 2025. An estimated 3,400 newly displaced subscribers relocated between different health zones (and an additional 3,600 subscribers may have relocated within their health zone - not shown). The main origins of displacements were Nyantende (also recorded by ACLED), Mulungu, and Ruzizi (also reported by country sources), while the main host zones were Nyangezi, Uvira, and Kabare. The largest detected displacement flow was from Nyantende to Nyangezi, involving approximately 700 newly displaced subscribers. No cash transfers were initiated for these inferred displacements from Vodacom CDRs. Statistics on cash transfers done via the operational pipeline with another MNO are provided in the abstract.

5. FURTHER DEVELOPMENTS

This study demonstrates the feasibility of using mobile phone data to rapidly identify and assist IDPs in conflict settings at scale. The developed pipeline could become a new asset for humanitarian crisis response: it can be used in remote places with mobile network coverage; and crucially it permits very rapid cash assistance. Currently the entire process can last as little as 6 or 7 days, which is a step change for emergency relief, particularly in remote areas.

Given the success of our pilot in the DRC, all partners are working on maturing the solution, focussing on improving the robustness and validation of results (rather than speed, as at least 6 days are required to cover steps 1-6), while also ensuring that potential operational delays are minimised through automation. Flowminder is refining its CDR analysis methods, inferring displacements from trajectory-based features and not only from outliers in group mobility, improving the accuracy of displacement detection and its sensitivity (enabling detection of smaller numbers of displacements per location and shorter displacement distance for example). We will validate and compare the performances of our CDR displacement detection methods by conducting a phone survey in North East DRC, currently in preparation.

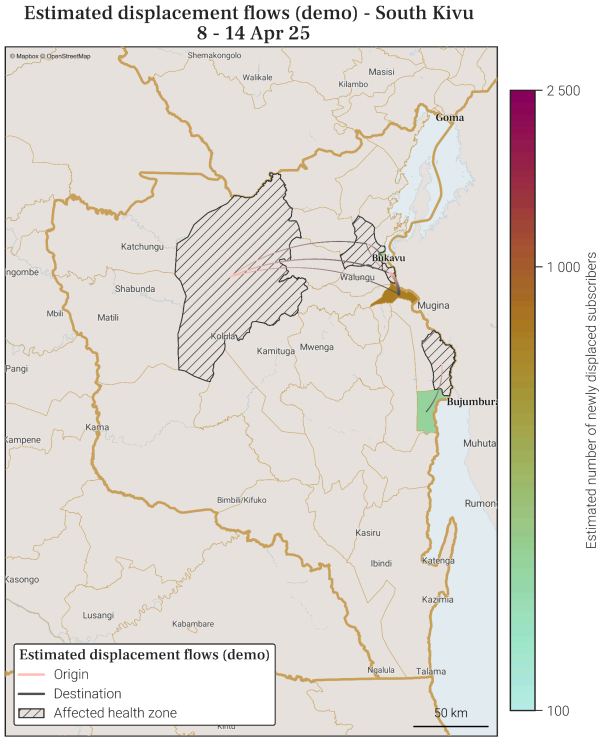

Fig. 1. Estimated new displacements (8-14 April 2025), derived in privacy conscious manner from Vodacom CDRs [3]

Funding. The work was funded by GiveDirectly, Inc., Sub-Grant Agreement No. 91174

Data Availability. Data underlying the results presented in this paper are not publicly available due to privacy.

Disclosures. The authors declare no conflicts of interest

References

1. Flowminder Foundation, August 2023: Gang violence in Port-au-Prince , Haiti. Haiti August 2023 Gang violence, Port-au-Prince | Report 1 (ENG)
2. Report: risks we faced delivering cash in 2024 | GiveDirectly
3. Flowminder Foundation, Showcase: Mobility estimates derived from Vodacom Call Detail Record: population mobility and displacements in the Democratic Republic of the Congo (DRC). https://www.flowminder.org/resources/publications-reports/mobility-estimates-in-drc-population-mobility-and-displacements

On the utility of Digital Contact Tracing on empirical contact network

CHUAN LI[1,2,5], VINCENT GAUTHIER[2], MIGUEL NUNEZ-DEL-PRADO[3,6], HUGO ALATRISTA-SALAS[4*], AND HASSINE MOUNGLA[1,2]

[1] *LIPADE, Université Paris Cité, France.*
[2] *SAMOVAR, Telecom SudParis, Institut Polytechnique de Paris, Palaiseau, France.*
[3] *The World Bank, USA.*
[4] *De Vinci Higher Education, De Vinci Research Center, Paris, France.*
[5] *Sorbonne Université, Paris, France.*
[6] *Universidad Peruana de Ciencias Aplicadas, Lima, Peru.*
[*] *hugo.alatrista_salas@devinci.fr*

1. INTRODUCTION

In response to the pandemic, which exposed health system vulnerabilities, the Peruvian government launched *Perú en tus manos* on April, 2020, during Peru's first COVID-19 lockdown. It was is designed to enable citizens to detect potential exposure hotspots and facilitate self-assessment and contact tracing [1]. The app mixes self-reported symptoms (usually via a digital triage form linked to national ID), real-time location data (GPS and potentially Bluetooth interactions), and reported case data to display an in-depth heat-map at the neighborhood scale. App users can be warned of entering zones of high risk and report their health status upon mandatory quarantine, aiding the contact tracing activities with the Ministry of Health and the *Te Cuido Perú* task force.

This study assesses both the effectiveness and limitations of digital contact tracing (DCT) technology in capturing human mobility during the pandemic. It uses a subset of the Perú en tus manos dataset, which comprises GPS locations, timestamps, and self-reported infection status from 1.66 million users. Of these users, 663,250 were active, and 80,068 reported a positive SARS-CoV-2 diagnosis.

A temporal analysis was conducted to validate the dataset against World Health Organization (WHO) statistics[1], revealing consistent trends in case numbers from March to June 2020. User activity fluctuated, peaking in August with the most registered devices and declining between April and July with the fewest active users. Mobility increased during relaxed restrictions and decreased during lockdowns, correlating with infection peaks. These results confirm the dataset's reliability in reflecting mobility and transmission patterns. The temporal analysis of the *Perú en tus manos* dataset reveals trends consistent with official COVID-19 statistics from the WHO, as shown in Figure 1(A). Both data sources, smoothed using a 7-day moving average, exhibit a marked increase in cases from March to June 2020, highlighting the dataset's reliability in tracking infection dynamics. Figure 1(B) illustrates fluctuations in user activity, where the proportion of registered devices peaked in August 2020, but the number of active users declined from about 0.5 million in April to its lowest in July. Mobility trends, displayed in Figure 1(C), indicate that daily trips per active device rose during periods of relaxed restrictions (July and September) and dropped during strict lockdowns, with a positive correlation observed between mobility and COVID-19 case peaks in April and June. These figures collectively demonstrate the relationship between human movement patterns, public health measures, and infection rates.

[1] World Health Organization (WHO) report: https://www.who.int/countries/per

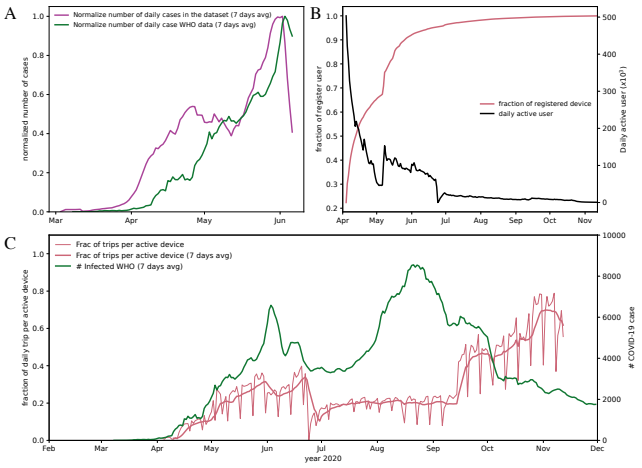

Fig. 1. **A**) Evolution of the Normalized number of COVID-19 Case (7 days average) in Peru WHO compare with normalized number of the case declared in our own dataset from April to July 2020. **B**) Fraction of daily active DCT devices from April to December 2020. **C**) Evolution of the number of daily trips per active device in Peru from April to November 2020.

2. MAIN FINDINGS

In order to measure the effectiveness of the *Peru en tus manos* DCT app, this work is divided into two major stages: 1) analyzing the mobility data captured by the DCT app, and 2) understanding how mobility patterns impact infection dynamics.

A. Analyzing the mobility data captured by the DCT application

An analysis of mobility data from the Digital Contact Tracing (DCT) application was conducted to assess changes in travel patterns during the COVID-19 pandemic, using the radius of gyration [2] to measure average distances traveled from individuals' residences. Covering April to November 2020, the results indicate that lockdown measures effectively limited overall mobility in Lima. Three phases were identified: a period of minimal movement under strict restrictions (April–June), followed by gradual increases as measures relaxed (July onward). Spatial variations revealed that city center residents traveled shorter distances due to nearby essential services, while suburban residents exhibited higher mobility, often commuting to the city for work despite restrictions.

In order to efficiently identify locations with the highest interaction density, the interaction duration and the temporal variability of user activity must be considered. We employed a bipartite

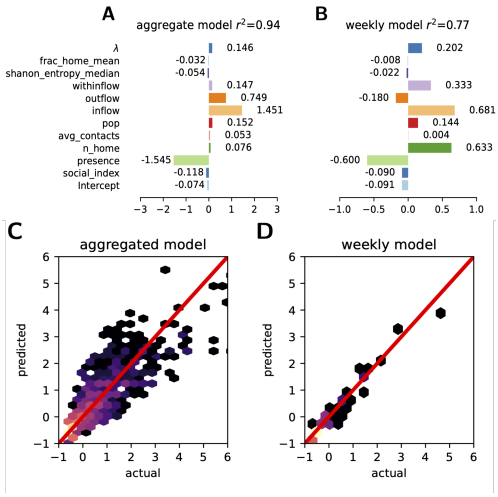

Fig. 2. Performance of spatial regression models incorporating both spatial adjacency and mobility features. (A, B) Regression coefficients of key predictors for the aggregated and weekly models, respectively, including OD flow components and social indicators. (C, D) Predicted vs. actual infection counts for the weekly and aggregated models. Color intensity reflects data point density. The red diagonal line represents the ideal fit. The aggregated model achieves higher predictive accuracy, while the weekly model captures finer-grained temporal dynamics.

stream graph model [3] to capture the temporal and structural aspects of interactions between DCT devices and points of interest (POIs), which are represented by H3 cells over time intervals. Analysis of the average and normalized temporal degree revealed that individuals tended to avoid crowded areas during lockdowns or in the early stages of an outbreak. Additionally, measures of centripetality and anisotropy showed that Lima's city center remains the main hub of urban mobility. POIs in this area are highly connected, which highlights its pivotal role in the city's mobility network.

A temporal contact analysis of DCT devices was later conducted by projecting the bipartite stream graph onto a stream graph to approximate the actual temporal contact network among users. Two assumptions guided the analysis: devices connected to the same POIs (H3 cells) during the same time period were considered to be in contact, and a device was considered to remain at a POI between consecutive location records. Users were classified into two groups: those who tested positive for SARS-CoV-2 and a control group. The results revealed that infected individuals were more likely to form small clusters of 1–30 people but were less likely to attend large gatherings than the control group. Additionally, the infected group exhibited a higher average temporal degree seven days prior to confirmed infection. Comparisons of clique size and temporal degree further supported the hypothesis that infected users exhibited distinct behavioral patterns with greater contact rates. These findings underscore the significance of group size and temporal degree as indicators of individual contact behavior, even when measured using spatial proximity data.

B. Understanding mobility patterns and infection dynamics

This study analyzes the influence of mobility and social strata on infection dynamics by leveraging a comprehensive Origin–Destination (OD) dataset and WHO reported infected population. A neighborhood graph based on the standardized H3 grid preserved spatial structure while capturing local interactions. Our objective is to understand the underlying mechanisms driving infection rates rather than to optimize prediction accuracy.

Figures 2(A) and 2(B) show the regression coefficients for the aggregated and weekly models, respectively. The aggregated model is more accurate ($R^2 = 0.95$) than the weekly model ($R^2 = 0.77$). Both models show that inflow and within-region mobility positively influence transmission. However, outflow appears to stabilize infections over time in the aggregated model. Socioeconomic status is negatively correlated with infection rates, indicating greater vulnerability in lower-income areas. Increased home presence significantly reduces infections, and both models exhibit spatial dependence. Figures 2(C) and 2(D) further confirm the superior predictive performance of the aggregated model.

Finally, we explore how socioeconomic factors influence disparities in mobility and infection risk during the pandemic. Using entropy-based metrics, we found that wealthier individuals exhibited more limited spatial movement and lower entropy values, indicating that they visited fewer distinct locations. The stay-at-home rate (SHR) revealed that higher-income groups had the highest SHR, reflecting a greater capacity for self-isolation [4], likely due to remote work or access to private transportation. Conversely, lower-income individuals demonstrated lower SHRs and higher mobility, often necessitated by factors such as commuting or accessing essential services. These findings underscore the socioeconomic constraints on mobility and isolation capacity.

3. CONCLUSION

This study evaluates the effectiveness of DCT by analyzing the *Perú en tus manos* dataset, which includes GPS data, timestamps, and self-reported infection statuses from 1.66 million users. Despite a low adoption rate of approximately 5.7% and limited individual tracing, findings indicate a strong correlation between mobility trends and public health interventions, suggesting compliance among DCT users with Peru's public health policies. The analysis of contact patterns revealed that individuals who contracted the virus had more contacts than those in the control group. While the data from DCT applications has proven valuable for policymakers in assessing social distancing strategies and estimating potential infections, concerns regarding user engagement—primarily due to privacy issues and app design—pose challenges to DCT efficacy. A key finding of this study is that meaningful insights can be drawn from coarser spatial data. The study highlights the need for alternative data collection methodologies, such as utilizing mobile phone metadata, to better understand contact and mobility patterns, especially in scenarios where user engagement with DCT applications is low.

REFERENCES

1. A. Alvarez-Risco, S. Del-Aguila-Arcentales, and J. A. Yáñez, "Telemedicine in peru as a result of the covid-19 pandemic: perspective from a country with limited internet access," The Am. J. Trop. Med. Hyg. **105**, 6 (2021).
2. M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, "Understanding individual human mobility patterns," nature **453**, 779 (2008).
3. M. Latapy, T. Viard, and C. Magnien, "Stream graphs and link streams for the modeling of interactions over time," Soc. Netw. Anal. Min. **8** (2018).
4. L. Lucchini, O. D. Langle-Chimal, L. Candeago, *et al.*, "Socioeconomic disparities in mobility behavior during the covid-19 pandemic in developing countries," EPJ Data Sci. **14**, 25 (2025).

Changes in social contact time following digital SARS-CoV-2 post-exposure notification: A retrospective case-control study at a US university.

Daniel T Citron,[1,*] Po-Shen Loh,[2] Shannon K Yee,[3] Anna Bershteyn[1]

[1]*Department of Population Health, NYU Grossman School of Medicine, NY, NY, USA*
[2]*Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA, USA*
[3]*GWW School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA, USA*
[†]*These authors contributed equally to this work.*
**daniel.citron@nyulangone.org*

BACKGROUND

In the first year of the COVID-19 pandemic, digital contact tracing smartphone applications were developed for the purpose of reducing transmission by encouraging individuals to voluntarily quarantine after exposure[1]. Privacy protocols used in many of these apps prevented longitudinal linkage of users' contact network data, making it difficult to observe whether users reduced social contact following digital exposure notification. While in theory such tools were believed to have great potential in reducing COVID-19 transmission[2], rigorous evaluation of such tools is necessary for deploying them in the service of public health. We used anonymous data from a digital exposure notification app, NOVID[3], which used alternative privacy measures and retained longitudinal interaction data, to evaluate changes in the users' time in proximity to other users after digital exposure notification.

METHODS

We analyzed a dataset of the anonymous app users in a US university community from January−March 2021, before COVID-19 vaccines were widely available. As part of the university's COVID-19 mitigation measures, students and staff were encouraged to use the app's notification features to facilitate timely testing and self-isolation. Data collection included all university-affiliated app users and their direct contacts, estimated to constitute between 6% and 26% of the on-campus population during that period[4]. Users provided consent for their anonymous data to be used for research. We analyzed the network of contacts and categorized strong contacts as lasting more than 15 minutes at a separation of under 10.1 meters distance.

We conducted pre-post analysis of contact time with other users before and after self-reporting a positive COVID-19 diagnosis, and before and after receiving exposure notification by the app. We focused on users who had at least three days active in the contact network prior to the report or notification and excluded users who deleted the app or went inactive within 7 days in order to focus on consistent app users with measurable baseline social interactions.

For each user who self-reported or received a notification, we calculated the cumulative contact time from 7 days prior through 7 days after the event. We performed a difference-in-differences analysis, comparing the contact time before and after a user received an exposure notification against users with similar baseline contact rates who had not received a notification or reported case in the same period. We repeated this analysis for users who reported a COVID-19 case.

RESULTS

Fig. 1 A: Timeline of weekly active users for the study period. B: Timeline of positive COVID-19 cases reported each week. C: Number of exposure notifications distributed to users each week. D: Total daily contact time between all pairs of users, broken down by weak and strong links.

There were 3221 users active in the Georgia Tech contact network over the time period of analysis. The number of active users peaked at 2766 during the week beginning January 22 (Fig. 1A), the week following the start of classes on campus. Cumulatively, 25 positive cases were reported to the app and 127 exposure notifications were transmitted by the app during the study period. No user reported a positive test more than once. On a weekly basis, cases (Fig. 1B) and exposure notifications (Fig. 1C) both peaked during the week beginning January 22, 2021. This was the week following the start of in-person classes, and also the week with the most positive cases were tested by the campus health services. Among active users, the app recorded 159,360 daily interactions (24,499 strong and 134,861 weak) totaling 191,257 hours of contact time (52,800 strong and 138,456 weak).

In the difference-in-differences analysis (Fig. 2, left panel), we found that users who self-reported a case to the app reduced their total contact time with others by 29% ($p = 0.002$) in the 7-day period immediately after the time of the report vs. the 7-day period before the time of the report. This is markedly different from the behavior of the matched control group, which did not appear to alter their contact time with others in the same way. By contract, we observed no statistically significant difference in contact time before vs. after exposure notification. The 7-day cumulative contact time of users who received exposure notifications did not differ significantly compared to matched controls who did not receive an exposure notification.

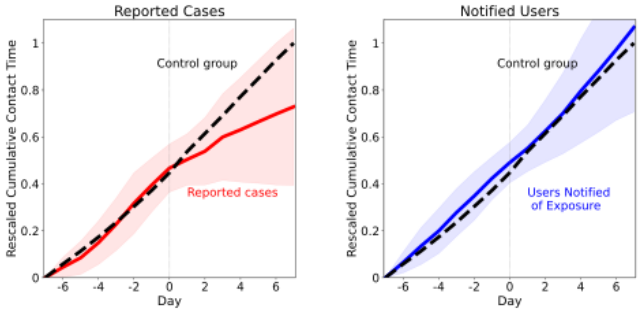

Fig. 2. Left: Cumulative contact time for all users who reported a positive case. Right: Cumulative contact time for users who received an exposure notification. Day 0 represents the time of the report/notification, such that the time before and after represent the days leading up to the report/notification and the days after, respectively. Colored lines show the mean cumulative contact time calculated across all users; shaded regions show ± one standard deviation from the mean; dotted black lines represent the control group average for the same period.

Reported cases were found to decrease their total contact time by 30 hours on average following self-reports of testing positive, compared to a 10-hour increase in total contact time for the control group ($p = 0.001$). This result is consistent for total contact time as well as for weak links only, suggesting that many who reported a positive case reduced their contact time with others in the 7 days following the time of the report. By contrast, notified users were found to increase total contact time by 6.2 hours on average, compared to a 0.9 hour increase for the control group, a difference that was not statistically significant. Consistent with the results of the pre-post analysis, this implies that users did not change their contact behavior following an exposure notification. In counterfactual analysis, artificially reducing weak interaction times by 11% following an exposure notification would have been enough to generate a statistically significant behavioral change.

DISCUSSION

Our findings suggest that the users of this app at this US university did not take measures to reduce their contact time following a digital exposure notification. By contrast, we also found a detectable change in contact time among users who self-reported a positive COVID-19 case.

Several factors may have contributed to the lack of measurable social distancing after exposure notification. First, contacts who received notifications were fully anonymous, and there was no mechanism for holding notified individuals accountable for behavior change. Second, users may have considered it too costly to comply with social distancing, missing work or class based on an app notification. Third, notified users were likely to be asymptomatic and may have perceived their risk of illness or further spreading COVID-19 to be relatively low. Fourth, notified users may have changed their behavior in ways undetectable to the app, such as remaining outdoors or masking while indoors. Lastly, our study involved a small population of users in a single university setting, and future studies must draw from larger and more diverse populations.

Our findings constitute one of the first empirical assessments of the impact of digital exposure notifications on user behavior in the United States which leveraged behavioral data collected through the app rather than self-reported survey data. Like any public health effort, digital exposure notification tools require evaluation of effectiveness, and our results invite an opportunity to redesign the development and deployment of such tools.

FUNDING

The authors received no financial support for the research, authorship, and/or publication of this article.

REFERENCES

1. Cebrian M. The past, present and future of digital contact tracing. Nature Electronics 2021; 4: 2–4.
2. Abueg M, Hinch R, Wu N, et al. Modeling the effect of exposure notification and non-pharmaceutical interventions on COVID-19 transmission in Washington state. npj Digital Medicine 2021; 4. DOI:10.1038/s41746-021-00422-7.
3. Expii Inc. NOVID. 2022; published online Dec 21. https://www.novid.org/ (accessed Dec 21, 2022).
4. Loh PS, Bershteyn A, Yee SK. Lessons Learned in Piloting a Digital Personalized COVID-19 "Radar" on a University Campus. Public Health Reports 2022. DOI:10.1177/00333549221112532.

PRIMO-Scrap: towards automating the cartography of smartphone apps for mobility

CYRILLE KOUMA[1], FATEMEH SALMANI[2], AND DAPHNE TUNCER[3,*]

[1] ESIEE Paris, France
[3] Universite Gustave Eiffel, France
[2] Ecole nationale des ponts et chaussees, Institut Polytechnique de Paris, France
[*] daphne.tuncer@enpc.fr

1. INTRODUCTION

Smartphone applications (apps) have become essential ingredients of mobility practices [1]. Mobility services are today commonly mediated through apps that enable users to plan, organize and optimize their mobility directly from their smartphones. Available functions offered by apps include navigation using geolocation, trip planning, online reservation and ticketing, as well as information service such as real-time announcement about traffic conditions. Understanding the ecosystem of these apps, including in terms of their technical and usage policy features, can provide new insights onto the evolution of people's mobility.

In this short paper, we present PRIMO-Scrap, a framework to automatically collect and format publicly available information about apps used for mobility. PRIMO-Scrap can be used to compile datasets for the study of the digital ecosystem of mobility in a given region, as well as the dynamics of this ecosystem over time. An example use case that we have been exploring using a dataset of 705 apps collected for France using PRIMO-Scrap is the risk that the use of these apps raises in terms of personal data collection. We provide details about the framework and report some observations regarding the collected data.

The source code of PRIMO-Scrap and the collected dataset are available at: http://github.com/DT16/PRIMO

2. PRIMO-SCRAP

An overview of the PRIMO-Scrap framework is depicted in Fig. 1. PRIMO-Scrap builds on three main functions: 1) a taxonomy parser that takes as input a taxonomy of keywords related to mobility practices that are combined and converted into scraping requests to apps download platforms (*e.g.,* Google App Play Store); 2) a web scraper that automatically retrieves information from the selected download platform(s) for the apps matching the inputs keywords in the taxonomy; 3) an app specifier that converts the collected information into a semi-structured format building on specification designed for these apps.

A key feature of PRIMO-Scrap is its modular design that enables to use different taxonomy inputs and integrate scraping libraries specific to different types of app download platforms.

A. Taxonomy parser

PRIMO-Scrap collects public data available on apps download platforms for apps whose description matches a set of input keywords. To select a set of keywords for apps that are relevant to the domain of people's mobility, we create a taxonomy of these apps, *i.e.,* a classification of apps based on the mode and means of transportation with which they are associated and on the service they offer, *e.g.,* rent a bicycle.

We implement our taxonomy in a tree structure where the root corresponds to the generic "mobility" keyword, the first level to the transport mode (individual, public, collective), the second to the mean used (bicycle, car, train, *etc.*) and the leafs to actions related to the mobility practice such as rent, purchase, navigate, find itinerary. The taxonomy parser implements a method based on the bigtree Python library [2] to recursively traverse the graph and extracts combination of keywords against which the description of apps on the download platform is compared. To explore a large space of possible and relevant matching apps, each action on the leaf node is associated with a set of synonyms, *e.g.,* rent, lease, hire, borrow.

B. Web scraper

We collect a dataset of apps from the Google app download platform, *i.e.,* Google Play App Store. We implement the PRIMO-Scrap scraper using the Google Search API tool SerpApi [3] that enables to scrape the results from app store.

Apps can be retrieved from Google Play Store using two approaches: 1) by typing the name of the application in the search engine of the Play Store or 2) by entering specific keywords related to the application. Our scraper uses the second approach. More specifically, we create a set of requests based on the list of keywords processed by our taxonomy parser. Each request is the concatenation of the keywords extracted from a node at level 2 in the tree, *i.e.,* a mean of transport, to a leaf, *i.e.,* an action, *e.g., rent car*, with one request per synonym of the action. For each request, we collect data for the twenty most relevant applications. We automatically eliminate redundancies between the set of apps collected via the different requests.

C. App specifier

Apps download platforms provide various types of information for each app, such as its popularity, genre, technical specs or developer details, *etc.* For each app, the retrieved information is stored in a json file as shown on the right side of Fig. 1. This includes general information about the app such as its full name, a summary of its functionality and a description of the service its provides; insights about its novelty, popularity and price; indications about the genre under which it is classified (*e.g.,* Maps and Navigation) and its rating; information about the company behind the app, including their address; the list of permissions users consent to give when installing this app and the privacy policy that is associated with the service that it provides.

3. COLLECTED MOBILITY APP DATASET

We use PRIMO-Scrap to collect a dataset of apps for mobility available for download in France. We initially collected a total number of 1000 unique apps but we noticed that some matching apps were classified under a genre not related to mobility (*e.g.,* Sports). After manual review of the genre and apps, we implemented a configurable filter to automatically remove from the dataset apps for which the genre is not related to the domain. This gives us a dataset of 705 apps with a total collection time of around 1 hour and 50 min.

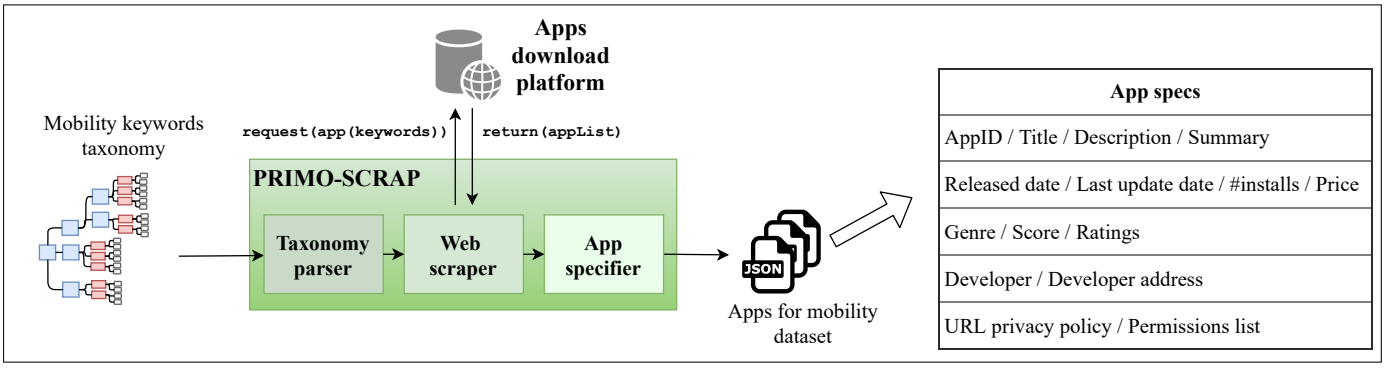

Fig. 1. Overview of the PRIMO-Scrap framework.

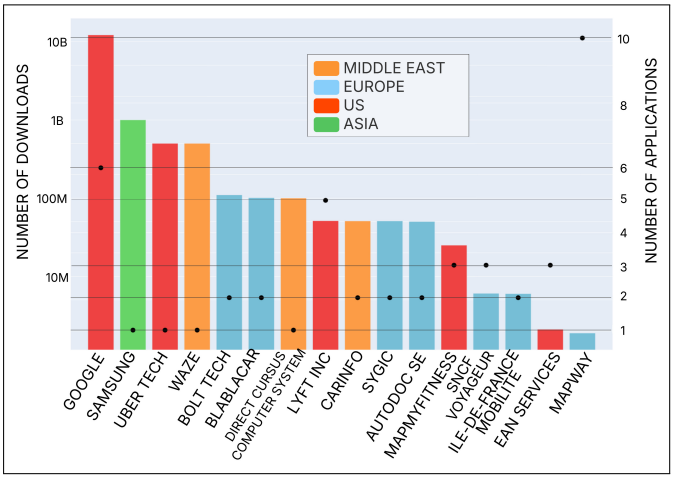

Fig. 2. Main developers of apps for mobility showing their popularity in number of downloads on Google Play Store and the number of different apps that they provide.

Fig. 2 shows statistics about the most common developers of these apps with an indication of their region of origin (based on the addresses provided on Google Play Store). Fig. 3 illustrates the types of permissions these apps require to operate. Some permissions grant access to data not directly related to the mobility actions such as access to photo data for Maps and Navigation apps, which raises questions with respect to personal data collection by these apps.

4. CONCLUSIONS

The use of smartphone applications (apps) has become central to people's mobility. Understanding the ecosystem of these apps is essential to investigate mobility practices. In this paper, we present PRIMO-Scrap, a framework to automatically collect online information pertaining to apps used for mobility. We use PRIMO-Scrap to collect and analyze a dataset of 705 apps available for download from France on the Google Play Store platform. We make the source code of PRIMO-Scrap and the dataset available to the community.

ACKNOWLEDGMENT

This work is funded by the French National Research Agency as part of the Mob Sci-Dat Factory project (ANR-23-PEMO-0004) under the France 2030 program.

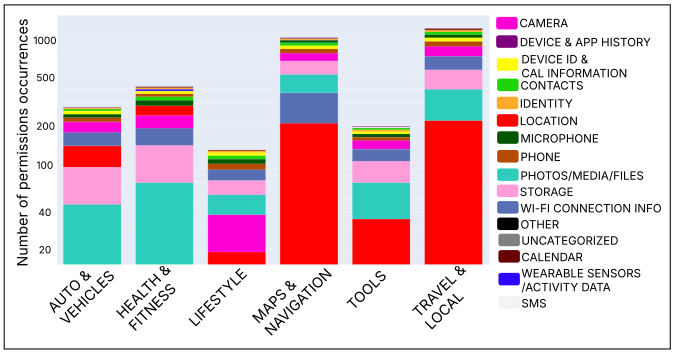

Fig. 3. Overview of the permissions required to use an app per genre (as defined on Google Play Store).

REFERENCES

1. N. A. Khan, M. A. Habib, and S. Jamal, "Effects of smartphone application usage on mobility choices," Transportation Research Part A: Policy and Practice, vol. 132, pp. 932–947 (2020).
2. bigtree, "Big Tree Python Package," https://pypi.org/project/bigtree/.
3. Google Search API, "SerpApi," https://serpapi.com/.

Economic spillover effects of electric vehicle charging station placement on local businesses: a staggered adoption study

M. Mavin De Silva[1,2,3], Callie Clark[1,4], Tadachika Nakayama[2], and Takahiro Yabe[1,5,*]

[1] Center for Urban Science and Progress (CUSP), Tandon School of Engineering, New York University, Brooklyn, NY, USA
[2] Extreme Energy-Density Research Institute, Nagaoka University of Technology, Nagaoka, Niigata 940-2188, Japan
[3] Department of Transport Management & Logistics Engineering, Faculty of Engineering, University of Moratuwa, Sri Lanka
[4] Marron Institute of Urban Management, New York University, Brooklyn, NY 11201, USA
[5] Department of Technology Management and Innovation, Tandon School of Engineering, New York University,, Brooklyn, NY, USA
[*] takahiroyabe@nyu.edu

1. INTRODUCTION

Electric vehicle (EV) infrastructure is rapidly transforming the urban transportation landscape, promising a cleaner and more sustainable future for cities worldwide. In New York City, where light-duty vehicles account for nearly 80% of transportation-related emissions, EVs represent a critical tool in the effort to decarbonize mobility and reduce local pollution. As EV technology becomes more accessible, through dropping vehicle prices and state incentives, public investment is now shifting toward building the charging infrastructure needed to support mass adoption. With over $7.5 billion allocated to installing 500,000 public EV charging stations (EVCSs) across the United States by 2030, questions about how, where, and to what effect these stations should be deployed are increasingly urgent [1]. Much of existing research has focused on how the placement of EVCSs affects EV adoption rates and access equity. However, relatively little is known about the economic ripple effects that charging infrastructure can create on the hyperlocal scale. In practice, many EVCSs are located near businesses and commercial corridors, implicitly pairing charging time with opportunities to shop, dine, or access services. These collocations raise key questions: do EVCSs create measurable increases in foot traffic to surrounding businesses? And if so, what types of business benefit the most and under what spatial conditions? To address these questions, we design a quasi-experiment using mobile phone GPS location and EVCS deployment history data to analyze the causal impact of EVCS placement on visitation patterns to businesses in New York City, home to one of the densest EVCS networks in the country. Leveraging a novel staggered difference-in-differences (DID), approach that capitalizes on the temporal roll-out of EVCS installations, we analyze how daily visit patterns to nearby businesses change before and after EVCS deployment.

2. POI MATCHING STRATEGIES

A major challenge in estimating the causal effect of EVCS installation is site selection bias, as providers often target areas with high EV adoption or business activity [2, 3]. To address this, we implement a staggered adoption framework using spatial and temporal variation in EVCS rollouts. As shown in Figure 1a, POIs within 400m of EVCSs are classified as treated or control, with at least a 3-month gap between installations to ensure temporal separation. Control POIs are matched by borough and business category but linked to a different EVCS installed at least three months later than the treated POI's station. From eligible controls, we select the one with the smallest pre-treatment difference, defined as the change in average visits from 3–6 months to 1–3 months before installation. This gap reduces confounding trends and supports causal inference. Figure 1b illustrates the staggered DID setup, where POI j serves as a control for treated POI i, which is affected by a later EVCS opening. While prior research has highlighted the strategic placement of EVCS, few studies have rigorously addressed the locational endogeneity in causal inference frameworks [2]. To tackle this, we compare the staggered adoption approach with two non-staggered alternatives: a non-staggered matching approach and a non-staggered propensity score matching (PSM) method. In both non-staggered approaches, control POIs are those beyond 400m of any EVCS installed during the same period. The non-staggered matching approach follows a two-step process: exact matching by borough and POI category, followed by nearest-neighbor matching to minimize pre-treatment differences. Similarly, non-staggered PSM involves exact matching and then nearest-neighbor matching based on propensity scores estimated from covariates, following Zheng et al. (2024) [2]. These covariates include built environment factors (e.g., population and road density, walkability), socio-demographics (e.g., income, unemployment, race), and POI-level activity (e.g., pre-installation visits). We also use logistic regression to examine how these covariates relate to treatment likelihood. As shown in Figure 1c, the staggered approach yields more treated-control matches and eliminates spatial clustering, ensuring geographically closer and more balanced pairs, especially valuable in EVCS dense areas.

For an effective evaluation of treatment effects, it is important to understand the covariate balance between treatment and control groups. For all three approaches, we assess the covariate balance between the treatment and control groups through graphical representation, as shown in Figure 1d. The staggered approach shows strong balance, with covariate differences converging toward zero ($-0.1 <$ covariate difference < 0.1), reflecting robust matching. In contrast, the non-staggered matching method performs less well, especially in dense areas like Manhattan and Brooklyn. The non-staggered PSM approach, which uses logistic regression to model treatment assignment, improves balance over non-staggered matching but still falls short of the staggered method. Figure 1e further shows that matched pairs under the staggered approach are geographically closer, with a tighter distribution of distances than the other two methods. Although PSM improves distance over standard matching, the staggered approach achieves the shortest matched-pair distances overall. These differences are statistically significant ($p < 0.001$), highlighting the spatial advantage of incorporating treatment timing.

Following the selection of control POIs using the three approaches, we incorporate the matched control POIs alongside the treated POIs and proceed with the staggered DID analysis.

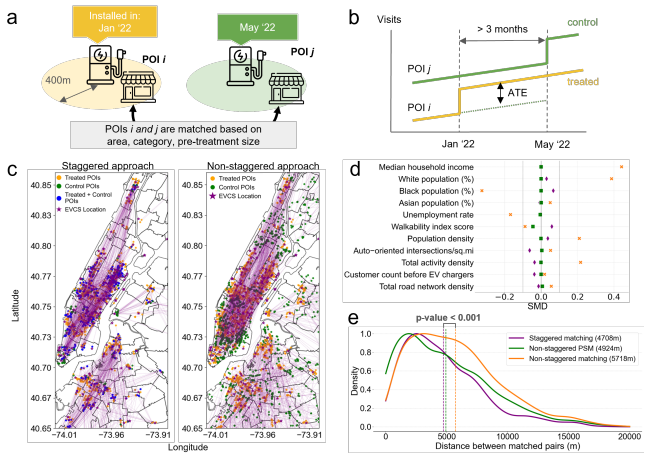

Fig. 1. Determinants of the staggered adoption framework

The model specification is defined as follows:

$$Y_{it} = \beta D_i \times T_t + u_i + \omega_{\text{year}} + \epsilon_{it} \quad \textbf{(1)}$$

In this equation, Y_{it} represents the outcome variable, which is the number of customers at POI i during time period t. β represents the treatment effect, quantifying the impact of adding a new EVCS on the number of customers at the respective POI. The term D_i is a binary treatment indicator that takes on the value of 1 for treated establishments when an EVCS opens for service nearby. Similarly, T_t is a binary time indicator that shifts from zero to one in the post-treatment period. u_i denotes the POI fixed effects. ω_{year} represents the year fixed effects. ϵ_{it} represents the error term.

3. HETEROGENEITY BY SPATIAL VARIATION, POI TYPES AND ACROSS DIFFERENT INCOME GROUPS

The staggered DID approach estimates that the average treatment effect of an EVCS placement on a business is 28.52 customers (p-value < 0.001). We then categorize distances between treatment POIs and nearby EVCS into four 100m bins (0–400m) to examine distance-varying treatment effects. Using the staggered adoption approach, the 0–100m range shows the strongest impact, with a significant increase of 39 customers ($p < 0.05$). However, as the distance from the EVCS increases to the range of 300–400m, the magnitude of the treatment effect diminishes, resulting in a smaller yet still significant increase of ~ 30 ($p - value < 0.001$) in customer counts. Accounting for the average of 3.0 ports per EVCS during the study, we estimate that each additional port increases visits by $\beta = 13.04$ ($p < 0.01$) within 100m. These results highlight the importance of accessibility and suggest that EVCSs are strategically sited to benefit nearby businesses, supporting more effective EVCS deployment strategies.

Next, we determine whether these effects are homogeneous across different POI categories. To measure this heterogeneity, we employ the following model specification to estimate treatment effects stratified by POI types: $Y_{it} = D_i \times T_t \times \sum_{c=1}^{11} \beta_c \text{Cat}_{ict} + u_i + \omega_{\text{year}} + \epsilon_{it}$. Here, Cat_{ict} represents customer visits for POI Category c at POI i during time period t, where $\boldsymbol{c}$ can take on one of eleven possibilities: Apparel, Beauty and Spa, Buildings and Entertainment Centers, Dining, Groceries, Home Improvements and Furnishings, Hotels and Casinos, Leisure, Medical and Health, Shopping Centers, and Shops and Services. According to the treatment effects obtained using staggered adoption approach, where Shopping Centers ($\beta = 117.83$, p-value < 0.001), Buildings and Entertainment Centers ($\beta = 85.73$, p-value < 0.001), Leisure ($\beta = 67.76$, p-value < 0.001), Hotels and Casinos ($\beta = 45.81$, p-value < 0.05) and Dining ($\beta = 14.94$, p-value < 0.05) are positive and significant, indicating that the presence of EV chargers may influence consumer behavior, encouraging visits to convenience type stores, restaurants, and recreational/rest areas during charging times. This suggests that newly attracted customers are more likely to be recreational or exploratory visitors than regular and loyal customers, as evidenced by the larger treatment effects observed for outdoor markets, entertainment venues, and shopping centers.

We also examine how EVCS installations affect customer counts across income groups by estimating the average visitor income for each POI based on their origin CBGs during the 3 months before EVCS placement. Visitors are grouped into four income brackets: $< \$60k$, $\$60$–$100k$, $\$100$–$150k$, and $> \$150k$. Each POI is assigned to the dominant income group among its visitors. The staggered adoption results show significant positive effects across all income levels, suggesting that EVCS accessibility may have supported broader EV adoption, including among lower-income users. However, the largest impacts are seen in higher-income groups, notably \$100–150k ($\beta = 42.43$, p-value < 0.001) and $> \$150k$ ($\beta = 37.32$, p-value < 0.001), indicating that businesses serving higher-income customers tend to experience larger effects, reflecting the broader trend that EV drivers are predominantly from higher-income demographics[1]. From a practical perspective, our research highlights the potential for policymakers and businesses to leverage EVCS as a catalyst for local economic growth, particularly in diverse commercial areas. Our findings suggest that EVCS providers can adopt a business model similar to the traditional paradigm of "gas station–convenience store", where fuel stations are strategically paired with retail establishments to maximize consumer spending[4]. Just as gas station owners capitalize on ancillary sales, EVCS operators and local businesses could collaborate to internalize the economic benefits of EVCS installations, transforming charging locations into hubs of commercial activity that attract and retain customers[5].

Acknowledgment. T.Y. acknowledges support by the National Science Foundation under grant number CMMI 2425021. M.M.D.S. acknowledges support by the Japan Society for the Promotion of Science (JSPS) Grant-in-Aid for JSPS Fellows [grant number 24KJ1180].

Disclosures. The authors declare no conflicts of interest.

Data availability. The data that support the findings of this study are available from Placer.ai data platform, but restrictions apply to the availability of these data, which were used under the license for the current study and are therefore not publicly available. Information about how to request access to the data and its conditions and limitations can be found in https://docs.placer.ai/reference/welcome-to-papi. Data access requests should be submitted through Placer.ai's marketplace customer page https://marketplace.placer.ai/, where the Sales team at Placer.ai may be contacted.

REFERENCES

1. H. A. U. Khan, S. Price, C. Avraam, and Y. Dvorkin, The Electr. J. **35**, 107096 (2022).
2. Y. Zheng, D. R. Keith, S. Wang, *et al.*, Nat. Commun. **15**, 1 (2024).
3. Y. Babar and G. Burtch, Manuf. & Serv. Oper. Manag. **26**, 797 (2024).
4. NACS, "Consumer behavior at the pump," NACS Report (2019).
5. M. L. Arlt and N. Astier, Energy Econ. **124**, 106777 (2023).

Understanding the Impacts of EV Charging Stations on Surrounding Businesses

CALLIE CLARK[1,2], ANNE DRISCOLL[1], XIYUAN REN[2], SALSABIL SALAH[2], MARTA C GONZÁLEZ[3], JOSEPH CHOW[2], AND TAKAHIRO YABE[1]

[1] Center for Urban Science & Progress, New York University, NY, USA
[2] Department of Civil and Urban Engineering, New York University, NY, USA
[3] Department of Civil and Environmental Engineering, University of California at Berkeley, CA, USA
* callieclark@nyu.edu.com

EXTENDED ABSTRACT

Electric vehicle charging stations (EVCS) constitute a critical component of modern urban infrastructure, and the rate of installation is remaining steady [1]. Evidence suggests that the installation of an EVCS can increase annual consumer spending at proximate businesses by up to 3.2% [2]. More broadly, the dependence between visits to POIs and the surrounding landscape of businesses has been demonstrated within the context of behavior-based dependencies [3]. We extend this idea to EVCS selection, where EV drivers may select their EVCS with consideration for the surrounding POI landscape. In this paper, we construct and evaluate a behavior-based dependency network between EVCS and surrounding POIs to investigate to what extent newly installed EVCSs attracts new visitation patterns into the neighborhood, creates spillovers to nearby businesses and amenities, and reshapes the local economic interdependence network. We hypothesize that the surrounding EVCS environment, including POI type and socio demographic information, is an important predictor for who will visit a POI during charging and in determining the overall change in the dependency network.

Prior to this work, the charging preferences of EV drivers have been investigated using surveys [4–7], field studies [8–11], analyses of EVCS charging locations [2, 12] and station-level charging session data [13, 14]. These sources are either limited by sample size and study longevity, or confined to the duration of charging session event. Our study provides a new lens into understanding EV driver charging behavior and it's impact on surrounding businesses— through large-scale individual mobility data. We utilize data from Cuebiq, a data aggregator that compiles device location data across multiple providers from anonymized, opted-in users, to gain insights into real-world charging behavior [15] and mobility patterns at an unprecedented level of granularity across a large-scale network.

Identifying EV Drivers from Individual Mobility Data

The first step in understanding charging behavior is identifying EV drivers. To this end, we integrate individual movement data with the geolocations of 6,759 public charging ports and 2,253 gas stations to identify how users are interacting with both POI types. The individual movement data provides latitude, longitude and timestamps for each observed ping. Our study focuses on the nine counties of the California Bay Area — home to approximately 7.6 million people during 2022. For the months of January to March 2022 145,320 Spectus users meet our criteria of having a home location within the study area and a median monthly observation rate of at least one ping every 30 minutes.

We use the ping data to capture the behavior that may occur at a timescale that would not be present in the "stops" data. We then define a "slow point" as any sequence of location pings in the raw data indicating a movement speed below 10 km per hour. Next, we define a "charging session" as at least two points

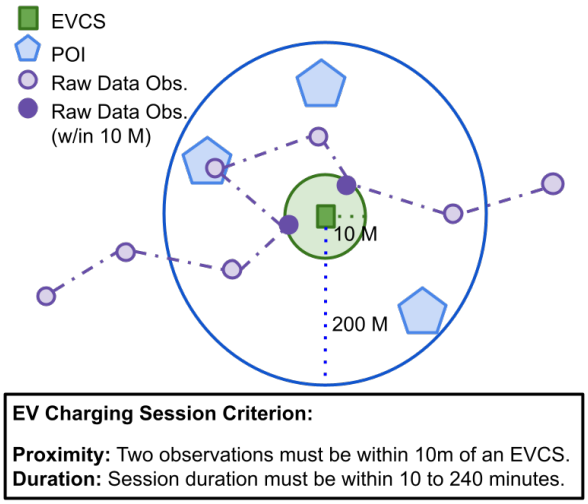

Fig. 1. EV Charging Session Criterion Diagram

within 10m of an EVCS, one of which must be a slow point. We calculate the duration of the charging session as the difference between the observation directly before the user is identified at the EVCS and the observation directly after the last seen EVCS observation (see Figure 1). In our sample, 14.8% of users have at least one charging session during the month.

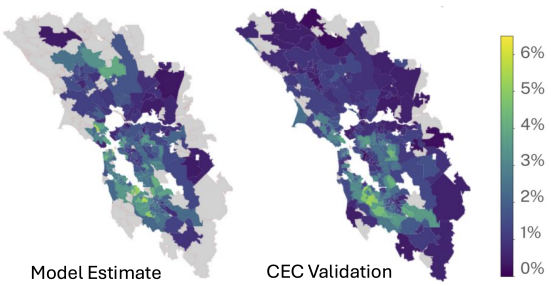

Fig. 2. EV Ownership Model Results and Validation

The EV ownership rate in the Bay Area is 4.9% [16], notably lower than the 14.8% of our users with at least one charging session. Given our sample only captures public charging, and will not capture the 26.2% [17] of EV drivers who never charge at a level-2 public charging station, this group is certainly too broad. To distinguish true EV charging sessions from instances where users simply park near EVCS, we design a model that identifies EV drivers based on their behavior over time.

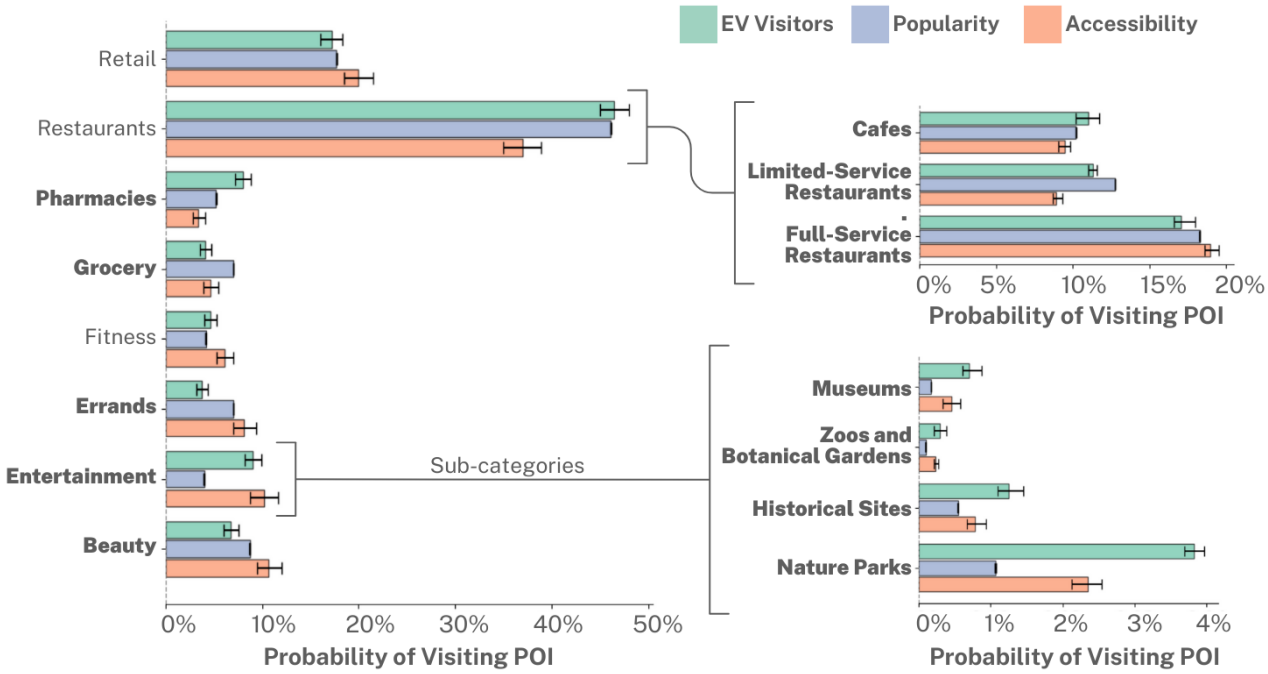

Fig. 3. Probability of visiting a POI by type with 95% Confidence Intervals for different scenarios: EV drivers includes visits during EV charging sessions, Popularity includes all POI visits within 250 meters of an EVCS, and Accessibility captures the spatial availability, or the count of each type of POIs that exist within 250 meters of an EVCS.

We use individual-level behavior over the course of three months to predict zipcode level EV ownership. We selected our model by iteratively testing different models and features, like gas stops per day, unique charging stations visited, charging duration, VMT, % multi-family home ownership. For each model and feature subset, we predict individual EV ownership, select the top 3.2% of users, aggregate to the zip code level and compare our inferred cohort to zip-level EV ownership rates reported by the California Energy Commission. Our final model is a linear regression model with a correlation of 0.58 (shown in 2). In this model, we find the strongest individual predictors for EV ownership are the number of unique EV stations visited and the number of daily gas stops. All subsequent analyses are performed on this inferred cohort of 5,400, or 3.2%, users out of our total sample.

Evaluating EV driver Behavior

Using our inferred EV driver cohort, we can evaluate driver behavior during charging sessions. We use the Spectus reported "stops" data to understand where drivers go during charging. We combine the Spectus stops data with eight Safegraph Points of Interest (POI) categories: restaurants and cafes, beauty, fitness, pharmacy, grocery, entertainment, errands and retail to explore the available POIs within 250 meters [18] . We then compare the probability of visiting a POI by category for EVCS drivers to both the entire study population and the spatial availability. We find that EV drivers visit restaurants, and entertainment POIs (which include museums, parks, etc) more frequently during EVCS sessions than the general public, with statistical significance. The error bars in Figure 3 are 95% confidence intervals derived using bootstrapping. We then further look at what categories drive this and find that parks and cafes have a significantly higher probability of visitation during an EVCS session than in general. These findings highlight the potential for businesses, like cafes and parks, located near EVCS to capture economic benefits from charging-related foot traffic.

Acknowledgment. T.Y. and C.C. acknowledge support by the National Science Foundation under grant number CMMI 2425021.

REFERENCES

1. A. Brown, J. Cappellucci, A. Heinrich, *et al.*, "Electric vehicle charging infrastructure trends from the alternative fueling station locator: First quarter 2024," Tech. rep., National Renewable Energy Laboratory (NREL), Golden, CO (United States) (2024).
2. Y. Zheng, D. R. Keith, S. Wang, *et al.*, Nat. Commun. **15**, 7437 (2024).
3. T. Yabe, B. García Bulle Bueno, M. R. Frank, *et al.*, Nat. human behaviour pp. 1–11 (2024).
4. J. H. Lee, D. Chakraborty, S. J. Hardman, and G. Tal, Transp. Res. Part D: Transp. Environ. **79**, 102249 (2020).
5. G. Tal, D. Chakraborty, A. Jenn, *et al.*, (2020).
6. F. A. Bhat, G. Y. Tiwari, and A. Verma, Transp. Policy **149**, 177 (2024).
7. A. A. Visaria, A. F. Jensen, M. Thorhauge, and S. E. Mabit, Transp. Res. Part A: Policy Pract. **165**, 120 (2022).
8. T. Franke and J. F. Krems, Transp. Res. Part F: Traffic Psychol. Behav. **21**, 75 (2013).
9. L. H. Björnsson and S. Karlsson, Appl. Energy **143**, 336 (2015).
10. X. H. Sun, T. Yamamoto, and T. Morikawa, Transp. Res. Part D: Transp. Environ. **46**, 26 (2016).
11. H. Yu and D. MacKenzie, Transp. Res. Rec. **2572**, 56 (2016).
12. G. Bauer, C. W. Hsu, M. Nicholas, and N. Lutsey, White Pap. ICCT (2021).
13. P. Huang and Z. Ma, Renew. Energy **236**, 121511 (2024).
14. S. Kim, D. Yang, S. Rasouli, and H. Timmermans, Transp. Res. Part C: Emerg. Technol. **82**, 248 (2017).
15. Cuebiq, "Cuebiq mobility data for research," https://www.cuebiq.com/about/data-for-good/. Accessed: 2025-02-01.
16. California Energy Commission, "Light-duty vehicle population in california," (2024). Data last updated [insert date last updated]. Retrieved November 15, 2025, from https://www.energy.ca.gov/zevstats.
17. N. Turner, "Ev driver survey," (2025).
18. DEWEY, "Global places - poi data," (2025). Data last updated [02/07/2025]. Retrieved November 11, 2024, from https://app.deweydata.io/products/63eeb9d8-8439-42b3-a75e-2d9cbdf64cb8/package.

Causal inference in the city: Improving Urban Policy Evaluation Through Mobility-Aware Methods

BIJIN JOSEPH[1,3,*], HAMISH GIBBS[3], TAKAHIRO YABE[2], AND ESTEBAN MORO[1,3,*]

[1] Department of Physics, Northeastern University, Boston, MA, USA
[2] NYU Tandon School of Engineering, New York, NY, USA
[3] Network Science Institute, Northeastern University, Boston, MA, USA
[*] bi.joseph@northeastern.edu,e.moroegido@northeastern.edu

1. INTRODUCTION

Urban policy interventions play a central role in shaping mobility, accessibility, and equity in cities. Rigorous evaluation of their impact requires quasi-experimental causal techniques that go beyond simple correlational analysis [1]. However, the complex interplay of behaviors across spatially connected areas creates unforeseen spillovers that often violate the assumptions of standard causal inference [2]. In this study, we introduce a novel approach to causal inference in urban contexts that explicitly accounts for human mobility and behavior flows between areas. Using large-scale, anonymized mobile phone data [3], we construct a mobility-flow network in urban areas to account for potential spill-overs and apply community detection algorithms to redefine treatment and control groups, moving beyond simplistic geographic-distance-based assignments to those groups. This Community Detection Method (CDM) enables more robust identification strategies by mitigating spatial spillovers. We demonstrate the efficacy of our approach across three distinct interventions: Boston's fare-free bus program, New York City's Open Streets initiative, and the shutdown of Boston's Orange Line. In all cases, CDM outperforms traditional distance-based designs in estimating treatment effects and passing placebo tests. Notably, our analysis reveals that standard analysis of Boston's fare-free program significantly underestimates its impact due to spillovers and concurrent confounding factors. These findings underscore the importance of incorporating mobility-aware network methods in policy evaluation to better understand the reach, efficacy, and design implications of urban interventions.

2. METHODS

We use geo-location data from Spectus, a mobility data provider. The dataset contains anonymized, privacy-preserving records of device stays across urban areas to different Points of Interest (POI) [3]. Socioeconomic data at the Census Block Group (CBG) level are sourced from the American Community Survey (ACS).

Randomized Controlled Trials (RCTs) are the gold standard for causal inference [4], but are often infeasible in urban settings due to logistical, ethical, and financial constraints. Researchers therefore turn to quasi-experimental designs that approximate random assignment, using methods such as Difference-in-Differences (DiD) [5] and the Synthetic Control Method (SCM) [6] to estimate treatment effects from observational data.

To operationalize these quasi-experimental designs to spatial interventions, we define a unit of analysis U, as either a CBG or a POI. The initial step in causal inference is to divide these units into treatment and control groups. In the conventional Distance Method (DM) (Fig. 1), the treatment group consists of all U that lie within a given radius from the intervention, while control units are selected from comparable areas located beyond this radius. To reduce potential contamination, a buffer zone is typically implemented around the treated area to exclude units

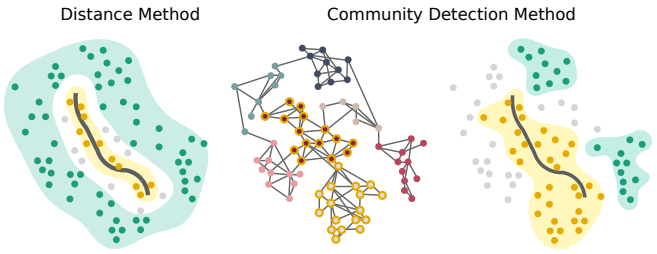

Fig. 1. Distance vs. Community Detection Method. (A) Traditional *Distance Method (DM)*: census block groups (CBGs) within 1 km of the intervention form the treatment group (green), while those farther than 2 km constitute the control group (orange). *Community Detection Method (CDM)*: we construct a mobility network among CBGs and apply a community-detection algorithm; CBGs in communities that contain the intervention are assigned to the treatment group.

that are neither clearly treated nor sufficiently distant to serve as controls. Units in the control group can be used to build the SCM and compared with units in the treatment group to estimate the impact of the intervention.

A key limitation of the Distance Method (DM) is its failure to account for behavioral spillovers in highly interconnected urban systems. To address this, we build a mobility-flow network between units in U using our large-scale data and apply a community detection algorithm to identify functional mobility areas—clusters of units closely linked by human movement. Units in communities intersecting the intervention are assigned to the treatment group, while control units are drawn from communities most distant in the network, thereby minimizing spillover. We refer to this approach as the Community Detection Method (CDM) (Fig. 1A).

After defining treatment and potential control units, we apply SCM to build a counterfactual for the treatment group, i.e., an estimate of how it would have evolved had the intervention not occurred. To do that we use both the time series of our outcome variable but also relevant demographic confounders to ensure the treated group and counterfactual are demographically similar. We then estimate the intervention's effect using DiD. To test robustness, we conduct a placebo test by randomly splitting the control group and repeating the SCM and DiD, assessing the likelihood of spurious effects under random assignment.

3. RESULTS

Although our study examines three distinct interventions, space constraints limit this abstract to reporting results from Boston's fare-free bus program [7]. In August 2021, the City of Boston eliminated fares on Route 28, a bus line serving neighborhoods with high social vulnerability that have been designated as priority areas for infrastructure investment. The program aimed to

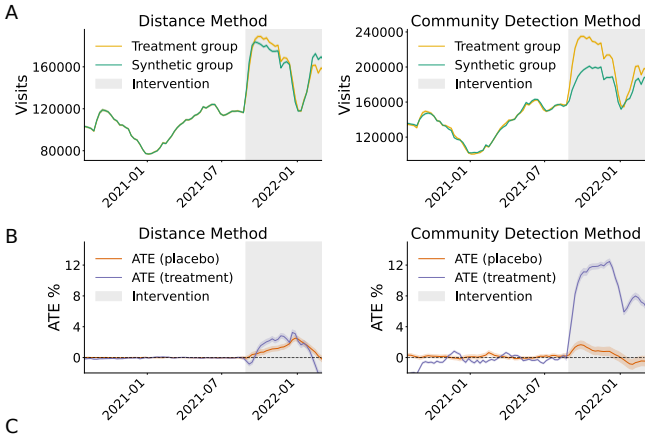

Method	Treatment Group	Placebo Group
Distance Method	2.71%***	0.85%***
	(0.01%)	(0.01%)
Community Detection Method	16.15%***	0.65%
	(0.28%)	(0.45%)

Note: $^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$.

Fig. 2. Causal estimates of Route 28 intervention. (A) Time series of total POI visits for the treatment group (orange) and its synthetic control (green); shaded area indicates the intervention period. (B) Average Treatment Effect over time: orange shows the estimated effect, purple shows a placebo ATE from randomly assigned treatment; shaded bands denote standard error. (C) Average Treatment Effect DiD estimates for the Distance Method (DM) and Community Detection Method (CDM), along with placebo results during the intervention period.

enhance mobility for residents in these communities. To evaluate its impact, we applied both the traditional Distance Method and our proposed Community Detection Method to assign CBGs to treatment and control groups. For each method, we constructed a synthetic control using SCM and then employed a DiD estimator to quantify the effect of the intervention. We evaluated multiple dimensions of impact, ranging from visits to transportation hubs to overall visits to points of interest (POIs) by individuals affected by the intervention. While the results were consistent across these measures, here we report only the findings related to the total activity levels of individuals exposed to the intervention.

As shown in Fig.2A, both the Distance Method and the Community Detection Method exhibit an increase in mobility in both treatment and control groups right after the intervention started. This is because COVID-19 mobility restrictions were lifted around the same time as the pilot program. Such temporal confounding complicates naive before-after comparisons [7], underscoring the need for a causal inference framework that adjusts for concurrent external shocks. By applying the Difference-in-Differences (DiD) estimator, we control for these common trends and isolate the effect attributable to the intervention. The DiD estimates or the Average Treatment Effect (ATE) on the total number of visits reported in Fig.2C indicate that the Distance Method yields only a modest ATE. In contrast, the Community Detection Method reveals a substantially larger impact of the intervention. This difference highlights the role of behavioral spillovers: the DM fails to account for indirect exposure to the intervention, thereby diluting the estimated effect. In contrast, CDM leverages mobility-based community structure to mitigate such spillovers, resulting in more accurate identification of treated and unaffected units. This conclusion is further supported by the placebo analysis shown in Fig. 2B. When the control group is randomly partitioned into pseudo-treatment and pseudo-control units, the CDM yields placebo estimates that are not statistically different from zero, while the DM still registers a significant effect, suggesting residual contamination due to spill-overs.

4. CONCLUSION

Evaluating urban interventions is difficult due to the impracticality of RCTs and the interconnected nature of cities, where policy effects often extend far beyond their geographic origin. Conventional distance-based methods are inadequate: *in urban areas, there is effectively no physical boundary beyond which units can be considered unaffected.* We address this challenge with a quasi-experimental framework that uses mobility data to define treatment and control groups based on *behavioral distance*, i.e., the likelihood that areas influence one another through human movement. Our Community Detection Method (CDM) groups units into mobility-based communities, reducing spillovers and improving causal estimates.

Our analysis of Boston's fare-free bus program highlights the importance of accounting for both confounding trends and behavioral spillovers. Because the intervention coincided with the lifting of COVID-19 restrictions, citywide mobility increased, confounding simple before-after comparisons. The Distance Method (DM) underestimates the program's effect by ignoring spillovers into neighboring areas. In contrast, CDM captures the broader reach of the intervention by forming behaviorally coherent control groups. Placebo tests confirm that CDM reduces contamination, yielding null effects when no true treatment is applied. We further applied CDM to New York City's Open Streets program [8] and Boston's Orange Line shutdown [9], finding consistently more robust and interpretable estimates than with distance-based methods.

These findings highlight the importance of using behavioral data in causal inference. It improves the accuracy of impact estimates and helps reveal how urban interventions spread through the city, leading to better-informed and more equitable policy decisions.

REFERENCES

1. A. Abadie, A. Diamond, and J. Hainmueller, J. Am. statistical Assoc. **105**, 493 (2010).
2. K. Butts, arXiv preprint arXiv:2105.03737 (2021).
3. E. Moro, D. Calacci, X. Dong, and A. Pentland, Nat. communications **12**, 4633 (2021).
4. E. Duflo, R. Glennerster, and M. Kremer, Handb. Dev. Econ. **4**, 3895 (2007).
5. J. J. Heckman, H. Ichimura, and P. E. Todd, The Rev. Econ. Stud. **64**, 605 (1997). Publisher: [Oxford University Press, Review of Economic Studies, Ltd.].
6. A. Abadie and J. Gardeazabal, Am. Econ. Rev. **93**, 113 (2003).
7. City of Boston, "City of Boston Fare Free Program Mid Program Report," https://www.boston.gov/sites/default/files/file/2023/03/Fare%20Free%20Mid%20Program%20Report.pdf (2023). Accessed 27 June 2025.
8. New York City Department of Transportation, "Open Streets," https://www.nyc.gov/html/dot/html/pedestrians/openstreets.shtml (2025). Accessed 27 June 2025.
9. City of Boston, "Orange Line Shutdown Resources," https://www.boston.gov/departments/mayors-office/orange-line-shutdown-boston (2025). Accessed 27 June 2025.

The Urban Impact of AI: Modeling Feedback Loops in Next-Venue Recommendation

GIOVANNI MAURO[1,2*], MARCO MINICI[3], AND LUCA PAPPALARDO[1,2]

[1] ISTI-CNR, via G. Moruzzi 1, Pisa, 56124, Italy
[2] Scuola Normale Superiore, Piazza dei Cavalieri,7, Pisa, 56126, Italy.
[3] ICAR-CNR, Via Pietro Bucci 8/9c, Rende, 87036, Italy.
[*] giovanni.mauro@sns.it

The integration of information technologies into cities has transformed citizen interactions with urban systems. Urban Recommender Systems, from routing to Points of Interest, improve user experiences, but their long-term impact on mobility remains unclear, necessitating an exploration of biases and broader urban dynamics [1]. Privacy and ethical concerns limit observational and experimental studies. To address this, we develop an Agent-Based Model (ABM) to examine algorithmic adoption and its urban implications.

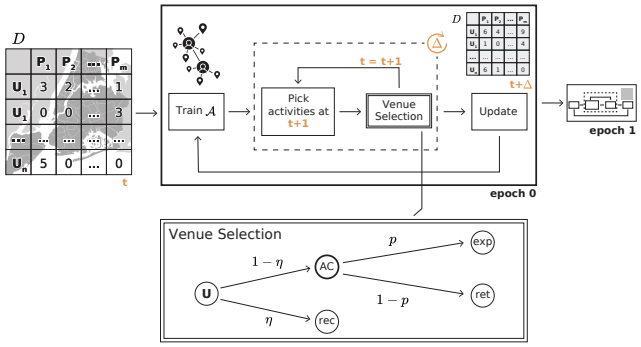

Fig. 1. Overview of the simulation framework.

Our ABM builds on a modified version of the Explorer-Returner paradigm (EPR) [2], which effectively models real-world mobility dynamics. We extend this framework by incorporating an AI-based RS to guide agent movement.

We use a Foursquare dataset containing mobility data from New York City, collected over approximately 10 months, from April 3, 2012, to February 16, 2013 (304 days) [3]. This real-world data provides the empirical basis for initializing agent behaviors and training recommender models.

Figure 1 provides a schematic overview of our simulation framework. At each time step, agents either move according to the classical EPR mechanism or follow a recommendation produced by a trained AI-based recommender system. The probability of selecting the AI-based option is controlled by a global adoption parameter η. The recommender is periodically retrained during the simulation at fixed intervals Δ, using an enriched dataset composed of the original training set D_{train} and the simulated visits up to that point, i.e., $D_{\text{train}} \cup \mathcal{S}$.

MEASURES

Venue perspective We quantify inequality in the distribution of visits across venues using the Gini index, which measures concentration: a value of 0 indicates perfect equality, while higher values reflect visits being concentrated on fewer venues. This aligns with findings that recommender systems can increase consumption inequality [4, 5]. Given a visit dataset D, the collective Gini index is computed as:

$$G = \frac{1}{|V|}\left(|V| + 1 - 2 \cdot \frac{\sum_{i=1}^{|V|}(|V| + 1 - i)x_i}{\sum_{v \in V} x_i}\right),$$

where x_i is the total number of visits to venue v_i, sorted in ascending order. To capture individual behavior, we compute a personal Gini index G_u for each user $u \in U$, based on their own visits D_u. The average individual inequality is:

$$\overline{G} = \frac{1}{|U|}\sum_{u \in U} G_u.$$

Individual Perspective We construct a co-location network where nodes are individuals and edges link those who visited the same venue within a time window $[t_1, t_2]$. Formally, $(u_m, u_n) \in E$ if

$$\exists\, (u_m, v_j, t_k),\ (u_n, v_j, t_l) \in D \quad \text{with} \quad t_k, t_l \in [t_1, t_2].$$

To examine hierarchical structure, we identify the top h highest-degree nodes U_{rich} and compute the rich-club density as the ratio of observed to possible links among them: $\frac{h(h-1)}{2}$ [6].

We also analyze the degree distribution $P(k)$, defined as:

$$P(k) = \frac{|\{u \in U' : \deg(u) = k\}|}{|U'|},$$

where $\deg(u)$ is the degree of node u in the network. Fitting a line to the log-log plot of $P(k)$, we report the slope α: higher α indicates more skewed, less uniform interactions.

RESULTS

We evaluate the impact of algorithmic adoption on urban mobility by simulating several types of recommender systems, ranging from simple collaborative filtering methods to more complex deep learning models. Among these, ItemKNN and MultiVAE emerge as the most contrasting cases in terms of their effects on individual and collective dynamics.

Figure 2 summarizes the outcomes across key mobility metrics as a function of adoption rate η. Detailed explanations of each recommender are provided in the main paper [7].

Individual diversity increases. The average individual Gini coefficient $\overline{G}$ (Figure 2a) consistently decreases with higher η, indicating more balanced visit patterns at the user level. While simple algorithms like ItemKNN yield modest changes, deep learning-based models (e.g., MultiVAE) lead to substantial drops—up to 60%—in $\overline{G}$. This suggests that as users rely more on recommender systems, they explore a wider range of venues.

Collective inequality rises. However, the collective Gini coefficient G (Figure 2b) often increases with η, revealing a growing concentration of visits on a limited set of venues. Models such as MultiVAE and UserKNN drive sharp increases in G, with MultiVAE reaching a 47% rise at full adoption. ItemKNN is the exception, slightly reducing G across all adoption rates.
Figure 2c–d extends the analysis to the co-location network.

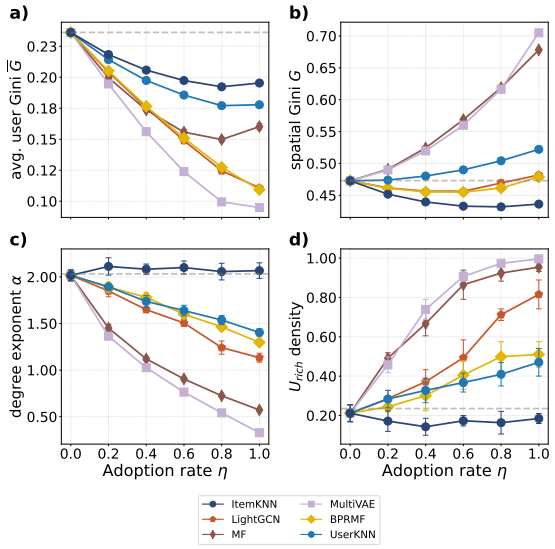

Fig. 2. Effects of recommender system adoption η. (a) Average individual Gini coefficient $\bar{G}$; (b) Gini coefficient of the distribution of visits across venues G; (c) Power-law exponent of the co-location network's degree distribution α; (d) Internal density of the rich club (top 15 nodes by degree) in the co-location network. Each curve is an algorithm; each point represents the average over five simulation runs.

Broader interpersonal mixing. Most recommender systems increase the exponent α of the degree distribution (Figure 2c), indicating a flatter, less skewed structure. This suggests co-location ties are more evenly distributed—i.e., users interact with a more diverse set of others. The effect is especially strong for MultiVAE and MF. ItemKNN, however, has negligible impact on α.

Tighter elite clustering. Rich-club density, measured among the top 15 most-connected users, rises sharply for most algorithms as adoption increases (Figure 2d), nearing full interconnection under MultiVAE. This implies that highly connected individuals increasingly co-locate with each other. Again, ItemKNN shows little change from baseline levels, consistent with its overall muted impact.

Visualization of the effect Figures 3a–b show venue visits by a representative user (User 900). Without recommendations ($\eta = 0$), visits are concentrated on a couple of venues in Midtown Manhattan. Under full adoption ($\eta = 1$), the visits of the same user become spatially dispersed. The individual Gini coefficient drops from $G_u = 0.41$ to $G_u = 0.19$—a 54% reduction (Figure 3c).

At the collective level, the opposite trend emerges. The aggregate visits of 250 users become more concentrated under $\eta = 1$, especially in central Manhattan (Figures 3d–e), with the collective Gini increasing from $G = 0.32$ to $G = 0.50$ (Figure 3f). These spatial patterns illustrate the general dynamic observed across most recommender systems: as adoption increases, individual mobility becomes more diverse, while collective behaviour becomes more concentrated.

Understanding the Individual–Collective Diversity Trade-off Recommender systems increase individual diversity by encouraging users to explore new venues. However, under full adoption these exploratory visits often concentrate on the same popular locations. With MultiVAE, the top 10% of venues attract 57% of new visits, up from 33% in the historical data, while less popular venues lose visibility. This suggests that individual diversity rises through shared attention to already popular places.

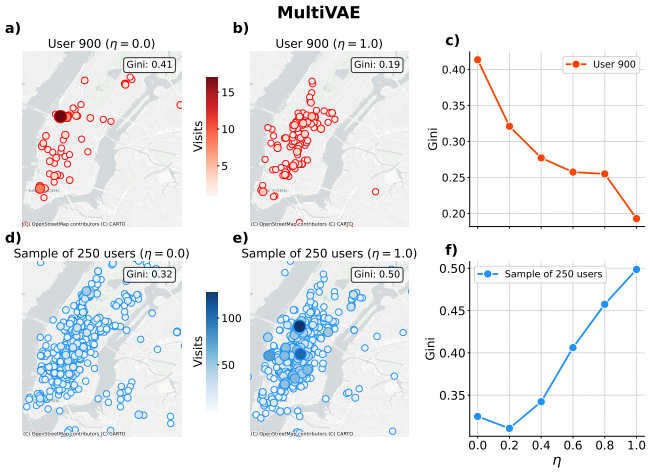

Fig. 3. Venue visitation patterns under MultiVAE for null ($\eta = 0$) and full ($\eta = 1$) adoption. (a–b) Spatial distribution of visits by User 900; visit intensity shown by point size and color. (c) Individual Gini coefficient declines with η. (d–e) Aggregated visits for 250 users show increasing spatial concentration. (f) Collective Gini coefficient rises from 0.32 to 0.50.

CONCLUSION

In this study, we modelled the human–AI feedback loop in next-venue recommendation and introduced a simulation framework to assess its impact on urban mobility. Using real-world check-in data from New York City, we tested several recommender system types and found a consistent pattern: as individuals rely more on algorithmic suggestions, their personal visitation behaviour becomes more diverse. However, this comes at the cost of reduced collective diversity, as many users are funnelled toward the same popular venues. This tension reveals how personalization can unintentionally amplify global concentration. This work has several limitations. Our simulation-based approach cannot fully substitute real platform data, and individual behaviour is modelled with simplified decision rules. In addition, the results are derived from a single urban context, which limits generalisability. However, the framework offers a foundation for future research. Key directions include studying the impact of retraining dynamics, designing recommender systems that self-correct for bias, and incorporating richer models of human mobility and social behaviour.

REFERENCES

1. D. Pedreschi, L. Pappalardo, E. Ferragina, *et al.*, "Human-ai coevolution," Artif. Intell. p. 104244 (2024).
2. L. Pappalardo, F. Simini, S. Rinzivillo, *et al.*, "Returners and explorers dichotomy in human mobility," Nat. communications **6**, 8166 (2015).
3. D. Yang, D. Zhang, V. W. Zheng, and Z. Yu, "Modeling user activity preference by leveraging user spatial temporal characteristics in lbsns," IEEE Trans. on Syst. Man, Cybern. Syst. **45**, 129 (2014).
4. L. Pappalardo, E. Ferragina, S. Citraro, *et al.*, "A survey on the impact of ai-based recommenders on human behaviours: methodologies, outcomes and future directions," arXiv preprint arXiv:2407.01630 (2024).
5. D. Lee and K. Hosanagar, "How do recommender systems affect sales diversity? a cross-category investigation via randomized field experiment," Inf. Syst. Res. **30**, 239 (2019).
6. S. Zhou and R. J. Mondragón, "The rich-club phenomenon in the internet topology," IEEE communications letters **8**, 180 (2004).
7. G. Mauro, M. Minici, and L. Pappalardo, "The urban impact of ai: Modeling feedback loops in next-venue recommendation," arXiv preprint arXiv:2504.07911 (2025).

Mobile Phone Data Sheds Light on the Localized Activity Patterns of Informal Workers

JISHAN DUAN[1,*], SUKANKANA CHAKRABORTY[1], AND NEAVE O'CLERY[1]

[1] *Centre for Advanced Spatial Analysis, University College London, London, UK*
[*] *jishan.duan.23@ucl.ac.uk*

1. INTRODUCTION

The informal economy constitutes a substantial share of global economic activity, especially in developing countries where it represents one-third to two-thirds of total economic output [1, 2]. Informality often emerges in response to burdens such as taxation, labor regulations, and business licensing [1, 3]. Operating outside governmental regulatory systems [1, 4], workers in these informal sectors are difficult to track and often invisible or misclassified in official survey-based records [1, 3], which produces biased insights and leaves them unrecognized and unprotected in social systems [5, 6].

In recent years, location-based data such as mobile phone data have become increasingly common in analyzing urban flows and daily activity patterns [7, 8]. These data offer fine-grained spatial and temporal resolution, which are useful in capturing granular mobility behavior of individuals. We combine large-scale mobile phone GPS data with official firm registration records and building footprint data to develop a novel approach that can distinguish the commuting patterns of informal and formal workers in Bogotá D.C.. The study address two research questions: How do commuting patterns differ between informal and formal workers in terms of spatial distribution, commuting distances and times? And are informal workers' commuting patterns related to amenity and transport hubs?

Our findings show that informal workers, concentrated on the city's outskirts, exhibit highly localized commuting patterns centered near their homes. In contrast, formal workers live and work in more central, well-connected areas near the Central Business District (CBD). Across all socioeconomic groups, low-income informal workers have the shortest commutes, while low-income formal workers travel the farthest. Public transport supply is significantly related to formal commuted between hubs, but for informal workers, this effect is only observed when destinations have a high concentration of informal jobs.

2. STUDY AREA AND DATA

A. Bogotá D.C.

Our study is conducted in Bogotá D.C. which is the capital city of Colombia. It is also the largest city in the country and one of the most populated urban areas in Latin America. In 2022, the employment rate of Bogotá was 59.3% with 42% of workers engaged in informal sectors [9].

We focus our analysis on Bogotá's urban zones, where the population density is 12,681 inhabitants per km². For our analysis, we divided the city into 756 hexagonal grids, each of size $500m \times 500m$.

B. Data

Mobility data Our analysis uses anonymized mobile phone GPS data from Quadrant which consists of location pings across the Bogotá metropolitan area collected between January, 2022 and December, 2022. The time-period was carefully chosen to reduce any effect of COVID-19 restrictions on the data. A pre-processing step was applied to prepare the data for analysis. Specifically, we dropped pings that had poor accuracy. We also dropped users who were active less than 10 times in a month. The cleaned dataset included 318 million pings from 1.3 million unique devices in the Bogotá metropolitan area. We validated this dataset against multiple other datasets (2021 Multipurpose Survey, 2019 and 2023 travel surveys, formal employment records, and WorldPop) for robustness checks.

Registered firm data We use the registered firm data from the Bogotá Chamber of Commerce (CCB) in our work. This dataset includes the name of the registered firm, its founding date, street address, UPZ and ISIC codes. Firms were geocoded using the Google Earth Engine API and then matched to the nearest building in the footprint. Of the 417,725 firms in the CCB dataset, 397,111 (94%) were successfully geocoded and matched to buildings.

Amenity and public transport data We use Points of Interest (POIs) data in our analysis. POI data for Bogotá was extracted from a 2022 local OpenStreetMap (OSM) dataset using the osmium Python package [10]. Each POI included geographic coordinates and the category of business the POI belonged to, namely: food, retail, entertainment, education, health, public services and industrial, yielding a total of 51,949 POIs.

We use public transport data to measure individuals' commuting times and public transport availability at home and work locations. The public transport data was from Bogotá's 2022 General Transit Feed Specification (GTFS), covering the TransMilenio, TransMiZonal and TransMiCable services. The dataset includes 8,406 stops linked to 1,196 trips, with mode identifiers, stop names and WGS 84 coordinates.

3. METHOD AND RESULTS

A. Home/work locations detection and informal/formal workers identification

We detect home and work locations of users using the Mobilkit algorithm [11]. We identify home locations by clustering pings detected between 19:00–7:00 hours (in local time) and work locations between 8:00–18:00 hours, Monday to Thursday (to minimize weekend bias)[12, 13]. Only users with distinct home and work locations are retained. Work locations are then matched to building footprints by maximum overlap within a defined spatial catchment. Users are classified as informal if their work location is not a registered firm, and formal otherwise, yielding 43,356 informal (38%) and 70,538 formal (62%) workers, which is also consistent with official data (42%) [9].

As shown in Fig 1, we find that informal workers largely live in the outskirts of the city and work locally (i.e. with short commutes, 28.9 minutes by public transport on average). Formal workers reside more centrally and travel longer for work on average (31.7 minutes by public transport). We also observe that these spatial divides are more pronounced in low-income groups, where the average commuting time of formal workers is 1.5 times (41.7 minutes) that of informal workers (26.5 minutes).

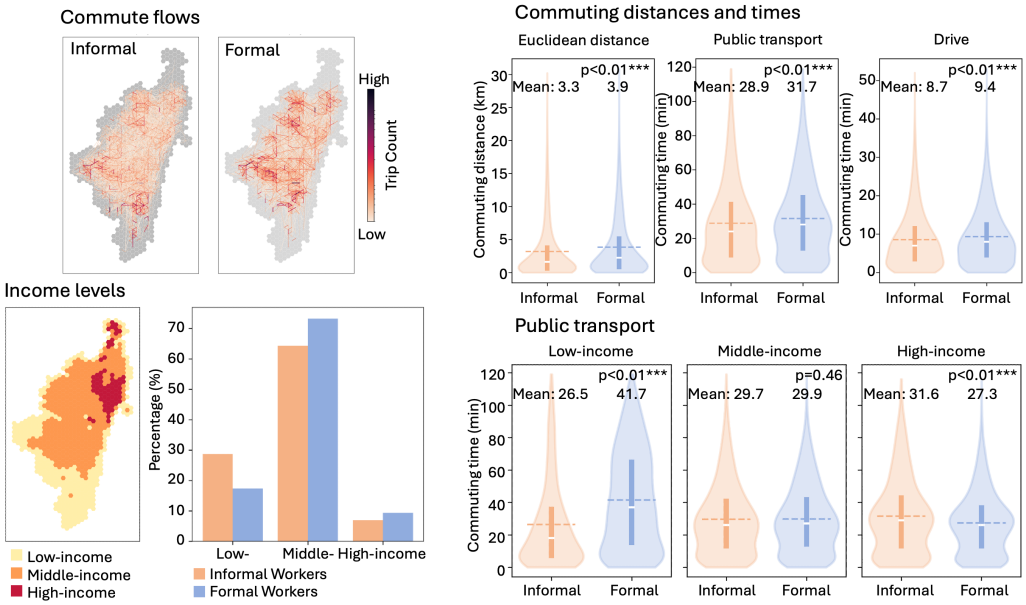

Fig. 1. Informal workers live in low-income peripheral areas and commute locally. Formal workers commute within the central areas of the city. Low-income informal workers have the shortest commutes, while low-income formal workers commute the longest distances.

B. Gravity model framework

Next, to explore how urban structures shape commuting patterns differently for informal and formal workers, we apply gravity-based regression models to analyze flows between home and work locations while focusing on the roles of amenities and transport infrastructure.

In the unconstrained gravity model [14], the flow (T_{ij}) from the origin i to the destination j is proportional to the product of the population at the origin (P_i) and the employment at the destination (E_j), divided by the distance between i and j (d_{ij}). However, as pointed out by Thorsen et al. in [15], the traditional gravity model can often underestimate short trips due to the distance deterrent in its model and this could therefore disadvantage informal workers in our analysis who typically have shorter commutes. Thus, following Thorsen et al. [15] we incorporate a Kronecker Delta (δ_{ij}) in our gravity model to explicitly capture intra-grid commuting:

$$\delta_{ij} = \begin{cases} 0, & \text{if } i \neq j \\ 1, & \text{if } i = j \end{cases} \tag{1}$$

Using a gravity model that captures spatial heterogeneity, we analyze the relationship between commuting flows and key features, including socioeconomic class, amenities, and public transport infrastructure, to assess their effects on informal and formal commutes:

$$\begin{aligned} \log T_{ij} = {} & \beta_0 + \alpha_1 P_i + \alpha_2 E_j + \beta_d d_{ij} + \mu\delta_{ij} \\ & + \sum_k \beta_{O_k} O_{ik} + \sum_l \beta_{D_l} D_{jl} \\ & + \beta_m MiddleIncome + \beta_h HighIncome \\ & + \gamma Informal \end{aligned} \tag{2}$$

where O_{ik} is the density of urban feature k at the origin i, D_{jl} is the density of urban feature l the destination j, $MiddleIncome$ and $HighIncome$ are dummy variables indicating whether a worker resides in a middle- or high-income area (1 if yes, 0 otherwise). $Informal$ is a dummy variable equal to 1 if the worker is informal worker, and 0 if formal, β_0 is the constant, βO_k, βD_l, β_m, β_h, and γ are parameters associated with different variables.

We find that formal commutes are positively associated with public transport supply. Informal commutes are positively associated with public transport supply, but only when the destination has a high supply of informal jobs. We find no significant results for amenities.

REFERENCES

1. C. Elgin, M. A. Kose, F. Ohnsorge, and S. Yu, "Understanding informality," Report C.E.P.R. Discussion Paper 16497, Centre for Economic Policy Research (2021).
2. G. Ulyssea, "Informality: Causes and consequences for development," Annu. Rev. Econ. **12**, 525 (2020).
3. Chen and M. Alter, "The informal economy: Definitions, theories and policies," Report WIEGO Working Paper No. 1 (2012).
4. G. Ulyssea, "Firms, informality, and development: Theory and evidence from brazil," Am. Econ. Rev. **108**, 2015 (2018).
5. J. P. Bocarejo and L. F. Urrego, "The impacts of formalization and integration of public transport in social equity: The case of bogota," Res. Transp. Bus. Manag. **42** (2022).
6. M. Suárez, M. Murata, and J. Delgado Campos, "Why do the poor travel less? urban structure, commuting and economic informality in mexico city," Urban Stud. **53**, 2548 (2016).
7. E. Moro, D. Calacci, X. Dong, and A. Pentland, "Mobility patterns are associated with experienced income segregation in large us cities," Nat. Commun. **12** (2021).
8. T. Yabe, B. García Bulle Bueno, M. R. Frank, *et al.*, "Behaviour-based dependency networks between places shape urban economic resilience," Nat. Hum. Behav. **9**, 496 (2024).
9. DANE, "Gran encuesta integrada de hogares (geih): Metodología," Report, Departamento Administrativo Nacional de Estadística (2023).
10. S. Hoffmann, "Pyosmium," .
11. E. Ubaldi, T. Yabe, N. Jones, *et al.*, "Mobilkit: A python toolkit for urban resilience and disaster risk management analytics," J. Open Source Softw. **9** (2024).
12. J. Reades, F. Calabrese, A. Sevtsuk, and C. Ratti, "Cellular census: Explorations in urban data collection," IEEE Pervasive Comput. **6**, 30 (2007).
13. C. Bergroth, O. Järv, H. Tenkanen, *et al.*, "A 24-hour population distribution dataset based on mobile phone data from helsinki metropolitan area, finland," Sci. Data **9** (2022).
14. G. K. Zipf, "The p1 p2/d hypothesis: On the intercity movement of persons," Am. Sociol. Rev. **11**, 677 (1946).
15. I. Thorsen and J. P. Gitlesen, "Empirical evaluation of alternative model specifications to predict commuting flows," J. Reg. Sci. **38**, 273 (2002).

Revealing Hidden Parental Mobility Behaviors Through Large-scale GPS Data

SHUSEI HAYASHI[1,*], TAKAHIRO YABE[2], AND NAOYA FUJIWARA[3]

[1] Graduate School of Information Sciences, Tohoku University, Sendai, Japan
[2] Department of Technology Management and Innovation, Tandon School of Engineering, New York University, Brooklyn, NY, USA
[3] Graduate School of Information Sciences, Tohoku University, Sendai, Japan
[*] hayashi.shusei.s7@dc.tohoku.ac.jp

1. INTRODUCTION

Urban poverty persistence remains a critical challenge. For example, the OECD estimates that it takes four generations for Japan's lowest-income families to reach average income levels. Children from low-income families who grow up in socially diverse cities earn 20% more as adults compared to their parents' income, highlighting the importance of cross-income interactions [1]. It is pointed out that neighborhood environments play important role in upward mobility for disadvantaged children [2], but it is unclear how local communities contribute to child development.

Human mobility data offers unprecedented opportunities to analyze detailed behavioral dynamics and understand people's urban experiences [3, 4]. It is difficult to directly extract their behavior from human mobility data, because small children do not possess mobile phones. However, it would be possible to understand mobility patterns of their parents, who mediate children's urban experiences, based on human mobility data. This study addresses this gap by developing methods to identify parents through childcare facility drop-off patterns and analyzing their distinct urban mobility behaviors. We focus on the exploratory visits of parents and compare with those of non-parents.

2. DATA AND METHODS

We utilized user-consented GPS location data from Blogwatcher Inc. [5], comprising 826,415 users in Tokyo who possessed at least one location record on every day of February 2022, and 108,789 users in Miyagi Prefecture with location records on at least 100 days in 2019. Point of Interest (POI) data from Zenrin's "Telepoint" database with 6.4 million entries was used to identify visit purposes. Additional geospatial datasets included road boundaries from the Geospatial Information Authority of Japan and administrative boundaries based on the 2020 Population Census.

GPS records of parents were estimated through drop-off behaviors at early childhood facilities. From the POI data, childcare facilities in Tokyo (3,656 facilities) and Miyagi Prefecture (750 facilities) as of 2019 were extracted, and the latitude and longitude of each facility were obtained. Circular polygons with a radius of 30 m centered on the coordinates of the extracted childcare facilities were then generated. Subsequently, using road boundary data, the portions of these polygons extending into roads were clipped, and representative polygons for the childcare facilities were constructed (Fig. 1).

Users, who had at least 10 GPS records in the morning (7-10 AM) in these polygons on at least five days, were estimated as parents. Note that facilities with more than 500 unique users, much larger than the typical size of Japanese childcare facilities, were excluded. Users with too many records or visiting multiple facilities were excluded, as they likely occupational related with these facilities. Numbers of users estimated as parents and non-parents are 2,684 and 68,319 in Tokyo, and 687 and 35,154 in Miyagi, respectively.

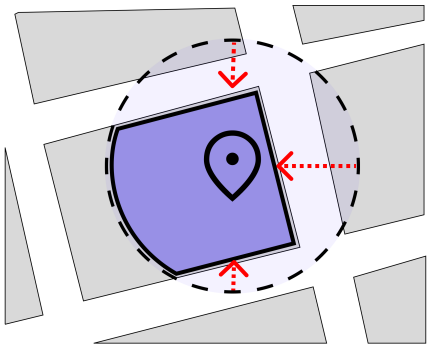

Fig. 1. Schematic diagram of polygons of childhood facilities.

Figure 2 displays typical visit patterns of parents to the childcare facilities; the visits occur only in the morning and the evening, suggesting drop-offs and pick-ups of parents. We analyzed residential distribution and frequency, total duration, average duration and visits home stays and visits. Visit purposes were determined by matching GPS records with the nearest POI within 50 meters. All statistical comparisons used Kolmogorov-Smirnov (KS) tests.

The probability of exploratory visits, unique locations visited only once during the observation period, serve as an indicator of individuals' propensity to explore new locations. This probability depends on the number of previously visited locations as $P_{\text{new}} = \rho S^{-\gamma}$, where $\rho \approx 0.6$ and $\gamma \approx 0.21$ [6]. This relationship is regarded as a universal law of individual behavior. Here, we examine whether this relationship holds in the present dataset.

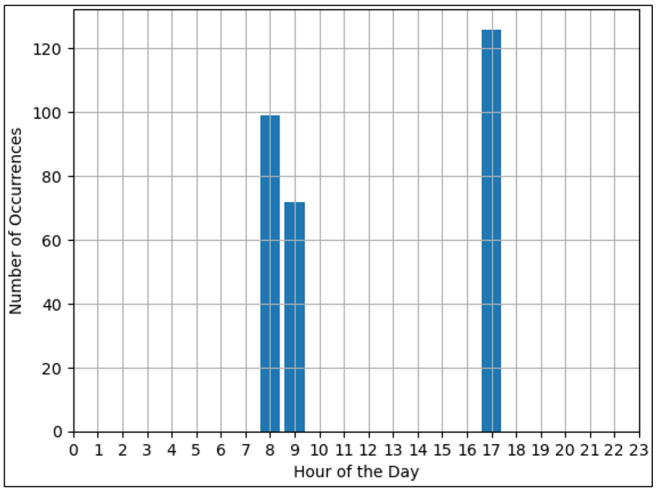

Fig. 2. Example of hourly number of GPS records within a childcare facility for a user estimated to be a parent.

3. RESULTS

Comparison of estimated home locations between parents and non-parents revealed that parents demonstrated distinct residential preferences, concentrating in suburban areas while avoiding central Tokyo's expensive districts. The number of home stays also differed between the two groups, with parents exhibiting significantly higher frequency of home stays compared to non-parents in terms of KS test. These findings would reflect the behavioral constraints associated with childcare responsibilities.

Differences between parents and non-parents in Miyagi prefecture regarding both the frequency of visits to each POI category and the distance from home were observed. Parents visit supermarkets more frequently and restaurants less frequently than non-parents. In Miyagi Prefecture, the frequency of visits to "Baby Goods Stores" was also significantly higher for parents. This is plausible because parents of a child attending a childcare facility may also be raising other similarly aged children, being necessary to buy baby-related goods. The travel distance analysis indicated that non-parents travel longer than parents to the majority of POIs. These patterns may reflect the temporal and financial constraints of parents associated with childcare responsibilities; they prefer affordable suburban residences, spend more time at home, and make fewer visits to restaurants and entertainment venues that are distant and financially burdensome.

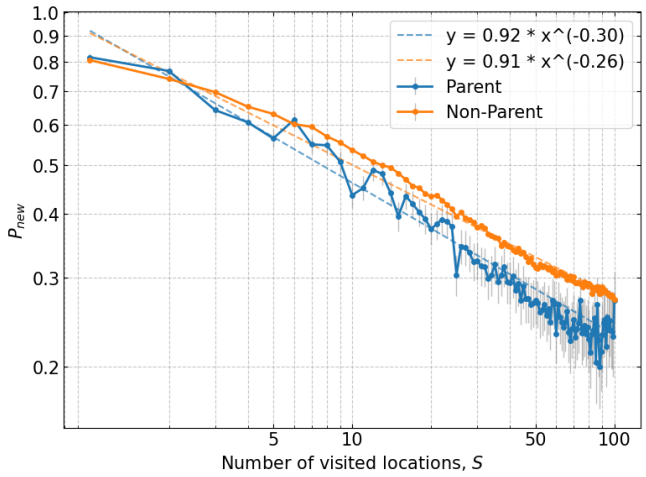

Fig. 3. Probability of visiting a new place, P_{new}, as a function of the number of already visited locations. Solid and dashed lines show the empirical distributions for parents and non-parents and their power-law fits, respectively. These distributions were obtained based on the method proposed in [4]

As shown in Fig. 3, we fitted the data to the power-law function $P_{\text{new}} = \rho S^{-\gamma}$. The estimated values of the parameters are $\rho \approx 0.92$ and $\gamma \approx 0.30$ for parents, and $\rho \approx 0.91$ and $\gamma \approx 0.26$ for non-parents. These estimated parameter values are generally consistent with those reported in the previous research [6].

Although the smaller sample size for parents leads to wider 95% confidence intervals, their P_{new} values are consistently lower than those of non-parents for most values of S. This suggests that parents tend to explore new locations less frequently.

4. DISCUSSION

We estimated parents through childcare facility drop-off patterns and revealed their distinct behavioral characteristics. It was suggested that parents demonstrate constrained mobility patterns including frequent home stays, smaller activity radius, and budget-conscious facility choices. These patterns would reflect the constraints of childcare responsibilities.

Furthermore, the analysis of exploratory visits reveals a crucial aspect of parental mobility. As indicated by the lower probability of visiting new locations, parents tend to be less exploratory than non-parents. This finding would be particularly significant, as it has been reported that a higher propensity for spatial exploration correlates with more economically diverse social contacts [4]. Such interactions with individuals from different backgrounds are known to foster upward income mobility [1]. Consequently, our results suggest that the status of being a parent may inherently place individuals in a situation where opportunities for diverse experiences are limited, not only for themselves but also for their children.

The findings provide insights for urban policy aimed at breaking poverty cycles, as constrained parental mobility may limit children's exposure to diverse urban experiences and cross-income social interactions. However, several limitations warrant consideration and methodological improvement. The parent estimation methodology based on drop-off and pick-up requires further refinement and validation, because it does not capture parents, e.g., who use late-night/early-morning care or school buses. Biases caused by the samples of the mobility data and POI data should be considered carefully. We should also take care of inaccurate assignment of POI, as it is based on the proximity between a GPS record and POI.

Future research should extend data collection periods, validate parent identification methods, improve visit location accuracy, and incorporate socioeconomic factors to examine cross-income interactions. Understanding how parental mobility constraints affect children's access to urban diversity is essential for designing interventions that facilitate social capital formation and break intergenerational poverty cycles.

Funding. Japan Society for the Promotion of Science (JSPS) KAKENHI (24K01108, 24K03007); Japan Science and Technology Agency (JST) PRESTO (JPMJPR21RA).

Acknowledgment. This study was conducted as part of the collaborative research with the Center for Spatial Information Science, the University of Tokyo (Joint Research No. 1273). The authors thank Zenrin Co., Ltd. for providing access to the geocoded telephone directory database "Telepoint Pack DB (Yellow pages)". We also thank Blogwatcher, Inc [5]. for providing the GPS location data used in this analysis.

Disclosures. The authors declare no conflicts of interest.

Data Availability. The GPS location data underlying the results presented in this paper are not publicly available and were shared under a strict research-use-only agreement with Blogwatcher, Inc., ensuring that no individual can be identified from the analyses. The POI data from Zenrin Co., Ltd. (Telepoint Pack DB) used in this study are also subject to licensing restrictions. Other geospatial datasets such as road boundaries and administrative boundaries are publicly available from the Geospatial Information Authority of Japan and the Statistics Bureau of Japan.

REFERENCES

1. R. Chetty, M. O. Jackson, T. Kuchler, *et al.*, Nature **608**, 108 (2022).
2. R. Chetty, N. Hendren, and L. F. Katz, Am. Econ. Rev. **106**, 855 (2016).
3. H. Nilforoshan, W. Looi, E. Pierson, *et al.*, Nature **624**, 586 (2023).
4. E. Moro, D. Calacci, X. Dong, and A. Pentland, Nat. communications **12**, 4633 (2021).
5. Blogwatcher, https://www.blogwatcher.co.jp/ (Accessed 10 September 2025).
6. C. Song, T. Koren, P. Wang, and A.-L. Barabási, Nature **6**, 818 (2010).

Measuring Bias in CDR Mobility Indicators Arising from Multiple SIM Ownership

SVETA MILUSHEVA[1,*], PAUL BLANCHARD[1], AND OSCAR BARRIGA CABANILLAS[1]

[1] World Bank, 1818 H St NW, Washington DC, 20043
[*] smilusheva@worldbank.org

1. INTRODUCTION

In many low- and middle-income countries, a significant share of mobile phone subscribers use two or more active SIM cards. This practice arises from subscribers' overlapping motives to reduce the impact of high off-net tariffs, maximize network coverage, and retain legacy numbers [1–3]. Nationally representative surveys carried out in 10 African countries show that the share of phone users with multiple SIM cards ranges from 19% to 51%, with consistently higher rates observed in urban areas . This highlights the widespread nature of the practice.

Multiple SIM ownership can distort the mobility indicators derived from Call Detail Records (CDRs) transactions, as these statistics rely on the number, frequency, and location of the transactions of each SIM. This paper explores the determinants of multiple SIM ownership, and its impact on the detected mobility and home location for a sample of mobile subscribers from a mobile network operator in Cote d'Ivoire. The paper sheds light on the potential biases of mobility patterns from CDR data and discusses its implications for its use in policy interventions as well as potential solutions.

2. DATA

Phone Survey Data We conducted a phone survey with a randomly selected sample of 2,025 individuals with phone numbers from Orange Cote d'Ivoire. Individuals over 18 years who consented to the survey and to link their survey responses with features generated using their CDR data were included. The survey includes sociodemographic information, assets and consumption, and mobility information. It also asks about multiple SIM ownership.

CDR Data We use CDR data for a set of 761 SIM cards that are linked to phone survey respondents.[1] Individual level transaction data was processed directly by the MNO using open source libraries [4, 5] which produce a set of monthly features on phone and data usage at the individual level aggregated per month. We use 6 indicators that provide measures of mobility and phone usage: number of antennas (representing number of unique places visited), percent of interactions the user had while at home, radius of gyration, frequent antennas (the number of locations that account for 80% of the locations where the user was), active days, and number of interactions. Additionally, we also use individual KPI indicators calculated by the phone provider monthly for billing purposes that include information on the home location of the phone user. We use data for the period from June 2023 to September 2023.

3. METHODS

We first analyze demographic characteristics associated with multiple SIM ownership. We use a dummy variable for owning more than one SIM.[2] We run a logit regression with heteroskedasticity robust standard errors, predicting the likelihood of owning multiple SIMs based on seven characteristics: gender; age category (18-24, 25-34, 35-44, 45-54, 55-64, 65+), the district where the person resided the last 30 days; the household size; a literacy dummy variable; education category (no education, primary, secondary 1, secondary 2, post-secondary and higher education); employment category (employed, unemployed but looking, unemployed not looking, seasonally unemployed, student, retired, housewife, agricultural worker, self-employed); and wealth index.[3] We also use SHAP to capture the importance of different features taking into account the contribution of other features.

We also analyze the relationship between multiple SIM ownership and mobility based on survey responses to the following: (1) whether in the last 3 months the person has spent more than 2 weeks away from home (2) how many nights a week they spend at home (3) how many different villages/towns in Cote d'Ivoire they have traveled to in the last 12 months where they have spent at least 2 weeks and (4) how many trips greater than 5km from home they typically do per day, per week and per month.

We then move into studying potential bias in measurement of mobility with CDR data due to multiple SIM ownership. We may expect this to be the case if multiple SIM ownership leads to non-random gaps in CDR observations that are correlated with mobility behavior. We analyze this first stage by regressing the number of active days and total observations on multiple SIM ownership. We then regress measures of mobility on multiple SIM ownership using linear regression with robust standard errors clustered at the user level. We use CDR features for the four months before and during the phone survey (June-Sept 2023). Given that there may be important differences between those with and without multiple SIMs, we also run the analysis using propensity score matching on demographic and mobility characteristics to identify matched pairs of individuals with and without multiple SIM cards based on nearest neighbor.[4]

Finally, we use the home location from the KPI indicators at the region level and compare to the home location provided by the respondent in the survey. We compare the information from the KPIs for September, which matches the month in which the survey was collected and analyze whether multiple SIM ownership is related to the accuracy of the home location measured by the CDR data. Again, we also limit the analysis using propensity score matching to only those observations where there is a matched pair between those with and without multiple SIMs based on their demographic characteristics.

4. RESULTS

Characteristics Associated with Multiple SIM Ownership Our logit regression shows that men, those aged 35-64, and those with

[1] A requirement from Orange to have clear recordings of the specific consent for each individual in order to link their data limited the sample that has linked data to only 761 individuals.The survey company experienced technical challenges that led to a large part of the sample lacking an audible recording that was accepted by Orange.

[2] Less than 7% of respondents own more than 3 SIM cards.

[3] The wealth index was calculated using Principle Component Analysis based on the ownership of 11 assets: motorcycle, bicycle, refrigerator, cook stove, air conditioning, washing machine, computer, TV, car, bed and radio.

[4] We allow for a maximum propensity score difference between matched pairs of .01 and we match without replacement. This significantly reduces the sample but allows us to compare individuals that are otherwise extremely similar in their demographic and mobility characteristics and only differ in the number of SIM cards they own.

more education are more likely to have multiple SIM cards.Those living in Abidjan, housewives and agricultural workers are less likely to own multiple SIM cards. There is a concave relationship between wealth and SIM ownership, with multiple SIMs having a higher concentration in the middle of the distribution. This aligns with users having multiple SIMs to maximize the usage of different providers' promotions, which the wealthiest are less likely to do. Looking at feature importance based on the SHAP values, we find that gender is the most important predictor, followed by being employed, wealth score, having no education, and being in the 45-54 age group (Figure 1).

Based on the survey data, we find a significant positive correlation between multiple SIMs and the number of locations where the person spent more than 2 weeks in the past 12 months. This is aligned with the hypothesis that individuals who may travel and spend more time in different parts of the country might be more likely to have multiple SIMs potentially due to differential provider coverage across the country . Once we control for gender, age category and education category, though, this relationship is no longer significant; therefore, that relationship may be driven by the characteristics of individuals who own multiple SIMs and who may also travel more.

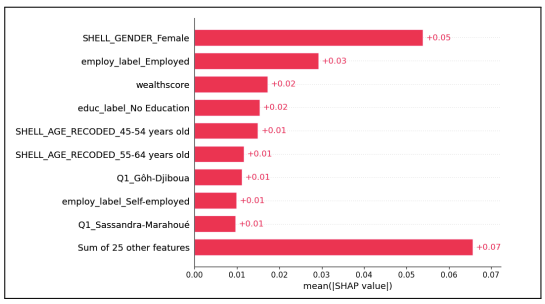

Fig. 1. Mean absolute SHAP values of demographic features predicting multiple SIM ownership. Bars show each feature's average contribution magnitude to the model's predictions across all individuals in the sample.

Bias in CDR Mobility Measures Due to Multiple SIM Ownership Regressions of number of voice, SMS, data and total interactions on a dummy for multiple SIM ownership all produce negative coefficients, though they are not significant. When we limit to the matched sample based on propensity score matching, all regression coefficients remain negative and they are significant for data and total observations. The difference is very large, with multiple SIM owners having 44% fewer data observations. Since data observations are collected passively, if users switch between SIM cards in their phone, a lot fewer observations will be collected since any one SIM is active for a shorter period of time compared to someone with only one SIM. The number of active days per month are not significantly different between those with and without multiple SIMs. This aligns with the fact that SMS and call usage does not differ significantly.

The large difference in the number of data observations could lead to bias in the amount of mobility measured since fewer observations would mean fewer opportunities for capturing the location of the SIM. Multiple SIM ownership is not significantly related to total number of antennas in a month, though when limiting the sample to only matched observations, the coefficient becomes more negative (Figure 2). Multiple SIM ownership is positively and significantly associated with the radius of gyration. The number of antennas that account for 80% of the user's locations is not significantly related to owning multiple SIMs. The percent of observations at home is significantly affected by owning multiple SIMs when we restrict the sample to matched observations. The percent of observations is higher for those with multiple SIMs, which would signal missing a higher percent of locations when away from home. Finally, the entropy of antennas is negatively related with owning multiple SIMs. This relationship is significant when we limit to the matched sample.

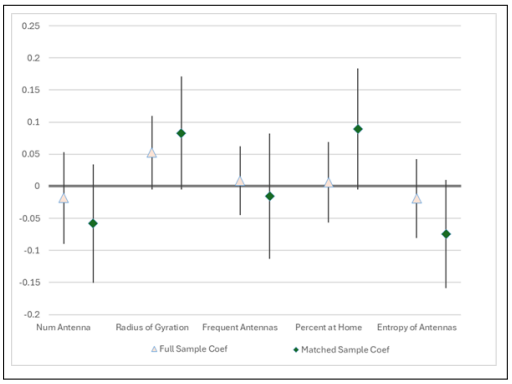

Fig. 2. Coefficients for Relationship Between CDR Mobility Features and Multiple Sim Ownership Using Full Sample and Matched Sample. The green diamonds are the coefficients on the multiple SIM dummy variable from linear regressions using the matched sample based on propensity score matching. The light orange triangles are the coefficients on the multiple SIM dummy variable from linear regressions using the full sample. Data is from June-Sept 2023. Standard errors are clustered at the SIM card.

Home Location Analysis We find that in 24% of SIMs, the home location in the KPI data is different from the home location provided by the respondent. Regressing the accuracy of the home location on multiple SIM ownership, there is no significant relationship. This is true even when using propensity score matching to limit the sample to a set of individuals that are matched on characteristics.

5. DISCUSSION

Our evidence shows that owning more than one SIM card is related to socio-economic characteristics, concentrated among middle-income, working-age men outside the capital. By combining individual survey data with CDR indicators of mobility, we show that ownership of multiple SIMs systematically thins the CDR trail, as multi-SIM users generate 44% fewer data events. This differential usage does not affect long-term location and mobility measures such as home location or most frequent antennas since these are still detected with the more limited observations. It can miss some locations and lead to bias in measures of short term mobility, which we see in the impact on percent of observations at home, the entropy of antennas and the radius of gyration. Moreover, our results point to the fact that ignoring multiple SIM ownership can lead to under-estimating mobility for populations living in areas with coverage gaps and more sensitive to price. Future work should therefore explore the mobility indicators when including information from multiple carriers, and weighting schemes to correct or attenuate the biases documented in the paper.

REFERENCES

1. K. Andersson and D. Göller, Rev. Netw. Econ. **19**, 189 (2021).
2. B. Ilembo and J. Walwa, Tanzan. J. Dev. Stud. **20** (2022).
3. E. Sutherland, Sim Cards & Cust. (September 5, 2009) (2009).
4. Y.-A. De Montjoye, L. Rocher, and A. S. Pentland, The J. Mach. Learn. Res. **17**, 6100 (2016).
5. E. Aiken, S. Bellue, D. Karlan, *et al.*, Nature **603**, 864 (2022).

Maintaining Connectivity in Subterranean Emergency Response: A 5G and Wi-Fi Mesh Mobility Dataset

NICOLAS PEUGNET[1], ANNE FLADENMULLER[1], THI-MAI-TRANG NGUYEN[1,2], BRIGITTE KERVELLA[1,3], AND JULIEN THOUROT[4]

[1] *Sorbonne Université, CNRS, LIP6, F-75005 Paris, France {nicolas.peugnet,anne.fladenmuller,thi-mai-trang.nguyen,brigitte.kervella}@lip6.fr*
[2] *Université Sorbonne Paris Nord, L2TI, 93430 Villetaneuse, France*
[3] *Université Picardie Jules Verne, UPJV, Amiens, France*
[4] *Ingénieur en Chef de l'Armement (ICA), Brigade des sapeurs Pompiers de Paris (BSPP)*

1. INTRODUCTION

Reliable communication is critical for firefighters during emergency interventions, especially in environments such as underground parking lots, where telecommunication infrastructures are often absent or unreliable. While private 5G networks offer a solution, their range is limited by physical obstacles. Wi-Fi mesh networks, used as mobile ad-hoc networks (MANETs), can help extend coverage through multi-hop relaying.

To explore the effectiveness of hybrid 5G/Wi-Fi mesh networks, we conducted a real-world experiment in collaboration with the Paris Fire Brigade (Brigade des Sapeurs-Pompiers de Paris, BSPP), during a live rescue training exercise. This paper presents the resulting dataset [1] of a realistic emergency response in a subterranean environment. The dataset provides mobility traces and connectivity logs and is made openly available to support further research in the field of resilient mobile networking for public safety applications.

2. INTERVENTION SCENARIO

The BSPP has a strictly defined operating mode. A team always consists of 6 members: a driver, a team leader (*chef d'agrès*) and two pairs of firefighters, each composed of a unit leader (*chef d'équipe*) and a servant. When multiple teams are deployed at the same location, a commanding officer (*chef des opérations de secours*, or *COS*) must also be present to coordinate the response.

The exercise scenario was designed in close collaboration with the firefighters to reflect realistic conditions, involving two operational teams and one commanding officer.

The experiment was carried out on one of the *Sorbonne Université* campuses underground parking (Figure 1), which comprises five interleaved levels, numbered -1 (topmost) to -5 (lowest). A single fire source was positioned in the left section of level -5.

Two firefighter teams participated in the drill. The movements of Team 1 (red arrows) and Team 2 (blue arrows) are annotated in Figure 1 with numbered steps corresponding to key events and their timestamps t_i (see Figure 4).

At around 16:12 briefing of both teams in front of the parking at ground level. Then at 16:19 (t_1, arrow 1.1) Team 1 went down in the staircase via emergency access 1 and at 16:30 (t_3, arrow 1.2) the first pair of Team 1 entered the exclusion zone and simulated an accident. At 16:38 (t_4, arrow 1.3), after receiving an emergency signal, the second pair of Team 1 responded to evacuate the first pair (Fig. 2). The wireless router of pair 1 is left on site. Meanwhile at 16:22 (t_2, arrow 2.1), Team 2 went to the emergency access 2. Shortly after (arrow 2.3), the first pair of Team 2 entered the exclusion zone and at 16:42 (t_5), located the fire source and attached the lifeline to mark the path to the fire source. At 16:50 (t_6, arrow 2.4) Team 2 returned and the exercise ended at 16:58.

[1] Accessible online at: https://www.youtube.com/watch?v=VL96Zk4tzs0

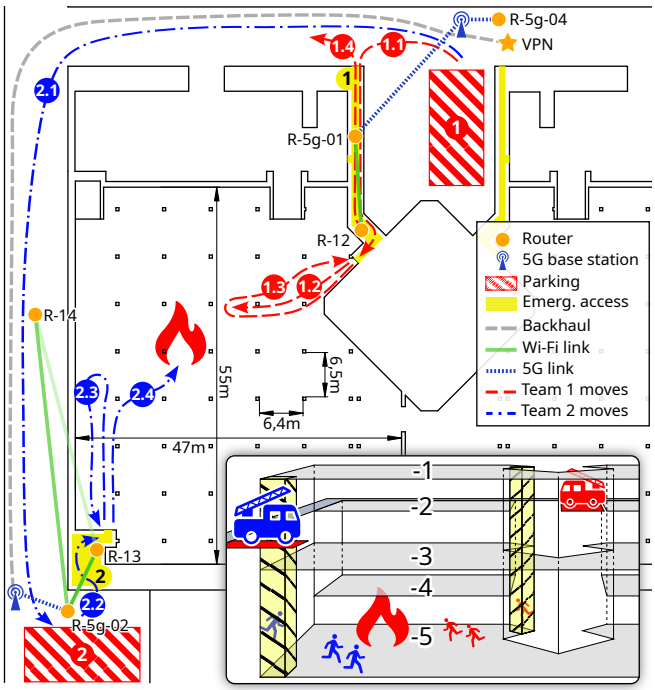

Fig. 1. Plan of the campus underground parking. The lower-right part is a 3D view of the parking, the rest is a top view showing the positions of the various elements of the intervention and reflects the different steps of the scenario.

Fig. 2. Screenshot from the demonstration video[1] capturing a moment of the exercise. The second pair of Team 1 is rescuing the first one at 16:41, which corresponds to arrow 1.3 in Fig. 1.

3. NETWORK TOPOLOGY

The demonstration showed that MANETs can be used to extend the range of private 5G networks and relay the information in remote places, thus motivating the use of a hybrid Wi-Fi and 5G

architecture. Each firefighter in the experiment carried a terminal device connected to a Wi-Fi router. One router was assigned for each pair and additional routers were carried individually by the other firefighters. Wireless routers provided by *Green Communications* created a wireless mesh routing infrastructure using QOLSR [2] based routing algorithm.

The firefighter trucks were connected in a tactical 5G bubble through a private 5G base station, with coverage extended by 5G/Wi-Fi and Wi-Fi routers. Some routers (`R-5g-xx`) featuring 5G capabilityare able to connect directly to the truck antennas, which were deployed at fixed locations as shown in Fig. 1.

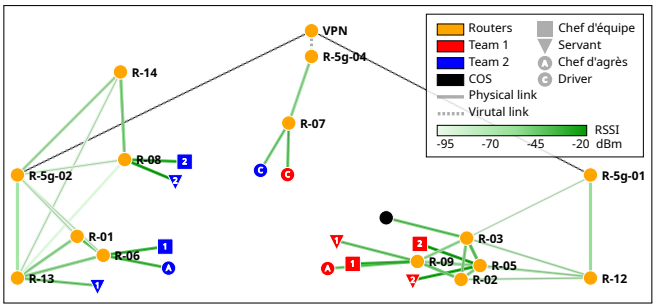

Fig. 3. Screenshot of the preview application[2] showing the network obtained with the VPN, omitting the 5G backhaul.

Due to limitations with the available hardware and software we had to use a QOLSR aware VPN to join together the mesh networks through the 5G core network. Fig. 3 shows the resulting topology, with the dashed lines being the virtual links between each client and the VPN server. The `5g-01` and `5g-02` routers are each connected to a different 5G cell, but as they are connected to the VPN server, behind the 5G core-network, they can act as a bridge between the disconnected meshes. From their point of view, each router sees only one big network that contains all of them.

4. COLLECTED DATA

Each router logged the network status every second, storing it in GraphML [3] files. These files contain directed graphs with 30 nodes: 15 routers, 13 terminal devices, and 2 5G base stations. Node and link attributes (listed in Table 1) vary over the 46-minute experiment, as does the number of links (see Fig. 4).

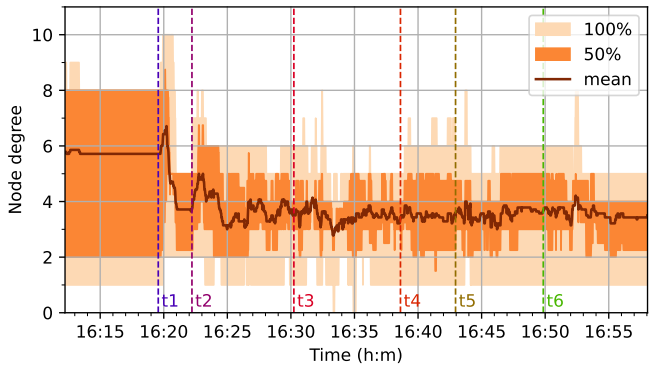

Fig. 4. Router nodes' degree (number of links from this node) evolution over time, only accounting router to router links. The timestamps t_i correspond to the events described in Section 2.

[2]Accessible online at: https://ene5ai.proj.lip6.fr/

The dataset [1] contains two archives, the raw data collected by each router (`raw.zip`) and the preprocessed data in multiple formats (`artifacts.zip`). This latter archive has been produced by merging the data from all routers into a single directory of GraphML files, providing a global state of the network, then exporting it in JSON and aggregating it in CSV spreadsheets.

Table 1. Main link and node attributes

Link attributes		Node attributes
`rx-bytes`*	`tx-bytes`*	`hostname`
`rx-packets`*	`tx-packets`*	`type`
`connected-time` (seconds)*	`tx-retries`*	`main-ip-addr`
`signal-avg` (RSSI)*	`tx-failed`*	`routes`**
`rssi`†		`imei`†

* Only for Wi-Fi links, obtained from nl80211 [4]
** Only for router nodes, obtained with `ip -d -j r` [5]
† Only for 5G links/nodes

5. CONCLUSION

This dataset captures firefighter connectivity and mobility traces during a live intervention scenario, showcasing a real-world use of a hybrid 5G/Wi-Fi network. This open-access dataset allows for simulations of realistic emergency response situations.

REFERENCES

1. N. Peugnet, A. Fladenmuller, T.-M.-T. Nguyen, and B. Kervella, "ENE5AI Dataset (5G and Wi-Fi Mesh Mobility)," https://doi.org/10.57745/DFMCYS (2025).
2. H. Badis and K. Al Agha, "QOLSR, QoS routing for ad hoc wireless networks using OLSR," Eur. Trans. on Telecommun. **16**, 427 (2005).
3. U. Brandes, M. Eiglsperger, J. Lerner, and C. Pich, "Graph Markup Language (GraphML)," in *Handbook of graph drawing and visualization*, R. Tamassia, ed. (Chapman & Hall, London, 2010), pp. 517–541.
4. J. Berg, "Linux wireless," https://wireless.docs.kernel.org/ (2024).
5. M. Litvak, *ip-route(8) — Linux manual page* (2012).

The Anatomy of Olympic Games: a Mobile Traffic Demand Perspective

MÁXIMO PIRRI[1, 2, *], DIEGO MADARIAGA[1], ZBIGNIEW SMOREDA[3], AND MARCO FIORE[1]

[1] *IMDEA Networks Institute, Spain*
[2] *University of Alcalá, Spain*
[3] *Orange Research, France*
[*] *maximo.pirri@networks.imdea.org*

1. INTRODUCTION

Summer Olympic Games are the major multi-sport event in the world, bringing together athletes, coaches, reporters, and fans from over 200 countries. During the 2024 Olympics, 9.5 out of the 10 million available tickets were sold as five different sports broke attendance records; fan zones attracted over 90,000 people per day; and 11.2 million visitors traveled to the Paris region [1].

Such a large-scale multi-week event clearly represents a circumstance of particular stress on the local transport, accommodation or communication infrastructures. Yet, our understanding of how a city-wide happening like the Olympics impacts the indigenous mobile traffic demands or how Mobile Network Operators (MNOs) prepare to accommodate them is limited.

In this paper, we provide first insights into the dynamics of mobile traffic demands during the Olympic Games. We benefit from the privileged vantage point of Orange –the MNO serving as the official communications partner to the event organization– and provide original figures about the overall data traffic usage in the Paris region during the Olympics and the specific demands at the venues hosting the different events.

2. OLYMPICS AND CITYWIDE DEMANDS

We first study the impact that Olympic Games have on the overall network load. The Games are held during the summer period, when seasonal migration typically induces a significant drop in mobile traffic consumption, as shown in Fig. 1 using anonymized traffic volumes to abide by the non-disclosure agreement with the MNO. There, the reduction in the demand during summer is computed with respect to a regressed line that captures the growing trend of consumption in normal (*i.e.*, non-vacation) times over the full observation interval: the drop is in the 30–40% range. However, the Olympics have a clear impact on summer mobile network utilization, as traffic is much higher than usual for the period and reduction in the demand is strongly mitigated.

While the total traffic load generated during the Olympic Games is not extraordinary compared to that normally absorbed by the local RAN, it may be that the demands during the event have peculiar spatiotemporal dynamics that do not align well with the existing RAN planning: for instance, the Olympics may cause network dimensioning issues and service degradation if they were to create moments of high demand in specific neighborhoods where the peak traffic is usually low. To verify if this is the case, we study the spatial and temporal patterns of the overall mobile traffic during the Olympics and compare it against that observed in normal non-holiday and summer conditions.

Interestingly, the geographical distribution of traffic looks almost indistinguishable across the three time periods, implying that the Games do not induce spatial unexpected patterns in the mobile demand. This is confirmed by measuring the Earth Mover's Distance (EMD) between maps of the three different periods, which is in the order of 0.01–0.02; for comparison, the

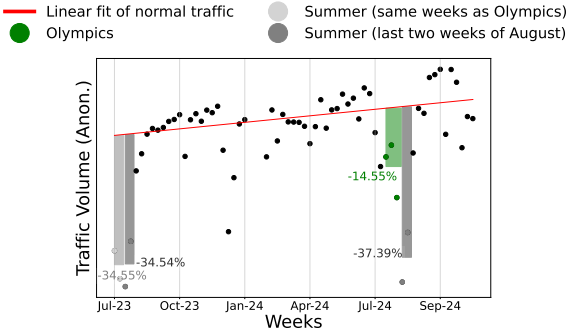

Fig. 1. Weekly mobile traffic consumption (dots) for the observation period between Aug-23 and Oct-24, with a linear fit over the normal, *i.e.*, non-vacation, periods (red solid line). Traffic changes are reported for the Olympic Games weeks (green region), and typical summer periods (dark gray).

EMD between morning and afternoon periods of a typical day is ten times higher at around 0.1, indicating that the traffic fluctuations normally observed in a working day are much higher than the differences in the geographical patterns of demands spawned by the Olympic Games.

When looking at time series, not only the volume but also the temporal fluctuations of traffic observed during the Olympics seem an interpolation of the behaviors in the summer and normal periods. We more formally prove our point by computing the dynamic time warping (DTW) distance between the time series, upon an individual min-max normalization that removes the bias of volume differences. We observe that the DTW is 30% lower between the Olympics and a normal period than between the latter and any summer time. This also holds for the same weeks of the Olympics in 2023, denoted by a lighter gray interval in Fig. 1, proving that the temporal patterns observed during the Olympic Games are atypical due to the event.

Overall, Olympics do not appear to strain the Paris citywide communication infrastructure and the majority of legacy RAN deployment seems sufficient to serve demands during the event.

3. A FOCUS ON OLYMPIC VENUES

As the footprint of the Games on the citywide mobile traffic demand is not remarkable, we study if more circumstantial effects emerge at the venues where the sports events take place. For each Olympic venue, we single out the RAN antennas providing coverage to the local events by considering all sites located within the venue surface plus a 500-meter surrounding buffer. This lets us associate to every venue a mobile traffic time series for the three weeks of the Olympics: Fig. 2 shows examples of such time-varying demand volumes at a venue. Time intervals during which qualifying events and finals take place are marked in gray and green, respectively. It is easy to spot a big increase in mobile traffic whenever an event is underway or close to begin.

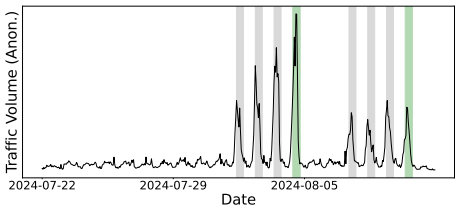

Fig. 2. Time series of the mobile traffic at Golf National venue during the three Olympic weeks. Shaded intervals pinpoint the qualifying rounds (gray) and finals (green) periods.

We quantitatively assess the impact of the Olympics on mobile demands at all venues in Fig. 3. We examine the increase in traffic relative to a non-holiday period as observed at times when Olympic events happen at each venue (blue). Formally, the relative increase R_e for an Olympic event e is computed as

$$R_e = \frac{r_e - 1}{r_e + 1}, \quad \text{with} \quad r_e = \frac{O(T_e)}{N(T_e)}, \tag{1}$$

where T_e is the time span of the event, $O(\cdot)$ is the measured traffic at the venue during the Olympics, and $N(\cdot)$ is the average traffic recorded in non-holiday days at the venue during the same weekday and hours of T_e. As a term of comparison, we report Eq. (1) metric for typical summer (yellow) replacing $O(T_e)$ with the average demand of the remaining August 2024 weeks.

Fig. 3 shows the distribution of the R_e values at each venue as an interquartile box plot reporting the median and the 5$^{\text{th}}$ to 95$^{\text{th}}$ percentile as whiskers. The plot reveals that activity during events is higher than normal (>0) at 24 out of the 25 sport venues, as well as in fan zones and Olympic Village. The effect is also present at 3 out of 4 airports but not at train stations. Control tests on typical summer show a negative relative increase in almost all venues in the rest of August 2024, hence the surge in mobile demands above is to be ascribed to the Olympics.

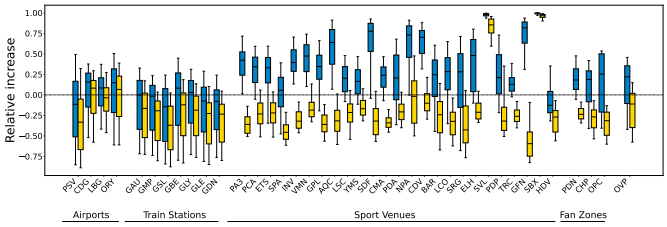

Fig. 3. Distribution of relative traffic volume increase during Olympic events (blue) and during typical summer (yellow).

4. OLYMPICS AND MOBILE SERVICES

The observed unusual traffic demands at Olympic venues may be linked to unconventional consumption of specific classes of mobile applications. To explore this possibility, we quantify the relative usage of popular applications at Olympic venues. More precisely, we compute the symmetric Revealed Comparative Advantage (sRCA) of each application a at every venue v as

$$\text{sRCA}_{a,v} = \frac{\text{RCA}_{a,v} - 1}{\text{RCA}_{a,v} + 1}, \quad \text{where} \tag{2}$$

$$\text{RCA}_{a,v} = \frac{\bar{O}_{a,v} / \sum_{a'} \bar{O}_{a',v}}{\sum_{\ell} \bar{O}_{a,\ell} / \sum_{\ell} \sum_{a'} \bar{O}_{a',\ell}} \tag{3}$$

and $\bar{O}_{a,v} = \sum_{T_e \in \mathcal{V}} O_{a,v}(T_e)$ is the traffic volume generated by a at v during all time intervals $\mathcal{V}$ when Olympic events were hosted at

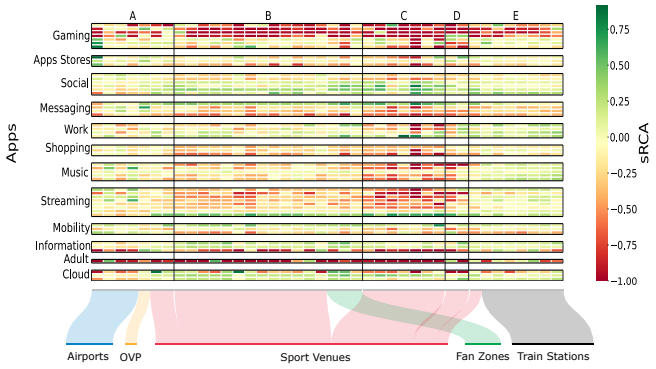

Fig. 4. Heatmap of the sRCA values for 50 mobile applications (rows) at clustered Olympic venues (columns). In the bottom part, a Sankey diagram maps venue clusters and categories.

v. We denote by ℓ all the locations in the Paris region. The RCA expression in Eq. (3) captures how prevalent the demand for application a is consumed at venue v (numerator) with respect to its usual consumption in the region (denominator) [2]. The sRCA equation in Eq. (2) simply normalizes the RCA in the $[-1, 1]$ range.

We summarize the sRCA values for all combinations of applications (rows) and venues (columns) as a heatmap in Fig. 4. We also group the venues based on their application demand distribution, by running an agglomerative hierarchical clustering and using as similarity metric the Euclidean distance between the vectors of sRCA values across applications. This results in five clusters (A to E) showing comparable patterns of sRCA values.

The most interesting behavior are observed in the two main clusters of sport venues (B and C). The two clusters appear to show similar patterns of application usage, but with different intensity: *e.g.*, both feature increased demands for many Social apps, whereas the consumption of streaming services tends to decreace. Yet, the effect is much stronger in cluster C than B. A closer inspection reveals that cluster C is composed of the seven Olympic venues with the highest mean relative increase in traffic volume in Fig. 3. This is not accidental: the demands at those venues are almost completely induced by the Olympic Games, while the usages in cluster B mix the significant habitual traffic at those locations with that produced by the Olympics. We conclude that usages in cluster C are more representative of a *pure* demand created by subscribers attending Olympic events.

5. CONCLUSIONS

We analyzed the mobile traffic served during the 2024 Olympics by the operational network of the official communications partner of the event, revealing previously unknown phenomena in the demands generated by such a unique happening.

Acknowledgment. This work was supported by grant CNS2023-143870 (6G-IRONWARE), funded by MICIU/AEI /10.13039/501100011033 and by the European Union NextGenerationEU/PRTR; and by CoCo5G (ANR-22-CE25-0016), funded by the French National Research Agency.

REFERENCES

1. P. je t'aime, "Paris 2024 olympic games: Preliminary tourism report," https://parisjetaime.com/eng/media/article/preliminary-report-olympic-games-a1755 (2024). [Accessed 11-05-2025].
2. B. Balassa and M. Noland, ""revealed" comparative advantage in japan and the united states," (1989).

Exploring urban event dynamics using mobile application measurements

DIEGO MADARIAGA[1], ANDRÉ F. ZANELLA[2,†], MÁXIMO PIRRI[1], ZBIGNIEW SMOREDA[3], AND MARCO FIORE[1]

[1] IMDEA Networks Institute, Madrid, Spain
[2] Telefónica Innovación Digital, Barcelona, Spain
[3] Orange Research, Châtillon, France
[*] diego.madariaga@networks.imdea.org
[†] Work done while at IMDEA Networks Institute

1. INTRODUCTION

Over the years, smartphones have become crucial tools for advertising, communicating, and reporting public events, bridging online and offline. Through mobile applications, interested parties can report and get informed in real time about these events, as well as instantly communicate with other peers. Overall, this facilitates the exchange of emotions and information among actors of public events [1]. This work expands on previous work by proposing a new framework to model the dynamics of multiple types of public events through the lenses of mobile networks, showcasing the spatiotemporal impact of those events, and offering insights into the relevance of different applications for each type of event. In this study, we focus on the impact of different types of community events on the consumption dynamics of mobile services. We refer to community events as any activity that is public in nature and can be reasonably anticipated to be attended by community members. Therefore, community events take place in a specific location and time, and involve the gathering of a group of people sharing a common interest. This definition includes, but is not limited to, trade fairs, music concerts, festivals, political protests, parades, and sports events.

2. DATA COLLECTION AND MODELING

Network traffic measurements are provided by Orange, a global telecommunications company with a leading position in France, in the deployment of its 4G and 5G infrastructures. Data was collected for a total of 239 days, from January 31 to May 31, 2023 and from May 20 to September 15, 2024.

For the purpose of this work, we focus on events within three core categories: (*i*) sports events, including football and rugby matches, as well as running competitions; (*ii*) social events such as public demonstrations related to diverse causes; and (*iii*) cultural events like concerts and festivals. All these events can also be classified into static, *e.g.*, those happening at a stadium or arena, or dynamic, *e.g.*, public marches. After investigating gatherings within these groups, we select the events with the highest attendance within the dates in our dataset, and collect information about their time and location. This set of target event days is used to generate ground truth, pointing out those antennas greatly affected by the target events. Due to the natural diversity of mobile services, different types of community events are expected to impact a different set of applications, as it depends on the activity carried out by the people gathered. Therefore, we leverage the Boruta feature selection algorithm to ensure that all and only the relevant mobile services are used to characterize the ground truth in each case. In our configuration, the Boruta algorithm is built around an XGBoost classifier (XGBC), which is trained using the original features in the ground truth as well as permuted copies of all these features. Thus, Boruta tells apart

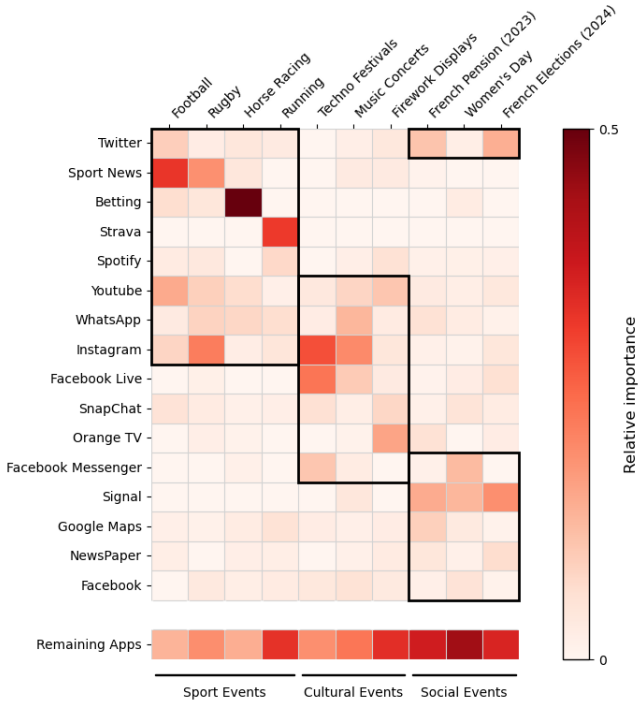

Fig. 1. Mobile traffic consumption patterns computed as the relative contribution of individual mobile services to the identification of different community events.

those mobile service features that are shown to be significantly more important than their randomized versions. Considering the application consumption features that remain after the feature selection stage, we train different XGBC models to learn the mobile traffic consumption patterns observed in the ground truth, obtaining an F1-score larger than 0.9 for each type of event.

From the modeling process, we find that each event category owns a well-defined mobile traffic usage signature that tells about the under- or over-consumption of individual services during a certain type of event. More precisely, we compute and inspect Shapley-based explanations of each community event model to pinpoint the mobile applications that better characterize each type of event. Figure 1 visualizes the relative contribution of key mobile services to the output of community event models. When comparing the most relevant applications inside individual core categories (areas surrounded by solid-line black rectangles), we find that events of the same nature tend to depend on a unified set of mobile services. In contrast, different core categories rely on almost disjoint sets of mobile services, evidencing that the nature of each event strongly delimits which are the most relevant applications. Despite the overall alignment

in terms of most relevant apps inside core categories, the relative contribution of these apps greatly varies across individual community events.

For example, even though sports events share some common patterns, such as the over-consumption of WhatsApp and the under-consumption of YouTube, they exhibit some significant differences. For instance, all types of sports events exhibit an over-consumption of Sport News websites, yet this consumption pattern is much more characteristic of football games and, to a lesser degree, rugby games. Similarly, the over-consumption of Instagram is much more related to rugby games, while the over-consumption of Twitter/X is more connected to football games. Remarkably, horse racing events and running races show a disproportional contribution of mobile applications, with a single application dominating the others. Indeed, the over-consumption of betting websites represents almost 50% of the contribution to modeling horse racing events, being much less relevant for football and rugby games. In the case of running races, the Strava application, used to track physical activities, constitutes over 30% of the relative contribution to the model.

Public demonstrations exhibit a generalized over-utilization of Signal for instant messaging, which contrasts with the leaning towards WhatsApp during sports and cultural events. Signal is an open-sourced and end-to-end encrypted application that has gained increasing attention as an alternative to WhatsApp because of its perceived privacy and security [2]. Indeed, previous research has indicated that protestors rely on Signal as a platform for coordination and self-defense against police surveillance [3]. Both waves of protest in response to the French pension reform in 2023 and the outcome of the French elections in 2024 also align on the relevance of higher consumption of Twitter/X and NewsPaper services, implying a stronger similarity in the behavior of participants across these instances of social uprising. Finally, protests related to the French pension reform in 2023, which consisted mainly of marches within city centers, are more linked to the over-consumption of Google Maps. Contrarily, the other two categories of social movements are mostly related to crowd gatherings in key areas of French cities, thus indicating a lower dependence on navigation services.

Regarding the core category of cultural events, music-related community events (techno festivals and music concerts) are largely characterized by higher consumption of Instagram and Facebook Live, which indicates the prevalence of live sharing on social media during these events. Nevertheless, both types of events still exhibit important differences in their mobile app consumption patterns, such as the predilection of WhatsApp for instant messaging and higher consumption of SnapChat during music concerts, in contrast to the preference for Facebook Messenger and the under-utilization of SnapChat in techno festivals. As for firework displays, they are mostly related to an increased usage of Orange TV for live TV broadcasting, followed by a generalized over-consumption of social media.

3. SPATIOTEMPORAL RECONSTRUCTION OF EVENTS

Following the previous stages, we obtain different event-specific models that can be used to explore any day and city in our dataset and identify occurrences of similar events. Thus, by running any of our classification models over an entire target day and city, we obtain a list of potentially affected carriers by a particular type of event during specific time intervals. Given that the community events considered in our study are expected to be continuous in space and time, we incorporate a secondary step to reduce the number of false positives and pinpoint any potential event as a singular entity. More precisely, we apply a density-based spatiotemporal clustering algorithm (ST-DBSCAN) over

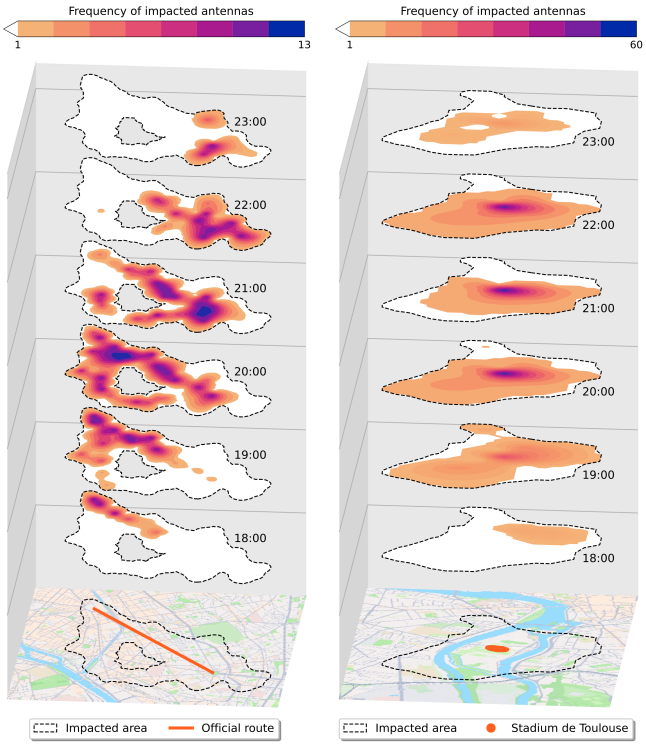

Fig. 2. Identification and reconstruction of two large-scale community events in France, a public demonstration in Paris (left) and a football match in Toulouse (right).

the pairs of carriers and time intervals previously classified as an event by the XGBC. Therefore, a potential event is associated with a dense spatiotemporal cluster of carriers classified as affected by the model. Notably, this method allows the isolation of multiple events of the same nature happening on the same day, as they will likely fall into different clusters.

After experimentation, we find that the combined XGBC and ST-DBSCAN consistently identify community events of varying magnitudes across different locations, based on traffic patterns observed in a small set of example events. These results confirm that events within the same category exhibit a stable traffic consumption signature, which can be leveraged to study and compare user behavior in diverse situational contexts. As representative examples, Figure 2 presents the resulting reconstruction of two large-scale events held in France, a public demonstration and a football match. As shown, the proposed methodology provides insights into the spatiotemporal dynamics that shape human behavior around the area where the events take place.

Acknowledgment. This work was supported by CoCo5G (ANR-22-CE25-0016), funded by the French National Research Agency; and by grant CNS2023-143870 (6G-IRONWARE), funded by MICIU/AEI /10.13039/501100011033 and by the European Union NextGenerationEU/PRTR. D. Madariaga's work was supported by the ORIGAMI project (GA 101139270) funded by SNS JU and the European Union.

REFERENCES

1. J. T. Jost, P. Barberá, R. Bonneau, *et al.*, "How social media facilitates political protest: Information, motivation, and social networks," Political psychology **39**, 85 (2018).
2. C. F. Griggio, M. Nouwens, and C. N. Klokmose, "Caught in the network: The impact of whatsapp's 2021 privacy policy update on users' messaging app ecosystems," in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, (2022), pp. 1–23.
3. T.-y. Ting, "From 'be water'to 'be fire': Nascent smart mob and networked protests in hong kong," Soc. Mov. Stud. **19**, 362 (2020).

Unequal Journeys to Food Markets: Continental-Scale Evidence from Open Data in Africa

ROBERT BENASSAI-DALMAU[1], VASILIKI VOUKELATOU[2], ROSSANO SCHIFANELLA[1], STEFANIA FIANDRINO[1], DANIELA PAOLOTTI[1], AND KYRIAKI KALIMERI[1,*]

[1] *ISI Foundation, Turin, Italy*
[2] *United Nations World Food Programme, Rome, Italy*
[*] *kyriaki.kalimeri@isi.it*

1. INTRODUCTION

Food market accessibility is a critical yet underexplored dimension of food systems, particularly in low- and middle-income countries. Here, we present a continent-wide assessment of spatial food market accessibility in Africa, integrating open geospatial data from OpenStreetMap and the World Food Programme. We compare three complementary metrics: travel time to the nearest market, market availability within a 30-minute threshold, and an entropy-based measure of spatial distribution, to quantify accessibility across diverse settings. Our analysis reveals pronounced disparities: rural and economically disadvantaged populations face substantially higher travel times, limited market reach, and less spatial redundancy. These accessibility patterns align with socioeconomic stratification, as measured by the Relative Wealth Index, and moderately correlate with food insecurity levels, assessed using the Integrated Food Security Phase Classification. Overall, results suggest that access to food markets plays a relevant role in shaping food security outcomes and reflects broader geographic and economic inequalities. This framework provides a scalable, data-driven approach for identifying underserved regions and supporting equitable infrastructure planning and policy design across diverse African contexts.

2. DATA AND METHODS

Markets. 107,760 OSM food-related amenities (supermarkets, butcheries, open-air stalls) are harmonised with 3.380 WFP market centres via iterative k-means, yielding unique market centroids.

Friction surfaces. 30-arc-second rasters provide motorised and walking travel speeds that incorporate roads, topography, waterways, and administrative borders. We evaluate which are the most adequate metrics to capture spatial market accessibility in data-restricted countries, and the impact of diverse accessibility definitions on the final insights. We shed light on market accessibility differences between urban and rural areas as well as the areas of diverse socioeconomic status. Finally, we explore the relation between spatial accessibility to food markets and food security in Africa.

Metrics. Shortest path routing (using `scikit-image`) produces two continental travel-time stacks (motorised, walking). Population weights come from WorldPop 2020 (https://www.worldpop.org); rurality from the URCA four-class scheme [1]; economic status from Meta's 1-km Relative Wealth Index [2]; food insecurity from five-year IPC averages(https://www.ipcinfo.org/).

We used a set of geospatial metrics validated in the literature [3–5], namely, (i) **Closeness**— travel time to the nearest market, (ii)**Availability**—count of markets reachable within a 30-min motorised window, (iii) **Spatial redundancy**— entropy of an exponential choice model ($\beta = 0.077, \text{min}^{-1}$) truncated at four hours. Larger entropy is associated with a better spatial distribution of the surrounding markets. All metrics have been computed at the 30 arc-second level .

3. RESULTS

Key findings of this study showed a clear urban advantage. 360M urban residents reach a median of 54 markets in <30min; 250M rural+hinterland residents often reach one market after 1–4h. Further, there is an economic stratification within each country related to the market accessibility, with RWI–access correlations are sizeable (median $|\rho| \approx 0.55$). The bottom national quintile consists of 7-h walk to 6 markets, while the top quintile, 1-h walk to 70 markets. Food insecurity is also related to market access, with administrative units in IPC Phase 3 (crisis) showing median motorised times of 1.6h and entropy <1.0 versus 0.7h and entropy 1.7 in Phase 1. Finally, all three metrics inter-correlate strongly ($R! > 0.8$ at country median level) yet capture distinct facets: travel-time emphasises remoteness, cumulative counts emphasise option breadth, entropy emphasises resilience to single-market failure (see Fig. 2 for the resulting maps).

4. CONCLUSIONS

Our work can be used as an analytical assessment tool for population vulnerability assessments. The main policy implications that emerge from this study highlight underserved corridors, high-population and low-entropy (e.g., northern Nigeria), that should be prioritised for improved market availability or better spatial access. A major limitation of this study lies at the data gaps: 12 countries lack any WFP markets; friction surfaces date to 2020; ambulant markets, price dispersion, and facility size are unobserved. Consequently, absolute values may understate true access in sparsely mapped regions. Still, relative patterns are robust.

Open data and reproducible GIS pipelines make it possible to quantify continental food-market access with granular social resolution. The resulting maps expose where geography, poverty, and food insecurity collide—and furnish a basis for equitable, evidence-driven policy.

REFERENCES

1. A. Cattaneo, A. Nelson, and T. McMenomy, "Urban-rural continuum," Figshare (2020).
2. G. Chi, H. Fang, S. Chatterjee, and J. E. Blumenstock, "Microestimates of wealth for all low- and middle-income countries," Proc. National Acad. Sci. **119** (2022).
3. D. J. Weiss, A. Nelson, H. S. Gibson, *et al.*, "A global map of travel time to cities to assess inequalities in accessibility in 2015," Nat. 2018 553:7688 **553**, 333 (2018).
4. W. G. Hansen, "How accessibility shapes land use," J. Am. Inst. planners **25**, 73 (1959).
5. S. Erlander, "Accessibility, entropy and the distribution and assignment of traffic," Transp. Res. **11**, 149 (1977).

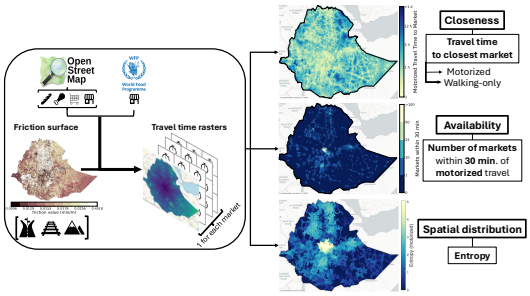

Fig. 1. Pipeline for the computation of the accessibility metrics. This figure illustrates an example from Ethiopia. First, food shops and markets are extracted from OSM and the WFP's market price and Market Functionality Index data. Market centroids are computed to harmonize both datasets. Using either a walking-only or motorized friction surface, we apply a shortest-path routing algorithm to generate a travel time raster for each market. These rasters are then used to derive accessibility metrics: travel time to the nearest market, the number of markets within a given travel time threshold, and entropy.

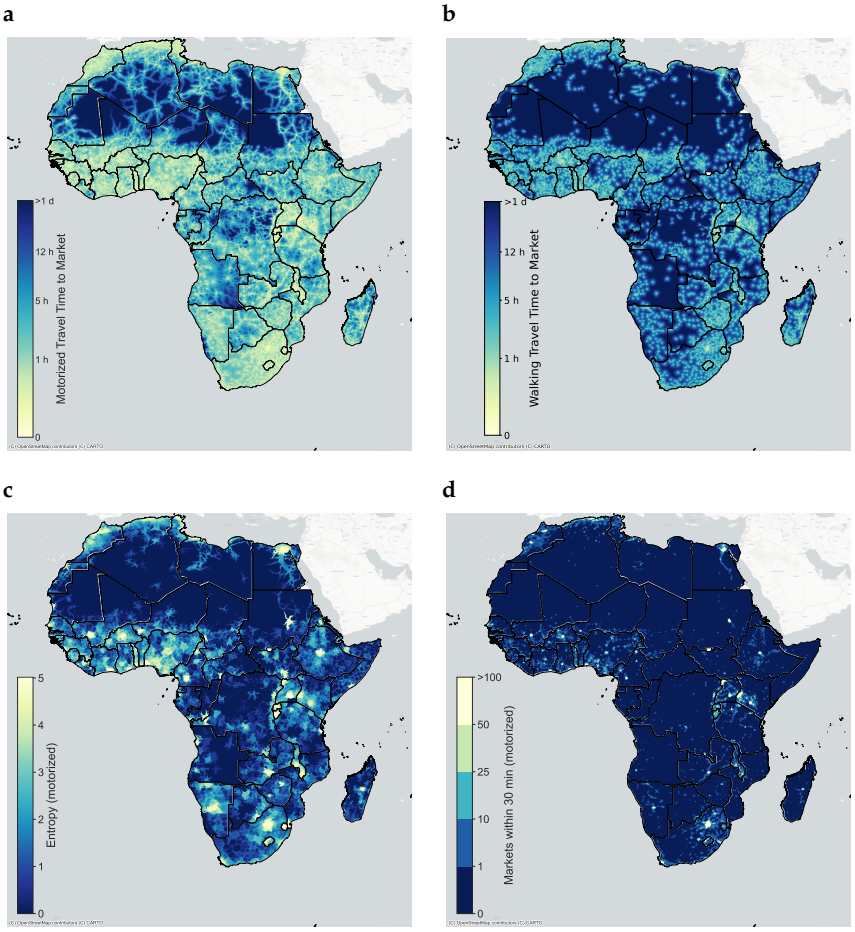

Fig. 2. Spatial accessibility to food markets in Africa. (a) Motorized and **(b)** walking travel time to the closest market show widespread disparities in accessibility across the continent, with rural and remote areas facing travel times of several hours or more.**(c)** Entropy illustrates the evenness of market distribution; higher entropy suggests a more balanced spread of accessible markets, typically concentrated in capital cities and urban centers.**(d)** Availability (cumulative accessibility) reflects the number of markets reachable within 30 minutes of motorized travel, favoring urbanized regions.

An Urban Geography of Mobile Traffic Demands

SACHIT MISHRA[1,2], DIEGO MADARIAGA[1], CEZARY ZIEMLICKI[3], DIALA NABOULSI[4], AND MARCO FIORE[1]

[1] *IMDEA Networks Institute, Spain*
[2] *Universidad Carlos III de Madrid, Spain*
[3] *SENSE / Orange Innovation, France*
[4] *École de Technologie Supérieure, Canada*
[*] *{sachit.mishra, diego.madariaga, marco.fiore}@imdea.org, cezary.ziemlicki@orange.com, diala.naboulsi@etsmtl.ca*

1. INTRODUCTION

In today's world, people rely on mobile applications to manage many aspects of their lives, including getting informed and forming opinions, organizing their daily personal and professional activities, moving with private and public transport, spending their free time, or even dating. This entanglement of mobile applications with real-world life is especially strong in dense urban areas, where individuals are known to make substantially higher use of such discussions [1].

Prompted by these considerations, we aim to investigate the existence of connections between the spatiotemporal consumption patterns of mobile applications and the underlying *urban fabrics*, *i.e.*, the entangled blend of physical infrastructure, human activities and socio-economic characteristics that combine to determine the functionality of an urban territory. Our aim is therefore that of (*i*) identifying recurrent utilization schemes of one or more mobile applications over space or time and (*ii*) linking such schemes to the urban fabric that explains their emergence. This approach enables us to establish causal relationships between the city's landscape and the complex mobile service demands it generates, ultimately revealing why people prefer specific mobile applications at different times and locations. The result has clear social implications [2] as well as important applications to networking, where it can, *e.g.*, drive a better urban-fabrics-informed planning of networks supporting application-specific operations such as slicing [3].

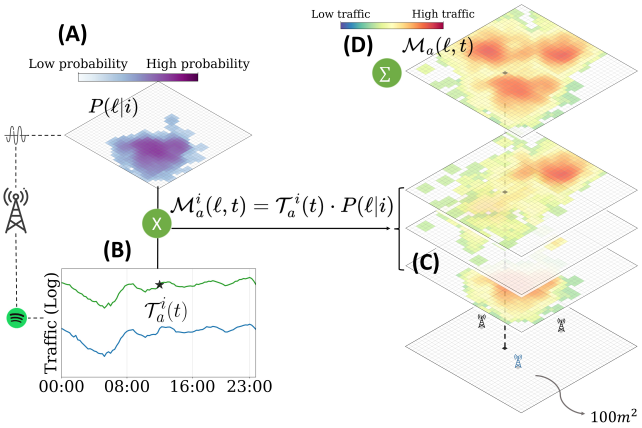

Fig. 1. Methodology for computing the traffic maps; (a) Coverage matrix for an base station i, (b) Traffic time series for Spotify, (c) Multiplication of the coverage matrix with the traffic time series, (d) Summation of all traffic maps.

2. DATA AND METHODS

We use hourly median mobile data traffic from 30 applications in 2019, aggregated at the eNodeB level across the 10 largest French cities. To map this traffic spatially, we employ eNodeB coverage data generated using a commercial RF signal propagation tool. Coverage is represented as a 600×600 matrix of 10,000 m^2 tiles, where each cell stores the probability $p(\ell \mid i)$ that user equipment (UE) in tile ℓ connects to base station i. Using this information, we follow the steps in Figure 1 to derive application-specific traffic distributions across IRIS zones, a geographical boundary defined by the French national statistics institute, each typically inhabited by 1,800 to 5,000 people. After following these steps, we obtain the hourly traffic demand for 30 different applications across various IRIS zones in the 10 cities.

To explore the underlying structure of mobile application usage, we rely on an Exploratory Factor Analysis (EFA) approach [4]. EFA assumes that the structure of a large set of observable variables can be modeled via a linear combination of unobservable common factors and error terms. We cast EFA so as to identify the *common factors* that explain specific spatial and temporal patterns in the mobile usage dynamics at service level. In our case, EFA is employed to summarize the traffic patterns of mobile applications observed at a large number of locations. Thus, the input for EFA is a set of zones (*i.e.*, samples), each characterized by different traffic features (*i.e.*, variables) that describe the consumption of mobile services. The objective is to identify a limited set of common factors that explain the largest possible portion of the total variance in the traffic demand data across targeted zones of the study.

We proceed by engineering the traffic features fed into the EFA model by constructing a concatenation of the following components: 1) The mean mobile traffic demand during weekdays and weekends for 30 applications, which captures their temporal usage patterns and highlights differences in weekday-weekend trends. 2) The weekday and weekend Symmetric Revealed Comparative Advantage (sRCA) values, serving as a signature to indicate the relative importance of each application compared to others during specific hours of the day in different IRIS zones.

3. RESULTS

Applying EFA to the service-level demands yields 20 common factors, which we related to distinct Urban fabrics. Table 1 presents a summary of the main 8 factors that collectively account for 55.2% of the total variance in the data. The remaining factors are associated with a minor fraction of the variance, and focus on less prominent urban fabrics, such as suburban areas or large parks in the city outskirts. We next discuss representative examples from such factors.

The digital divide of wealth. EFA yields two key components: factor loadings (L^2) and factor scores. The loadings capture how strongly each factor is associated with the typical hourly traffic demand and SRCA of application a on weekdays and weekends, according to our signature definition,i.e., values closer to 1 or -1 indicate a stronger correlation with the application. Figure 2 (first row) illustrates example heatmaps for Factor 1: the top two rows correspond to the SRCA of app a on weekends and weekdays, respectively, while the bottom two rows depict the normalized traffic demand for the same periods. The factor scores quantify

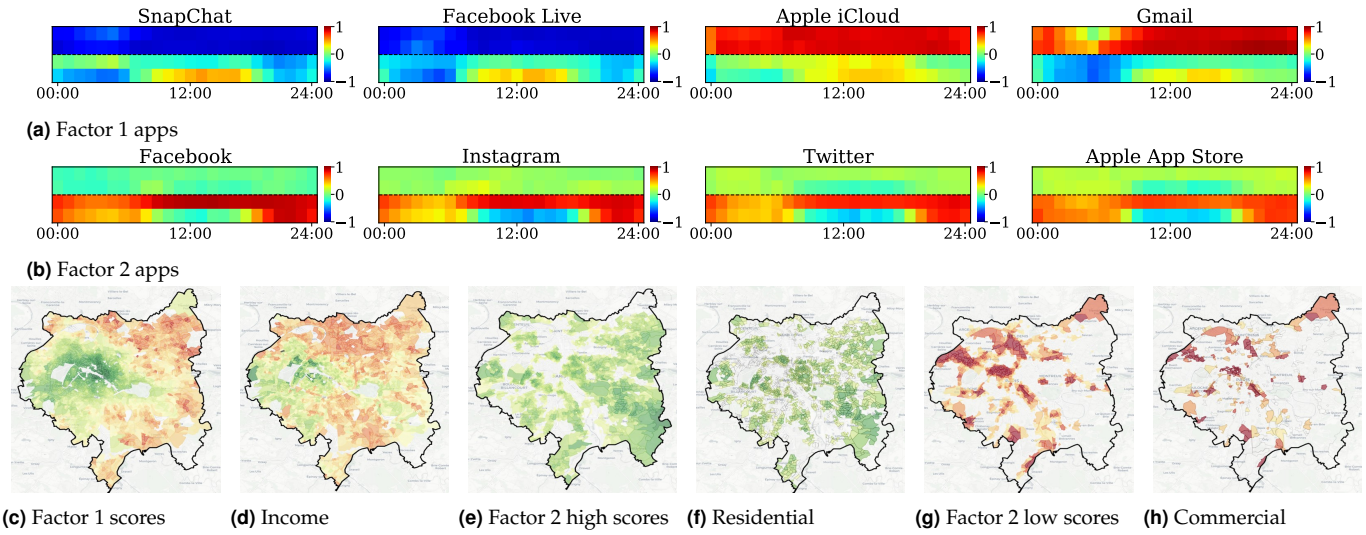

Fig. 2. Application-level contributions and geographical score distributions for Factor 1 and Factor 2.

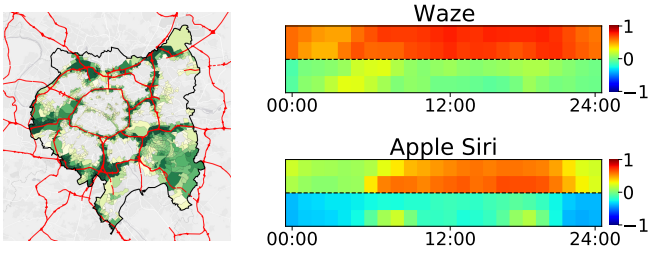

Fig. 3. Factor 8 scores in Paris, only high scores are shown. Factor 8 loadings for Waze and Siri: in each plot, SRCA (resp., traffic) loadings are top (resp., bottom) two rows.

the prevalence of each factor k across spatial zones. These are shown in the third row of Figure 2, revealing the geographic distribution of the factor throughout the studied cities.

The heatmap of L^2 for Factor 1 shows high loadings (in red) for productivity-focused apps like mail and cloud services, indicating elevated relative usage. In contrast, messaging and social media apps exhibit low SRCA loadings (in blue), reflecting below-average usage. These patterns remain consistent across weekdays and weekends. The normalized traffic loadings (bottom rows) are less distinct in the case of Factor 1. Figure 2c shows the factor 1 scores across Paris, dividing the city into two distinct regions. When compared with the income distribution in Figure 2d, the similarity is evident, with a high Pearson correlation of 0.7 between the scores and INSEE's average income data per IRIS zone. Similar strong correlations are observed in other cities as well.

The work-life dynamics of mobile traffic. Factor 2 L_a^2 values reflect a different dynamic than Factor 1. sRCA loadings of Factor 2 hover around zero, indicating no strong relative preference for any specific app. Instead, the defining pattern lies in normalized traffic, with a marked drop during weekday working hours and peaks in evenings and weekends. This behavior holds across all monitored applications, as shown in the second row of Figure 2.

Figure 2e shows Paris areas with high evening and weekend traffic, while Figure 2g highlights zones with higher weekday work-hour activity. These patterns align closely with residential and commercial/industrial areas shown in Figures 2f and 2h. By overlapping high and low factor 2 scores with land use maps, we compute F1 scores of 0.68 (residential) and 0.54 (commercial) for Paris. Similar results appear in other cities, confirming Factor 2's strong link to urban function.

Factor	Urban Fabric	Applications	Time
1	Income level	All	All
2	Residential areas	All	8 am – 6 pm
4	Budget nightlife zones	All	10 pm – 6 am
6	University areas	Instagram, Twitter/X	All
8	Main roadways	Waze, Siri	All
10	Train commuting zones	All	7–8 am, 6–7 pm
14	Airport regions	All	All
15	App-driven locations	Pokémon Go	All

Table 1. Summary of 8 prominent common factors identified by EFA in the traffic generated by 30 mobile applications across 5,097 IRIS zones in 10 French cities.

Connecting traffic on roads and networks. Factor 8 yields high L_a^2 for mobility-oriented applications such as Waze and Google Maps, as well as for Apple Siri. The loadings, exemplified in Figure 3. These services are used more than usual either throughout the day (e.g., Waze) or mainly during daylight hours (Siri). Figure 3 reveals the root cause of this factor: high-score IRIS zones closely follow the layout of the city's major highway networks.

4. CONCLUSIONS

Our results demonstrate the effectiveness of the approach to reveal a number of tangled patterns in space and time that are potentially associated to specific applications. Our work provides both a promising methodology for in-depth traffic characterization and novel insights on how urban fabrics control the demands for mobile services.

5. ACKNOWLEDGMENT

This work was supported by the CoCo5G (ANR-22-CE25-0016), funded by the French National Research Agency; and CHIST-ERA ECOMOME, funded by the Fonds de Recherche du Québec.

REFERENCES

1. S. Mishra, Z. Smoreda, and M. Fiore, "Second-level digital divide: A longitudinal study of mobile traffic consumption imbalance in france," in *The World Wide Web Conference*, (2022), pp. 2532–2540.
2. D. A. Rohlinger and S. Sobieraj, *The Oxford Handbook of Digital Media Sociology* (Oxford University Press, 2022).
3. K. Park, S. Sung, H. Kim, and J. il Jung, Comput. Networks **234**, 109908 (2023).
4. M. W. Watkins, J. black psychology **44**, 219 (2018).

Exploring the Digital Life of Urban Parks Through Mobile Traffic

ANDRÉ FELIPE ZANELLA[1,*], LINUS W. DIETZ[3], SANJA ŠĆEPANOVIĆ[2], KE ZHOU[2], ZBIGNIEW SMOREDA[4], AND DANIELE QUERCIA[2,3]

[1] *Telefónica Innovación Digital, Barcelona, Spain*
[2] *Nokia Bell Labs, Cambridge, UK*
[3] *King's College London, London, UK*
[4] *Orange Research, Châtillon, France*
[*] *andrefelipe.zanella@telefonica.com - Work done while at IMDEA Networks*

1. INTRODUCTION

Understanding how parks are used is crucial so urban planners can optimize the design and management of these areas. This is usually achieved by considering their form (e.g., the amount and type of greenery) and functions (e.g., do they promote specific sport activities, socializing, or nature-appreciations). With the widespread use of smartphones, mobile application traffic patterns, which act as a promising proxy for the presence and movement of populations in specific areas, have already proven useful in revealing detailed and nuanced usage patterns across different urban areas, offering a promising approach for understanding the patterns and types of urban park usage across different times and contexts.

Previous studies that explored urban green space utilization through mobile phone data either focused on datasets that showcase movement without mobile traffic usage—such as GPS [1, 2] and traffic signaling [3]—examined usage patterns without context, like CDR data [4], or analyzed traffic from a single application such as Strava [5]. Consequently, while these provided insights into park access, they did not delve into the specifics of mobile app usage within parks, leaving gaps in understanding beyond just the timing of visits. For this work, we leverage a large-scale app-level mobile traffic dataset (collected from Jan. 31 to May 31, 2023), selecting urban parks in Paris we're able to distinguish traffic consumed inside from outside of parks. Our results showcase how people not only interact with their smartphones differently in these spaces, but also how different types of parks lead to different smartphone behaviors.

2. RESULTS

People use different apps in parks compared to the the city. With a final set of 45 parks located in Paris and its suburbs that we're confident that we can isolate its traffic from the neighborhood, we explore and characterize how these parks influence smartphone utilization versus the rest of the city, as well as if there's uniformity of smartphone usage across parks. To compare mobile app usage within parks to that in other urban spaces, we select 41 relevant mobile apps, which together account for 80% of the non-encrypted traffic, and organize them into nine categories. Fig. 1A shows the application preference (represented by the Relative Symmetric Comparative Advantage (RSCA) metric) for these app categories in both parks and city areas, with $*$ representing categories that have a statistical difference between the distribution of app usage between parks and the city (obtained from a t-test across distributions).

Park-specific applications: Apps in the music and news categories stand out as being used significantly more in parks compared to other areas of the city. This difference in mean usage is statistically significant, as confirmed by a t-test (p-value < 0.005 for music, and 0.049 for news). While categories such as fitness, video, shopping, and travel also tend to be more prevalent in parks on average, their differences with the rest of the city are less pronounced and do not reach statistical significance.

City-specific applications: On the other hand, app categories like social media and games are used significantly more in the rest of the city compared to parks, as indicated by a notable difference in the mean of the distributions (p-values of 0.005 and 0.027, respectively). Social media apps have negative RSCA values in parks, while most of the city shows positive values, indicating under-utilization in parks. Gaming apps are also under-utilized in parks, but are used more evenly across the city.

Despite the differences across apps, the confidence intervals (represented by lines) for RSCA values in parks show fluctuations between under- and over-utilization across some app categories (e.g., games). This variability suggests that app usage in parks is not uniform, indicating that we should next to explore patterns across parks.

Most parks see higher weekday use. To explore temporal aspects of parks, we divide the total traffic volumes into weekdays and weekends. We found that approximately 75% of parks experience higher traffic on weekdays than on weekends (Fig. 1B). The variation in the traffic ratio, ranging from 0.72 to 1.86, is quite significant. The three parks with the *highest weekend traffic* are primarily oriented towards nature activities: Parc de la Fosse Maussoin, Parc de la Vallée-aux-Loups and Parc Georges Brassens, and Parc de la Vallée-aux-Loups. The first two are located in the Parisian suburbs, feature dense green spaces. Meanwhile, the last is located more centrally, offers a pond, meadows, tree groves, and various gardens, including a vineyard. All are situated mostly in residential areas.

Conversely, from the three parks with the *highest weekday traffic*, Jardins Abbé Pierre and Parc Monceau are more central (with many business and universities nearby), while Parc Suzanne Lenglen is located on the southern outskirts of Paris (before the Suburbs). The latter park is a multi-sport facility park, used for sports classes, close to the Ministry of Defense and the digital business district in Issy-les-Moulineaux — which may explain its higher weekday usage.

Uncovering three functional uses of parks. We used each park's per-application RSCA and weekday-to-weekend traffic ratio as features of a spectral cluster algorithm, finding 3 significant clusters through the Silhouette score. Besides an apparent spatial distribution across clusters (Fig. 1C), we note that the parks in each cluster show a variability in temporal traffic patterns (Fig. 1D) and differences in app preferences (Fig. 1E). Based on these patterns, we name clusters based on their functional uses: recreational, lunchbreak and cultural.

Recreational Parks (in blue) are located in the suburban areas of metropolitan Paris. According to the 2017 Census, 68% of the

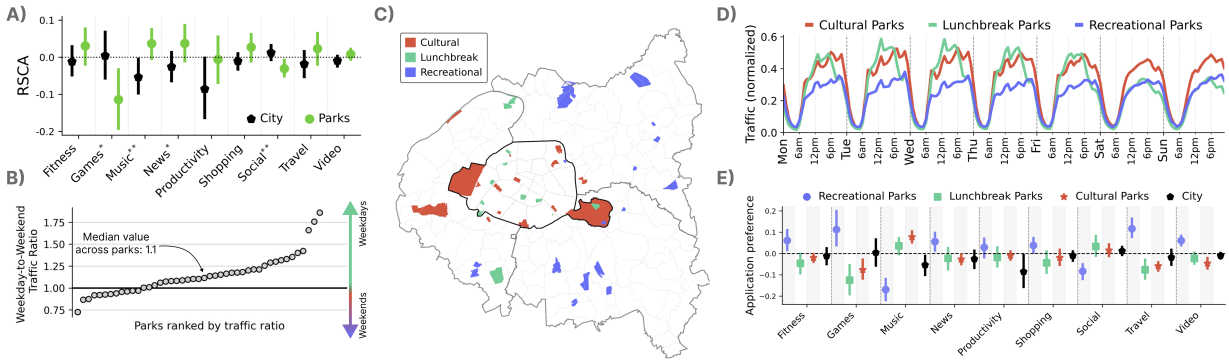

Fig. 1. (A) App preference between parks and city; (B) Weekday-to-weekend traffic ratio; (C) Park clusters across Paris; (D) Median week traffic and (E) App preference across clusters;

inhabitants of metropolitan Paris live in these less densely populated, mostly residential areas. Notably, parks in this category have the *lowest traffic consumption* compared to the those in the other two categories, using on average 55% less traffic.

Lunchbreak Parks (in green) are located in densely built areas near offices and universities. Their temporal traffic patterns are highly concentrated within work days, with the *median traffic being 58% higher on a workday compared to the weekend* (these are mostly to the right of Fig. 1B). These are comparatively smaller, with their area being only 36% of the median area of the set. Their traffic patterns, proximity to business areas, and the types of apps used on workdays suggest that office workers frequent these parks during their work breaks.

Cultural Parks (in red) are mainly known for attractions such as the Champ de Mars with the Eiffel Tower, Jardin des Tuileries and Parc de la Villette. Interestingly, those parks have an *even distribution of traffic across days*, with working days having only 7% more traffic than weekends. This could be explained by those parks being *popular throughout the week*, relating to the regions they are in and their tourist purpose in Paris.

How apps are used differently across park types. To understand how mobile app are used in each cluster, we calculate the application preference in each (Fig. 1E). *Cultural parks* experience a steady flow of visitors throughout the week, given their central locations and their role as cultural hubs. Social media and music streaming apps stood out with higher-than-average usage, likely driven by tourists who prioritize social engagement over apps related to shopping, games, productivity, or news while exploring the city. *Lunchbreak parks* follows a similar pattern, with social and music apps having a higher preference. This highlights how these two park types in the city center exhibit similar app usage patterns to the rest of the city (marked in black in Fig. 1E), indicating app consumption may not be influence by these parks. *Recreational parks* in contrast have distinct app usage patterns from all. Unlike cultural and lunchbreak parks, recreational parks exhibited higher usage across most app categories, except for music and social media. Notably, they were the only parks with increased usage in the fitness category, alongside higher engagement with news and travel apps.

These findings suggest that these spaces serve a unique role by offering settings for relaxation and everyday activities while promoting physical activity and reducing social media usage. When visitors used their phones, they tended to favor apps for reading and planning instead of social and music apps.

Acknowledgment: Supported by the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation program under Grant Agreement No 101192369 (6G-MIRAI).

REFERENCES

1. W. Liu, Q. Chen, Y. Li, and Z. Wu, "Application of gps tracking for understanding recreational flows within urban park," Urban For. & Urban Green. **63**, 127211 (2021).
2. Y. Zhai, P. K. Baran, and C. Wu, "Spatial distributions and use patterns of user groups in urban forest parks: An examination utilizing gps tracker," Urban forestry & urban greening **35**, 32 (2018).
3. Y. Zhai, H. Wu, H. Fan, and D. Wang, "Using mobile signaling data to exam urban park service radius in shanghai: methods and limitations," Comput. Environ. Urban Syst. **71**, 27 (2018).
4. C. Guan, J. Song, M. Keith, *et al.*, "Delineating urban park catchment areas using mobile phone data: A case study of tokyo," Comput. Environ. Urban Syst. **81**, 101474 (2020).
5. V. Heikinheimo, H. Tenkanen, C. Bergroth, *et al.*, "Understanding the use of urban green spaces from user-generated geographic information," Landsc. Urban Plan. **201**, 103845 (2020).

Spatial Heterogeneity in Human Mobility Responses to London's Ultra-Low Emission Zone Expansion

YIKANG WANG[1,*] AND CHEN ZHONG[1]

[1] Centre for Advanced Spatial Analysis, University College London, London, UK
[*] yikang.wang.21@ucl.ac.uk

1. INTRODUCTION

Cities worldwide are implementing sustainable transport policies like Low Emission Zones (LEZs) and Ultra Low Emission Zones (ULEZs) to improve air quality and promote sustainable transport. While the impact of London's ULEZ on air quality is well-documented [1], its effects on human mobility patterns and socio-spatial inequalities remain less understood. Most studies examine correlations without establishing causality or lack a holistic view of diverse transport modes and populations [2].

This study addresses this gap by quantifying the causal impact of the first ULEZ expansion on human mobility. Using large-scale mobile phone location data and an extended Interrupted Time Series (ITS) model, we investigate two key questions: 1) To what extent did the ULEZ expansion impact human mobility? 2) How did this impact vary across different geographic areas and socio-demographic groups? Our findings offer crucial insights for designing equitable and sustainable urban transport policies.

2. RESEARCH AREA AND DATA

Launched in April 2019, the ULEZ initially covered Central London. It expanded to Inner London on October 25, 2021, and again to the Greater London Area (GLA) in August 2023 (Fig. 1). This phased approach provides a unique opportunity to examine the evolving impact of the ULEZ on mobility across different geographic zones and over time.

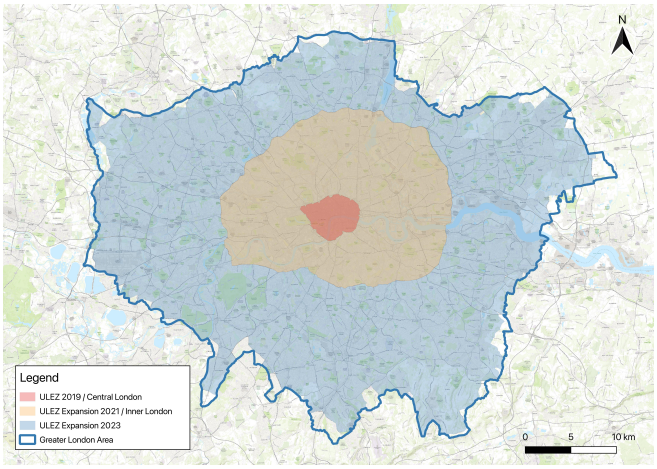

Fig. 1. The phased expansion of London's ULEZ.

Our primary data source is anonymised mobile phone location data from approximately 1.35% of London's population over 73 days (1 October - 12 December 2021), covering the period before and after the ULEZ expansion on 25 October 2021. From this, we extracted trips and activities, inferring home/work locations to derive mobility indicators at the Middle Layer Super Output Area (MSOA) level [3]. These indicators include (1) the number of activities, (2) travel distance, and (3) travel duration. We also incorporated MSOA-level 2021 Census data (e.g., car ownership, education levels) and Public Transport Accessibility Level (PTAL) data to control for socio-demographic and environmental factors.

3. METHODOLOGY

Given the absence of a suitable control group for a city-wide policy in London, we employ an Interrupted Time Series (ITS) model. ITS is a robust quasi-experimental method that uses pre-intervention trends to establish a counterfactual scenario, thereby isolating the policy's causal effect [4]. The basic ITS model is specified as:

$$Y_t = \beta_0 + \beta_1 \text{Time1}_t + \beta_2 \text{Intervention}_t + \beta_3 \text{Time2}_t + \epsilon_t \quad \textbf{(1)}$$

where Y_t is the mobility indicator at time t. Intervention_t is a dummy variable for the policy's implementation (0 before, 1 after), capturing the immediate effect (β_2). Time1_t and Time2_t are number of days since the beginning of observation and the intervention, capturing the baseline trend (β_1) and the slope change after the intervention (gradual effect, β_3), respectively.

To explore spatial heterogeneity, we extend the model by incorporating MSOA-level geographic and socio-demographic variables and their interactions with the policy intervention [5, 6]:

$$\begin{aligned} Y_{it} = {} & \beta_0 + \beta_1 \text{Time1}_t + \beta_2 \text{Intervention}_t + \beta_3 \text{Time2}_t \\ & + \sum \beta_{4+m} \text{DayOfWeek}_t^m + \beta_{10} \text{PublicTransport}_i \\ & + \sum \beta_{10+m} \text{Census}_i^m \\ & + \text{Intervention}_t (\beta_{17} \text{PublicTransport}_i + \sum \beta_{17+m} \text{Census}_i^m) \\ & + \text{Time2}_t (\beta_{24} \text{PublicTransport}_i + \sum \beta_{24+m} \text{Census}_i^m) + \epsilon_{it} \end{aligned} \quad \textbf{(2)}$$

where Y_{it} is the outcome for MSOA i at time t. DayOfWeek_t^m are dummy variables controlling for weekly seasonality. PublicTransport_i is the MSOA's public transport accessibility level (PTAL). Census_i^m represents MSOA-level variables including residence density, gender, education, car ownership, driving to work, and share of residents working from home. The interaction terms quantify how neighbourhood context (e.g., census and public transport access) moderates the abrupt and gradual effects of the ULEZ expansion.

4. RESULTS

A. Aggregate and Temporal Effects

The ULEZ expansion prompted an immediate behavioural response. Across the GLA, daily trips dropped by 8.15% on the day of expansion, recovering to pre-expansion levels within two weeks (Fig. 2).

However, recovery patterns differed. As shown in Table 1, cross-boundary travel (e.g., between Inner and Outer London) showed no recovery in activity numbers. Across all zones, travel duration increased more significantly and recovered more slowly than travel distance, suggesting a modal shift to slower, more sustainable options.

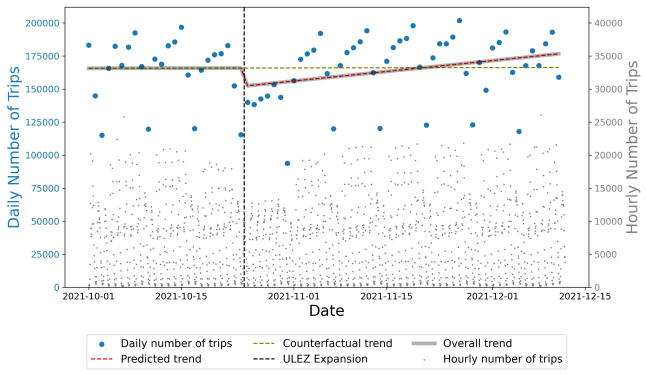

Fig. 2. Trend in the daily number of trips in the Greater London Area, showing the abrupt drop at the ULEZ expansion and subsequent recovery.

Table 1. Abrupt and gradual effects on travel behaviours between origins and destinations (OD).

OD	Indicator	Abrupt	Gradual Effect
GLA ↔ GLA	Activities	-8.15% **	Recovered in 2 wks
	Duration	+5.49% ***	Slow recovery
	Distance	+3.01% *	Recovered in 1 wk
Inner London ↔	Activities	-9.23% ***	Slow recovery
Inner London	Duration	+6.80% ***	Slow recovery
	Distance	+3.76% ***	Recovered in 3 wks
Inner London ↔	Activities	-5.53% **	No recovery
Outer London	Duration	+3.63% ***	Slow recovery
(cross-boundary)	Distance	+1.85% *	Recovered in 1 wk, then decrease
Central London ↔	Activities	-4.05%	No recovery
Inner London Ring	Duration	+5.69% ***	Slow recovery
(cross-boundary)	Distance	+2.04% **	Slow recovery

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Recovery categories are defined by the gradual effect's (β_3) statistical significance and speed. "Recovered in X wks" indicates a return to pre-expansion levels within that timeframe. "Slow recovery" signifies a statistically significant trend toward recovery that did not reach pre-expansion levels within our 5-week post-intervention study period. "No recovery" indicates no statistically significant recovery trend was observed.

B. Spatial and Socio-demographic Heterogeneity

The policy's impact was not uniform across London. The most significant changes occurred in Outer London, an area outside the expanded zone but highly car-dependent. These areas experienced the largest abrupt increases in commuting duration and the most pronounced decreases in casual (non-work) activities, highlighting a strong spatial spillover effect (Fig. 3).

Socio-demographic factors were key moderators. MSOAs with higher car ownership experienced a more pronounced decrease in casual activities and a larger increase in commuting duration, suggesting these residents were forced to cut non-essential trips. Conversely, areas with better public transport accessibility (higher PTAL) were more resilient, showing a smaller reduction in activities. This indicates the ULEZ expansion disproportionately burdened residents in car-dependent areas with limited access to sustainable transport alternatives.

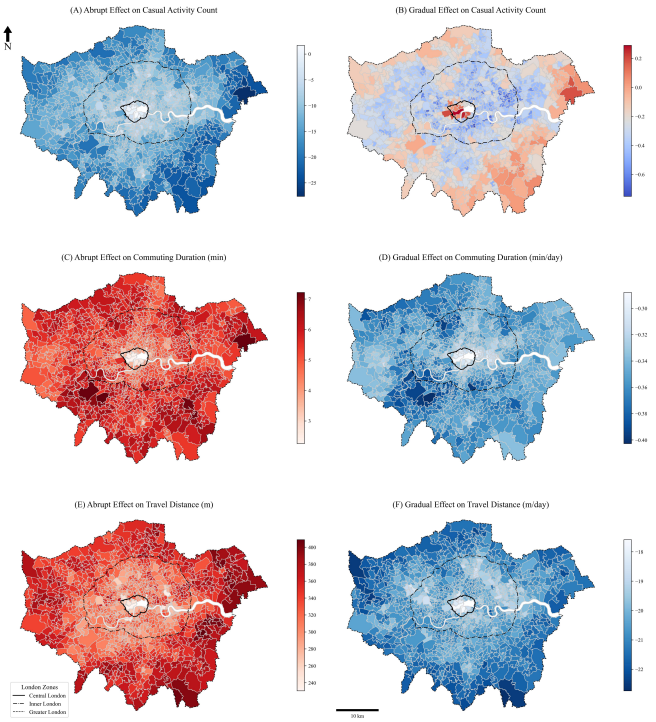

Fig. 3. Spatial heterogeneity of the ULEZ expansion's abrupt and gradual effects on casual activity counts, commuting duration, and travel distance.

5. CONCLUSION

This study provides a nuanced, causal analysis of the 2021 ULEZ expansion's impact on human mobility. Our findings show that the policy successfully induced a shift towards more sustainable (and slower) transport modes but that its effects were highly heterogeneous. The significant impacts on car-dependent Outer London and the lack of recovery for cross-boundary trips highlight potential disruptions to established activity patterns and the risk of deepening socio-spatial inequalities. These results underscore the critical need for policymakers to implement complementary measures, such as improving public transport in underserved areas, to ensure a just and equitable transition to sustainable urban mobility.

REFERENCES

1. H. Hajmohammadi and B. Heydecker, "Evaluation of air quality effects of the London ultra-low emission zone by state-space modelling," Atmos. Pollut. Res. **13**, 101514 (2022).
2. Greater London Authority, "Inner London ultra low emission zone expansion one year report," (2023).
3. C. Zhong, N. S. Aslam, Y. Wang, *et al.*, "Anonymised human location data for urban mobility research," CASA Working Paper 240 (2024).
4. J. L. Bernal, S. Cummins, and A. Gasparrini, "Interrupted time series regression for the evaluation of public health interventions: a tutorial," Int. J. Epidemiol. **46**, 348–355 (2017).
5. A. Freni-Sterrantino, R. Ghosh, D. Fecht, *et al.*, "Bayesian spatial modelling for quasi-experimental designs," Environ. Int. **128**, 109–115 (2019).
6. W. Zhang and K. Ning, "Spatiotemporal heterogeneities in the causal effects of mobility intervention policies during the COVID-19 outbreak," Ann. Am. Assoc. Geogr. **113**, 1112–1134 (2023).

The temporal dimension of experienced segregation in cities

JOSHUA ROSEN[1,2,8], ERIK WEIS[3], HAMISH GIBBS[3], SARAH A. SANCHEZ[1,2,8], TAKAHIRO YABE[4,5,6], DANIEL T. O'BRIEN[1,2], BRENNAN KLEIN[3,7], AND ESTEBAN MORO[3,7,8*]

[1] School of Public Policy & Urban Affairs, Northeastern University, Boston, Massachusetts, USA
[2] Boston Area Research Initiative, Boston, Massachusetts, USA
[3] Network Science Institute, Northeastern University, Boston, Massachusetts, USA
[4] Institute for Data, Systems, and Society, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA
[5] Center for Urban Science and Progress (CUSP), Tandon School of Engineering, New York University, Brooklyn, New York, USA
[6] Department of Technology Management and Innovation, Tandon School of Engineering, New York University, Brooklyn, New York, USA
[7] Department of Physics, Northeastern University, Boston, Massachusetts, USA
[8] Social Urban Networks Lab, Northeastern University, Boston, Massachusetts, USA
[*] e.moroegido@northeastern.edu

Cities have been traditionally viewed as a great equalizer of economic prosperity, the core mechanism of which is the interactions between people of different demographic characteristics [1–3]. However, a long history of work shows that segregation in cities undercuts this possibility [4]. Recent work has sought to demonstrate the (in)equality in cities using high-resolution mobility data, finding that experiences of segregation extend beyond residential configurations of cities to the everyday interactions between individuals [5]. However, these comprehensive studies overlook another dimension recognized to be an important aspect of experience inequality: time [6, 7].

Our study builds on this foundation by empirically measuring space-time inequality using high-resolution mobility data. We leverage anonymized mobility data from Spectus to quantify how inequality manifests dynamically across cities, within cities, and between points of interest (POIs). Unlike static segregation metrics, our approach captures fluctuations in socioeconomic integration at different times of day. For instance, POIs that appear integrated at one moment may become highly segregated at another due to shifting activity patterns. The COVID-19 pandemic highlighted the importance of these temporal dynamics, as changes in work schedules and public space usage fundamentally altered mobility-driven inequality [8]. Moreover, recent research suggests that large metropolitan areas, despite their diversity, often exhibit higher exposure segregation due to differentiated space usage across socioeconomic groups [9]. Our work extends these insights by integrating both spatial and temporal dimensions to offer a more comprehensive framework for understanding mobility-driven inequality.

Across 11 major US CBSAs, we systematically map the structure of space-time inequality over a six-month period, capturing both weekdays and weekends—in doing so, we document a consistent shape of urban inequality throughout a day. In particular, our results confirm the existence of substantial variation in space-time inequality, beyond typical residential segregation (Fig. 1c). This variation exists on three geographic scales: (1) between cities, (2) within cities, and (3) across POI types. We find that across all hours of the day, low-income individuals experience the greatest level of segregation compared to other income groups. Furthermore, we find that the POI categories that experience the most change in daily segregation are amenity spaces being used by different demographic groups at different times—thus revealing how temporal sorting contributes to the patterns described above.

Furthermore, we find that temporal sorting in mobility itself may both reinforce and disrupt segregation. To understand why some individuals are more likely to encounter integrated environments, we introduce a concept we call *time-richness*—the degree to which a person has both the control over their schedule and the capacity to use it meaningfully. We derive this from mobility patterns by measuring two behavioral dimensions: *predictability*, which reflects how structured or constrained someone's daily routine is, and *exploration*, which reflects how widely and diversely they move through the city. Here, we measure predictability using permutation entropy of hourly visitation sequences, capturing the regularity of an individual's mobility patterns, and exploration as the fraction of newly visited places within a rolling 14-day window. We operationalize our measure of time-richness using the first principal component between an individual's normalized exploration and predictability. We show that this time-richness is largely uncorrelated with income, and is a stronger predictor of experienced segregation than income: people with more temporal autonomy are better able to access integrated spaces, thereby "breaking the cycle" of segregation that structures many urban routines. These findings position time, not just money, as a critical axis of inequality that shapes who we encounter and where.

We expand on the analysis of time-richness by presenting several stylized facts regarding this newly measured group. We find that people who are *time-rich* tend to live further from the urban core and, like the income-rich, travel greater distances on average to reach amenities. We also document hourly variation in amenity visitation likelihoods between the time-rich but low income, and time poor but high income—thereby again suggesting that time poses as a critical axis for how we understand experienced segregation. Finally, we exploit the shift to Daylight Saving Time as a natural experiment, offering initial causal evidence that positive shocks to time reduce experienced segregation.

In documenting the temporal dimension of experienced segregation in cities, our analyses show that urban segregation depends not only on where people live, work, and gather, but on *when* they do so. In other words, the periodic rhythms of segregation we observe reflect a pattern of temporal sorting—one that cannot be understood by examining income alone. This temporal inequality emerges at the intersection of residence, employment, and leisure, offering fresh insights into how urban form and zoning shape both spatial and temporal access to resources. A more nuanced understanding of these dynamics suggests that policy and design interventions—such as flexible work and transit schedules, staggered public programming, and multi-use community venues—can actively counter the cyclic peaks of segregation by redistributing temporal as well as spatial opportunities. In sum, by revealing the typical shape of space-time

inequality and identifying time-richness as its driving mechanism, our work calls for a deeper commitment to crafting urban environments that allocate hours and places equitably, so that everyone can fully participate in the life of the city.

REFERENCES

1. J. Jacobs, *Cities and the Wealth of Nations: Principles of Economic Life* (Random House, 1984).
2. J. Jacobs, *The Economy of Cities* (Knopf Doubleday Publishing Group, 1969). Business & Economics, 288 pages.
3. R. Florida, City & Community **2**, 3 (2003).
4. D. S. Massey and N. A. Denton, *American Apartheid: Segregation and the Making of the Underclass* (Harvard University Press, Cambridge, MA, 1993).
5. E. Moro, D. Calacci, X. Dong, and A. Pentland, Nat. Commun. **12**, 4633 (2021).
6. D. Massey, *Politics and Space/Time* (University of Minnesota Press, 1994).
7. V. M. Netto, J. V. Meirelles, M. Pinheiro, and H. Lorea, Cybergeo: Eur. J. Geogr. (2018).
8. C. Santana, F. Botta, H. Barbosa, *et al.*, Nat. Hum. Behav. **7**, 1729 (2023).
9. H. Nilforoshan, W. Looi, E. Pierson, *et al.*, Nature **624**, 586 (2023).

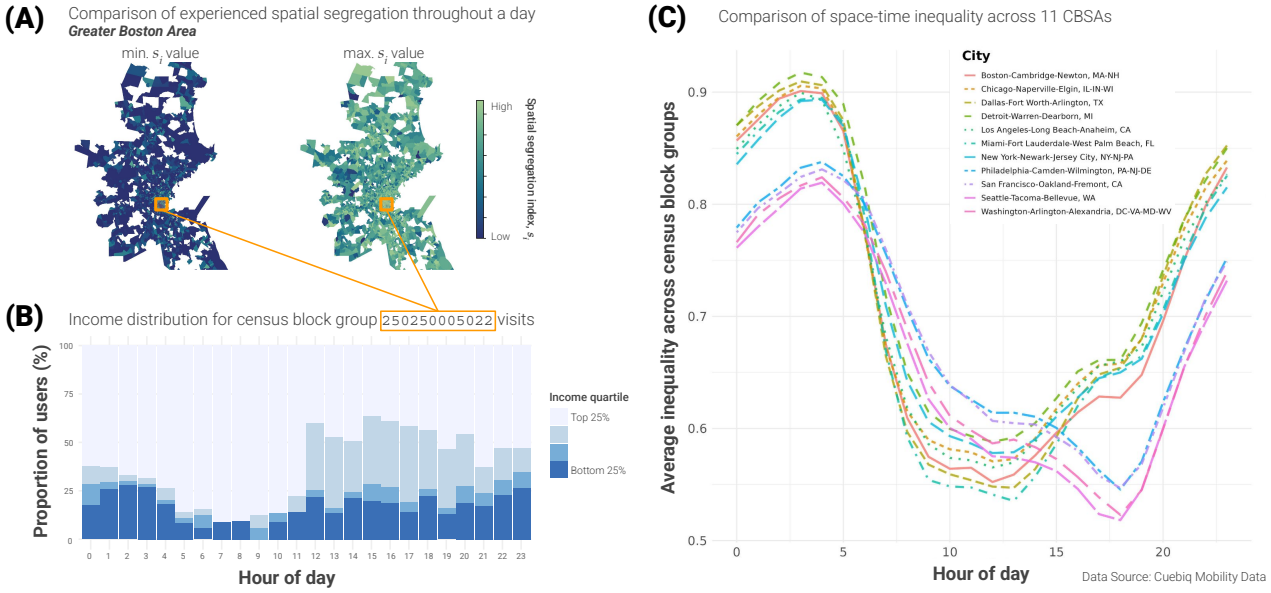

Figure 1. (A) For the Boston CBSA, each census block group (CBG) is shaded to indicate its maximum (lighter) and minimum (darker) experienced inequality across the day. **(B)** The diurnal cycle of population composition (by four income groups) for a single CBG, illustrating how income mix varies hourly. **(C)** The average hourly experienced inequality for each CBG, aggregated across 11 CBSAs, highlighting the broader shape in experienced inequality over the course of a day.

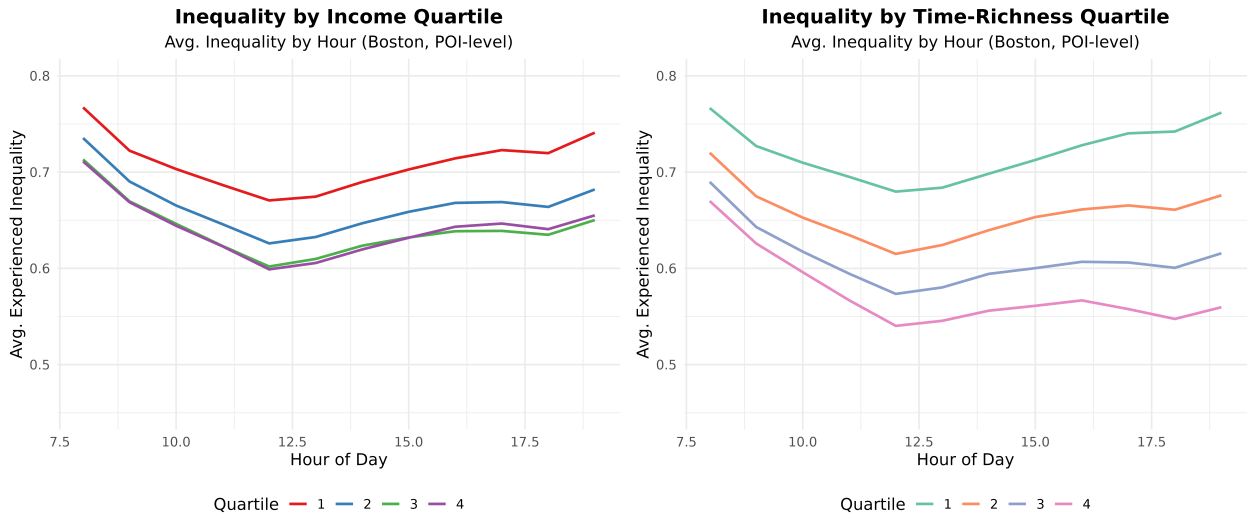

Figure 2. Average experienced segregation from 8 AM to 8 PM in the Boston, MA CBSA by quartile. *Left:* Income quartile (Group 1 = lowest income; Group 4 = highest income). *Right:* Time-richness quartile (Group 1 = most time-poor; Group 4 = most time-rich).

Mobility behavior moderates residential partisan segregation in US metropolitan areas

MARCO TONIN[*1,2], MICHELE TIZZONI[†1], BRUNO LEPRI[‡2], AND ESTEBAN MORO[§3]

[1] Department of Sociology and Social Research, University of Trento, Via Verdi 26, I-38122, Trento, Italy
[2] Mobile and Social Computing Lab (MobS), Fondazione Bruno Kessler (FBK), Via Sommarive 18, I-38123, Trento, Italy
[3] Network Science Institute, Department of Physics, Northeastern University, Boston, 02115, MA, USA

1. INTRODUCTION

The American electorate today is more politically polarized than in recent decades, with distinct opinions on various important issues. This polarization is reflected in the political homogeneity within social networks, which is associated with several political outcomes, including voter turnout, political preferences, and affective polarization. Beyond ideological or opinion-based differences, people are also segregated in daily life, with limited exposure to those with opposing political views. The social and political divide is primarily associated with residential patterns, as Democrats tend to cluster in metropolitan areas, while Republicans are more segregated in rural communities [1], a trend increasingly driven by generational change in urban contexts and ideological conformity in rural areas [2]. However, similar to the dynamics of income segregation [3], individuals' mobility behavior can shape their experienced political segregation even within urban or rural areas, either increasing or reducing the diversity of social encounters. Such diversity is influenced both by individuals' lifestyles [4] and by the characteristics of the places they visit [5].

From this perspective, this study aims to explain experienced partisan segregation by examining its relationship with mobility behavior. First, we compare between- and within-area variation in experienced partisan segregation. Second, we assess the moderating role of mobility behavior in the relationship between residential and experienced partisan segregation.

2. MATERIALS AND METHODS

Data sources

Our research integrates high-resolution anonymized individual mobility data provided by Cuebiq, collected between October 2016 and April 2017, with presidential vote results at the precinct level from the 2016 election, as well as socioeconomic data from the Census at the census block group (CBG) level. Specifically, data includes 545,741 individuals in 11 major US Metropolitan Statistical Areas (MSAs) (i.e., Boston, Chicago, Dallas, Detroit, Los Angeles, Miami, New York, Philadelphia, San Francisco, Seattle, Washington DC).

Individual and place segregation

First, we compute the political composition of visitors for each point of interest (POI) based on the time individuals spent at each location and their inferred political affiliation. Political affiliation is inferred from the results of the 2016 presidential election in the precinct where each individual resides, defined as the difference between the share of votes for Republicans and Democrats. This measure ranges from -1 (if the entire precinct voted for Democrats) to 1 (if the entire precinct voted for Republicans).

[*] marco.tonin@unitn.it
[†] michele.tizzoni@unitn.it
[‡] lepri@fbk.eu
[§] e.moroegido@northeastern.edu

Second, we calculate each user's experienced segregation based on the political composition of the POIs they visit. Similar to political affiliation, both place- and individual-level experienced segregation range from –1 (Democratic segregation) to 1 (Republican segregation). Finally, we compute the relative experienced and residential partisan segregation taking the absolute value of the difference between each segregation measure and the median value within the corresponding MSA.

3. RESULTS

Between and within-area variation in experienced political segregation

People living in the same area can experience different levels of political segregation in their daily lives. Although contextual factors such as county and neighborhood account for a substantial share of the variation, they do not fully determine individuals' exposure to political groups. In this context, we model the relationship between experienced partisan segregation and individuals' characteristics and behaviors using multilevel regression analyses with random intercepts with two nested geographic levels (County and Census Tract). Finally, we compute Variance Partition Coefficients (VPCs) to assess how much of the variance in segregation is attributable to differences between geographic areas and how much remains within them at the individual level.

Our findings highlight that the within-area variance in experienced partisan segregation not explained by county and tract differences ranges from 10.9% in Washington, DC to 54% in Miami (Fig. 1a). This means that people who live in the same area can be either segregated or exposed to the political outgroup. The remaining variance is mainly explained by differences between counties, while only a small portion is explained by differences between tracts.

Mobility patterns moderate partisan segregation

Mobility patterns moderate residential partisan segregation and can explain part of the within-area variance. To understand how human behavior and mobility impact experienced segregation, we model the relationship between relative experienced segregation and interaction terms combining relative residential sorting with mobility patterns. Specifically, we characterize mobility behavior using three different metrics, which are the exploration rate, defined as the tendency to visit new places, the category diversity, defined as the variety of activity categories visited, and the average distance from home, which measures the average distance traveled to reach an activity space. We also control for socioeconomic characteristics, including ethnicity, educational attainment, and income.

Overall, as shown in Fig. 1c, including mobility patterns improves model performance by explaining part of the variance in relative experienced segregation. Mobility is particularly predictive in Boston (+11.2%), Miami (+11.9%), Los Angeles (+13.2%), and Seattle (+14.4%).

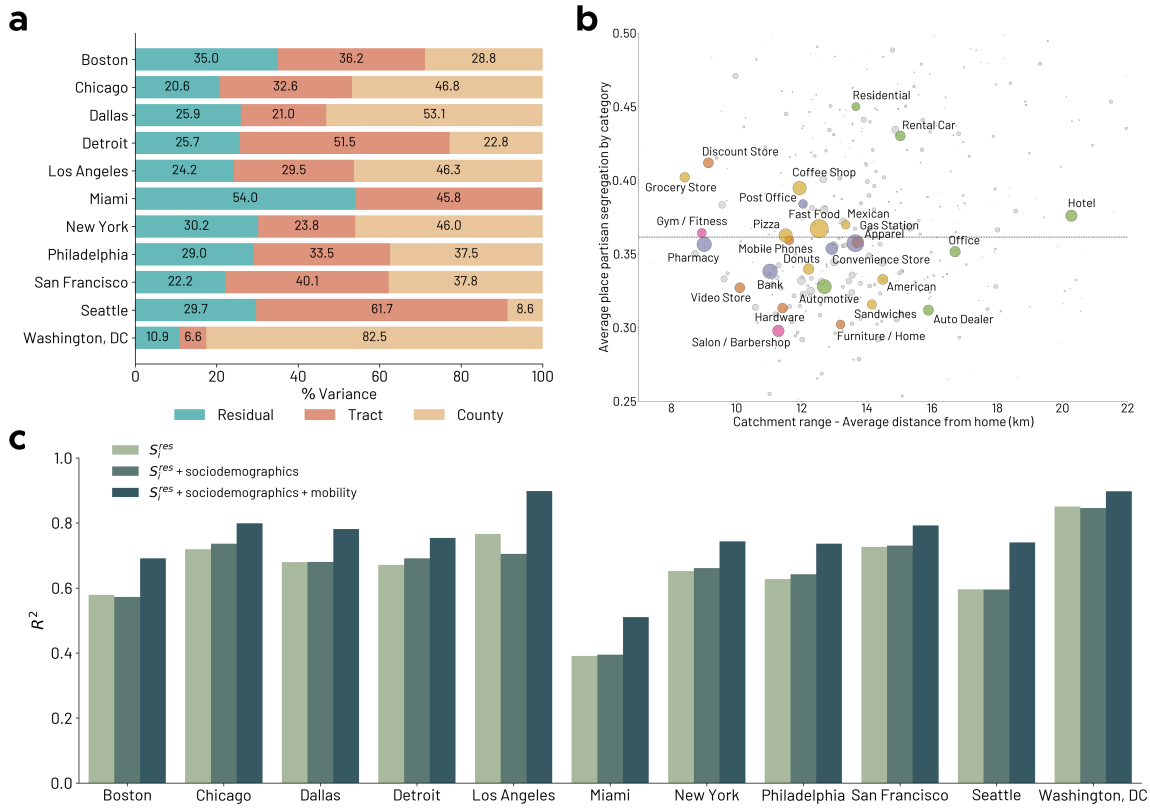

Fig. 1. (a) Variance Partition Coefficients (VPCs) from the multilevel regression models show that substantial share of experienced partisan segregation remains within areas, even after accounting for differences between Counties and Tracts; **(b)** Considering the points of interest of all 11 MSAs, place partisan segregation by category persists with catchment range. The horizontal dashed line represents the average place segregation, while dot size reflects the number of activities in the category; **(c)** Models performances (R^2) of the linear regressions with and without mobility patterns.

However, the effect of mobility behavior varies depending on where individuals reside. Specifically, higher exploration rates, diversity of category visited, and larger distances from home are generally associated with reduced partisan segregation, increasing exposure to politically diverse environments. These effects are strongest for individuals residing in highly segregated areas, while mobility has a more limited impact for those living in less segregated neighborhoods. Overall, the interaction between mobility patterns and residential segregation highlights how individual movement can moderate the extent of partisan segregation in daily life.

Mobility behavior and place segregation

The average place partisan segregation by category persists with catchment range, defined as the average distance between POIs and individuals' home locations, with a Pearson's correlation of -0.196. However, as shown in Fig. 1b, this pattern varies across categories. While daily necessities such as grocery stores, discount stores, and pharmacies, tend to have relatively high average partisan segregation and short catchment ranges, categories with a medium catchment range (around 12 kilometers from home) show a wide spread in segregation levels. For example, fast food restaurants, coffee shops, and post offices display higher segregation compared to salons, barber shops, or automotive services, despite similar distances from home. At the upper end of the catchment range, categories such as hotels and offices tend to show average segregation levels. Overall, the figure highlights how both distance and activity type shape the partisan composition of places individuals visit in daily life.

4. DISCUSSION

This study leverages high-resolution individual mobility data to investigate partisan segregation patterns and the moderating role of mobility behavior. Our findings show that, although the residential context greatly influences the composition of individuals' networks, it does not fully determine their partisan exposure. Even people who live in the same area can have different segregation patterns, as highlighted by the consistent within-areas variance. From this perspective, mobility behavior can moderate residential segregation, mitigating the level of experienced segregation, especially in more segregated areas. Finally, place-based segregation persists with distance, but varies widely across activity types, further shaping individuals' partisan exposure. In this context, future work aims to understand the relationship between individuals' experienced partisan segregation and political voting behavior in the US.

REFERENCES

1. J. R. Brown and R. D. Enos, Nat. Hum. Behav. **5**, 998 (2021).
2. J. R. Brown, E. Cantoni, R. D. Enos, *et al.*, Work. paper (2022).
3. E. Moro, D. Calacci, X. Dong, and A. Pentland, Nat. communications **12**, 4633 (2021).
4. S. Talaifar, D. Jordan, S. D. Gosling, and G. M. Harari, J. Pers. Soc. Psychol. (2025).
5. Y. Zhang, S. Cheng, Z. Li, and W. Jiang, Sci. reports **13**, 9768 (2023).

Tracing Behavioural Changes of High Street Visitors in London Using Mobile Location Data

XINGLEI WANG[1,*] AND TAO CHENG[1,]

[1] SpaceTimeLab, Dept. of Civil, Environmental & Geomatic Engineering, UCL, Gower Street, London, UK, WC1E 6BT
[*] xinglei.wang.21@ucl.ac.uk

1. INTRODUCTION

High streets serve as critical arteries of local communities where economic vitality, social interaction, and everyday convenience converge [1]. Prior research has primarily focused on quantifying and analysing high street performance through metrics such as vacancy rates [2] or by tracking longitudinal changes in footfall [3, 4], as well as examining fine-grained temporal patterns [5]. However, we argue that high streets also provide a powerful lens through which to observe patterns of human mobility and behavioural change, particularly in response to large-scale disruptions. To this end, our study adopts a human-centred perspective. Specifically, we investigate how the COVID-19 pandemic altered visitation patterns to 616 high streets across London, leveraging high-resolution mobile phone data to uncover shifts in both the intensity and purpose of human presence.

The spatial scope of this research is notable: analysing hundreds of high streets across one of the world's most densely populated and socially diverse metropolitan regions allows us to capture unique behavioural changes. Temporal granularity is equally important; by examining data snapshots from February 2020 (pre-pandemic), February 2021 (lockdown), and February 2022 (post-restriction), we trace the pandemic's arc and its longer-term behavioural impacts.

This study contributes to urban mobility research by quantifying changes in human behaviour through mobile location data. Our findings offer valuable insights into behavioural adaptation, hybrid work patterns, and the partial reversion to pre-pandemic rhythms, and demonstrate how anonymised mobile phone data can serve as a powerful tool for urban and behavioural analysis. A key methodological contribution is the introduction of a novel role-based classification of high street visitors. By categorising individuals as residents, workers, or visitors, we are able to uncover more nuanced and detailed behavioural shifts across different user groups—an aspect that has been largely overlooked in the existing literature.

2. STUDY AREA AND DATA DESCRIPTION

Our study focuses on the Greater London area, home to over 9 million residents and a complex network of urban high streets. A high street is defined as a commercially active corridor of mixed-use urban development, typically characterised by retail shops, services, hospitality venues, and significant pedestrian traffic. We analyse 616 designated high streets across London, providing comprehensive spatial coverage.

The primary dataset is a mobile location dataset comprising millions of anonymised mobile phone location pings from users. These data were collected via a range of location-based service apps (from over ten different providers) in compliance with GDPR regulations. The dataset spans over three years, from which we extract data for three specific monthly windows: February 2020, February 2021, and February 2022.

3. METHOD

The workflow consists of the following key steps:

1. Data cleaning: Raw mobile phone data were first cleaned to remove anomalous records such as GPS positioning errors.

2. Home location detection: We inferred each user's home location at the Lower Layer Super Output Area (LSOA) level, denoted as *home_lsoa*. A person's home was defined as the LSOA where they spent the most time during nighttime hours (22:00–07:00). If no nighttime records were available, daytime data were used as a fallback.

3. Stay detection: We identified stationary periods (or "stays") where a user remained at the same location for more than a specific duration — set at 5 minutes in this study. Each stay record includes a start time and an end time, from which the duration of stay is derived.

4. Identity inference: We classified individuals visiting a given high street into one of three categories: residents, workers, or visitors.

- If a person's *home_lsoa* intersected with the high street's LSOA boundary, they were classified as a resident.
- Non-residents were further divided into: Workers: those who spent more than 5 hours at the high street during standard working hours (09:00–18:00). Visitors: all other non-residents who did not meet the worker threshold.

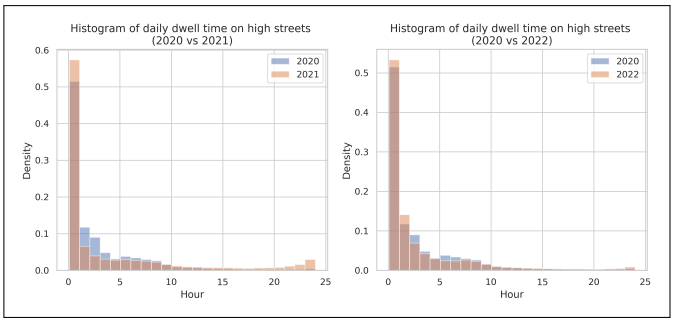

Fig. 1. Comparison of daily dwell time in three years

With stays identified and visitor types inferred, we derived the following key behavioural indicators:

a. Total Duration of Stay per Person per Day. From the detected stays, we measured how long each person spent on the high street each day. Time spent at urban locations serves as a valuable proxy for behavioural intent and engagement [6]. We then compared daily dwell time distributions across the three years using histograms to reveal temporal shifts, highlighting behavioural compression or expansion under pandemic restrictions and recovery.

b. Proportions of Different Footfall Types. We aggregated the number of residents, workers, and visitors per day across each day of the week. For each high street, we computed the

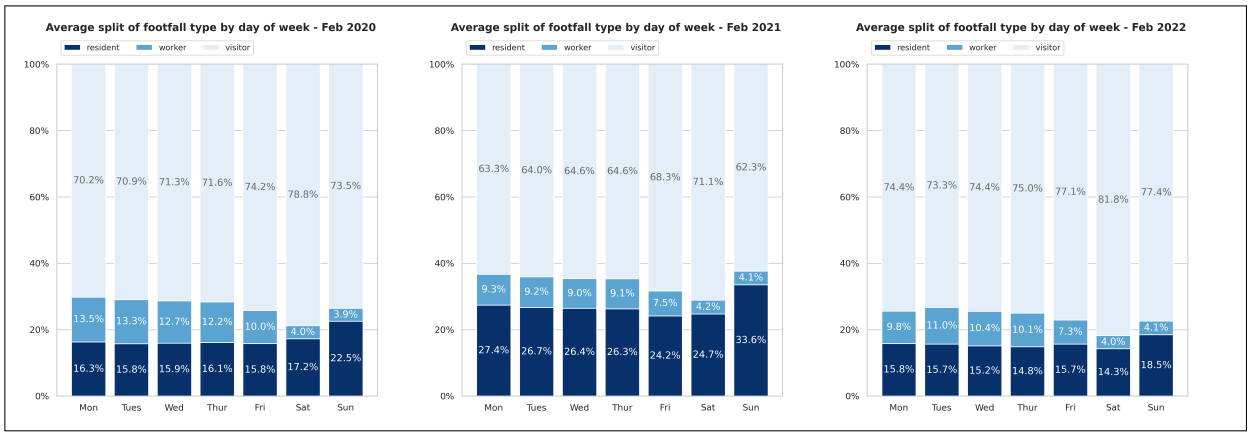

Fig. 2. Comparison of daily dwell time in three years

daily proportions of each footfall type, then averaged these proportions over all weeks in February to generate final results. This allowed us to compare footfall composition across time and uncover structural changes in usage patterns.

4. RESULTS AND FINDINGS

A. Duration-based Visitation Patterns

Figure 1 displays histograms of dwell time distributions for February 2020, 2021, and 2022.

From 2020 to 2021, the data reveal a marked increase in short visits (<1 hour), likely reflecting essential shopping under strict lockdown measures. Simultaneously, moderate-length visits (1–4 hours) dropped significantly, corresponding to the shutdown of non-essential retail, hospitality, and social venues. Notably, we also observed a rise in long-duration presence (>20 hours), indicative of residents remaining in their local areas, often due to work-from-home mandates.

By February 2022, the dwell time distribution showed partial normalisation. Discrepancies in the 0–5 hour range became less pronounced compared to 2020, suggesting some behavioural reversion. However, the lack of full convergence implies enduring changes in how people use high streets post-pandemic.

B. Footfall Typology and Weekly Patterns

Figure 2 presents the average daily proportion of each footfall type across the week for the three study periods.

In 2021, residents comprised a larger share of high street visitors (e.g., weekday proportion increased from 16% in 2020 to 27%). This increase aligns with lockdown-related immobility and localised activity. In contrast, workers' share declined sharply during the same period, then partially rebounded in 2022. However, 2022 figures remained below 2020 levels, highlighting the persistent shift toward hybrid or remote work.

Visitor proportions followed a V-shaped trajectory: decline in 2021 and increase in 2022. Interestingly, the 2022 share of visitors surpassed 2020 levels, not necessarily due to more visitors in absolute terms, but likely reflecting a decline in worker share.

A salient trend emerges when examining weekday worker patterns. In 2022, the proportions of workers on Mondays and Fridays were significantly lower compared to midweek days (Tuesday through Thursday). This pattern supports emerging evidence of the "TWaTs" (Tuesday, Wednesday, and Thursday) hybrid work schedule, where employees are more likely to be on-site midweek [7].

5. CONCLUSIONS

This study provides a high-resolution empirical account of how the COVID-19 pandemic disrupted and reshaped urban high street usage in London. By combining spatial-temporal mobile phone data with behavioural classification, we reveal both temporary and structural shifts in mobility patterns.

Our key findings include compressed visitation durations during lockdown, a persistent shift toward hybrid work, and a partial but incomplete return to pre-pandemic rhythms by 2022. Moreover, we highlight how anonymised mobility data can offer real-time, scalable insights into urban behavioural dynamics.

These insights have implications for urban planners, retailers, public health authorities, and policymakers aiming to understand and adapt to new patterns of urban life. As cities continue to evolve in the post-pandemic era, such data-driven approaches will be vital in designing resilient, responsive, and inclusive urban environments.

Funding. X. W. is jointly funded by UCL Dean's Prize and China Scholarship Council [No. 202106270039].

Data Availability Statement. Data underlying the results presented in this paper may be obtained from the authors upon reasonable request.

REFERENCES

1. M. Carmona, "London's local high streets: The problems, potential and complexities of mixed street corridors," Prog. planning **100**, 1 (2015).
2. L. Dolega and A. Lord, "Exploring the geography of retail success and decline: A case study of the liverpool city region," Cities **96**, 102456 (2020).
3. X. Wang, X. Zhang, and T. Cheng, "The ups and downs of london high streets throughout covid-19 pandemic: Insights from footfall-based clustering analysis (short paper)," in *12th International Conference on Geographic Information Science (GIScience 2023)*, (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2023), pp. 80–1.
4. P. Ballantyne, A. Singleton, and L. Dolega, "Using unstable data from mobile phone applications to examine recent trajectories of retail centre recovery," Urban informatics **1**, 21 (2022).
5. X. Wang, J. Shi, and T. Cheng, "Capturing high street dynamics on a finer scale: a case study in the context of covid-19," in *the 32nd GIS Research UK Conference (GISRUK 2024)*, (2024).
6. F. Zhou, X. Zhang, X. Wang, and T. Cheng, "Customer profiling based on mobile apps gps data: A case study on westfield shopping malls," in *2023 30th International Conference on Geoinformatics*, (IEEE, 2023), pp. 1–5.
7. The Telegraph, "Londoners told to take the Tube on two days as more work from home," https://www.telegraph.co.uk/news/2023/01/25/londoners-told-take-tube-two-days-work-home/.

When *Proximity Falls Short: Inequalities in Commuting and Accessibility in Santiago, Chile*

CESAR MARIN-FLORES[1], LEO FERRES[2,3,*], HENRIKKI TENKANEN[1]

[1] *GIST Lab, Department of Built Environment, Aalto University, Finland*
[2] *Data Science Institute, Universidad del Desarrollo, Chile*
[3] *ISI Foundation, Turin, Italy.*
*lferres@udd.cl

1. INTRODUCTION

Urban mobility has traditionally been studied using data sources like travel surveys. While these have provided valuable insights, they have significant limitations: surveys are expensive, infrequent, and often fail to capture seasonal variations in travel patterns. Mobile phone data offers a practical alternative, providing continuous information on movement patterns. Unlike traditional surveys, this data captures both temporal and spatial variability. One such data source is the eXtended Detail Records (XDR), which comprise all billable records per user [1] providing cost-effective, scalable, and up-to-date information applicable to domains, such as transportation [2] and human migration [3].

Understanding individual movement between places is fundamental in mobility studies, yet mobility alone does not fully explain access to opportunities. The concept of accessibility provides a more comprehensive approach, operationalizing the interrelation between land use, transport, and socio-economic conditions [4]. Defined by Hansen (1959), accessibility measures the potential opportunities individuals can reach, often expressed in terms of travel time, cost, or distance. Combining mobility data with accessibility metrics enables a novel perspective on urban dynamics. This approach is particularly relevant for cities in the Global South, where social inequalities are well documented, yet studies utilizing big data to explore these dynamics remain scarce.

This study examines how commuting patterns in Santiago relate to socio-economic disparities using up-to-date mobile phone records. Specifically, we investigate whether systematic differences exist in commuting times and accessibility across economic levels. To achieve this, we analyze global commuting times and accessibility metrics to identify spatial patterns of advantage and disadvantage neighborhoods. Our approach integrates XDR data with accessibility methods to provide a detailed assessment of how different social groups interact with key opportunities in the city.

2. METHODS

Our methodology consists of three main steps (Fig.1). First, we derive commuting patterns using one month of XDR data from Telefónica Chile, covering approximately 28–31% of the national mobile phone market. This dataset includes all billable records from over 1,300 Base Transceiver Stations (BTS) in urban Santiago. Residential locations are inferred by identifying the most frequently used BTS during nighttime, while daytime locations serve as work-time anchors, following the approach of Ahas et al. (2010). These locations are referred to as key opportunities. Second, we compute optimal commuting routes and travel times using r5py [7], a routing tool that integrates public transportation and walking. The analysis relies on OpenStreetMap for the street network and General Transit Feed Specification (GTFS) data for public transport schedules, covering a total of 84 million origin-destination pairs. Third, we incorporate the socioeconomic dimension using the Social-Material Index [8], which provides fine spatial granularity and accounts for factors beyond economic status, such as education levels and overcrowding.

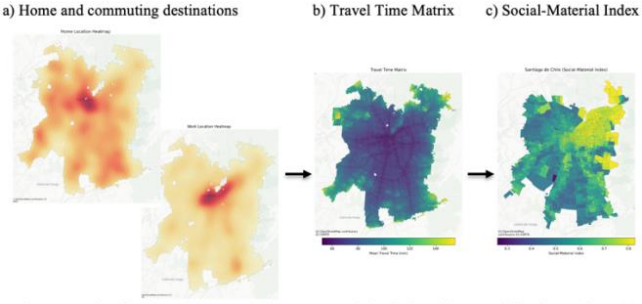

Fig. 1. Visualization of home and key opportunities derived from mobile phone data (a). The travel time matrix of the city for each H3 Level 9 hexagon (b). Social-material index obtained from open-source data (c).

3. RESULTS

The analysis, based on 610,000 recurrent commuters, reveals an average commuting time of 46.8 minutes for the combination of public transport and walking. This aligns with previous findings indicating that Santiago has the longest commuting times in Chile, ranging between 40 and 60 minutes [9]. We assess equity in commuting times using the Palma ratio [10], obtaining a value of 1.1, which suggests minimal differences in average commuting times between the most and least privileged areas. However, spatial patterns indicate significant disparities. A bivariate map (Fig. 2) combining the Social-Material Index with commuting times reveals that northeastern Santiago, home to the city's wealthiest neighborhoods, does not exhibit a clear clustering of low travel times, even in areas with a high density of key

opportunities. In contrast, disadvantaged areas in the southern periphery—such as La Pintana, Puente Alto, Lo Espejo, and San Bernardo—form more defined clusters, with some showing both low socio-economic status and high commuting times. These extreme cases of dual disadvantage are largely absent near metro lines, highlighting the role of efficient public transport in mitigating inequalities.

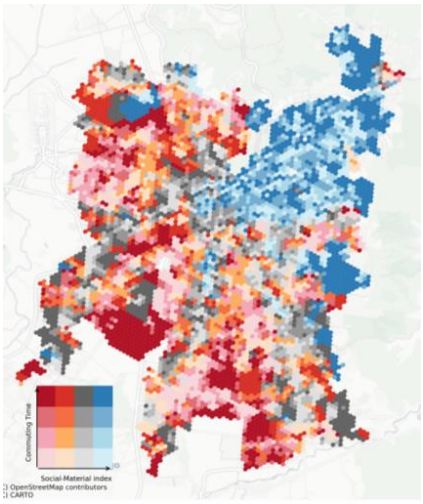

Fig. 2. Bivariate map showing the SMI (color variations) and average commuting time (clarity variations)

Figure 3 further illustrates stark disparities in access to opportunities, with high-accessibility areas concentrated in economically privileged neighborhoods. Rather than a direct link between more opportunities and shorter commutes, we observe a cone-like distribution converging toward the city's average commuting time. This suggests that commuting inequalities stem less from differences in travel time and more from people's freedom to choose among numerous reachable opportunities within a reasonable commuting time.

4. CONCLUSIONS

In this paper, we investigated commuting patterns in Santiago and their relationship with socio-economic factors. Our results indicate that there are no significant differences in average commuting times by public transport between privileged and disadvantaged areas. Our methodology allows for a spatial evaluation of these variables, helping to identify areas that are both economically and accessibility disadvantaged with high granularity. These areas could become key targets for more effective land use and transport policies. By combining the accessibility approach with mobility data, we were able to compare the potential access to key opportunities with actual commuting behavior simultaneously. This approach reveals inequalities that are hidden when using traditional, separate methods. For instance, we found that proximity to key opportunities alone does not necessarily translate into shorter commuting times. However, a limitation of this study is that we do not account for modal choice. We do not evaluate travel times by car, whose ownership is unevenly distributed across the city and tends to reduce commuting times in wealthier areas. As a result, the apparent equality in global metrics may mask disparities that need to be identified in future research.

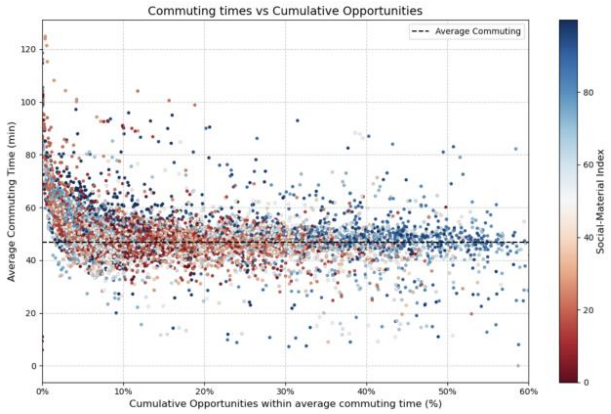

Fig. 3. Cumulative key opportunities accessible within the city's average commuting time (46.8 min) from each hexagon (X-axis), compared to the average commuting time of the corresponding area

Data availability

The mobile phone data used in this study came from a private provider, one of the major telecom operators in Chile. Due to privacy agreements, the data cannot be shared publicly.

References

[1] F. Calabrese, L. Ferrari, and V. D. Blondel, "Urban Sensing Using Mobile Phone Network Data: A Survey of Research," *ACM Comput. Surv.*, vol. 47, no. 2, pp. 1–20, Jan. 2015, doi: 10.1145/2655691.

[2] E. Graells-Garrido, D. Caro, and D. Parra, "Inferring modes of transportation using mobile phone data," *EPJ Data Sci.*, vol. 7, no. 1, p. 49, Dec. 2018, doi: 10.1140/epjds/s13688-018-0177-1.

[3] B. Aydoğdu and A. A. Salah, "Mobile Phone Data Based Indicators of Migration." Κατάλογος Δεδομένων SoDaNet, 2024. doi: 10.17903/FK2/DCJEMR.

[4] R. H. M. Pereira, T. Schwanen, and D. Banister, "Distributive justice and equity in transportation," *Transport Reviews*, vol. 37, no. 2, pp. 170–191, Mar. 2017, doi: 10.1080/01441647.2016.1257660.

[5] W. G. Hansen, "How Accessibility Shapes Land Use," *Journal of the American Institute of Planners*, vol. 25, no. 2, pp. 73–76, May 1959, doi: 10.1080/01944365908978307.

[6] R. Ahas, S. Silm, O. Järv, E. Saluveer, and M. Tiru, "Using Mobile Positioning Data to Model Locations Meaningful to Users of Mobile Phones," *Journal of Urban Technology*, vol. 17, no. 1, pp. 3–27, Apr. 2010, doi: 10.1080/10630731003597306.

[7] C. Fink, W. Klumpenhouwer, M. Saraiva, R. Pereira, and H. Tenkanen, *r5py: Rapid Realistic Routing with R5 in Python*. (Sep. 09, 2022). [object Object]. doi: 10.5281/ZENODO.7060437.

[8] Observatorio de Ciudades, Pontificia Universidad Catolica de Chile, 2022. [Online]. Available: https://ideocuc-ocuc.hub.arcgis.com/maps/c83a1ea2c31b4850b65a481b21e4919f/about

[9] Herrera and Razmilic, "De la casa al trabajo: Análisis de un tiempo perdido," 2018.

[10] A. Karner, R. H. M. Pereira, and S. Farber, "Advances and pitfalls in measuring transportation equity," Aug. 05, 2023. doi: 10.31235/osf.io/y246u.

Hybrid work has increased income segregation at third places in US cities

XIAOTONG XU[1,2], CALLIE CLARK[2,3], WEI MA[1], AND TAKAHIRO YABE[2,4,*]

[1] Department of Civil and Environmental Engineering, The Hong Kong Polytechnic University, 11 Yuk Choi Rd, Hong Kong SAR, 999077
[2] Center for Urban Science and Progress, Tandon School of Engineering, New York University, 370 Jay Street, New York, 11201
[3] Marron Institute of Urban Management, New York University, 370 Jay Street, New York, 11201
[4] Department of Technology Management and Innovation, Tandon School of Engineering, New York University, 2 Metrotech, New York, 11201
[*] takahiroyabe@nyu.edu

1. INTRODUCTION

The concept of the "third place" introduced by Ray Oldenburg [1], refers to public spaces distinct from the home (the first place) and the workplace (the second place). These are familiar communal environments, such as cafés or parks, where individuals can relax, engage in informal conversation, and foster social connections with both friends and strangers through shared interests or activities. Numerous studies [2, 3] highlight that third places are vital for sustaining social bonds, fostering belonging, and sparking creativity in today's increasingly fragmented society.

However, the COVID-19 pandemic significantly increased the adoption of remote and hybrid work in certain industries, altering the long-established relationship between the first (home) and second (work) place. By combining multiple survey sources [4, 5], such as the American Heritage Time Use Study (AHTUS), the American Time Use Survey (ATUS), and the Survey of Working Arrangements and Attitudes (SWAA), Barrero et al. [6] estimate that the share of full workdays performed from home among individuals aged 20 to 64 in the United States was just 7% in 2019. This figure surged to 60% in 2020 and stabilized at 28% by June 2023. The pandemic catalyzed a widespread and irreversible shift toward hybrid work, with its continued prevalence in the post-pandemic era indicating increased recognition and acceptance of this work arrangement.

As the workforce shifts toward hybrid work arrangements, patterns of third place visitation are inevitably reshaped. Despite extensive speculation [7, 8], there is limited behavioral evidence on how third place usage has evolved post-pandemic and the forces driving these shifts. In this study, we leverage anonymized, high-resolution mobile phone location data from June 2019 and June 2022 across eight major U.S. Metropolitan Statistical Areas (MSAs) to identify individual hybrid work behaviors and visits to third places. Our findings suggest that the growing prevalence of hybrid work has driven individuals to prefer third places nearer to their homes, signaling a structural transformation in work patterns that exacerbates income-based segregation at these venues, with meaningful implications for social inclusion and inequality.

2. METHOD

Using anonymized, high-resolution mobile phone location data provided by Cuebiq from 2019 and 2022 across eight major U.S. MSA regions, we identify individual-level hybrid work patterns. By further integrating the SafeGraph POI dataset, we examine their visitation behaviors to third places.

A. Estimation of hybrid workers

We use mobility data from June 2019 (pre-pandemic) and June 2022 (post-pandemic) to estimate workers conducting hybrid work. First, we use individuals' daytime and nighttime stay locations over the course of a month to identify each user's workplace and home locations, respectively. Then we classify their commuting patterns on weekdays to distinguish between in-person and hybrid workers, with those commuting on less than 80% of weekdays categorized as hybrid workers. Based on this classification, we can calculate the proportion of hybrid workers who work and live in these MSA regions.

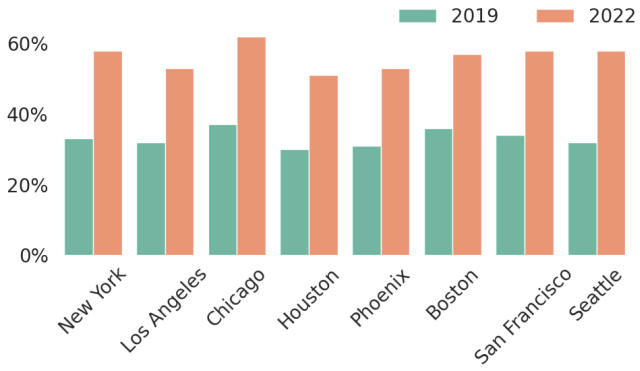

Fig. 1. The proportion of hybrid workers among employees working in the CBD areas across the different MSA regions.

B. Estimation of third places

For users with identified home and work locations, we analyze their visits to third places over the month. Specifically, we match each non-home, non-work stop to nearby POIs using k-nearest neighbors algorithms [9] with a maximum distance of 50 meters. We then identify third places based on criteria such as repeated visits, dwell time, and distance from home and work. This allows us to capture each user's third-place visitation behavior.

3. INCREASED HYBRID WORK IN US CITIES

To examine the trend of changes in hybrid work, we compare the estimation results from 2019 and 2022, and find that hybrid work became significantly more prevalent in the post-pandemic period. It reveals that the proportion of hybrid workers increased by 10 percentage points across all MSA regions following the pandemic, with the proportion in Central Business Districts (CBDs) experiencing an almost twofold increase, as shown in Figure 1.

Although the proportion of hybrid workers has shown a widespread increase, the adoption of hybrid work exhibits clear occupational disparities. These remote-compatible occupations tend to be more knowledge-intensive, highly digitalized, and allow for greater autonomy, such as roles in software engineering, finance, and marketing. They not only support task execution via digital platforms but also offer flexibility in both time and

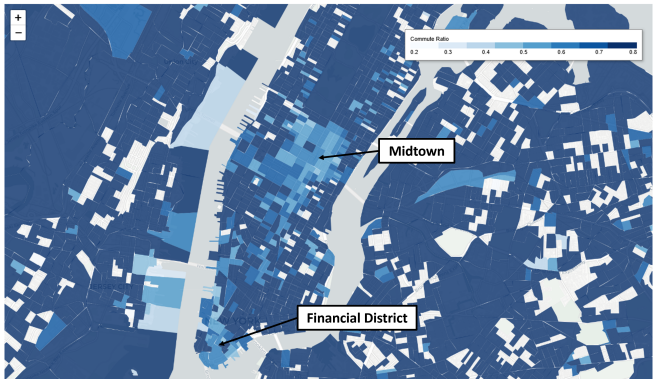

Fig. 2. Spatial heterogeneity in Friday office attendance, NY-NJ MSA (June 2022).

location. Notably, these occupations account for a disproportionate 40% of total U.S. wages [10], underscoring their dominant role in the hybrid work transition.

Taking the New York–Newark–Jersey metropolitan statistical area (NY-NJ MSA) as an example, our analysis reveals clear spatial heterogeneity in hybrid work arrangements across occupations. As shown in Figure 2, we show the proportion of on-site workers on Friday at the Census Block Group (CBG) level for June 2022. Notably, the office attendance on Fridays is significantly lower in Manhattan than in other areas. This pattern aligns with Manhattan's occupational structure, particularly in the Financial District and Midtown Manhattan, which concentrate a large number of finance and tech-sector workers with higher adoption rates of hybrid work.

4. CHANGE IN THIRD PLACE VISITATION PATTERNS

We categorize weekdays into in-office days and work-from-home (WFH) days based on individuals' commuting patterns, and analyze how third place visitation behaviors differ between these two work modes. Our analysis remains centered on individual-level data from June 2022, focusing on those with workplaces in New York's CBD area, as this allows us to best evaluate the impact of hybrid work adoption on engagement with third places.

First, we compare the distance from third places to individuals' home and workplace on in-office days versus WFH days. The results are shown in Figure 3. It shows that people tend to visit third places that are closer to their home on WFH days. Specifically, the median distance to home on WFH days is 2.87 km, compared to 4.79 km on in-office days ($t = -42.98$, $p < 0.001$). Conversely, the median distance to the workplace on WFH days is 12.02 km, while it is 4.14 km on in-office days ($t = 80.28$, $p < 0.001$).

We further analyze the income-based diversity of third places visited during in-office and WFH days. Specifically, using the American Community Survey (ACS), we compute the Gini index based on the median household incomes of visitors' home CBGs, representing the income heterogeneity of individuals frequenting each location. A higher Gini value indicates a more socioeconomically diverse visitor group, whereas a lower value suggests a more homogeneous group. The distribution of the income Gini index is given in Figure 4, the median Gini index of third places visited on WFH days is 0.212, compared to 0.237 on in-office days, suggesting that individuals tend to experience income segregation when they visit third places near home on WFH days. A t-test confirms the difference is statistically significant ($t = 24.72$, $p < 0.001$).

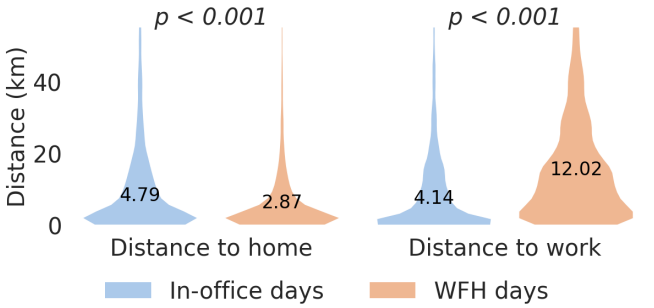

Fig. 3. Distribution of distances from third places to home (left) and workplace (right).

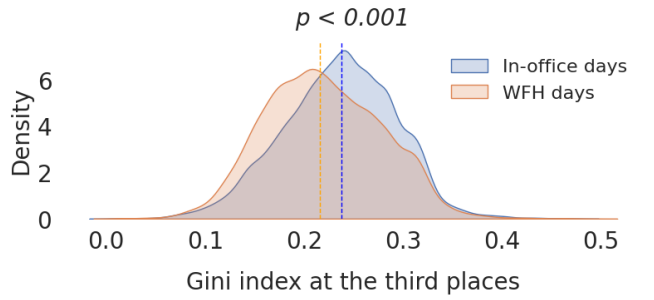

Fig. 4. The distribution of income Gini index at third places.

5. CONCLUSION

This study uses mobility data from eight U.S. MSA regions to confirm the post-pandemic rise of hybrid work. By examining shifts in visitation patterns to third places, it provides behavioral evidence of emerging income segregation. The findings carry important implications for understanding post-pandemic social interaction and inequality in access to communal spaces.

6. BACK MATTER

Funding. T.Y. and C.C. acknowledge support by the National Science Foundation under grant number CMMI 2425021.

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

REFERENCES

1. R. Oldenburg, *The great good place: Cafes, coffee shops, bookstores, bars, hair salons, and other hangouts at the heart of a community* (Da Capo Press, 1999).
2. D. M. Littman, Am. J. Community Psychol. **69**, 436 (2022).
3. N. Lee, Cities **131**, 104049 (2022).
4. K. Fisher, J. Gershuny, S. M. Flood, *et al.*, "American heritage time use study extract builder: Version 1.2," (2018). Dataset.
5. S. M. Flood, L. C. Sayer, D. Backman, and A. Chen, "American time use survey data extract builder: Version 3.2," Dataset (2023). Accessed via IPUMS.
6. J. M. Barrero, N. Bloom, and S. J. Davis, J. Econ. Perspect. **37**, 23 (2023).
7. W. Li, E. Zhang, and Y. Long, Sustain. Cities Soc. **103**, 105258 (2024).
8. Cody Cook, The Rev. Econ. Stat. (2024).
9. S. M. Omohundro, "Five balltree construction algorithms," Tech. Rep. TR-89-063, International Computer Science Institute (1989).
10. J. I. Dingel and B. Neiman, J. Public Econ. **189**, 104235 (2020).

Beyond Distance: Mobility Neural Embeddings Reveal Visible and Invisible Barriers in Urban Space

GUANGYUAN WENG[1,4], MINSUK KIM[5], YONG-YEOL AHN[5], AND ESTEBAN MORO[1,2,3*]

[1] Network Science Institute, Northeastern University, Boston, MA 02115, USA
[2] Department of Physics, Northeastern University, 02115 Boston, USA
[3] Media Lab, Massachusetts Institute of Technology, Cambridge, MA 02139
[4] Khoury College of Computer Sciences, Northeastern University, Boston, MA 02115, USA
[5] Luddy School of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN 47405, USA
[*] e.moroegido@northeastern.edu

1. INTRODUCTION

Cities are social and opportunity hubs that foster diversity, social interaction, innovation, and economic growth [1–3]. However, cities are also characterized by persistent inequality in accessibility and mobility for their residents [4]. The urban landscape is scarred by large visible physical barriers like natural obstacles, transportation infrastructure, or amenities, which constrain mobility and social connections [4, 5]. At the same time, invisible demographic boundaries defined by class, race, income, or lifestyles, create social barriers between groups that restrict everyday mobility between neighborhoods, perpetuating segregation and isolation [6–8].

Despite their impact on urban mobility, there is no comprehensive study of visible and invisible barriers in cities due to their complex, heterogeneous, and uneven effects across sociodemographic and geographical contexts. Physical barriers create significant mobility constraints, particularly affecting pedestrians at short spatial scales, while enabling connectivity at larger distances [9]. Intricate interactions between barriers often amplify their collective impact on urban mobility [10]. For example, some highway barriers in the U.S. were deliberately designed to reinforce racial segregation, linking their visible and invisible role in constraining mobility [11]. In parallel, sociodemographic factors such as race and income strongly affect mobility, restricting visitation and accessibility to amenities [6, 12]. Some barriers are not easily recognized: while parks are traditionally viewed as inclusive public spaces, evidence suggests they can function as socio-demographic barriers [4, 13]. Traditional studies of mobility barriers have relied mainly on small case studies [4, 14, 15]. More recent work has used large-scale mobility data to identify barriers at broader spatial scales [9, 16, 17], but these studies have focused primarily on visible, physical obstacles, often overlooking invisible, demographic, or institutional factors. Without a more comprehensive perspective, conclusions about urban mobility and accessibility risk being incomplete or misleading.

2. METHODS

To address this gap, we combine large-scale mobility data with neural embedding techniques that infer mobility barriers between urban areas from observed movement patterns. Our mobility data consists of anonymized trajectories of 25.4 million devices across 11 metro areas in the United States from 2019 to 2022 provided by Spectus. It is complemented by 2.5 million verified venues from Foursquare and Safegraph, as well as spatial data on major physical obstacles from Open Street Map.

To systematically represent mobility patterns across cities and at scale, we use a neural embedding-based approach that infers mobility barriers by measuring effective distances between locations from observed trajectories. We adopt the Word2Vec model due to its formal equivalence with the mobility gravity law [18]. Mobility trajectories are treated as sequences of visited locations at Census Block Group (CBG) level, allowing the model to learn vector embeddings $\mathbf{w}_i$ for each location i based on co-occurrence patterns. The resulting embeddings capture behavior-based distances $d_{ij}^{(e)}$ (i.e., cosine distance between $\mathbf{w}_i$ and $\mathbf{w}_j$) between CBGs that reflect actual movement patterns better than geographic distances, $d_{ij}^{(p)}$. To validate our approach, we model normalized flux using a gravity model and find that embedding distance explains mobility better than geographic distance across cities and years (e.g., $R^2 = 0.61$ vs. $R^2 = 0.33$ in Boston).

To quantify mobility barriers, we identify CBG pairs where observed mobility is significantly lower than expected given their geographic distance. To do that, we compare embedding-based distances $d_{ij}^{(e)}$ to geographic distances $d_{ij}^{(p)}$ and define a barrier when the embedding distance is substantially greater than the value predicted from physical distance using a log-linear model. We focus on CBG pairs within 20 km to account for CBSA size and minimize broader spatial effects. To understand factors associated with barrier presence, we examine four groups of factors known to shape human mobility. First, we consider physical barriers measured by highways, railways, parks, and waterways crossed between areas. Second, we account for amenity structure, using the number of intervening POIs and the Jensen–Shannon distance in POI category composition. Third, we incorporate sociodemographic differences in income, racial composition, and public transit usage. Finally, administrative boundaries are included as a binary variable indicating whether CBGs belong to different counties.

3. RESULTS

Our results show that amenity structure is the primary determinant of mobility barriers across all cities, years, and distance ranges, operating through two opposing mechanisms. First, the number of intervening opportunities between CBGs is positively associated with barrier formation, as nearby alternative destinations reduce direct movement likelihood. In contrast, greater POI Jensen–Shannon distance—indicating higher functional complementarity—reduces barrier likelihood. Shopping, College, and Sport venues exhibit the most substantial barrier effects at all distances, while Transportation, Service, and Grocery facilities facilitate mobility between short-distance areas.

Administrative boundaries are the second strongest predictor of mobility barriers. CBG pairs in different counties are significantly more likely to have barriers, consistent with prior work showing that jurisdictional lines can disrupt movement [19]. Demographic differences—especially in income, race, and transit use—also strongly correlate with barriers across cities and distances, reflecting persistent segregation and activity space isolation [6, 20].

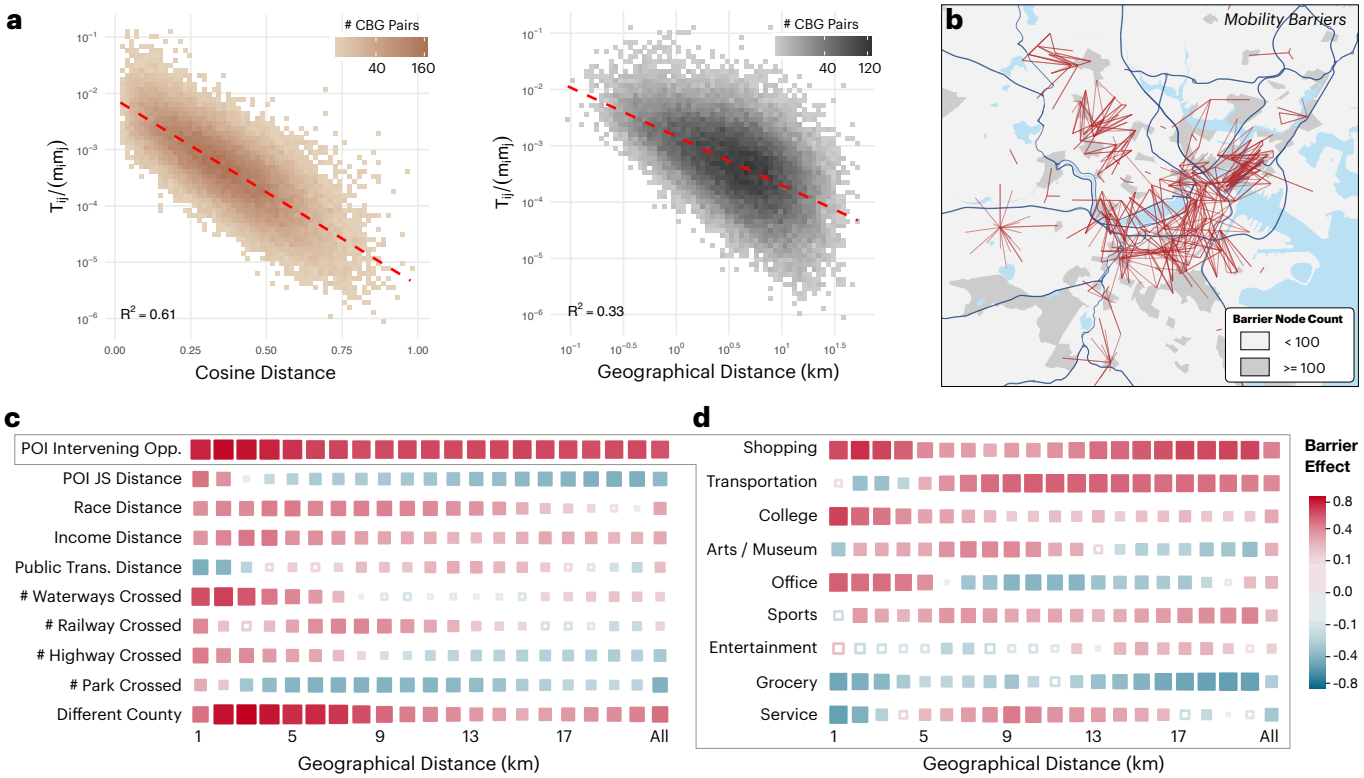

Fig. 1. Understanding mobility borders through human movement patterns and spatial embeddings. (a) Binned density plots for the Boston CBSA showing the relationship between normalized mobility flows between CBGs and embedding cosine distance (left) and geographic distance (right) in 2019. Red dashed lines show linear fits; embedding distance better explains flows ($R^2 = 0.61$ vs. $R^2 = 0.33$). **(b)** Sample of detected mobility barriers in Boston. Lines connect CBG pairs with barriers; CBGs involved in over 100 barriers are highlighted. Some align with visible infrastructure, while others occur between adjacent areas without physical separation. **(c)** Heatmap of logistic regression coefficients for barrier predictors across distances in 11 CBSAs, including POI features, physical infrastructure, demographic differences, and county boundaries. Empty symbols correspond to statistically non-significant coefficients (p-value < 0.05). **(d)** Breakdown of the intervening opportunities effect shown in (c), disaggregated by POI category.

Physical barriers are the least predictive of mobility barriers when compared to amenity, administrative, and demographic factors. While these features contribute to restricted movement at short distances, their effect diminishes and even reverses at larger spatial scales. Highways can impose constraints on local mobility but facilitate movements at longer distances. Similarly, parks function as barriers at short distances, contrary to their intended role, limiting mobility between nearby neighborhoods.

4. DISCUSSION

Neural embedding-based models like ours offer a powerful tool to measure complex functional access, uncover spatial inequality, and inform more inclusive approaches to urban planning and policy. As cities confront rising inequality, the need for frameworks that reflect real accessibility —grounded in behavior rather than geography— becomes increasingly urgent. Access is not defined by meters or minutes, but by the ability to move across the social, economic, and institutional barriers that structure urban life.

REFERENCES

1. M. Schläpfer, L. M. Bettencourt, S. Grauwin, *et al.*, J. Royal Soc. Interface **11**, 20130789 (2014).
2. A. Pentland, *Social physics: How social networks can make us smarter* (Penguin, 2015).
3. L. M. A. Bettencourt, Science **340**, 1438 (2013). Publisher: American Association for the Advancement of Science.
4. J. Jacobs, *The death and life of great American cities* (Modern Library, New York, 1993).
5. D. Appleyard, The ANNALS Am. Acad. Political Soc. Sci. **451**, 106 (1980).
6. Q. Wang, N. E. Phillips, M. L. Small, and R. J. Sampson, Proc. National Acad. Sci. **115**, 7735 (2018).
7. B. L. Levy, N. E. Phillips, and R. J. Sampson, Am. Sociol. Rev. **85**, 925 (2020). Publisher: SAGE Publications Inc.
8. F. Xu, Q. Wang, E. Moro, *et al.*, Nat. Hum. Behav. pp. 1–11 (2025). Publisher: Nature Publishing Group.
9. L. M. Aiello, A. Vybornova, S. Juhász, *et al.*, Proc. National Acad. Sci. **122**, e2408937122 (2025).
10. T. Wessel, Econ. Geogr. **98**, 464 (2022).
11. J. Trounstine, *Segregation by design: Local politics and inequality in American cities* (Cambridge University Press, 2018).
12. H. Barbosa, S. A. Hazarie, B. Dickinson, *et al.*, Sci. Reports **11** (2020).
13. T. Fraser, T. Yabe, D. P. Aldrich, and E. Moro, Comput. Environ. Urban Syst. **113**, 102173 (2024).
14. R. Grannis, Am. J. Sociol. **103**, 1530 (1998). Publisher: The University of Chicago Press.
15. E. O. Ananat, Am. Econ. Journal: Appl. Econ. **3**, 34 (2011).
16. M. Jin, L. Gong, Y. Cao, *et al.*, Comput. Environ. Urban Syst. **87**, 101625 (2021).
17. G. Pintér and B. Lengyel, arXiv preprint arXiv:2312.11343 (2023).
18. D. Murray, J. Yoon, S. Kojaku, *et al.*, Proc. National Acad. Sci. **120**, e2305414120 (2023).
19. G. Pintér and B. Lengyel, "Neighborhoods and Boundaries of Urban Mobility," (2023). ArXiv:2312.11343 [physics].
20. E. Moro, D. Calacci, X. Dong, and A. Pentland, Nat. communications **12**, 4633 (2021).

Correcting temporal bias in mobility data using time-use surveys

SARAH A. SANCHEZ[1,2,3*], HAMISH GIBBS[3,4], TAKAHIRO YABE[5], DANIEL T. O'BRIEN[1,2], AND ESTEBAN MORO[3,4]

[1] School of Public Policy and Urban Affairs, Northeastern University, Boston, Massachusetts, USA
[2] Boston Area Research Initiative, Boston, Massachusetts, USA
[3] Social Urban Networks Lab, Boston, Massachusetts, USA
[4] Network Science Institute, Northeastern University, Boston, Massachusetts, USA
[5] Tandon School of Engineering, New York University, New York, USA
[*] sa.sanchez@northeastern.edu

Introduction

Mobility data generated by mobile phones have been routinely used to measure patterns of social and economic activity in cities [1–4]. Despite its potential, mobility data suffers from well-documented limitations that can reduce the accuracy of these insights. These include **demographic biases**, where certain sub-populations are over- or under-represented in mobility data, and **temporal biases**, such as sample 'drift' over time or changes in individuals' mobile device usage, which may affect the behaviors captured in these datasets [5–7]. In this paper, we address temporal bias by benchmarking observed behavior against the American Time Use Survey (ATUS), a high-quality, demographically representative survey of self-reported activities in the U.S. Using this approach, we quantify the mismatch between activities recorded in mobility data with those reported in ATUS across time of day and income groups. We then apply these differences to re-weight mobility data and apply it to a specific metric, namely, measures of experienced income segregation at thousands of Points of Interest (POIs) in 11 metro-areas in the U.S. Our findings show that, while overall levels of segregation remain similar after adjusting for the temporal bias, there are still some significant differences at specific types of POIs and times of day. Our results then highlight the importance of correcting for temporal bias when using mobility data to measure urban inequality.

Methods

We analyze anonymized, high-resolution mobile location pings from Cuebiq between October 2016 and March 2017 across 11 cities in the U.S.: Boston, Chicago, Dallas, Detroit, Los Angeles, Miami, New York, Philadelphia, San Francisco, Seattle, and Washington. Mobility data processing and the conversion of individual pings to POI visits are based on previously established methods [2]. ATUS uses computer-assisted telephone interviews to record 24-hours of activities for a representative sample of U.S. residents, making it the most reliable data source analogous to mobility data [8]. To ensure compatibility, we analyzed ATUS respondents located in major cities (n = 2,502 respondents; excluding rural ones).

Mobility data is location-based while ATUS is activity-based. As such, ATUS has information on the type, location, start, and end-time of each activity that respondents reported doing in one day. In contrast, mobility data infers from cellphone pings where a user was and for how long, but we do not know for certain what they were doing there. Due to these differences, the POI categories of our mobility data and the location options provided by ATUS do not crosswalk perfectly. We focus on three POI categories which we stitched together across the 2 data sets: Food and Coffee / Tea (referred to as Food establishments), Grocery stores, and Gyms. We assigned an income quartile to each ATUS respondent to match the same bounds calculated for the mobility data quartiles [2].

We quantify temporal bias by first measuring the ratio for representation between ATUS and the mobility data (MD) with the equation:

$$R_{h,q,cat} = \frac{N_{cat}^{ATUS}(h;q) \ / N_{total}^{ATUS}(h;q)}{N_{cat}^{MD}(h;q) \ / N_{total}^{MD}(h;q)} \quad \textbf{(1)}$$

where N is the number of observed minutes (normalized), for each of the POI-categories (cat): (1) Food establishments, (2) Grocery stores, and (3) Gyms, for each hour (h) and income quartile (q). Secondly, we use that ratio to re-weight mobility data visits to generate a more representative sample regarding income quartile and time of day. Third, we estimate the impact of that bias by comparing the segregation calculated with unweighted data against weighted data. We examine these measures across POIs and categories. Initial baseline segregation for each POI was computed with the equation :

$$S_\alpha = \frac{2}{3}\sum_q \left| \tau_{\alpha,q} - \frac{1}{4} \right| \quad \textbf{(2)}$$

where $\tau_{\alpha,q}$ represents the proportion of time spent at place (α) by income quartile (q). POI-level segregation is quantified as a measure between 0 and 1, where $S_\alpha = 0$ is a fully income integrated POI signifying that the total time spent at that venue is evenly split across users from each of the four income quartiles. And where $S_\alpha = 1$ is a fully income segregated POI, signifying that it is only visited by users from one income quartile [2].

Results

We explore temporal differences and how they affect POI-level income segregation, Eq. (2), by comparing visits captured by mobility data to recorded time spent in the American Time Use Survey (ATUS) (Figure 1). We calculate a ratio (Eq. 1) between the two data sources, for every hour, income quartile, and POI category. First, we examine whether users from different income quartiles produced data in different ways throughout the day. We compare the cumulative time spent by ATUS respondents (Figure 1a) to cumulative time spent by mobility data users in the Boston metro-area (Figure 1b) at three categories of POIs. By quartile, mobility data has parallel curves, with the lowest income users spending the least amount of time and the highest income users spending the most. A similar pattern is seen in ATUS, although with more variation throughout the day. For all three POI types the mobility data curves are smoother compared to the ATUS curves. The greater fluctuations seen in ATUS, particularly for Grocery stores and Gyms are due to increased noise from fewer visits.

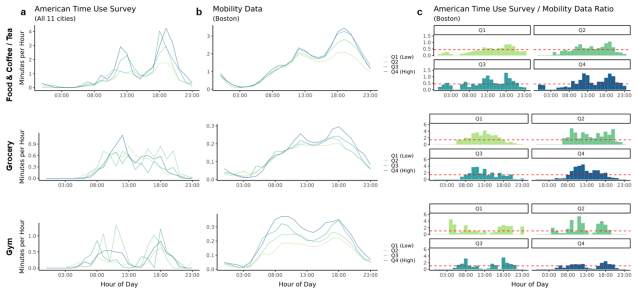

Fig. 1. Comparison of total time spent per hour shows temporal differences between quartiles in ATUS and mobility data. a. Cumulative time spent for the 2,502 ATUS respondents from 11 cities. **b.** Cumulative time captured from mobility data users in Boston. **c.** Ratio between ATUS and mobility data. Hours above the mean line (red dotted) show when visit type is under-represented in mobility data while hours below show over-representation. Hours without a ratio occur when no ATUS respondents from that quartile recorded spending time at that POI category during that hour.

We multiplied these ratios by the total minutes each quartile spent at each POI, to calculate original and weighted segregation measures. When they are averaged across POIs and across time by city, it resulted in a small, but frequently significant change between original and weighted income segregation (Figure 2). The absolute change in average weighted- and un-weighted income segregation across cities reveals the effect of differences in time-use recorded in each data source. We find the lowest difference in segregation in food POIs ($\Delta S_\alpha = 0.003$), and highest in gyms ($\Delta S_\alpha = 0.02$), with grocery showing relatively little change ($\Delta S_\alpha = 0.008$). These results are promising for the use of mobility data to measure income segregation, as temporal adjustments based on ATUS time-use produce marginal differences in segregation for Food and Grocery POIs.

Although the magnitude of the average changes in segregation were small across most POI categories and cities, the re-weighting provided significant adjustments to many of these city's segregation measures. For Food POIs, re-weighting income segregation resulted in a significant change in segregation for 5 cities (45%), while for Grocery stores and Gyms, significant changes occurred in 6 (55%) and 9 (82%) cities, respectively. Overall these results point to relative stability in estimates of place-based income segregation, as even significant differences in re-weighted income segregation produce a relatively small effect on the estimated segregation of places within a city.

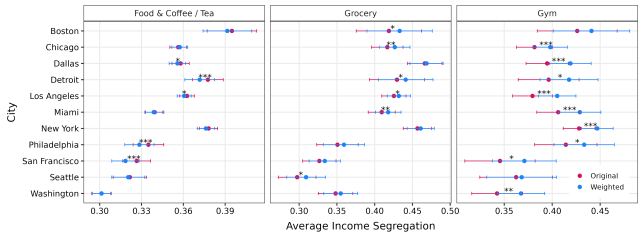

Fig. 2. Average income segregation decreases for Food establishments but increases for Grocery stores and Gyms. Comparing average income segregation by city before and after re-weighting with ATUS / MD ratios for each POI category. Significance codes: 0 '***' 0.001 '**' 0.01 '*' 0.05

As a means to understand the underlying factors which may influence differences between weighted and unweighted segregation, we ran a multiple linear regression analysis (Table 1). We assessed the influence of (1) number of POIs per census block group (CBG), a proxy for POI density which may cause misattribution of visits to POIs, (2) number of unique visitors to a POI, a proxy for POI popularity, and (3) the fraction of users by income quartile, proxies for whether POIs predominantly serve specific income segments, potentially indicating that businesses targeting particular ends of the income spectrum are especially sensitive to re-weighting.

Explanatory Variable	*Food & Coffee / Tea*		*Grocery*		*Gym*	
	Std. β	Std. Error	Std. β	Std. Error	Std. β	Std. Error
Segregation	-0.17 ***	0.01	-0.24 ***	0.02	-0.27 ***	0.02
# POIs per CBG	0.01	0	-0.01	0.02	0	0.02
# Unique visitors	-0.03 ***	0.01	-0.06 ***	0.02	-0.04 *	0.02
Fraction of Q1 visitors	-0.42 ***	0.01	-0.12 ***	0.02	-0.3 ***	0.02
Fraction of Q4 visitors	0.17 ***	0.01	0.1 ***	0.02	-0.44 ***	0.02
Model Results						
Conditional R^2	0.348		0.115		0.233	
Marginal R^2	0.346		0.115		0.233	

Table 1. Number of unique mobility data visitors along with the fraction of visitors from quartiles 1 (lower income) and 4 (higher income) have significant relationships with the change in segregation after re-weighting. Multiple linear regression results by POI category across all 11 cities. Significance codes: 0'***'0.001'**'0.01'*'0.05

Discussion

Behavioral biases need to be considered as they have the potential to skew our understanding of human behavior in cities, particularly when analyzing large-scale patterns such as experienced segregation. Our results show that, although aggregate measures of income segregation are relatively stable, re-weighting mobility data using external benchmarks, such as ATUS, can reveal significant shifts in specific places and income groups, underscoring the value of correcting for behavioral and temporal bias.

Our data and study have a few important limitations. By comparing and creating a ratio between the mobility data and ATUS, it implies that ATUS is a form of proxy of the ground truth, like any other time diary, its accuracy and reliability are dependent on what respondents report [9, 10]. Our ongoing work will focus on investigating and analyzing the changes in measured income segregation before and after re-weighting on an hourly basis.

REFERENCES

1. F. Xu, Q. Wang, E. Moro, *et al.*, "Using human mobility data to quantify experienced urban inequalities," Nat. Hum. Behav. pp. 1–11 (2025). Publisher: Nature Publishing Group.
2. E. Moro, D. Calacci, X. Dong, and A. Pentland, "Mobility patterns are associated with experienced income segregation in large US cities," Nat. Commun. **12**, 4633 (2021). Number: 1 Publisher: Nature Publishing Group.
3. T. Yabe, B. G. B. Bueno, X. Dong, *et al.*, "Behavioral changes during the COVID-19 pandemic decreased income diversity of urban encounters," Nat. Commun. **14**, 2310 (2023). Number: 1 Publisher: Nature Publishing Group.
4. S. Athey, B. Ferguson, M. Gentzkow, and T. Schmidt, "Estimating experienced racial segregation in US cities using large-scale GPS data," Proc. National Acad. Sci. **118**, e2026160118 (2021). Publisher: Proceedings of the National Academy of Sciences.
5. M. Salganik, *Bit By Bit: Social Research in the Digital Age* (Princeton University Press, 2017).
6. F. Keusch, A. Wenz, and F. Conrad, "Do you have your smartphone with you? Behavioral barriers for measuring everyday activities with smartphone sensors," Comput. Hum. Behav. **127**, 107054 (2022).
7. S. Bähr, G.-C. Haas, F. Keusch, *et al.*, "Missing Data and Other Measurement Quality Issues in Mobile Geolocation Sensor Data," Soc. Sci. Comput. Rev. **40**, 212 (2022). Publisher: SAGE Publications Inc.
8. U.S. Bureau of Labor and Statistics, "American Time Use Survey: Overview," https://www.bls.gov/opub/hom/atus/home.htm.
9. J. Gershuny, T. Harms, A. Doherty, *et al.*, "Testing self-report time-use diaries against objective instruments in real time," Sociol. Methodol. **50**, 318 (2020).
10. P. A. Phipps and M. K. Vernon, *Calendar and Time Diary* (SAGE Publications, Inc., 2009), chap. Twenty-Four Hours: An Overview of the Recall Diary Method and Data Quality in the American Time Use Survey, pp. 108–128.

Towards Assessing Accessibility Resulting from Integrating Demand-Responsive and Conventional Public Transport with Travel Time Uncertainty

OURAHOU MOHAMED[1,*], ARALDO ANDREA[1], ZIGRAND LOUIS[2], AND VIANA ALINE CARNEIRO[3]

[1] Télécom SudParis, Institut Polytechnique de Paris, 19 Place Marguerite Perey, 91120 Palaiseau, France
[2] Research and Development Team, Padam Mobility, 11 Tronchet Street, 75008 Paris, France
[3] INRIA Paris-Saclay, 1 Honoré d'Estienne d'Orves street, 91120 Palaiseau, France
[*] mohamed.ourahou@telecom-sudparis.eu

1. INTRODUCTION

Transportation networks provide people with the means to reach opportunities (i.e., accessibility), e.g., schools and workplaces [1, 2]. Conventional Public Transportation (CPT), e.g., fixed-route buses and trains, is most effective in densely populated areas, where it offers extensive and affordable coverage [3]. In low-density areas (e.g., suburbs), CPT provides insufficient accessibility, as few opportunities can be reached within reasonable travel times. Operators limit frequency and coverage to prevent high costs per passenger [4], leaving such areas car-dependent, with private vehicles often the only viable mode [5]. Demand Responsive Transportation (DRT) offers a flexible alternative, as routes and schedules adapt to users' requests. With growing interest in integrating DRT with CPT, it is essential to evaluate its impact on accessibility, particularly when DRT operates as stop-to-stop service, both as a feeder and as a standalone mode. Yet, rigorous methods to assess accessibility under these conditions remain lacking. To address this gap, we propose a method that accounts for travel time uncertainty in DRT, rather than relying on deterministic or oversimplified assumptions.

2. METHODOLOGY

A. Accessibility

We use an isochrone-based measure to quantify how many additional opportunities individuals can access within a given time threshold due to the introduction of DRT. Accessibility is computed by calculating the shortest travel time from each location to all others using available transport modes (DRT, CPT, walking), and counting the opportunities reachable within the threshold. The study area is divided into hexagons, with centroids $C \subseteq \mathbb{R}^2$ used as trip origins and destinations. The accessibility measure is expressed as:

$$A_u(t) = \sum_{v \in C} O_v \cdot \mathbf{1}_{T_{uv}(t) \leq \tau} \quad (1)$$

where $A_u(t)$ is the accessibility of hexagon u at time $t \in [00{:}00, 24{:}00]$, O_v is the opportunities in hexagon v, $T_{uv}(t)$ is the shortest travel time from u to v, departing at time t, τ is the travel time threshold, and $\mathbf{1}_{T_{uv}(t) \leq \tau}$ is an indicator function that equals 1 if $T_{uv}(t) \leq \tau$, and 0 otherwise.

B. Travel time prediction

$T_{uv}(t)$ is a random variable because it depends on factors such as detours imposed by other passengers, which may change from day to day. Therefore, $A_u(t)$ is also a random variable. The distribution of $T_{uv}(t)$ can be estimated using ML models trained on observed trips.

Let $\hat{A}_u(t)$ and $\hat{T}_{uv}(t)$ be the estimated travel time and accessibility. Via such models we extrapolate $T_{uv}(t)$ for (u,v) pairs with no observed trips; ML models exploit spatial similarity to do so. To account for uncertainty, we do not rely on a single deterministic value of $\hat{T}_{uv}(t)$; instead, we encode percentiles of its distribution. In this study, we consider only the 95th percentile (P95). Given a pair (u,v) and departure time t, the ML model is a function $f(u,v,t) = [\hat{\mu}_{uv}(t), \hat{\sigma}_{uv}(t)]$, the estimated mean and standard deviation of $T_{uv}(t)$. We compare three ML models: Linear Regression (LR), Random Forest Regressor (RFR), and Bayesian Neural Network (BNN). We tried three probability distributions—Normal, Gamma, and Lognormal— to compute percentiles and assessed point prediction accuracy and uncertainty quantification. We used Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) for accuracy, and Negative Log-Likelihood (NLL) for probabilistic estimate quality. The best-performing model was selected for the rest of the study.

C. Unified modeling of demand responsive Transportation with conventional public transportation

Accessibility is typically calculated using a graph structure, where shortest paths can be computed and travel times to opportunities are the time of those paths. To account for the dynamics of transport networks, the time dimension is incorporated into this graph, resulting in a Time-Expanded Graph (TEG) [6]. In a TEG, nodes are defined as *stoptimes* $(u, t_{\text{arr}}, t_{\text{dep}})$, representing arrival and departure of a vehicle at location $u \in \mathbb{R}^2$ at times $t_{\text{arr}}, t_{\text{dep}} \in \mathbb{R}$, respectively. For CPT lines with scheduled trips, it is feasible to use a TEG to compute accessibility, since *stoptimes* are scheduled in advance. Each CPT trip is represented by a sequence of *stoptimes* connected by edges. However, no scheduled stoptimes exist in DRT and such a graph cannot be computed. To address this, we adopt an approach similar to [7], augmenting the CPT TEG with "virtual lines". These represent potential DRT trips not necessarily observed in the past. While [7] model DRT solely as a feeder service—connecting areas to selected CPT stops—our approach generalizes this by also supporting complete DRT origin-to-destination journeys, without requiring transfer to CPT. We enrich the CPT TEG G with virtual DRT lines to enable centroid-to-centroid multimodal journeys (Figure 1). Virtual lines can start or end in CPT stops or in "DRT barycenters" in each hexagon. A DRT barycenter is the barycenter of pick-up or drop-off locations in the historical dataset (B1–B4 in Figure 1). For hexagons without observed pick-up or drop-off, we do not add virtual lines, restricting DRT modeling to areas with observed demand. The virtual lines enable five types of multimodal trips between origin u and destination v, as shown in Figure1 1. All walking connections between centroids, barycenters, and CPT stops are handled by CityChrone [8]. We consider that all trips start and end with walking, reflecting the fact that users need always to walk a short distance to reach CPT or DRT services, or to arrive at the destination centroid. We also model line changes between services (e.g., CPT–CPT or CPT–DRT) via walking links, also managed by CityChrone.

3. RESULTS

A. Study area and data

We evaluate our methodology on three French regions and present results for Orléans Métropole. We use a large real world historical DRT

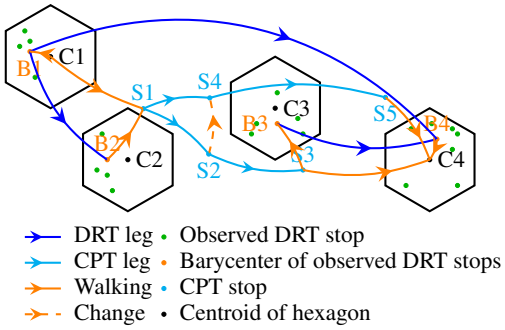

Fig. 1. Illustration of DRT integration with CPT

trip dataset (over 1 million trips, 2018–2024), collected by vehicles operated by [9]. Additionally, we incorporate publicly available datasets, including the 2025 CPT GTFS, street networks from OpenStreetMap (OSM) and locations of opportunities extracted from OSM.

B. Prediction of travel time distributions

Table 1 presents the evaluation scores for both point prediction accuracy (using MAE and RMSE) and probabilistic performance (using NLL under three distributions: Lognormal, Normal, and Gamma). RFR achieves the best point prediction accuracy overall, with the lowest MAE and RMSE in most cases. In Lille, for instance, RFR yields an MAE of 3.80 min and RMSE of 4.43 min, slightly outperforming BNN. BNN, however, records the lowest RMSE in Lille (3.96), and shows competitive MAE in other regions as well. The differences in MAE and RMSE between RFR and BNN are small and may be negligible in practice. For uncertainty quantification, BNN significantly outperforms both RFR and LR. Under the Normal assumption, BNN obtains the lowest NLL across all territories (e.g., 0.71 in Lille).

Given BNN's competitive point prediction performance and its clear advantage in uncertainty estimation, we select it as the final model for estimating travel time percentiles used to enrich the CPT TEG with virtual lines.

Table 1. Performance comparison of travel time distribution prediction models across three territories

Territory	Model	MAE (min)	RMSE (min)	NLL		
				Lognorm.	Normal	Gamma
Lille M.	LR	4.41	5.95	39.21	32.49	7.29
	RFR	**3.80**	4.43	117.82	142.46	7.49
	BNN	3.99	**3.96**	7.24	**0.71**	7.16
Orléans M.	LR	3.31	4.58	38.96	32.71	6.82
	RFR	**2.75**	**3.87**	108.50	125.57	6.80
	BNN	2.87	3.97	6.68	**0.65**	6.69
Yvelines	LR	3.29	4.31	51.88	45.45	6.86
	RFR	2.77	**3.85**	174.76	192.02	6.94
	BNN	**2.71**	3.98	6.66	**0.50**	6.68

C. Accessibility evaluation

The accessibility of each hexagon was calculated on the weekday of April 25, 2024, using a maximum travel time threshold of 30 minutes. As explained in the methodology, this was done by simulating a trip from that hexagon to every other hexagon in the region, departing at 7:00 a.m. The travel time threshold was based on the average travel time from peripheral areas to the center of Orleans, making it an appropriate cutoff for evaluating regional accessibility.

Figures 2 demonstrate the impact of integrating DRT with CPT. In the figures, the CPT-only map (left) reveal substantial areas with low accessibility scores (light blue) in peripheral regions. The combined CPT+DRT map (center) show markedly improved accessibility scores across almost the entire territory, with some areas reaching dark blue (high accessibility scores > 200). The relative improvement maps (right) quantify these enhancements with percentage differences calculated as (CPT+DRT - CPT)/CPT × 100. The green hexagons, represent areas newly made accessible by DRT that were previously unserved by public Transportation (No access to CPT within a considered maximum walk time of 15 Minutes).

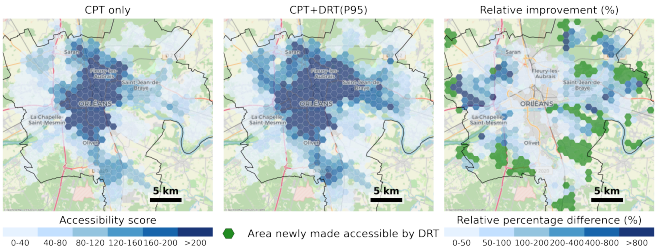

Fig. 2. Accessibility scores and relative percentage improvement between CPT only and P95 (CPT+DRT) scenarios

4. CONCLUSION

In this study, we first investigated the impact of DRT on accessibility using the proposed method. We also examined how travel time variability in DRT affects accessibility by using probabilistic machine learning models across three French territories. Our results demonstrate that the BNN with a Normal distribution provides the best overall performance, effectively balancing point prediction accuracy and the ability to model travel time uncertainty. However, our study has several limitations that suggest directions for future research. We assumed unrestricted availability of DRT vehicles, without considering fleet size constraints, vehicle capacity, or operational costs.

ACKNOWLEDGEMENTS

This work is part of the Mob Sci-Dat Factory project (ANR-23-PEMO-0004) under the France 2030 program. It was funded by Région Île-de-France and partially supported by the French ANR research project MuTAS (ANR-21-CE22-0025-01). The data were provided by Padam Mobility.

REFERENCES

1. E. Miller, OECD (2020).
2. W. G. Hansen, Journal American Planning Association **25** (1959).
3. S. Alamri, K. Adhinugraha, N. Allheeib, and D. Taniar, ISPRS International Journal Geo-Information **12**, 180 (2023).
4. A. Badeanlou, A. Araldo, and M. Diana, "Assessing transportation accessibility equity via open data," in *hEART 2022 - 10th Symposium of the European Association for Research in Transportation,* (2022).
5. P. Saeidizand, K. Fransen, and K. Boussauw, Cities (2022).
6. C. L. Hasif, A. Araldo, S. Dumbrava, and D. Watel, "A graph-database approach to assess the impact of demand-responsive services on public transit accessibility," in *15th ACM SIGSPATIAL International Workshop on Computational Transportation Science,* (ACM, 2022).
7. S. Diepolder, A. Araldo, T. Chouaki, *et al.*, "On the computation of accessibility provided by shared mobility," in *11th Symposium of the European Association for Research in Transportation (hEART),* (Zurich, Switzerland, 2023). https://hal.archives-ouvertes.fr/hal-04154712v2.
8. I. Biazzo, B. Monechi, and V. Loreto, Royal Society Open Science **6**, 191014 (2019).
9. Padam Mobility, "Use case," https://www.padam-mobility.com/ (2022).

Beyond Proximity: Rethinking Urban Accessibility Through Human Mobility Patterns in United States

SAUMITRA KULKARNI[1 *] AND ESTEBAN MORO[1,2 **]

[1] Network Science Institute, Northeastern University, Boston, MA 02115, USA
[2] Department of Physics, Northeastern University, Boston, MA 02115, USA
[*] kulkarni.saumi@northeastern.edu
[**] e.moroegido@northeastern.edu

Urban accessibility is typically measured by proximity to home, but this overlooks how people navigate cities and use their time. Using mobility data from 1.86 million individuals across 11 U.S. metro areas, we show that access to daily necessities is structured around a small set of key locations—containers—such as home, work, and other routine destinations. We introduce the offer-meets-demand ratio to measure alignment between available POIs to visit and actual visited POIs. Most individuals rely on multiple containers, with the second container often linked to job locations. Our findings highlight a dual nature of accessibility and call for urban policies grounded in human behavior.

1. INTRODUCTION

Traditional measures of accessibility to daily necessities in cities are based on proximity, either in terms of physical distance or time, from home. Planning and policy-making paradigms like the "15-minute-city" [1] or "food-desserts" [2] are based on this notion, assuming that close access to home ensures effective use. However, such models often overlook how people actually organize their routines and daily necessities across urban spaces. For example, only 14% of daily consumption trips in the US happen within 15-minute walking distance from home [3], and residents in large urban areas in the US travel a median of 6.7km from home to consume fast food [2]. In this study, we use high-resolution mobility data to investigate whether and how people access daily necessities beyond their immediate residential environment, and which locations actually serve as anchors for their daily routines.

2. DATA AND METHODS

In this study, we use individual-level mobility data from 1.86 million anonymized smartphones across 11 U.S. urban areas over a six-month period (October 2016 - March 2017). We identified visits to daily necessities by extracting the stays of people from mobility data around a large collection of points of interest (POI) obtained from Foursquare. We only focus on eight types of POIs as daily necessities that belong to the following categories: Banks, Coffee Shops, Grocery Stores, Gyms, Laundries, Pharmacies, Post Offices, and Salons.

We define the offers available to an individual as the set of the k-nearest POIs of a given category surrounding specific locations—referred to as containers (here we use $k = 10$) [4, 5]. Containers are locations that play a meaningful role in an individual's daily routine. We identify them by ranking all visited locations based on the total time an individual spends there, with the home location always designated as the primary (1^{st}) container. Subsequent containers may include workplaces, shopping centers, or other visited places in the city (see Fig. 1.A). To quantify how well these containers and their associated offers account for actual behavior, we define the offer-meets-demand ratio (ϕ^n) as the proportion of visits made by an individual to POIs within the combined offer sets of their first n containers, relative to their total visits to that POI category. A value of $\phi^n = 1$ indicates that all of an individual's visits are to nearby offers, while lower values suggest that many visits occur outside these offers. The case ϕ^1 corresponds to traditional accessibility measured only around home. We use the 2016-2020 5-year American Community Survey (ACS) to obtain population estimates, and the LEHD Origin-Destination Employment Statistics (LODES) data on the number of jobs at a census block group (CBG) level across 11 metro areas under consideration.

3. RESULTS AND DISCUSSION

We begin by examining whether individuals actually visit the POIs available near their homes. While the median distance to the nearest POI of a specific category from home is just 700 meters (IQR: 0.39–1.36 km), the median distance people travel to visit a POI of the same category is significantly greater—about 7 km (IQR: 3.07–15.19 km; see Fig.1.B). This indicates that, on average, individuals routinely bypass nearby options for their daily activities. To explore this further, we compute the offer-meets-demand ratio ϕ^1 around the 1^{st} container (home) for each individual. Under the assumption that individuals spend most of their time at home, we would expect high ϕ^1 values if proximity from home was the primary driver of behavior. Instead, we find three distinct groups: those who exclusively visit POIs near home ($\phi^1 = 1$, around 10% of them), those who never do ($\phi^1 = 0$, almost 45% of them), and a broad spectrum of individuals in between ($0 < \phi^1 < 1$, the rest; see Fig.1.C).

If not around home, then where do people do most of their daily activities? We define an individual's 2^{nd} container as the

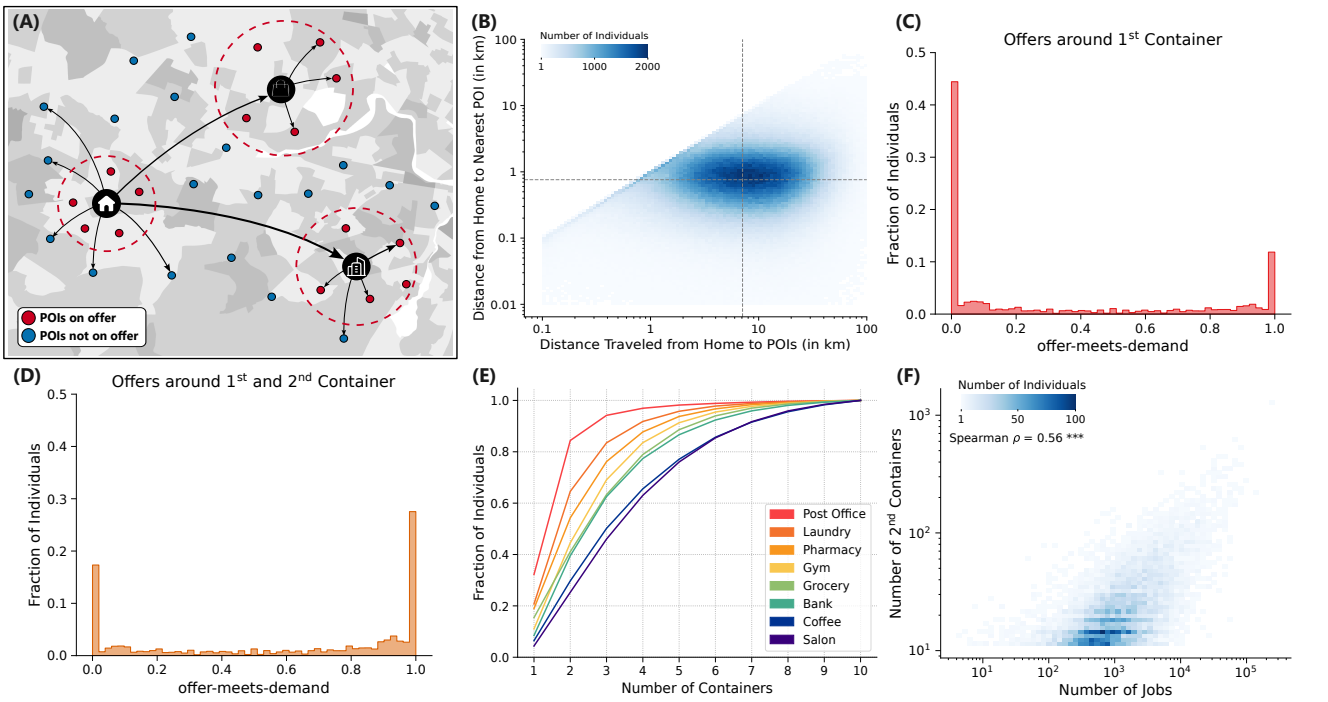

Fig. 1. **(A)** Illustration of the concept of offers around different containers ($k = 5$ and $n = 3$ here) for an individual. Blue points represent all POIs of a specific category in the city, while red points indicate POIs included in the offer set around each container. The offer-meets-demand ratio ϕ^n measures the proportion of visits to these offers out of total visits. Thickness of arrows roughly indicate frequency of visits. **(B)** Heatmap showing the distance from individuals' homes to the nearest POIs of categories under study versus the actual distance traveled to visit POIs of same categories. Grey dotted lines indicate median values. **(C)** Distribution of offer-meets-demand ϕ^1 for individuals with only offers around their 1st container (i.e. homes). **(D)** Distribution of offer-meets-demand ϕ^2 for same individuals with offers around their 1st containers and 2nd containers (i.e. place where they spend most of their time outside home). **(E)** Fraction of individuals with $\phi^n = 1$ as a function of number of containers included, shown separately for different POI categories. **(F)** Relationship between the number of 2nd containers and the number of jobs in each census block group.

location where they spend the most time outside of home. When we include the 10 nearest POIs around this 2nd container to the set of offers, the offer-meets-demand ratio ϕ^2 increases, and for many individuals, it approaches 1 (see Fig. 1.D). This suggests that individuals who do not visit nearby POIs around home often access them near other (next) key locations in their routines. We extend this analysis by progressively adding offers around each individual's ranked containers and computing ϕ^n for increasing values of n. As shown in Fig. 1.E, the fraction of individuals with $\phi^n = 1$ grows with the number of containers, but the rate of increase depends on the POI category. For instance, only 3 containers are needed to capture $\geq 75\%$ of individuals visiting Post Offices, Laundries, or Pharmacies—activities with relatively consistent patterns. In contrast, categories like Salons and Coffee Shops require 5 containers to reach similar coverage, reflecting more varied behavior. Grocery Stores, Gyms, and Banks fall in between, typically requiring 4 containers.

What are the containers that people organize their time around? Our results show that access to daily necessities is shaped not just by residential location but by a small set of recurring places in people's routines. For example, people seem to shop for groceries around four different containers (on average). The nature of these containers varies across individual lifestyles, but common patterns emerge. Notably, areas that contain a large number of individuals' 2nd containers tend to have higher concentrations of jobs (see Fig. 1.F), indicating that workplaces are key anchors in organizing access to daily necessities beyond the home. Ongoing work investigates the demographic and urban drivers of this behavior and whether individuals reorganize their offer-meets-demand in response to both exogenous shocks, such as crisis response policies, and endogenous shocks, such as changes in home or work location.

Our results reveal a dual nature of accessibility in cities—shaped by both where people live and where they routinely spend their time. Measuring access solely from home overlooks this important behavioral aspect. Importantly, identifying where these containers cluster across the city can inform decisions on where to deploy services, amenities, or policy interventions to maximize their impact. By incorporating the actual patterns of people's routines into accessibility metrics, we move toward more realistic, equitable, and effective urban planning.

REFERENCES

1. C. Moreno, Z. Allam, D. Chabaud, *et al.*, Smart cities **4**, 93 (2021).
2. B. García Bulle Bueno, A. L. Horn, B. M. Bell, *et al.*, Nat. Commun. **15**, 2291 (2024).
3. T. Abbiasov, C. Heine, S. Sabouri, *et al.*, Nat. Hum. Behav. **8**, 445 (2024).
4. L. Alessandretti, U. Aslak, and S. Lehmann, Nature **587**, 402 (2020).
5. L. Pappalardo, F. Simini, S. Rinzivillo, *et al.*, Nat. Commun. **6**, 8166 (2015).

Understanding Urban-Rural Disparities in Mobility Inefficiency for Colombia, Mexico, and India

NANDINI IYER[1,*], MASSIMILIANO LUCA[2], AND RICCARDO DI CLEMENTE[1,3]

[1] Complex Connections Lab, Network Science Institute, Northeastern University London, UK.
[2] Bruno Kessler Foundation, Via Sommarive 18, 38123 Povo (TN), Italy.
[3] ISI Foundation, Via Chisola 5, Turin 10126, Italy.
[*] nandini.iyer@nulondon.ac.uk

As rates of urbanisation increase, data-driven methods offer timely insights that can complement local knowledge, enabling more responsive and equitable decisions around access to resources and opportunities in both urban and rural areas.[1, 2]. However, urbanisation patterns differ between the Global South (characterised by increase in smaller urban settlements) and the Global North (driven by megacity growth) [3, 4]. This highlights the need for a unified framework to analyse mobility disparities between urban and rural areas, accounting for different forms of urbanisation. Previous studies often rely on static boundaries based on population density or built environment contiguity to define urban and rural boundaries [5], but these fail to capture interdependencies between rural and urban regions. Additionally, recent works generally compare urban and rural areas from the perspective of migration [6, 7], leaving a significant gap in understanding the emergence of inequalities in smaller-scale mobility (i.e. routine travel rather than residential migration). This study addresses these limitations by investigating mobility inefficiency inequalities in Colombia, India, and Mexico. Employing aggregated mobility data for three Low and Middle-Income Countries, we explore how these inequalities manifest differently between urban and rural areas.

To define urban and rural boundaries, we apply network percolation to a country's mobility network, which has been constructed from an origin-destination matrix. The origin-destination matrix is an aggregated version of Location Based Services (LBS) data from Cuebiq, assembled by the World Bank [8] at a 25 km^2 spatial scale from November 1, 2019 to December 31, 2019. In doing so, we derive a rural-urban classification that is based on how people move through space, rather than solely relying on static characteristics such as population density. Building on this definition, we examine mobility inefficiency (ϕ), which highlights the discrepancy between ideally efficient travel and real-world mobility patterns that can be impacted by congestion, limited infrastructure, or travel preferences. To estimate mobility inefficiency, we use stochastic methods to determine the average optimal shortest path length, comparing it to empirical trip lengths that are provided by the data [8]. Our findings show that, in each country, rural areas generally exhibit higher mobility inefficiency compared to urban areas, particularly for longer trips, later in the day (Fig 1).

The key contribution of this work lies in interpreting mobility inefficiency inequalities in the context of its rural and urban environment. In urban areas, we determine four distinct patterns (Figure 2), in terms of how aligned public transit access is with mobility efficiency throughout the week. Transit access is defined as the average between the coverage area and the density of transit stops, which are identified using OpenStreetMap. The coverage area is the ratio of the convex hull area and the geohash area, where the convex hull includes all transit infrastructure in a geohash. To portray the misalignment between transit access and mobility efficiency in different cities, we construct transit access

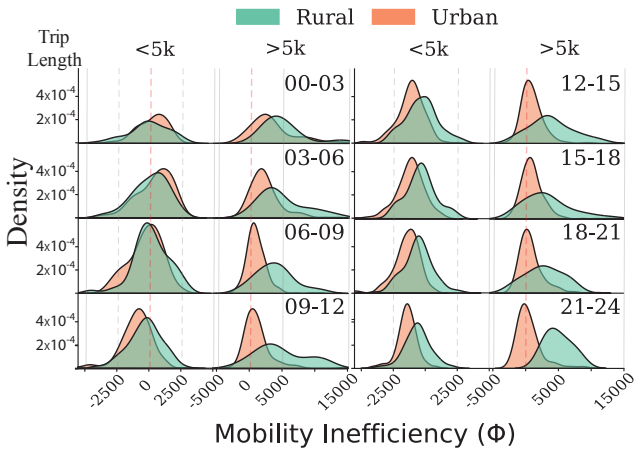

Fig. 1. Mobility inefficiency for short (< 5km) and long (> 5km) trips in Colombia, in urban (orange) and rural (green) areas at different time intervals. Higher values of ϕ indicate longer, less efficient trips. This finding hold true across all three countries.

($\mathbf{F}^{PT}$) and mobility efficiency ($\mathbf{F}_t^{-\phi}$) vector fields using the PT score and the negated mobility inefficiency metric, respectively. It is important to note that the mobility efficiency vector field is time-dependent while the transit access one remains static.

We calculate the force between a geohash, g, and each of its eight neighbours, representing the interaction strength of adjacent geohashes based on either PT score or mobility inefficiency. For each city, we compute the resultant transit access vector and efficiency vector, for all geohashes in a city, to obtain a final vector field for transit access ($\mathbf{F}^{PT}$) and mobility efficiency ($\mathbf{F}_t^{-\phi}$). Finally, we compare transit access and efficiency by computing the average angular difference between each of the geohashes in an urban area. We can then observe how each of the 15 analysed cities transit-efficiency alignment changes throughout the weekdays and weekends (Figure 2).

Rural areas are often overlooked in human mobility studies, underscoring the need for a quantitative approach to understand which types of rural areas experience particularly high levels of inequality, such as mobility inefficiency. Since rural regions exhibit much lower rates of transit access coverage and mobility data, the vector field approach used in urban settings would not be relevant in the rural context. We highlight the highly contextual nature of improving rural-urban connectivity, as factors such as city size, sprawl, natural barriers influence where rural areas can emerge and how well connected they are to nearby urban areas. Our findings reveal spatial disparities in mobility inefficiency that are unique to rural areas surrounding each city, using Ordinary Least Squares (OLS) regression to analyse the

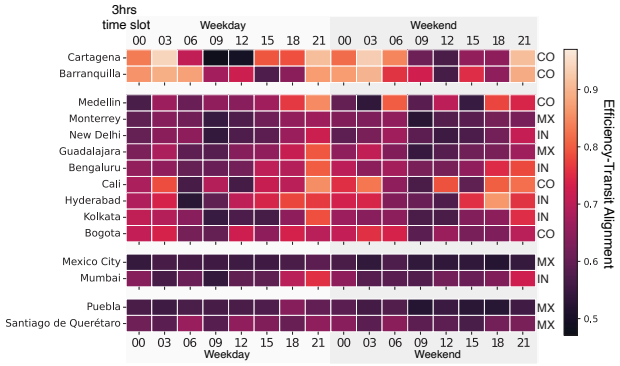

Fig. 2. Alignment between urban mobility efficiency and transit access for 15 different cities (rows) at different time intervals (columns). Cities are split into clusters using spectral clustering to understand overall trends. Cities in Colombia tend to exhibit more alignment between transit access and mobility efficiency, while three of the five Mexican cities have low alignment for the majority of weekdays and weekends.

relationship between various features and inefficiency. These features include distance from each city's centroid, population of the rural area, orientation, and its proximity to cities. The orientation is decomposed into its north-south and east-west distance, where positive values indicate a relationship between inefficiency and northern and eastern rural areas, respectively. The proximity to cities represents the absolute difference in distance between a rural area and its two nearest cities. It quantifies how much closer or farther the rural area is to one city compared to the other. This difference could provide insight into the area's relative isolation or proximity to urban hubs, which might affect its access to resources, services, or infrastructure.

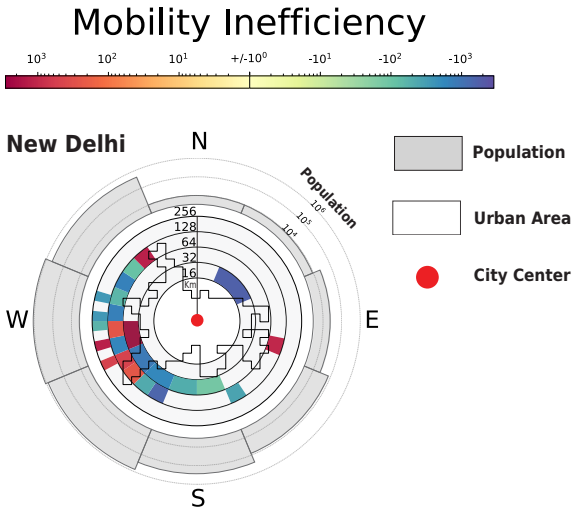

Fig. 3. Mobility inefficiency for rural populations surrounding New Delhi, conveying orientation around the city, distance from the urban centre (concentric rings), and rural population (outer grey bars)

Figure 3 focuses on rural areas that surround New Delhi (IN) and their respective features where orientation is shown around the circle, distance from the city centroid is expressed through concentric circles, light grey bars depict population. In the case of New Delhi (Fig 3), rural inefficiency is related to larger populations and east-west orientation, with areas further west experiencing higher inefficiency. Our findings for rural areas closest to the 15 cities in Figure 2 suggest that factors associated with orientation, such as natural barriers and socioeconomic disparities, population, and proximity to city centre, influence rural mobility inefficiency to different extents.

This study introduces a unified framework for understanding mobility inequality across urban and rural contexts, tailored to settings where mobility and infrastructure data are limited. In urban areas, we quantify how well public transit aligns with efficient mobility patterns, while in rural regions, we identify how factors such as population, orientation, and proximity to the nearest cities shape inefficiency. By grounding our analysis in observed movement data rather than relying on static definitions or complete transit datasets, such as GTFS, which remains scarce in many LMICs, we offer a scalable approach for identifying where and how mobility inequalities emerge. This enables more targeted, context-aware interventions to improve connectivity and access across diverse geographies.

REFERENCES

1. P. de Souza, S. Anenberg, C. Makarewicz, *et al.*, Environ. science & technology **58**, 280 (2023).
2. T. Abbiasov, C. Heine, S. Sabouri, *et al.*, Nat. Hum. Behav. **8**, 445 (2024).
3. T. Edensor and M. Jayne, "Introduction: urban theory beyond the west," in *Urban theory beyond the West*, (Routledge, 2012), pp. 18–44.
4. G. F. Randolph and C. Deuskar, J. Plan. Educ. Res. **44**, 279 (2024).
5. B. Chen, S. Wu, Y. Song, *et al.*, Nat. Commun. **13**, 4636 (2022).
6. H. Selod and F. Shilpi, Reg. Sci. Urban Econ. **91**, 103713 (2021).
7. M. Busso, J. P. Chauvin, and N. Herrera, Reg. Sci. Urban Econ. **91**, 103658 (2021).
8. W. Zhang, M. N. del Prado, V. Gauthier, and S. Milusheva, arXiv preprint arXiv:2410.00453 (2024).

Potential route diversification in road networks

GIULIANO CORNACCHIA[1,2,*], LUCA PAPPALARDO[1,3], MIRCO NANNI[1], DINO PEDRESCHI[2], AND MARTA C. GONZÁLEZ[4,5,6]

[1] ISTI-CNR, Pisa, Italy
[2] University of Pisa, Pisa, Italy
[3] Scuola Normale Superiore of Pisa, Italy
[4] Department of City and Regional Planning, University of California, Berkeley, CA, USA
[5] Energy Technologies Area, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
[6] Department of Civil and Environmental Engineering, University of California, Berkeley, CA, USA
[*] giuliano.cornacchia@isti.cnr.it

The structure of road networks shapes urban dynamics, influencing traffic congestion and urban sustainability. Recent studies reveal inefficiencies in road network usage: 98% of roads are underutilized [1], 16% of highway-based routes have faster alternatives [2], and traffic is typically concentrated into a few major corridors [3]. Route diversification, which disperses traffic across multiple paths, offers a promising solution to reduce congestion and improve overall network efficiency, yet its effectiveness depends on the road network's structure. Despite its importance, no metric currently quantifies a road network's capacity for route diversification. While several measures, such as detour index [4], inness [5], and route redundancy [6], capture aspects of network efficiency or resilience, they typically overlook either the multiplicity of alternative routes or the extent of their spatial overlap. To this end, we introduce DiverCity, a measure that quantifies the extent to which traffic can potentially be distributed across multiple minimally overlapping routes. The DiverCity of a trip (u, v) from origin u to destination v is defined as $\mathcal{D}(u, v) = \mathcal{S}(\text{NSR}(u, v)) \cdot |\text{NSR}(u, v)|$, where $\text{NSR}(u, v)$ is the set of near-shortest routes, i.e., routes no more than 30% longer than the shortest path, ensuring they remain practical for drivers, and $\mathcal{S}(\text{NSR}(u, v))$ measures their spatial spread as the average pairwise weighted Jaccard similarity. To identify the near-shortest routes, we generate up to $k = 10$ alternative routes using path penalization [7]. DiverCity ranges from 0 (a single route) to k (k disjoint near-shortest routes). Figures 1a–b illustrate low and high DiverCity examples.

Results. We apply DiverCity to 56 cities worldwide, each within a 30 km radius from the city center, spanning all six inhabited continents, varying population densities, and diverse road network structures. Using a radial sampling method, we systematically generate trips that cover various distances from city centers to peripheries. Each city's DiverCity value is derived as the median of the trip-level DiverCity values.

DiverCity reveals substantial variability (see Figure 1c). Tokyo exhibits the highest route diversification, while Mumbai has the lowest (38% lower). Notably, DiverCity is strongly linked to average edge load ($r = -0.644$), suggesting that cities with higher DiverCity achieve more balanced traffic distribution (see Figure 1d). DiverCity also correlates with key road network characteristics: cities with larger networks, more intersections, and more direct connections exhibit greater route diversification. Additionally, grid-structured cities exhibit higher DiverCity than non-gridded ones.

Within cities, DiverCity reveals a universal pattern: it is low in the city center, with an average value of 5 within the first kilometer (Figure 1e), and increases sharply with distance from the center, stabilizing at around 7.5 beyond 10 km. Interestingly, DiverCity declines in proximity to mobility attractors (e.g., highways and ring roads), which concentrate traffic flow due to their convenience, thereby restricting alternative route options and reducing DiverCity for nearby trips. Locations in the bottom 10% of DiverCity relative to their city's distribution are, on average, 1.33 km from the nearest mobility attractor, closer than the global average of 2 km and the 2.56 km for the top 10%. However, when strategically spaced and well-distributed, mobility attractors may enhance DiverCity at a city-wide scale. Cities with higher DiverCity have a higher density of attractors and more evenly distributed attractors compared to cities with lower DiverCity.

Speed limit tuning. To counteract the suppressive effects of mobility attractors on DiverCity, we propose speed limit tuning as a targeted intervention. Simulations show that moderate speed reductions ($\leq$50%) on mobility attractors significantly increase route diversification with minimal impact on travel times. Further reductions provide diminishing returns as attractors lose their competitive advantage. The diminishing returns occur because when the attractor's speed drops below that of the surrounding roads, it ceases to function as a shortcut, leading to a natural redistribution of drivers. This pattern is consistent in most cities, although the extent of improvement can vary based on the network's layout and the initial distribution of alternative routes.

Controlled Simulations. We test the robustness and generality of our findings through controlled simulations on a grid-based synthetic city with a 30 km radius, matching real-world cases. Within this framework, we vary the spatial distribution and speed limits of mobility attractors to analyze their impact on DiverCity. Our results closely align with real-world observations, confirming both the local and global effects of mobility attractors and demonstrating the effectiveness of speed limit reductions in mitigating their suppressive impact.

DiverCity relies exclusively on road network data, making it applicable to any city. Our study equips city planners, policymakers, and transportation authorities with valuable tools to promote more efficient road network use. By quantifying route diversification and identifying critical areas for intervention, DiverCity can guide infrastructure investments by pinpointing areas that would benefit from new road connections or speed limit adjustments. We provide an interactive platform (https://divercitymaps.github.io) to explore DiverCity across all considered cities.

REFERENCES

1. P. Wang, T. Hunter, A. M. Bayen, *et al.*, "Understanding road usage patterns in urban areas," Sci. Reports **2**, 1001 (2012).
2. P. Wang, L. Liu, X. Li, *et al.*, "Empirical study of long-range connections in a road network offers new ingredient for navigation optimization models," New J. Phys. **16**, 13012 (2014).

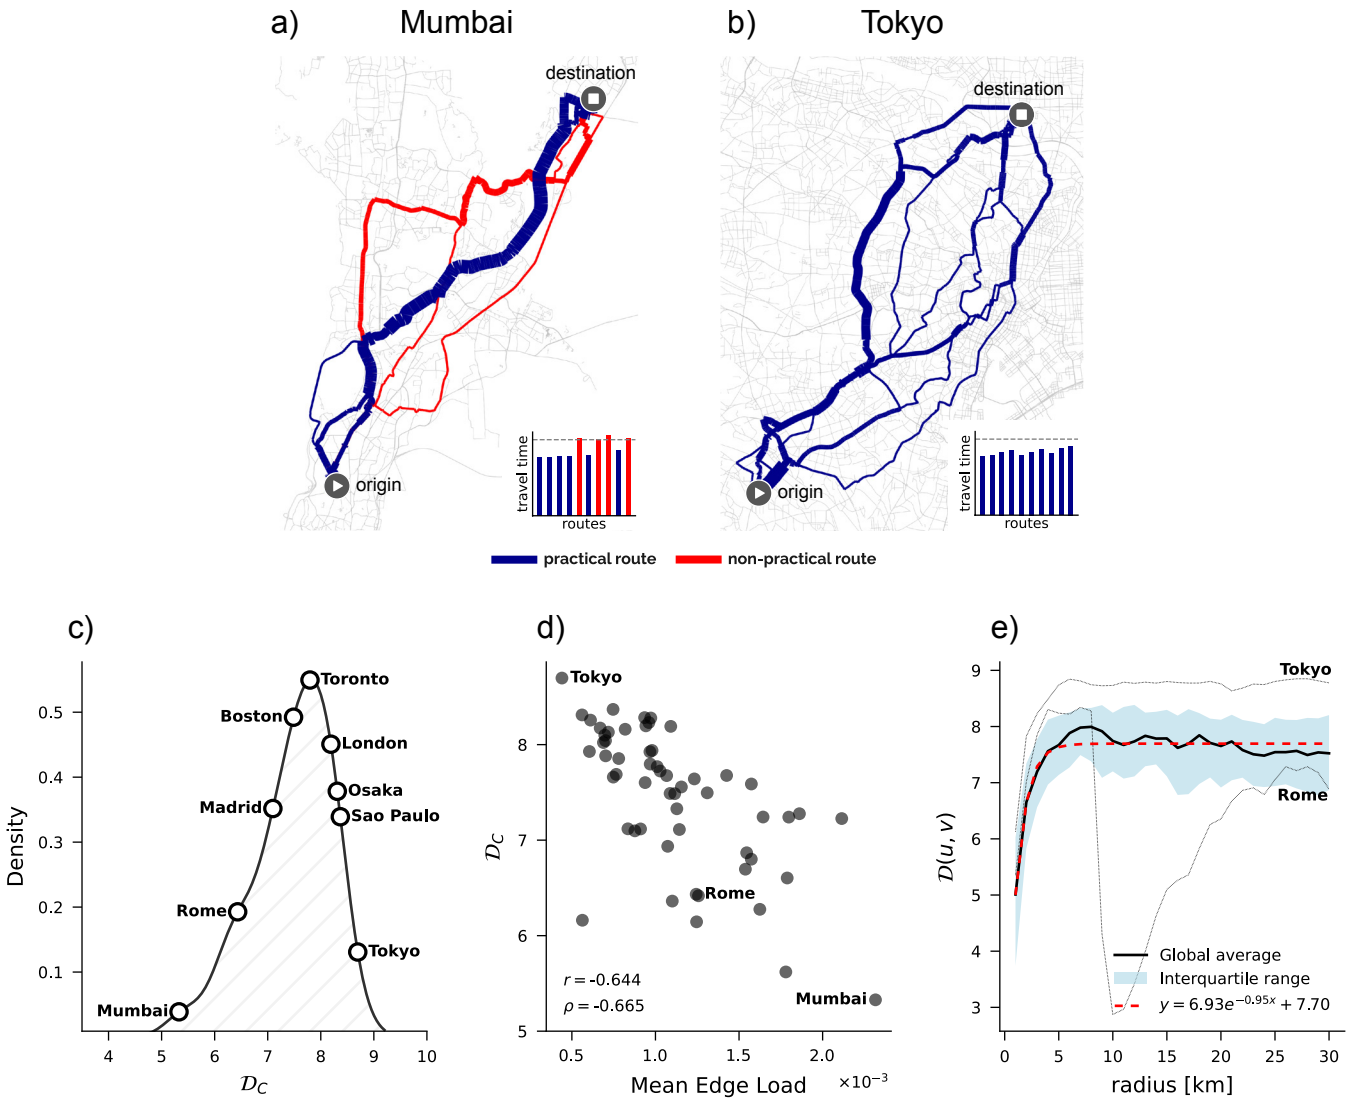

Fig. 1. **(a)** Example of a trip with low DiverCity (2.18) in Mumbai. **(b)** Example of a trip with high DiverCity (9) in Tokyo. For panels (a) and (b), inset bar plots show the travel time of each alternative route, with NSRs in blue and non-feasible routes in red. **(c)** Distribution of median DiverCity across 56 cities. **(d)** Median DiverCity versus average normalized edge load. **(e)** DiverCity for trips at varying radial distances from the city center. The black line shows the global average, the light blue area represents the interquartile range, and the red line is an exponential fit ($R^2 = 0.93$). Deviations include Tokyo's high values and Rome's localized drop near its ring road.

3. S. Çolak, A. Lima, and M. C. González, "Understanding congested travel in urban areas," Nat. Commun. **7**, 1 (2016).
4. D. Levinson and A. El-Geneidy, "The minimum circuity frontier and the journey to work," Reg. Sci. Urban Econ. **39**, 732 (2009).
5. M. Lee, H. Barbosa, H. Youn, *et al.*, "Morphology of travel routes and the organization of cities," Nat. communications **8**, 2229 (2017).
6. W. Jing, X. Xu, and Y. Pu, "Route redundancy-based network topology measure of metro networks," J. Adv. Transp. **2019**, 4576961 (2019).
7. D. Cheng, O. Gkountouna, A. Züfle, *et al.*, "Shortest-path diversification through network penalization: A Washington DC area case study," in *Proceedings of the 12th ACM SIGSPATIAL International Workshop on Computational Transportation Science,* (2019), pp. 1–10.

Urban semantic: dynamic land use through human mobility patterns

NANDINI IYER[1,*], MASSIMILIANO LUCA[2], OLLIN LANGLE-CHIMAL[3], MARTA C. GONZÁLES[3,4], AND RICCARDO DI CLEMENTE[1,5]

[1] Complex Connections Lab, Network Science Institute, Northeastern University London, UK.
[2] Bruno Kessler Foundation, Via Sommarive 18, 38123 Povo (TN), Italy
[3] Department of Civil and Environmental Engineering, University of California at Berkeley, Berkeley, 94720, CA, USA
[4] Department of City and Regional Planning, University of California at Berkeley, Berkeley, 94720, CA, USA
[5] ISI Foundation, Via Chisola 5, Turin 10126, Italy
[*] nandini.iyer@nulondon.ac.uk

Recent advances in human mobility research have leveraged large-scale digital traces and machine learning to reveal the hidden structure of cities. Pioneering work has uncovered universal patterns in human mobility [1], while later studies have used mobility flows to infer functional zones [2, 3], or have proposed models that learn urban representations from heterogeneous mobility datasets (e.g., taxi, bike-sharing) [4].

More recently, a growing body of work has applied embedding techniques to spatial data to uncover latent urban structures. For example, methods have been developed to identify functional similarity between places by applying Word2Vec to sequences of POIs ordered by shortest paths [5] or by embedding locations based on the co-occurrence of visits within individual mobility trajectories. [6]. Moreover, urban functions have been defined for the Greater London by using neural network embeddings on POI co-location data [7].

These methods demonstrate the power of representation learning to capture underlying spatial semantics. However, they typically treat space as static and either ignore or aggregate away temporal information. For example, Place2Vec [6] focuses on spatial co-occurrence but does not encode the time of day or sequence in which places are visited, making it difficult to infer how place function changes over time. Building on these foundations, we propose a novel embedding-based framework that captures the dynamic interplay between space and time in human mobility. By encoding individual trajectories as sequences of place-time tokens and applying Word2Vec, we learn representations that uncover contextual similarity between neighbourhoods at different times.

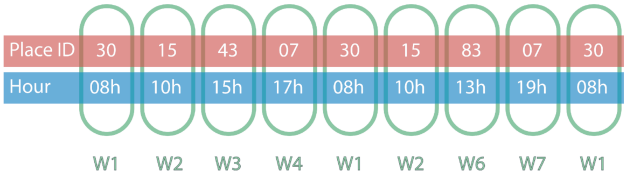

Fig. 1. Mobility trajectories can be encoded as sentences. We create string representations of place-hour combinations to represent words. We apply Word2Vec on trajectory sentences to apply a spatio-temporal perspective on human mobility patterns, understanding how similar different neighbourhoods are at each hour.

Our methodology begins with encoding mobility trajectories as spatio-temporal sequences, where each token represents a location visited at a specific time (Figure 1). This data is provided by Cuebiq for Dallas County, Texas, covering the period from October to December 2016. Space is encoded with respect to statistical administrative boundaries (i.e., Census Block Groups — CBGs — in the US). In this manner, an individual's trajectory is analogous to a sentence, while each place-hour combination resembles a word in the sentence (Figure 1). We apply Word2Vec, a technique commonly used in natural language processing, to this transformation to learn vector representations of each place-time combination based on its co-occurrence patterns in individual mobility trajectories. We cluster the learned word embeddings using spherical K-Means to capture location-hour combinations that belong to similar mobility contexts. Figure 2 illustrates how CBGs shift between clustered embeddings, revealing meaningful relationships between locations and hours that go beyond simple spatial proximity.

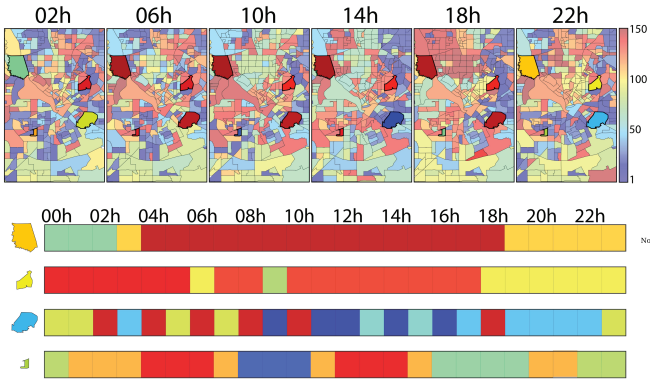

Fig. 2. CBGs can be interpreted based on how they change mobility contexts throughout the day. We define these mobility contexts by clustering Word2Vec word embeddings, which reveals neighbourhood-time combinations that appear in a similar context. Accordingly, each distinct colour reflects an embedding cluster.

To ensure these clusters meaningfully correspond to real-world urban functions, we validate and interpret them by integrating land-use data from the Global Human Settlement Layer (GHSL-BUILT-C) satellite imagery, which classifies areas as Residential, Non-Residential, and Open Space at a 10m resolution. Figure 3 highlights how we combine the raw land-use distribution (left-most, grey values) and the global land-use distribution (right-most, black values) to measure the relative lift (middle, black values), which captures how over or under-represented each land-use category is in relation to the land-use distribution for the entire Dallas County. Then, we quantify each cluster by aggregating the land-use distribution of all associated CBGs, estimating for each land-use type: (1) **average**: the average relative land use, (2) **tendency**: the percentage of CBGs that are

over-represented in a given land-use category and (3) **consistency**: how similar, the relative lifts are across all CBGs, using the coefficient of variance. This approach enables us to characterise clusters as a 9-element vectors (3 features for the 3 land-use types).

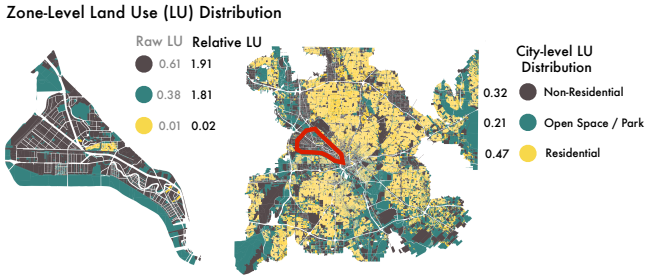

Fig. 3. Satellite imagery provides insight into the physical environment behind mobility context clusters. Using GHSL-BUILT-C land use (LU) data, we characterise each Census Block Group (CBG) by the proportion of open space, residential buildings, and non-residential buildings it contains. To contextualise these **raw LU patterns** (left), we divide each CBG's LU proportions by the **city-level LU distribution** of Dallas County (right). This yields a measure of **"relative land-use,"** (middle) in which values greater than 1 indicate that a particular land use type is more prominent in the CBG than it is countywide, highlighting its distinct spatial character.

To understand how urban areas behave over time, we analyse the sequence of mobility context clusters that each zone transitions through across the day. Each context cluster is defined by a 9-dimensional vector summarising the relative land-use composition (residential, non-residential, and open space) of its associated CBGs. By tracking a zone's cluster assignments at each hour over weekdays and weekends, we construct a temporal profile that captures its changing mobility context. We then cluster these temporal profiles to identify groups of zones that follow similar patterns of daily activation. In doing so we reveal, for instance, zones that exhibit consistently high residential relative land use during both weekday and weekend daytime, with non-residential land use near parity (values around 1) during the weekend daytime; or zones that shift from work-related activity (non-residential activations during weekdays) to residential use on weekends.

Figure 4 illustrates two such zone clusters, respectively, by showing the average relative lift of each land-use category across hourly time steps. Importantly, we distinguish between embedding clusters, which group place-time "words" with similar mobility contexts learned through Word2Vec, and zone clusters, which group zones based on how they transition across those embedding clusters over time.

This framework provides a powerful perspective for understanding how socioeconomic inequality manifests in urban space, not just in terms of where people go, but when and within what types of functional environments they move. By linking CBGs to demographic data from the US Census American Community Survey, we can investigate how different population groups, especially those in highly segregated neighbourhoods, interact with the city's mixed-use, residential, and non-residential spaces across the day.

To examine spatial inequality, we use the Index of Concentration at the Extremes (ICE) [8], defined as $ICE_{cbg} = \frac{p_{cbg,hi} - p_{cbg,lo}}{p_{cbg}}$, where $p_{cbg,hi}$ and $p_{cbg,lo}$ represent high- and low-income population counts within each CBG. Meanwhile p_{cbg} captures the

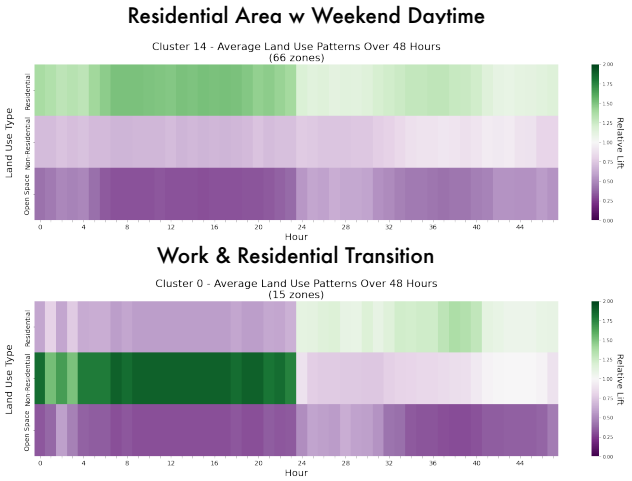

Fig. 4. Temporal land-use profile for a zone cluster, shown as a heatmap of average relative lift values. Each row represents a land-use category (residential, non-residential, open space), and each column corresponds to an hourly time step from weekday (0–23h) to weekend (24–47h). Values reflect the average relative lift of a given land-use category across all zones in the cluster, computed relative to the countywide land-use distribution. A value of 1 (light gray) indicates parity with the county average; values below 1 (purple) indicate under-representation; values above 1 (green) indicate over-representation. This highlights how zones within the cluster shift in land-use emphasis across the day and week.

total number of individuals in a given CBG. Mapping ICE scores onto zones (CBGs) allows us to explore which temporal urban contexts align with concentrated affluence or disadvantage. For instance, whether affluent areas remain stable and residential, while disadvantaged ones shift across mixed or fragmented roles.

Crucially, we can refer back to the trajectory sequences, to analyse how individuals from different ICE backgrounds navigate urban contexts over time. Even when visiting the same zone, high- and low-income groups may do so at different times, reflecting different functions: residential, commercial, or leisure. Ultimately, this framework presents a behaviourally grounded, temporally aware typology of urban space that moves beyond static land-use categories to reveal how the functional character of places evolves across the day and can be applied to reveal demographic divides in mobility.

Acknowledgment. We deeply thank Enrico Ubaldi and Bernardo Monechi for early discussions in defining the initial framework.

REFERENCES

1. M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, nature **453**, 779 (2008).
2. T. Louail, M. Lenormand, M. Picornell, *et al.*, Nat. communications **6**, 6007 (2015).
3. J. Reades, F. Calabrese, and C. Ratti, Environ. planning B: Plan. design **36**, 824 (2009).
4. Y. Fu, G. Liu, Y. Ge, *et al.*, IEEE transactions on knowledge data engineering **31**, 535 (2018).
5. Y. Yao, X. Li, X. Liu, *et al.*, Int. J. Geogr. Inf. Sci. **31**, 825 (2017).
6. W. Zhai, X. Bai, Y. Shi, *et al.*, Comput. environment urban systems **74**, 1 (2019).
7. H. Niu and E. A. Silva, Comput. Environ. Urban Syst. **88**, 101651 (2021).
8. D. S. Massey, Demography **33**, 395 (1996).

Working behaviour and the spatial complexity of economic activity

ZSÓFIA ZÁDOR[1 *], BALÁZS LENGYEL[2,3,4], AND RICCARDO DI CLEMENTE[1,5]

[1] Northeastern University London, United Kingdom
[2] MTA-ELTE KRTK Momentum Agglomeration Networks and Innovation Research Group, Budapest, Hungary
[3] Corvinus University of Budapest, Budapest, Hungary
[4] Institute of Economics, ELTE KRTK, Budapest, Hungary
[5] ISI Foundation, Torino, Italy
[*] zsofia.zador@nulondon.ac.uk

1. INTRODUCTION

Remote work, accelerated by the COVID-19 pandemic, has become widely adopted [1, 2]. Although remote work offers benefits to individuals and cities such as time savings, increased productivity, and reduced pollution, it also poses challenges to cities that require a rethinking of urban planning [3, 4]. Traditionally, the distance to the centre played a key role in urban planning, since urban travel was driven mainly by commuting to the centre. However, remote work reduces the demand for central workplaces and surrounding businesses [5] and increases suburban economic activity, contributing to the "Donut Effect" [6].

Remote work adoption is not equally accessible across industries and socio-economic factors, such as income and gender [7–9], therefore, exploring how these factors are interrelated with remote working is crucial to better understand its impact on cities. Prior studies have estimated remote work potential of certain jobs using occupational data [10] or professional platform trends [11], but the former approach focuses on hypothetical potential rather than actual behaviour, and the latter only captures limited segments of all job seekers. Finally, mobile phone data have been widely used to track mobility and socio-economic segregation [5], but this does not provide job related information. This study bridges that gap by combining mobile phone data with administrative company records to reveal how remote work adoption varies across industries, income groups, and neighbourhoods.

This study aims to better understand how remote work changes cities. To this aim, we examine the socio-economic and company-level determinants of remote work adoption in Budapest, Hungary. We estimate remote work prevalence from mobile phone data, measuring the change in workplace activity between "opening time" (September–October 2020) and "curfew time" (November–December 2020), when strict COVID-19 restrictions were in place. By linking these mobility trends with administrative company records, we identify the socio-economic and industry characteristics of locations with known remote work prevalence.

2. DATA AND METHODS

Using geolocated data from 1.4 million mobile devices in Budapest, we analyse hourly population density across approximately 0.01 km² hexagonal spatial units. We cluster the dynamics of work activity into three categories (Figure 1), which exhibit distinct temporal patterns: Cluster 1 aligns with traditional office schedules, Cluster 2 corresponds to multi-use spaces with extended working hours, including weekends, and Cluster 0 is concentrated on the urban periphery and characterised by non-standard activity patterns, such as nighttime and weekend work and increased activity during the mobility-restricted “curfew” period.

Based on the spatial distribution of these clusters (Figure 2a), we would expect distance to the city centre to drive remote work adoption, as central areas function as integrated workspaces. However, we find that spatial proximity alone provides limited explanatory power. Although Cluster 0 industrial areas on the periphery display consistently different working patterns, there is considerable variation in activity even within central zones. Spatially, Clusters 1 and 2 exhibit similar characteristics, but qualitative analysis reveals important distinctions: large commercial and institutional hubs with higher remote work adoption (Cluster 1) contrast with mixed-use zones that display

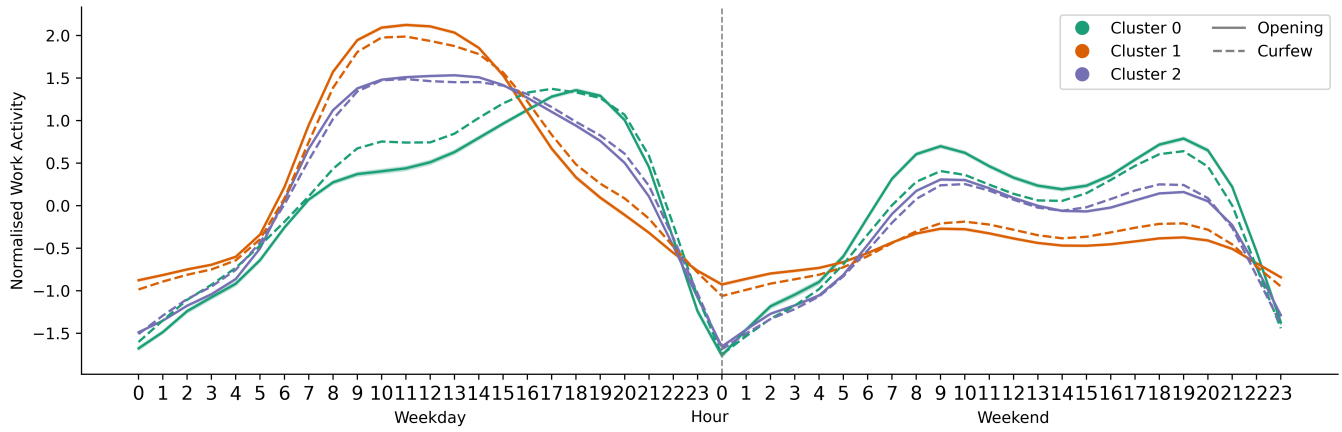

Fig. 1. Average work activity of the clusters during Opening and Curfew time.

(a) (b) (c)

Fig. 2. The interaction of working dynamics and company statistics. (a) Snapshot of Budapest's working activity clustered. (b) Scatter plot of the log working activity against the log number of companies. Colours represent the 3 working clusters; size = average employees per location. (c) Comparison of the mean socio-economic and company variables within each working cluster.

more diverse activity rhythms (Cluster 2). We find a correlation between log work activity and the log number of companies (Figure 2b), demonstrating the concurrent validity of the mobile phone data and administrative company records. Additionally, cluster membership moderates this relationship, with industrial areas (Cluster 0) showing a weaker association.

To explore the economic drivers of remote work adoption, we link working behaviour clusters with company-level data to examine where remote work is most prevalent, what socio-economic characteristics describe these locations, and how the presence of different industries relates to varying adoption of work-from-home policies. We estimate remote work for each location i as $\Delta\text{working activity}_i = N_{\text{curfew},i} - N_{\text{open},i}$, where $N_{\text{curfew},i}$ denotes the number of workers present at working hours during the curfew period at the location i, and $N_{\text{open},i}$ denotes the number of workers present at the same time during the opening period at the same location. This approach allows us to capture the locations from which workers shifted during the curfew period. We then estimate an OLS regression using $\Delta\text{working activity}_i$ as the outcome, with predictors capturing location, company, and socio-economic composition.

3. RESULTS AND DISCUSSION

We find that industry composition, local socio-economic composition and distance to the centre all play a role in explaining remote work adoption (Figure 2c). Remote work adoption is highest in the service sector and lowest in the primary sector, followed by the secondary sector. The OLS regression finds that remote work adoption is highest in finance-related industries ($\beta = -0.187$, $p < 0.05$), in contrast with a low adoption in commerce-related industries ($\beta = 0.072$, $p < 0.01$). Furthermore, remote work prevalence is negatively associated with the distance of workplaces from the city centre ($\beta = -0.214$, $p < 0.01$), as well as Cluster 1 areas that constitute large commercial and institutional hubs ($\beta = -0.454$, $p < 0.01$). Areas with more productive companies also exhibit higher levels of remote work adoption ($\beta = -0.078$, $p < 0.01$), suggesting that large, productive firms in the city centre were most likely to allow working-from-home.

On the other hand, the socio-economic composition of work locations during the opening time highlights the impact of unequal remote work on social mixing in cities. The OLS regression shows a strong association between locations with a stronger female presence ($\beta = 0.200$, $p < 0.01$), income equality ($\beta = 0.0499$, $p < 0.01$) and lack of remote work. This suggests that locations that previously had the potential to serve as social mixing hubs may now experience a decline of high-earning service workers, leading towards the homogenisation of areas and increased segregation.

By integrating high-resolution mobility data with administrative company records, this study offers novel insights into the socio-economic, industrial, and spatial dimensions of remote work. Understanding evolving urban workforce dynamics has critical implications for policymakers and urban planners seeking to adapt city structures to the post-pandemic environment. Our study highlights that with reduced commuting, the traditional importance of proximity to the city centre may diminish, reshaping surrounding businesses [5] and weakening the 'mixing pot' function where people of diverse backgrounds interact [12].

REFERENCES

1. J. M. Barrero, N. Bloom, and S. J. Davis, J. Econ. Perspect. **37**, 23 (2023).
2. J. M. Barrero, N. Bloom, and S. Davis, "Why Working from Home Will Stick," Tech. Rep. w28731, National Bureau of Economic Research, Cambridge, MA (2021).
3. C. G. Aksoy, J. M. Barrero, N. Bloom, *et al.*, Brookings Pap. on Econ. Activity **2022**, 281 (2022).
4. M. O. Faruque, T. Talukder, M. N. Pranto, *et al.*, Am. J. Ind. Bus. Manag. **14**, 869 (2024).
5. T. Yabe, B. G. B. Bueno, X. Dong, *et al.*, Nat. Commun. **14** (2023). Publisher: Springer Science and Business Media LLC.
6. A. Ramani, J. Alcedo, and N. Bloom, Proc. National Acad. Sci. **121**, e2408930121 (2024).
7. N. Gozzi, M. Tizzoni, M. Chinazzi, *et al.*, Nat. Commun. **12**, 2429 (2021).
8. X. Huang, J. Lu, S. Gao, *et al.*, Ann. Am. Assoc. Geogr. **112**, 286 (2022).
9. A. Kar, H. T. K. Le, and H. J. Miller, Ann. Am. Assoc. Geogr. **112**, 1023 (2022).
10. J. I. Dingel and B. Neiman, J. Public Econ. **189**, 104235 (2020).
11. "LinkedIn Data: Women, Gen Z More Likely to Apply to Remote Jobs," .
12. J. Jacobs, *The death and life of great American cities*, Pelican books (Penguin, Harmondsworth, 1972).

Bias estimates for MPD-derived mobility statistics in low-and middle-income countries

Roland Hosner*, Zachary Strain-Fajth, Véronique Lefebvre

Flowminder Foundation, Norrsken House, Birger Jarlsgatan 57C, 113 56 Stockholm, Sweden
*roland.hosner@flowminder.org

Received 5 July 2025; revised XX Month, XXXX; accepted XX Month XXXX; posted XX Month XXXX (Doc. ID XXXXX); published XX Month XXXX

ABSTRACT

Based on original and recent household and telephone surveys in the Democratic Republic of the Congo (DRC), Haiti and Ghana, we estimate biases of mobility and migration statistics derived from mobile phone data (MPD). We show that indicators of short- and long-term internal mobility differ greatly between mobile phone users and non-users, where phone users show higher mobility on average. These differences translate into biases when statistics are derived from mobile phone data and population inferences are made without adjustments. These findings are not only relevant for mobility research but also have wider implications for the use of Mobile Phone Data and other Big Data for analyses of mobility and migration.

INTRODUCTION

The interest in and use of Big Data for the analysis of human mobility has grown rapidly in recent years. Particularly in low- and middle-income countries (LMICs), the promise of utilising highly granular data for official statistics, public service provision and for humanitarian and development purposes is driving analyses of relatively new data sources such as MPD for statistics. Use cases range from routine mobility reports (1) to disaster response (2), epidemiological studies (3) and the study of rare events such as mobility in the wake of earthquakes and other natural hazards (4). There is also growing awareness of and consensus among researchers that biases inherent in such novel data, particularly stemming from the non-probability nature of these samples, need to be analysed, and where feasible, adjusted for (5-7). However, the sample selection mechanism cannot be identified from MPD alone, socio-demographic information is often incomplete or missing, so that standard data sources – survey, census and register data – have to be used in combination with MPD to assess and mitigate biases.

The work we present here covers the assessment and is aimed at highlighting the differences in mobility and migration indicators between mobile phone users and non-users that translate into biases when only MPD are used for target statistics. In LMICs, direct measures of the target statistics are often lacking, and MPD-derived statistics are often more granular in time and space than available statistics at higher temporal and spatial levels, making it impossible to calculate exact bias statistics. Therefore, we present bias estimates based on survey data.

DATA

For the DRC, the data of the 2021 microcensus were collected through face-to-face household interviews in seven provinces (8). For Ghana, we were able to submit questions to an ongoing household survey, the Annual Income and Household Expenditure Survey (AHIES) 2022 (9), covering all regions of Ghana. For Haiti, the household survey data used for analyses stem from the DHS data collection 2016/2017 (10).

RESULTS

For all three countries we find significant differences in mobility between mobile phone users and non-users, across a range of mobility indicators. In Ghana, the differences are generally larger than in Haiti and the DRC. Short-term mobility, such as spending one or several months during the year away from home, is more common among mobile phone users than non-users in both Ghana and Haiti. A more permanent and rarer type of mobility - having changed the place of residence in the previous six months - showed smaller differences between mobile phone users and non-users in both the DRC and Ghana. Long-term mobility, measured as any change of place of residence over the life course, is more prevalent overall, and significantly more common among mobile phone users in all three countries (see table 1). A smaller difference can be seen for the DRC, while in Haiti, the difference is much larger, with half of mobile phone users ever having internally migrated, but only about four out of ten non-users. The largest difference is seen for Ghana, with half of mobile phone users ever having migrated internally, but only a third of non-users.

Table 1. Mobility differences between mobile phone users and non-users in DRC, Haiti and Ghana

	Lifetime internal migrant		
Mobile phone use	DRC	HTI	GHA
Non-users	40.1%	38.3%	34.7%
Mobile phone users	42.9%	51.3%	50.9%
p-value[a]	0.031	0.000	0.000
Observations	9,833	25,033	30,774

[a] Wald tests for differences of means
Sources: DRC Microcensus 2021, Haiti DHS 2017, Ghana AHIES 2022 Q2

Across the three countries, we analysed a total 15 indicators of mobility and migration and found significant effects ($p<0.1$) of mobile phone use for 13 of these indicators in bivariate regression models in overall models, of which 10 were highly significant ($p<0.01$). All 13 indicators show higher mobility of mobile phone users in these overall models, where all respondents per country are included. The bivariate effects were partly explained when controlling for socio-demographic variables and region (administrative level 1), but 12 indicators remained significant and 11 highly significant also in the multivariate analyses.

Moving from mean tests to bias estimation, whenever general population parameters are estimated directly from samples of mobile phone users (excluding non-users in general or also users of MNOs not covered), the bias of such estimates is the difference between these estimates and the true population parameters. For the mean, this can be expressed as:

$$bias\left(\mu_{x,phoneusers}\right) = \hat{\mu}_{x,phoneusers} - \mu_x$$

If and when the parameter of interest differs between mobile phone users and the general population, there is bias in the direct MPD-derived estimate.

Taking Ghana as an example and using bias in % points as a measure for dichotomous indicators, biases across five mobility indicators and 16 regions (admin1) ranged between 0% and +12%. In other words, the mobility of mobile phone users was often higher than for the general population, resulting in a positive bias. The spread of biases on the district (admin2) level is much larger than that of biases on the regional level. Across all indicators and 148 districts, biases range from -5 % points to +26 % points (fig. 2). On lower sub-national levels, biases vary more substantially, and smaller biases at higher geographical levels may mask variation on lower levels.

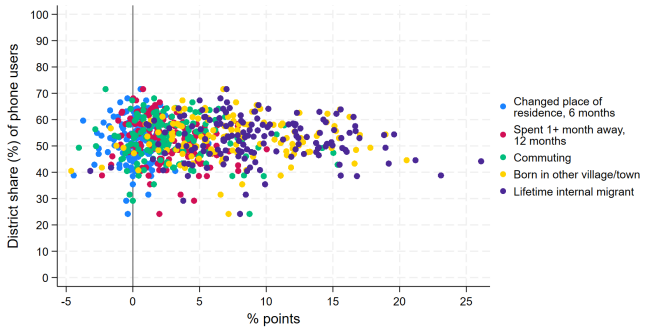

Fig. 1. Bias estimates for five mobility indicators across 148 Ghana districts. Data Source: Ghana Annual Income and Household Survey 2022, quarter 2; respondents aged 6+ years

Biases for more permanent forms of mobility show substantially larger variation in our data. However, this might be due to the so-called floor effect: when proportions are small, variation of these proportions between groups is statistically curtailed.

DISCUSSION

In summary, mobile phone use has a significant impact on mobility and migration statistics derived from MPD, which translates into biased estimates. Such biases may be small on the national level but become larger on sub-national levels. While acknowledging that it may not be possible to eliminate but only mitigate bias, future research should address these statistical quality issues by triangulating MPD and standard data sources such as surveys and providing estimates that are bias-adjusted, so that MPD-derived statistics can be used for analyses and practical applications, for example by leveraging integrated estimates (11) or modelled estimates (12).

In practical terms, when a statistic is biased by 10% points or more, this may have serious consequences for the interpretation of results and the applied use case. Health workers may allocate scarce resources in places that are more or less populated than indicated by the MDP-derived statistics, for example. As such, accuracy of statistics for the service provision, for humanitarian and development sectors is equally important as for official statistics.

References

1. Ghana Statistical Service. Population Mobility and Distributions in Ghana in 2024. 2025.
2. Bengtsson L, Lu X, Thorson A, Garfield R, von Schreeb J. Improved response to disasters and outbreaks by tracking population movements with mobile phone network data: A post-earthquake geospatial study in Haiti. PLoS Med. 2011; 8(8).
3. Kostandova N, Prosperi C, Mutembo S, Nakazwe C, Namukoko H, Nachinga B, et al. Comparing and integrating human mobility data sources for measles transmission modeling in Zambia. PLOS Global Public Health. 2025; 5.
4. Li T, Dejby J, Albert M, Bengtsson L, Lefebvre V. Detecting individual internal displacements following a sudden-onset disaster using time series analysis of call detail records. 2019.
5. Milusheva S, Bjorkegren D, Viotti L. Assessing Bias in Smartphone Mobility Estimates in Low Income Countries. In: Proceedings of 4th ACM SIGCAS Conference on Computing and Sustainable Societies. 2021. p. 364–78.
6. Pestre G, Letouzé E, Zagheni E. The ABCDE of Big Data: Assessing Biases in Call-Detail Records for Development Estimates. World Bank Economic Review. 2020;34: 89–97.
7. Gallotti R, Maniscalco D, Barthelemy M, De Domenico M. Distorted insights from human mobility data. Commun Phys. 2024; 7(1).
8. Flowminder. Methodology report for the DRC microcensus 2021. 2021.
9. Ghana Statistical Service. Quarterly Labor Force Report. 2022.
10. Institut Haïtien de l'Enfance (IHE), ICF. Enquête Mortalité, Morbidité et Utilisation des Services (EMMUS-VI 2016-2017). 2018.
11. Kim JK, Tam SM. Data Integration by Combining Big Data and Survey Sample Data for Finite Population Inference. International Statistical Review. 2021; 89(2): 382–401.
12. Zhang LC, Necula M, Filipponi D, Simeoni G. Methodologies and open source tools for integrating MNO and non-MNO data sources. Preliminary report on methodologies. 2024.

Funding. The work was funded by the William and Flora Hewlett Foundation

Data availability. The survey microdata underlying the results presented in this paper are available from Ghana Statistical Service upon request, from the DHS website and from the DRC Institut National de la Statistique upon request.

Disclosures. The authors declare no conflicts of interest

Predicting Vehicle Trajectories based on Latitude and Longitude Positioning Historical Data

BRENO KROHLING[1], GIOVANNI COMARELA[1], CLAYSON CELES[1], PAULO H. L. RETTORE[2], AND VINÍCIUS F. S. MOTA[1,*]

[1] Universidade Federal do Espírito Santo (UFES), Vitória, ES, Brazil
[2] Department of Communication Systems, Fraunhofer FKIE, Bonn, Germany
[*] vinicius.mota@inf.ufes.br

1. INTRODUCTION

Understanding vehicle movement dynamics in large cities by analyzing and predicting how vehicles move and where they are likely to go next enables the development of various smart city services and applications [1]. Indeed, trajectory prediction supports applications in urban planning, traffic management, mobile network handover, and vehicle communication within the Internet of Vehicles (IoV) paradigm [2–5]

In the context of urban mobility, Deep-Learning has been applied to predict trajectory tasks [6, 7]. However, urban environments are inherently complex and heterogeneous, with traffic patterns influenced by diverse temporal, spatial, and behavioral factors. Additionally, data collected from GPS and sensors often contains noise and inaccuracies, further complicating the learning process. These challenges limit the effectiveness of trajectory prediction models in real-world applications. Furthermore, most prediction models focus solely on estimating the next position of a given past trajectory. It is important to note that while trajectory prediction has been extensively studied in scenarios involving Points of Interest (POIs), predicting future trajectories based solely on raw coordinate data remains an open challenge.

In this paper, we propose an architecture that, given a sequence of position tuples $(time, latitude, longitude)$ of a trajectory, outputs the vehicle's position up to h unity of time ahead. The proposed architecture considers the outputs of a Long short-term memory network (LSTM) stacked as input to another LSTM, and so on. Considering the low variation of the latitude and longitude for short distances, we used two metrics to evaluate trajectory prediction: (i) Haversine distance between real and predicted positions; and (ii) Region of Interest, which divides the area into a rectangular grid and maps both actual and predicted positions to grid cells. This metric aims to evaluate whether the model can hit a vehicle's future region of interest. We analyzed prediction efficiency over time, using taxi trajectory datasets from Porto (Portugal) and San Francisco (USA).

2. PROPOSED ARCHITECTURE: STACKED LSTM

A. Problem Definition

Before defining the problem, we first present some definitions.

Position: Let $x_i = (lat_i, long_i)$ be a GPS position in Latitude and Longitude (WGS84) of a vehicle at time i.

Trajectory: A trajectory $T = \{x_1, x_2, ..., x_t\}$ represents a complete path traveled by a vehicle over t time units.

Sub-trajectory: Given a trajectory T of size t, the sub-trajectory $\tau_{j,w}$ of T is the w consecutive positions in T starting from position j, i.e., $\tau_{j,w} = \{x_j, x_{j+1}, ..., x_{j+w-1}\}$ for $w \leq t$ and $1 \leq j \leq t - w + 1$.

Thus, T has $t - w + 1$ sub-trajectories of size w. Now, we define the Trajectory Prediction Problems as short-term and long-term below.

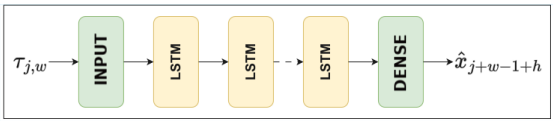

Fig. 1. Output of an LSTM layer stacked as input to the next.

Short-term prediction: Given a sub-trajectory $\tau_{j,w}$, the goal is to predict the vehicle's position in the next time unit ($h = 1$).

Long-term prediction: Given a sub-trajectory $\tau_{j,w}$, the goal is to predict the vehicle's position at times reflecting medium and/or long term ($h \gg 1$).

B. LSTM for Trajectory Prediction

LSTM networks are commonly used for sequential and temporal data due to its memory mechanism, enabling storage of past state information over long periods [8]. Given a sub-trajectory $\tau_{j,w}$, the model predicts the future position $\hat{x}_{j+w-1+h}$. To extract complex temporal patterns, we considered stacked LSTMs, where the output of one LSTM layer is the input to the next. Preliminary experiments indicated that stacked LSTMs outperform single-layer LSTMs. The model architecture is shown in Fig. 1.

The output $\hat{x}_{j+w-1+h}$ is a location (latitude and longitude). We map model outputs to grid cells. The mapping is performed to a grid M of dimensions $a \times m$ covering the region; each cell represents a region of interest. The mapping is defined as:

$$M_{r,c}(x_i) = a_r(x_i), m_c(x_i) \quad (1)$$

$$a_r(x_i) = \left\lfloor \frac{a(lat_{x_i} - lat_{min})}{lat_{max} - lat_{min}} \right\rfloor, r = 0, ..., a-1 \quad (2)$$

$$m_c(x_i) = \left\lfloor \frac{m(long_{x_i} - long_{min})}{long_{max} - long_{min}} \right\rfloor, c = 0, ..., m-1 \quad (3)$$

where $lat_{min}, lat_{max}, long_{min}, long_{max}$ are the grid's bounding coordinates. Model performance can thus be analyzed by region-of-interest hits. The grid's dimensions a and m determine result granularity: smaller cells mean higher precision and greater prediction difficulty.

C. Evaluation Methods

Traditional prediction models use metrics like RMSE or R^2 to quantify error, but for this problem, small metric errors can mean large real distances. For example, position $x = (22.905300, -43.133600)$ predicted as $\hat{x} = (-22.9067626, -43.1355232)$ yield RMSE $= 0.00171$ and $R^2 = 1$, but are about 200 meters apart. We use two main metrics:

Distance: The first metric is the real distance in meters between the real point (x_i) and the predicted point ($\hat{x}_i$), calculated by the Haversine distance as shown in 4.

$$D_{x_i,\hat{x}_i} = 2 \cdot r \cdot \arcsin \sqrt{\sin^2 \frac{\phi_{x_i} - \phi_{\hat{x}_i}}{2} + \cos \phi_{x_i} + \cos \phi_{\hat{x}_i} + \sin^2 \frac{\omega_{x_i} - \omega_{\hat{x}_i}}{2}} \quad (4)$$

where r is Earth's radius, ϕ is latitude, and ω is longitude. The second metric is based on region-of-interest (ROI) prediction:

Region of interest: Points x and $\hat{x}$ are mapped to grid cells M; given a test set X and predictions $\hat{X}$, the ROI metric is defined as shown in 5 and 6:

$$ROI = \frac{\sum_{x \in X, \hat{x} \in \hat{X}} F(M_{r,c}(x), M_{r,c}(\hat{x}))}{|X|} \cdot 100 \quad (5)$$

where

$$F(M_{r,c}(x), M_{r,c}(\hat{x})) = \begin{cases} 1, & M_{r,c}(x) = M_{r,c}(\hat{x}) \\ 0, & M_{r,c}(x) \neq M_{r,c}(\hat{x}) \end{cases} \quad (6)$$

Thus, ROI is an accuracy metric. Grids with larger granularities mean higher position precision.

3. EVALUATION

Given a sequence of past GPS positions, we aim to predict the short- and long-term ($h \geq 1$) vehicle's future positions. Evaluation measures are (1) Haversine distance between predicted and actual points, and (2) ROI accuracy, where a hit occurs when both positions fall in the same grid cell.

Regarding the datasets, the **Porto Taxi** dataset [9] contains GPS data from 442 taxi drivers between 2013 and 2014. It includes ~1.7 million trajectories, recorded at 15-second intervals throughout the trips. The **San Francisco** dataset [10] contains data from 500 taxi drivers. During a 30-day period, it includes ~900,000 trajectories, with positions recorded at intervals of approximately 10 seconds. Figure 2 depicts the trajectories on both datasets. San Francisco has a grid-like street pattern, while Porto presents a more irregular pattern, making prediction more complex. For each dataset, we evaluated the effect of the size of the sub-trajectory by dividing each trajectory into sub-trajectories of size $w = \{10, 20, 30, 40, 50\}$.

STACKED LSTMS PARAMETRIZATION. A previous search over LSTM hyperparameters hidden units (hu) and stacked layers (sl) was performed to find the best prediction model.

A. Results

Fig. 3 shows the results of short-term (next-step) and long-term (up to 10 steps ahead) predictions, with the corresponding bar charts summarizing the key ROI accuracy metrics for both cities. The model is more accurate in San Francisco (83% short-term, 52% long-term) than in Porto (64% short-term, 42% long-term), due to San Francisco's grid-like street layout (as shown in Fig. 2). Most short-term predictions have Haversine errors below 300m.

4. CONCLUSION

Despite the challenges for short and long-term trajectory prediction using only geospatial data, the proposed architecture achieved promising results. Two evaluation metrics were proposed: one for high-precision position estimation and another for region-of-interest hit rate. As future work, the prediction model can be applied to assist other problems, such as increasing efficiency in mobile network handover techniques.

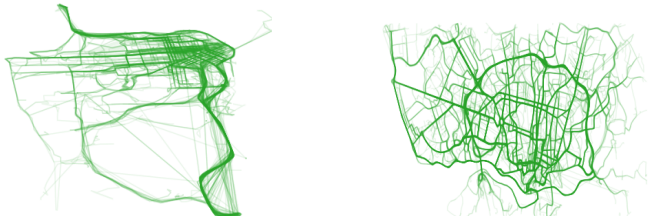

(a) San Francisco. **(b)** Porto.

Fig. 2. GPS trajectories from San Francisco and Porto.

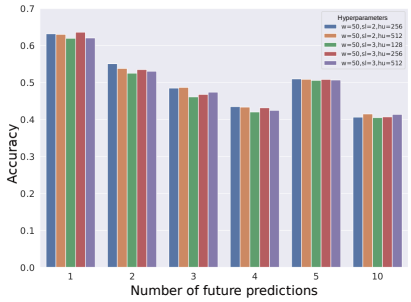

(a) Porto. Each cell ≈ 240m × 220m.

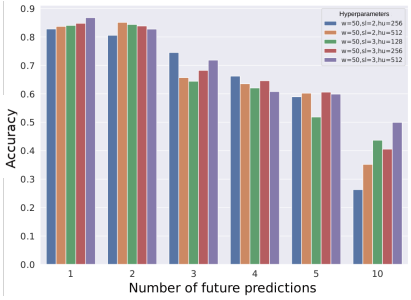

(b) San Francisco. Each cell ≈ 450m × 220m.

Fig. 3. ROI accuracy for Porto and San Francisco: short-term ($h = 1$) and long-term ($h = 10$). 20×50 grid cells.

Acknowledgment. This work was supported by CAPES (Brazil) - Code 001, the National Council for Scientific and Technological Development – CNPq, FAPES (#2024-9G1WP), and São Paulo Research Foundation (FAPESP) (#2020/05121-4, #2023/00148-0, #2025/00513-5).

Data Availability Statement. Data available in Refs. [9, 10].

REFERENCES

1. Y. Mehmood, F. Ahmad, I. Yaqoob, *et al.*, "Internet-of-things-based smart cities: Recent advances and challenges," IEEE Commun. Mag. **55**, 16 (2017).
2. N. Khansari, A. Mostashari, and M. Mansouri, "Impacting sustainable behavior and planning in smart city," Int. journal sustainable land Use Urban planning **1** (2014).
3. E. Maggi and E. Vallino, "Understanding urban mobility and the impact of public policies: The role of the agent-based models," Res. Transp. Econ. **55**, 50 (2016).
4. S. Sharma and B. Kaushik, "A survey on internet of vehicles: Applications, security issues & solutions," Veh. Commun. **20**, 100182 (2019).
5. Rettore, P. H. L., G. Maia, L. A. Villas, and A. A. F. Loureiro, "Vehicular data space: The data point of view," IEEE Commun. Surv. & Tutorials **21**, 2392 (2019).
6. P. Zhang, W. Ouyang, P. Zhang, *et al.*, "Sr-lstm: State refinement for lstm towards pedestrian trajectory prediction," in *Conference on Computer Vision and Pattern Recognition*, (2019), pp. 12085–12094.
7. A. D. Pazho, G. A. Noghre, V. Katariya, and H. Tabkhi, "Vt-former: An exploratory study on vehicle trajectory prediction for highway surveillance through graph isomorphism and transformer," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, (2024), pp. 5651–5662.
8. S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural computation **9**, 1735 (1997).
9. ECML/PKDD, "Porto taxi dataset," https://www.kaggle.com/c/pkdd-15-predict-taxi-service-trajectory-i (2015). Acesso: 2024-01-22.
10. M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "Crawdad data set epfl/mobility (v. 2009-02-24)," (2009).

Temporal Convolutional Network Forecasting of Tourism Demand Indicators Based on GPS Data: A Case Study in Japan

KETIAN CHEN[1,*], JUE MA[1], CHENBO ZHAO[1], AND YOSHIHIDE SEKIMOTO[1]

[1] Center for Spatial Information Science, The University of Tokyo
[*] ktchen@csis.u-tokyo.ac.jp

1. INTRODUCTION

Accurate forecasting of tourism demand is crucial for the tourism industry, as it enables tourism site planners and local authorities to allocate resources effectively in response to changing demand. This capability is especially important in single tourist sites, such as theme parks or resorts, since visitor volume often experiences more significant short-term fluctuations compared to broader tourist regions. The operations of these tourism sites are highly sensitive to factors such as weather changes, holidays or promotional events, which can cause rapid and pronounced changes in the volume of visitors. Consequently, precise forecasting of tourism demand is essential for timely staff scheduling and facility management, helping to prevent overcrowding and ensure a positive visitor experience. Traditional forecast models, such as time-series or econometric models have been widely applied with relatively low-frequency data, such as monthly, quarterly, or annual public statistics [1]. Recently, researchers have started to utilize diverse and more high-frequency data sources, such as online search engine trends, online review data, social media data, with advanced modeling techniques [2–5]. Online search engine data have proven especially effective due to their timeliness and a direct reflection of tourists' planning behaviors. GPS-based mobility data further expands the scope of tourism demand forecasting by offering more detailed insights with higher spatial and temporal resolution [6–8]. Specifically, GPS data enables the extraction of more detailed tourism behavior indicators beyond tourist volume, such as average stay duration, visitor origin, overnight stay ratio, and precise arrival and departure times. Those GPS data-based metrics reflect diverse tourist profiles and behavior patterns, facilitating targeted planning and marketing strategies. This data source is also more cost-effective over time compared to traditional surveys, especially beneficial for tourist sites without ticketing systems or formal entry monitoring.

Our main research objective is to develop a framework for forecasting multiple tourism demand indicators in a single tourist site using GPS data. Specifically, we first extract multiple tourism demand indicators from the GPS trajectory data. Then we employ a temporal convolutional network (TCN), a deep learning architecture, to forecast future tourism demand based on a sequence of historical tourism demand and related features, including time factors, meteorological conditions, and site-specific events or activities. To examine the practical application of our framework, we present a case study of Gotemba Kogen Resort Tokinosumika, a resort located in Shizuoka Prefecture, Japan.

2. METHODOLOGY

To model daily tourism demand, we developed a predictive model incorporating a range of features capturing key influencing factors. These include: (1) historical visitation metrics extracted from GPS data, including the number of visitors on the previous day and the same day of the previous week; (2) time-related dummy factors, including the day of the week and public holiday in Japan; (3) meteorological factors, including daily weather and average temperature; (4) local promotional event dummies to reflect tourism-driven campaigns; and (5) Google Trends indices for keyword search to capture online interest levels. All numerical features are normalized using Min-Max scaling to ensure consistent contribution during model training.

We choose the temporal convolutional network (TCN) model for forecasting tourism demand. Compared to traditional recurrent architectures such as Long Short-Term Memory (LSTM) or Gated Recurrent Unit (GRU) models, TCN typically requires shorter training time and uses less memory [9, 10]. Its architecture—based on causal and dilated convolutions—prevents future-to-past information leakage while efficiently capturing both short- and long-term temporal dependencies with interpretable receptive fields. Our TCN model takes 7 days of historical features as input (`window_size = 7`) and outputs the next day's prediction. The architecture consists of two 1D convolutional layers (kernel size = 1, filters = 9) for feature extraction, followed by a temporal convolution layer (kernel size = 28, causal padding) to capture sequential patterns. The output is flattened and passed through a dense layer to generate the forecast. The model is trained using the Adam optimizer (learning rate = 0.0001) with mean squared error (MSE) loss.

We use mobile GPS data provided by GeoTechnologies. GeoTechnologies offers anonymous and high-resolution GPS records from approximately 4 million monthly active mobile phone users across Japan, which generates about 1 billion location logs per day. Each log contains a unique user ID, geographic coordinates (latitude, longitude, and altitude), and a timestamp. To analyze visitor behavior at tourist destinations, we apply stay point detection techniques to estimate the arrivals and stays at specific locations.

3. CASE STUDY

We examine our model at Gotemba Kogen Resort Tokinosumika, a resort located in Gotemba City, Shizuoka Prefecture, Japan. To identify users who have visited the Tokinosumika area, we segment these logs into individual trajectories. We apply a stay-point detection algorithm, using a 100-meter distance and 10-minute duration threshold, to extract stops from raw trajectories. Users with at least one stay point intersecting the Tokinosumika area are labeled as site visitors. For the identification of overnight visitors, we further require that the stay satisfies both of the following conditions: the stay spans across calendar dates and includes the nighttime period between 23:00 and 6:00. We also identify local visitors from Shizuoka prefectures based on daily origin locations from individual trajectories. Each user is then weighted using the expansion factor1 provided by the data source to approximate the actual number of visitors.

We identified 282,833 overnight visitors in 2023 and 256,925

in 2024. According to information provided by the local resort operator, the actual number of overnight visitors in fiscal year 2024 was approximately 260,000. This close correspondence suggests the robustness and reliability of the data processing methodology. Moreover, we conducted a cross-validation using a separate set of mobility data provided by another company (Agoop). We applied the same extraction methods to this independent dataset and compared the daily distribution of all site visitors counts. Pearson's correlation coefficient (r) was 0.95, indicating strong consistency and supporting the validity of the GPS-based estimates used in our study.

We derive daily tourist visit records from 1 January 2023 to 31 December 2024. Then, we split the data into training (80%) and test (20%) sets. We evaluated the performance of the TCN model against several baseline models, including ARIMAX, SARIMAX, and Long Short-Term Memory (LSTM). Model performance was assessed using mean absolute error (MAE), coefficient of determination (R^2), weighted mean absolute percentage error (wMAPE), and root mean squared error (RMSE). Table 1 presents the performance of TCN and the other three models on the test set. In general, TCN outperformed all other models in predicting all three indicators of tourism demand: all site visitors, local site visitors, and nightstay visitors. It achieved the lowest MAE and RMSE for all site visitors and nightstay visitors, and maintained strong performance for local visitors, slightly ahead of LSTM. TCN also achieves the highest R-squared scores for all three indicators, highlighting its stronger ability to account for variability compared to baseline models.

Table 1. Performance Evaluation

All Visitors				
Model	MAE	RMSE	wMAPE	R^2
ARIMAX	1135.85	1324.59	33.99%	0.52
SARIMA	1118.96	1305.98	33.49%	0.53
LSTM	802.22	1143.96	24.44%	0.64
TCN	**764.27**	**1026.80**	**23.34%**	**0.70**
Local Visitors				
Model	MAE	RMSE	wMAPE	R^2
ARIMAX	458.61	631.81	33.00%	0.19
SARIMAX	438.33	599.12	31.54%	0.27
LSTM	334.17	465.29	**24.00%**	0.56
TCN	**344.81**	**464.99**	24.89%	**0.57**
Overnight Visitors				
Model	MAE	RMSE	wMAPE	R^2
ARIMAX	212.06	278.99	34.50%	0.28
SARIMAX	210.54	277.53	33.28%	0.29
LSTM	205.87	272.06	29.76%	0.31
TCN	**188.56**	**239.09**	**26.73%**	**0.40**

[a]Bold values indicate the best-performing model for each evaluation metric.

Figure 1 compares the predicted daily flow of all site visitors by the TCN model with the actual observations on the test dataset. TCN demonstrated the ability to accurately predict the timing of peak visitation days. This capability suggests a practical application for tourism site management: by identifying potential high-demand days in advance, operators could allocate resources, adjust staffing levels, or implement crowd management measures.

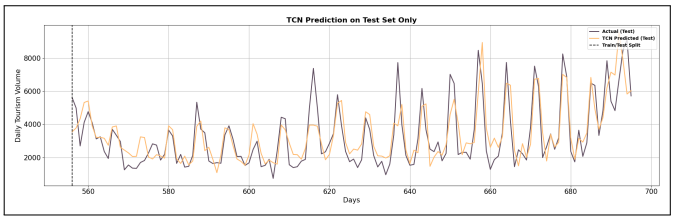

Fig. 1. Test Set Forecasting Results for All Site Visitors

4. CONCLUSIONS AND FUTURE WORKS

This study proposed a TCN-based deep learning approach for forecasting daily tourist visits using high-resolution GPS data. By applying stay point detection, we extracted visits from raw location logs and constructed time series representing daily visitor volumes Our experimental results show that the TCN model consistently outperformed traditional time series models. This forecasting capability can help tourist sites adjust facility and staff operations in advance to better accommodate expected visitor flows. In future work, we aim to extend our approach in two directions: (1) forecasting hourly visitor counts to support more precise tourism management, and (2) predicting more advanced behavioral metrics such as the proportion of long-stay visitors and the share of arrivals from specific municipalities. These enhancements would provide deeper insights into tourist profiles and support more targeted promotional campaigns and services.

ACKNOWLEDGMENTS

We thank GeoTechnologies, who kindly provided the mobility data for this study. This study partially includes the results of the Cross-ministerial Strategic Innovation Promotion Program (SIP) 3rd Phase, "Development of Smart Mobility Platform" promoted by the Council for Science, Technology and Innovation, Cabinet Office (Project Management Agency: New Energy and Industrial Technology Development Organization (NEDO) (Project Code JPNP23023)).This study was supported by JST SPRING, Grant Number JPMJSP2108.

REFERENCES

1. M. Akal, Tour. Manag. **25**, 565 (2004).
2. M. Hu, H. Li, H. Song, *et al.*, Tour. Manag. **90**, 104490 (2022).
3. P. F. Bangwayo-Skeete and R. W. Skeete, Tour. Manag. **46**, 454 (2015).
4. M. Hu, M. Li, Y. Chen, and H. Liu, Tour. Manag. **106**, 105004 (2025).
5. J.-W. Bi, H. Li, and Z.-P. Fan, Ann. Tour. Res. **90**, 103255 (2021).
6. W. Zheng, X. Huang, and Y. Li, Tour. Manag. **59**, 267 (2017).
7. A. D'Antonio, C. A. Monz, B. Crabb, *et al.*, Appl. Geogr. **147**, 102771 (2022).
8. D. Schmücker and J. Reif, Ann. Tour. Res. Empir. Insights **3**, 100061 (2022).
9. S. Bai, J. Z. Kolter, and V. Koltun, arXiv preprint arXiv:1803.01271 (2018).
10. C. Lea, M. D. Flynn, R. Vidal, *et al.*, "Temporal Convolutional Networks for Action Segmentation and Detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition,* (2017), pp. 156–165.

Parameter-free robust stop-detection with HDBSCAN

FRANCISCO BARRERAS[1,*] AND ANDRES MONDRAGON[1]

[1] Department of Computer and Information Science, University of Pennsylvania, 3333 Chestnut Street, Philadelphia, PA 19104

The study of human mobility has advanced greatly in recent years due to the availability of commercial large-scale GPS trajectory datasets [1] and, while there has been a rich variety of applications leveraging these data, the validity of research findings depends highly on the robustness of the processing algorithms used to go from "raw" GPS pings—(`lon, lat, time`)—to the mobility metrics used in each study. A large number of applications perform an initial clustering algorithm known as *stop-detection*, in which pings close in space and time—which might reveal that an individual is stationary—are grouped into *detected stops*. However, choosing the right algorithm and parameterization is not an easy task—the optimal choice might depend on how dense and 'bursty' a user's signal is, or how broad of an area they are roaming, and, while there are many open-source algorithms implemented, there is little guidance on how to choose these parameters optimally.

In particular, for stop-detection algorithms like the ones in [2–6], a 'coarse' parameterization—i.e. with looser criteria to cluster nearby pings—might merge together consecutive visits to small neighboring establishments. However, attempting to address this by using a 'fine' parameterization—i.e. with stricter criteria for clustering nearby pings—could result in 'splitting' a true stop into multiple shorter stops, or in not clustering the pings altogether, thus 'missing' the stop (see 1 for definitions). This problem is further aggravated by the fact that users with otherwise similar mobility could have much lower pinging frequencies, longer gaps in their signals—ranging from hours to several days—as well as measurement noise ('horizontal accuracy'). These factors affect the perceived density of pings—hence, the clustering—in a given area at a given time while being *a priori* unrelated to true mobility.

To address the issue of robustness to parameter choices in stop detection, we present a parameter-free algorithm which adapts `hdbscan` [7] to handle noisy and sparse spatiotemporal data. We evaluate our algorithm in a synthetic sandbox environment—Garden City [8]—which simulates realistic agent-based trajectories with sparsity, burstiness, GPS noise, and, crucially, with ground-truth data that allows comparing detected stops with true stops. Our proposed method outperforms different parameterizations of widely-used sequential and density-based stop detection algorithms across different tasks.

Table 1. Definitions of stop-detection terms and parameters.

Term	Definition
True stop	Part of a trajectory with stationary behavior.
Detected stop	Pings clustered by a stop-detection algorithm.
Missing stop	A true stop not detected by the algorithm.
Split stop	A true stop with separate parts correctly detected.
Merged stop	A true stop partly detected as a wrong location.
Invalid stop set	A group of stops with overlapping time intervals.
$\Delta_{\texttt{max}}$	Maximum diameter allowed for a stop in Lachesis.
ϵ	Max. distance between neighbor pings in DBSCAN.
`max_gap`	Max possible time gap between neighbor pings.
`min_pts`	Min neighbors for a ping to be 'dense' (core).

1. RELATED WORK

Stop-detection algorithms can broadly be classified into *sequential* or *density-based*. The former—e.g. as in 'project Lachesis' [4]—process pings chronologically, assimilating pings into a stop until a stopping condition is triggered—such as the stop exceeding a diameter $\Delta_{\texttt{max}}$ or having a time gap larger than `max_gap`. While simple and efficient, these methods are highly sensitive to measurement noise: a single outlier can end a stop prematurely or cause unnecessary splits.

Density-based methods, typically variations of DBSCAN [9], group pings based on spatial proximity—regardless of time—based on the reachability graph of the set of 'core' points—those with at least `min_pts` neighbors within a radius ϵ. It is well-known that these methods are more robust to noise, and can produce clusters of arbitrary size. They, however, depend on choosing a good radius ϵ, which is a more abstract parameter and dependent on signal density—which is heterogeneous over users and over time. Moreover, applying DBSCAN directly to GPS data neglecting the time dimension would certainly result in invalid stops that have temporal overlaps—a problem known as the 'many-visits-to-one-location' problem. Some variants address this problem through pre- and post-processing of the data, such as restricting the graph of neighbors— based on a time threshold `max_gap` as in `ta-dbscan` [3] or dynamically T-DBSCAN [6]—or simply splitting overlapping stops ex-post. Alternatively, some sequential algorithms implement 'DBSCAN-like' features to increase robustness, for instance those algorithms used by companies like SafeGraph [2] or Cuebiq—which uses a variant of the SOC algorithm [5]. Unlike DBSCAN, which requires a fixed ϵ, `hdbscan` [7] builds a hierarchy of clusters for all values of ϵ simultaneously. Each level of the hierarchy corresponds to a DBSCAN clustering at a particular ϵ. A final clustering is selected by choosing an 'antichain' in the hierarchy—a set of non-overlapping clusters from different levels. This approach effectively clusters points "as if" using different values of ϵ for different regions of the domain. Our method adapts `hdbscan` to spatiotemporal data by modifying the cluster selection step and adding post-processing to avoid invalid stops, as described next.

2. DATA AND METHODS

TA-HDBSCAN. `hdbscan` implements a greedy algorithm that selects the 'antichain' of clusters with the maximum sum of *cluster stabilities* [7]—a metric of the 'range of persistence' of a cluster as epsilon shrinks. The 'true' modes of the density are clusters that persist over a wide range, while 'spurious' clusters disappear quickly in comparison. In `ta-hdbscan`, we adapt this procedure to incorporate time, defining (and discarding) clusters as spurious when they do not have a minimum duration, and using a definition of cluster stability that gives less weight to the stability of 'redundant' pings in a short burst, compared to more 'isolated' pings which better support cluster.

Baseline algorithms. We compare our method against `lachesis` [4] and `ta-dbscan`[3], each evaluated with two parameterizations: one 'coarse' (better suited for sparse signals) and one 'fine' (better for dense signals). For reference, we also implement an `oracle` baseline with perfect knowledge of which pings in a

Fig. 1. Percentage of stops with errors of each type. `lachesis_coarse`:($\Delta_{max} = 40$m, `max_gap` $= 240$min), `lachesis_fine`:($\Delta_{max} = 25$m, `max_gap` $=$ 120min); `ta-dbscan_coarse`: ($\epsilon = 30$m, `max_gap` $= 240$min, `min_pts` $= 2$), and `ta-dbscan_fine`: ($\epsilon = 20$m, `max_gap` $= 120$min, `min_pts` $= 3$).

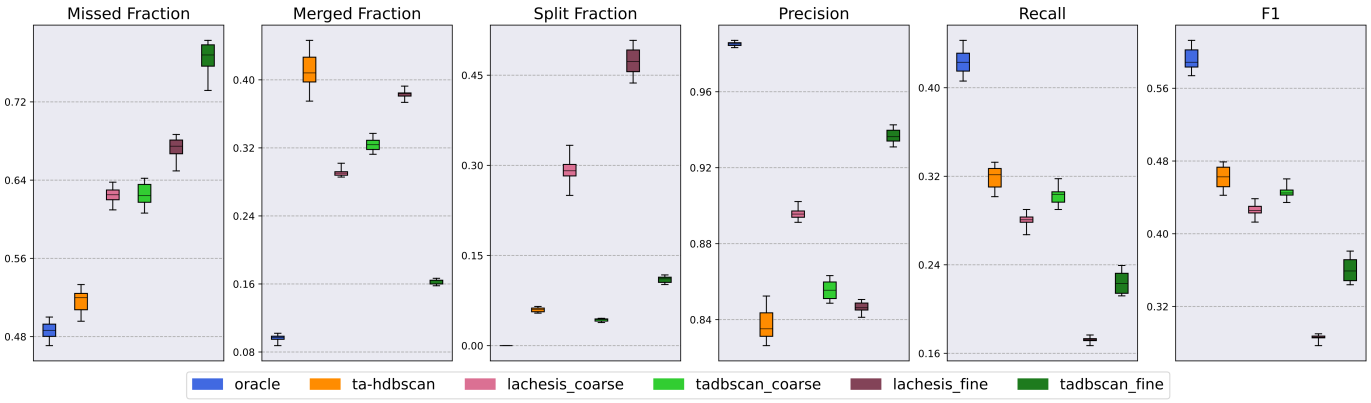

Table 2. Performance of `hdbscan` against baseline algorithms (best parameterization) for different (true) stop types.

Stop category	Value	Algorithm	F1	F1 (% of Oracle)
building_type	home	lachesis	0.50	72.97
		ta-hdbscan	**0.55**	**81.28**
		tadbscan	0.53	80.02
	park	lachesis	0.21	64.83
		ta-hdbscan	**0.24**	**79.28**
		tadbscan	0.17	55.97
	retail	lachesis	0.15	60.14
		ta-hdbscan	**0.15**	**56.66**
		tadbscan	0.13	46.15
	work	lachesis	0.40	80.50
		ta-hdbscan	**0.44**	**88.42**
		tadbscan	0.41	81.94
dwell_length	high (>3h)	lachesis	0.48	73.91
		ta-hdbscan	**0.53**	**82.46**
		tadbscan	0.51	80.57
	low (<1h)	lachesis	0.15	58.72
		ta-hdbscan	**0.14**	**56.46**
		tadbscan	0.11	43.76

sparse trajectory input belong to the same stop, to reflect the limits arising from missing data. The F1-scores in Table 2 are expressed as a percentage of `oracle`'s F1.

Post-processing. `ta-dbscan` and `ta-hdbscan` apply a simple recursive post-processing to break-apart invalid stops: for any set of overlapping stops, we find a minimal segment in the middle of the trajectory such that the 'tails' do not have invalid stops, and re-apply the algorithm on that middle and on each tail. Recursing if the middle also has invalid stops.

Error metrics. The aforementioned problems of "merging", "splitting" and "missing" (Table 1) are measured by counting the number of true stops that (temporally) intersect, respectively, a stop attributed to another location; multiple stops attributed to the correct location; and no detected stops at all. Furthermore, we can frame stop detection as a multi-class prediction task, assigning each minute of a trajectory to a stop (or gap) then assigned to a building ID (or 'street'). This allows us to compute per-minute precision, recall, and F1-scores.

Data. Refer to Data Availability section.

3. RESULTS AND CONCLUSION

Figure 1 shows that `ta-hdbscan` achieves the lowest missed-stop fraction and highest F1 score among all algorithms, recovering more true locations without the need for parameter tuning. As expected, sequential methods like `lachesis` show more splitting (due to sensitivity to noise), while density-based methods tend to merge stops. `ta-hdbscan` avoids this trade-off and performs comparably to the `oracle` baseline.

Table 2 breaks down F1 scores by stop type. `ta-hdbscan` consistently outperforms other methods in homes, parks, and workplaces, and in long-dwell stops, all of which contribute substantially to mobility metrics. The exception is short stops and visits to retail locations—also characterized by short dwell times. Interestingly, `lachesis` outperforms `ta-dbscan` in longer stops and buildings with large footprint (like 'park'), showing that perhaps the density constraints come at a cost in regimes in which noise is less of a problem.

Conclusions. We propose a parameter-free stop-detection algorithm based on `hdbscan` that overcomes the limitations of parameter setting and lack of robustness. The results are limited by the use of synthetic data and its possible idiosyncrasies, but they represent a proof of concept and possible improvement over existing processing pipelines.

Funding. This work was possible with the support of the National Science Foundation, NSF award 2341932.

Data availability. Our dataset consists of 3 week-long trajectories from 200 users generated using the Garden City trajectory generator [8]. 'Macro' trajectories follow an Exploration and Preferential Return model, and 'Micro' trajectories follow a constrained brownian motion. Available at https://github.com/Watts-Lab/nomad/tree/main/examples/research.

REFERENCES

1. F. Barreras and D. J. Watts, Nat. Comput. Sci. pp. 1–14 (2024).
2. SafeGraph, "Determining points of interest visits from location data: A technical guide to visit attribution," https://www.safegraph.com/guides/visit-attribution-white-paper (2025). Accessed: 2025-06-30.
3. D. Birant and A. Kut, Data & knowledge engineering **60**, 208 (2007).
4. R. Hariharan and K. Toyama, "Project lachesis: parsing and modeling location histories," in *International Conference on Geographic Information Science*, (Springer, 2004), pp. 106–124.
5. L. Xiang, M. Gao, and T. Wu, ISPRS Int. J. Geo-Information **5**, 29 (2016).
6. W. Chen, M. Ji, and J. Wang, Int. J. Online Eng. **10** (2014).
7. R. J. Campello, D. Moulavi, and J. Sander, "Density-based clustering based on hierarchical density estimates," in *Pacific-Asia conference on knowledge discovery and data mining*, (Springer, 2013), pp. 160–172.
8. T. H. Li and F. Barreras, arXiv preprint arXiv:2412.00913 (2024).
9. M. Ester, H.-P. Kriegel, J. Sander, *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise." in *kdd*, , vol. 96 (1996), pp. 226–231.

Tackling the Ambiguity of Stop Detection in Human Mobility Data: A Model-Agnostic Evaluation Framework

KAMIL SMOLAK[1,*] AND KATARZYNA SILA-NOWICKA[2]

[1] Institute of Geodesy and Geoinformatics, Wrocław University of Environmental and Life Sciences, Wrocław, Poland
[2] School of Environment, University of Auckland, Auckland, New Zealand
[*] kamil.smolak@upwr.edu.pl

1. INTRODUCTION

As human mobility research expands, calls for methodological standardisation grow [1]. Without clear benchmarks and common practices, studies can be difficult to compare or replicate [2]. Stop detection is a key preprocessing step, identifying locations where individuals remain stationary and extracting valuable information from raw trajectories. Numerous algorithms have been proposed to handle noise and extract relevant places [3–5], yet they are difficult to compare without ground-truth, since mobility patterns differ between individuals and environments. Methods can define a 'correct' stop differently, making it difficult to select the best approach. This variability and ambiguity mean that no single algorithm can be universally optimal.

We propose a model-agnostic framework with interpretable metrics to address this foundational challenge in evaluating stop detection methods. Specifically, these methods can produce predictable and simplified trajectories or rich and chaotic ones. By balancing individual-level predictability with trajectory complexity, our framework identifies Pareto-optimal solutions, allowing analysts to tailor stop detection algorithms to their needs. Furthermore, we show that selecting such a solution yields near-optimal stop detection settings.

2. DATA AND METHODS

We used a high-resolution human mobility dataset provided by the Near company, collected via smartphone applications. The data provider complies with privacy regulations and collects data with informed user consent. The dataset contains individual mobility trajectories in New Zealand, including location (longitude and latitude), timestamp, and a unique device identifier. It spans a 31-day period, from 2021-01-30 to 2021-03-01. Data preprocessing involved removing low-accuracy data points and filtering out inactive devices.

To establish a ground-truth for stop detection, we manually annotated stops in a subset of trajectories. We randomly selected 18 individuals, ensuring broad spatial coverage throughout the country. For each trajectory, we used a geographic information system to label points as either "moving" or "stopped." A stop was defined as at least three points within proximity for at least 5 minutes and validated against the map background to exclude artefacts (e.g., traffic jams). This also captured short stops such as bus waits or quick shop visits. Each stationary period was assigned a unique stop identifier, and recurring stops at the same location were grouped under a common label to reflect spatial consistency. This manual annotation process resulted in a total of 1,249 identified stops.

We evaluated the performance of three stop detection algorithms against this ground-truth: (1) a two-step method known as Project Lachesis [4], which filters stationary points based on rule-based thresholds before clustering them with DBSCAN; (2) Spatio-Temporal DBSCAN (ST-DBSCAN) [5]; and (3) Infostop [3].

Each algorithm has its own set of parameters that influence the detection outcomes. We applied each algorithm to manually labelled trajectories across a range of parameter settings. The performance of each setting was evaluated using a time-based Intersection-over-Union (IoU) metric, which quantifies the temporal overlap between predicted and ground-truth stop intervals. We defined the final score as the harmonic mean of IoU precision and IoU recall:

$$\text{Score} = \frac{2 \cdot \text{IoU}_{\text{pred}} \cdot \text{IoU}_{\text{gt}}}{\text{IoU}_{\text{pred}} + \text{IoU}_{\text{gt}}} \tag{1}$$

where

$$\text{IoU}_{\text{pred}} = \frac{\sum\limits_{p \in P} d(p) \cdot \max\limits_{g \in G} \text{IoU}(p, g)}{\sum\limits_{p \in P} d(p)} \tag{2}$$

$$\text{IoU}_{\text{gt}} = \frac{\sum\limits_{g \in G} d(g) \cdot \max\limits_{p \in P} \text{IoU}(p, g)}{\sum\limits_{g \in G} d(g)} \tag{3}$$

Here, P and G denote the sets of predicted and ground-truth stop intervals, respectively, while $p \in P$ and $g \in G$ refer to individual intervals. The function $d(\cdot)$ gives the duration of an interval. The term $\text{IoU}(p, g)$ measures the temporal overlap between a predicted stop p and a ground-truth stop g, provided that the two also overlap spatially.

Our framework guides parameter selection for stop detection algorithms by balancing individual mobility predictability with data richness. It penalises both trajectory oversimplification, which yields spuriously high predictability, and oversegmentation, which inflates entropy. To achieve this, we employ the real entropy S_{real} as a richness measure and the real predictability Π_{max} as defined in [6].

These measures require representing each trajectory as a sequence of visited locations, encoded as unique symbols (labels). Subsequent studies show that temporal aggregation strongly affects entropy and predictability estimates [2]. To account for this, we apply two common processing approaches: next place (NP) and next time-bin (NTB). In NP, the trajectory is encoded as a sequence of transitions between locations. In NTB, the trajectory is discretised into regular intervals of length Δt, with each interval assigned the symbol of the visited location. In this work, we set $\Delta t = 30$ min.

We assess candidate solutions by their balance between predictability Π_{max} and real entropy S_{real}. We identify the Pareto front, which is the set of solutions where neither metric can be improved without reducing the other. We evaluate solution quality by measuring how often the Top-1 and Top-10 solutions (ranked by the combined Score) lie on the Pareto front and by the proportion of candidate solutions it contains.

For each Pareto front, we also select a representative solution. This is defined as the solution closest to the ideal point,

where both metrics are maximised. For comparison, we report scores for stop detection methods with parameters chosen via sensitivity-based analysis. In this approach, parameters were chosen from neighbouring groups where the variability in the number of detected stops was minimal, indicating well-separated and stable stop detection.

3. RESULTS AND DISCUSSION

Tables 1 and 2 present the final scores, calculated according to Equation 1. Table 1 reports how often the Top-1 and Top-10 solutions appear on the Pareto front for trajectories represented with NP and NTB. The Pareto front reflects the trade-off between predictability Π_{max} and real entropy S_{real}. Table 2 presents the average scores of stop detection methods with parameters chosen by the distance-based criterion, along with the scores from the sensitivity analysis and the maximum achievable scores. Figure 1 illustrates an example of the candidate solution space for an individual using the Infostop algorithm.

Table 1. Proportion of Top-1 and Top-10 solutions included in the Pareto front and the corresponding front size, calculated for each method and trajectory representation.

Method	Rep.	Top-1 (%)	Top-10 (%)	Pop. (%)
Project Lachesis	NTB	**22%**	56%	20%
	NP	**22%**	**67%**	**17%**
ST-DBSCAN	NTB	50%	**100%**	45%
	NP	**67%**	83%	**16%**
Infostop	NTB	**50%**	**78%**	**17%**
	NP	28%	56%	**17%**

The proposed optimisation approach performs best with ST-DBSCAN and Infostop algorithms, achieving Top-1 solutions in more than 50% of cases and Top-10 solutions in nearly 80% of cases when using the NTB representation. Performance for the Project Lachesis algorithm is slightly lower in terms of Top-1 solutions but comparable for Top-10 coverage. Overall, performance is higher with the NTB representation. In most cases, the Pareto front reduces the candidate solution space to less than 20%.

Table 2. Scores from distance-based and sensitivity-based parameter selection, with maximum achievable scores for each method. Values in parentheses show the average distance to the optimal solution.

Method	Rep.	Distance-based	Sensitivity-based	Max
Project Lachesis	NTB	**0.895 (3.5)**	0.893	0.963
	NP	0.865 (17.4)		
ST-DBSCAN	NTB	**0.716 (162.8)**	0.472	0.811
	NP	0.559 (155.8)		
Infostop	NTB	**0.810 (24.2)**	0.796	0.933
	NP	0.810 (123.7)		

Solutions selected with the distance-based approach consistently achieve the best scores across all methods, particularly with the NTB representation. They also outperform sensitivity-based selection. Overall, Project Lachesis emerges as the best-performing and most robust method, also achieving the highest maximum score.

These results demonstrate that balancing predictability and entropy enables the selection of objectively good trajectory segmentation solutions. This balance ensures that the method captures meaningful and structured mobility patterns while properly reflecting both movement complexity and regularity. However, optimal solutions appear at different points along the Pareto front, indicating variability in this balance.

Previous studies show that processing mobility data without accounting for predictability and entropy often leads to over-simplification or over-segmentation, yielding trivial or noisy movement representations [2]. Our method addresses this issue by providing a simple and practical tool that mitigates such biases in mobility data processing. Importantly, it can be used to select appropriate parameters for stop detection algorithms even without ground-truth labels. We hypothesise that the method that best balances predictability and entropy also produces the most meaningful and accurate trajectory segmentations. This is suggested by the relationship between the shortest distance to the ideal solution and the algorithm scores in 2.

It is important to note that the metrics introduced in [6] are purely trajectory-based and do not incorporate contextual information. The consistently better performance of the NTB representation suggests that the temporal structure adds richness to the trajectories. It also highlights the importance of context in method performance.

This study is limited by manual annotations on only 18 trajectories, which constrains ground-truth diversity but suffices to demonstrate feasibility. Larger annotated datasets, more extensive experiments, and the incorporation of contextual information with refined solution selection criteria are important directions for future work to enhance robustness and practical utility.

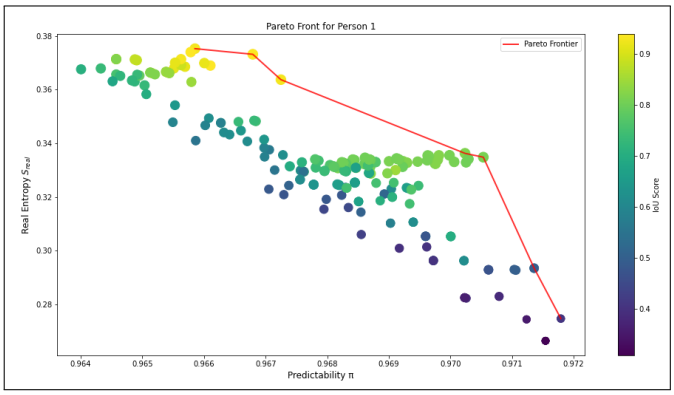

Fig. 1. Candidate solutions for one individual trajectory, calculated using the Infostop algorithm. The color indicates the score, red line is the Pareto front.

Funding. This work was supported by the Wrocław University of Environmental and Life Sciences (Poland) as part of the research project no. N060/0003/24 and the National Science Centre, Poland, grant no. 2021/41/N/HS4/03084.

Data Availability. Mobility data cannot be shared for privacy reasons.

REFERENCES

1. L. Pappalardo, E. Manley, V. Sekara, and L. Alessandretti, "Future directions in human mobility science," Nat. Comput. Sci. **3**, 588 (2023).
2. K. Smolak, K. Siła-Nowicka, J.-C. Delvenne, *et al.*, "The impact of human mobility data scales and processing on movement predictability," Sci. Reports **11**, 15177 (2021).
3. U. Aslak and L. Alessandretti, "Infostop: Scalable stop-location detection in multi-user mobility data," arXiv pp. 1–5 (2020).
4. R. Hariharan and K. Toyama, "Project lachesis: Parsing and modeling location histories," Lect. Notes Comput. Sci. (including subseries Lect. Notes Artif. Intell. Lect. Notes Bioinformatics) **3234**, 106 (2004). ISBN: 9783540235583.
5. D. Birant and A. Kut, "ST-DBSCAN: An algorithm for clustering spatial–temporal data," Data & Knowl. Eng. **60**, 208 (2007).
6. C. Song, Z. Qu, N. Blumm, and A. L. Barabási, "Limits of predictability in human mobility," Science **327**, 1018 (2010).

TransGM: Transferable Gravity Models for Adaptive Urban Policy

ADHAM ENAYA,[1,*] CHEN ZHONG,[1] MICHAEL BATTY[1]

[1] *Centre for Advanced Spatial Analysis, University College London*
*adham.enaya.22@ucl.ac.uk

1. INTRODUCTION

Cities adopt transport policies—such as bus rapid transit, congestion charging, and mixed-use regeneration—by learning from global exemplars, yet policy transfer rarely works "off the shelf" because variations in land use, governance, and behaviour alter impacts (Dolowitz & Marsh, 2000). Gravity models remain central for simulating the impact of policy interventions due to their simplicity and computational tractability, but their fixed coefficients often fail to generalize across contexts. (Batty, 2010; Lopane et al., 2023; Wilson, 2018). Conversely, machine learning advances in transfer learning offer strong predictive across diverse domains (Feng et al., 2024; Gong et al., 2025). We propose **TransGM**, a framework that balances interpretability, generalisability, and adaptability across cities by (i) encoding structural spatial differences, (ii) training a regularised gravity model in a data-rich source city, and (iii) adapting it to a data-scarce target city using divergence metrics.

2. METHODS

Spatial divergence encoding maps land-use distributions from OpenStreetMap onto 8×8 grids, quantifying local similarity via Jensen–Shannon Divergence (JSD) between the unified density probabilities across two cities, considering the local windows in the measurement as per Eq. (1).

$$\mathrm{JSD}_{\mathrm{overall}} = \frac{\frac{1}{N}\sum_{i=1}^{N} \mathrm{JSD}_{\mathrm{local}}(P_i \parallel Q_i)}{\mathrm{JSD}_{\mathrm{global}}(P, Q)} \tag{1}$$

Regularized gravity model incorporates POI-type–specific attraction weights (γ) and applying a log-linear transformation combined with L2-regularized regression to estimate parameters using gradient-based optimisation, as given by Eqs. (2) and (3).

$$T_{ij} = O_i^{\alpha} \cdot \sum_{k=1}^{n} \left(\gamma_k \cdot \mathrm{POI}_{j,k}\right) \cdot d_{ij}^{-\beta} \tag{2}$$

$$\mathcal{L} = \frac{1}{N}\sum_{i,j} \left(P_{ij} - y_{ij}\right)^2 + \eta(|\gamma|_2^2 + \alpha^2 + \beta^2) \tag{3}$$

Adaptive transfer uses a divergence-weighted regularization that penalizes deviation from source-city parameters as formulated in Eq. (4). This adaptive mechanism preserves transferable knowledge while allowing flexibility to adjust to the target city's spatial characteristics, updating model coefficients by preserving aligned features and flexibly adapting those that diverge across cities.

$$L_{\mathrm{transfer}} = \underbrace{\mathrm{MSE}(y_t, X_t\gamma_s)}_{\text{Prediction Loss}} + \underbrace{K \cdot f_{\mathrm{DA}}\big(\delta(\mathrm{POI})\big) \cdot |\gamma_t - \gamma_s|^2}_{\text{Adaptation Penalty}} \tag{4}$$

3. RESULTS

Model performance showed TransGM achieving R^2 = 0.82 compared with 0.76 for naïve transfer and 0.79 for the baseline in Coventry, with RMSE lowest across subsamples (25–100% flows). Attraction profiles indicated Birmingham favoured finance (0.71), leisure (1.27), and shopping (0.53), while Coventry baseline emphasized transport (0.97) and industry (1.23); adaptation shifted Coventry toward services, increasing finance by 0.49 and leisure by 0.22, reducing industrial and transport influence while increasing the weight of service-related POIs (see **Fig. 1** and **Table 1**).

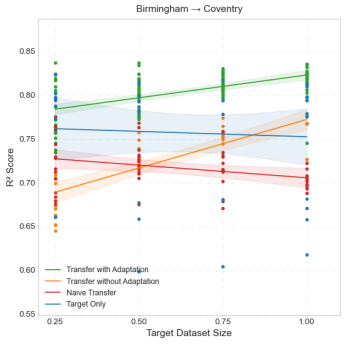

Fig. 1 R^2 scores for four models for transferring model from Birmingham to Coventry.

Travel decay parameters were β = 0.63 for Birmingham, β = 0.56 for Coventry baseline, and β = 0.37 for adapted Coventry, indicating a higher tolerance for longer trips.

Table 1 *Transfer* and *Adaptation* vs. Birmingham baseline.

POI	Birm	Cov.Trans	Cov.Adapt	Cov.Base
BUS	0.493	0.378	0.478	0.461
EAT	0.368	0.181	0.606	0.354
EDU	0.897	0.744	0.836	0.969
LEI	1.267	0.497	1.365	1.143
FIN	0.706	0.203	0.756	0.268
HIS	0.412	0.337	0.409	0.543
IND	0.776	0.542	0.503	1.231

MED	0.537	0.327	0.638	0.963
REL	0.392	0.306	0.454	0.783
SHP	0.53	0.129	0.589	0.613
TRN	0.636	0.792	0.438	1.097
Distance	**0.63**	**0.5**	**0.37**	**0.56**
R2	0.8	0.76	0.83	0.79

In **Fig. 2**, synthetic experiments showed that POI diversity (entropy 0.1–1.0) was strongly correlated with inflows to the centre after adaptation, indicating a replication land-use effect while adhering to local travel patterns, whereas naive transfers failed to alter the dynamics.

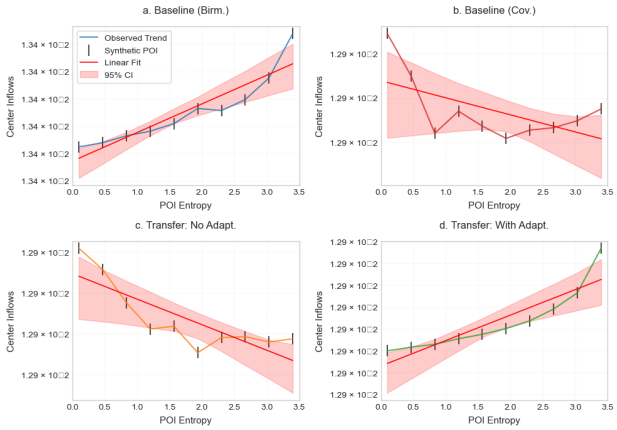

Fig. 2 The correlation between Point of Interest (POI) Entropy and Centre Inflows

4. Discussion and Conclusion

The study shows travel behaviour is highly context-dependent, making direct transfers of urban models or policies between cities unreliable. Comparing Coventry and Birmingham reveals differences in POI and distance decay coefficients that reflect their distinct urban functions: Coventry focuses on transport and industrial POIs, reflecting its industrial heritage, (Healey & Clark, 1984), while Birmingham emphasizes service-oriented POIs, consistent with its diversified, service-driven economy. (Flynn & Taylor, 1986). Adapting Birmingham's POI coefficients for Coventry highlights the importance of context: reweighting leisure and finance shifts Coventry toward a service-oriented profile while reducing industrial and transport influence. A flatter distance decay indicates greater tolerance for longer trips, suggesting coordinated transport and mixed-use planning could enhance accessibility and city-centre vitality. Policy implications include POI reweighting, economic incentives, and travel enhancements (e.g., free transport or parking) to replicate urban dynamics. The study shifts from traditional gravity models to adaptive, data-driven approaches using machine learning regularization to improve generalizability and policy transfer. The modular framework supports applications in transportation, housing, and economic planning, with future work focusing on validating robustness across diverse urban environments.

Funding European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (No. 949670)

Acknowledgment We thank Nathan Wycoff, Beatrice Taylor, and colleagues for feedback.

Disclosures Authors declare no conflicts of interest.

Data Availability Available upon request.

References

Batty, M. (2010, March 11). *Urban Modelling: Algorithms, Calibrations, Predictions*. https://www.semanticscholar.org/paper/Urban-Modelling%3A-Algorithms%2C-Calibrations%2C-Batty/ee16cf2a1099cc02f20611c3560fdfd8867d01b3

Dolowitz, D. P., & Marsh, D. (2000). Learning from Abroad: The Role of Policy Transfer in Contemporary Policy-Making. *Governance*, *13*(1), 5–23. https://doi.org/10.1111/0952-1895.00121

Feng, S., Yao, R., Hess, S., Daziano, R. A., Brathwaite, T., Walker, J., & Wang, S. (2024). Deep neural networks for choice analysis: Enhancing behavioral regularity with gradient regularization. *Transportation Research Part C: Emerging Technologies*, *166*, 104767. https://doi.org/10.1016/j.trc.2024.104767

Flynn, N., & Taylor, A. P. (1986). Inside the Rust Belt: An Analysis of the Decline of the West Midlands Economy. 1: International and National Economic Conditions. *Environment and Planning A: Economy and Space*, *18*(7), 865–900. https://doi.org/10.1068/a180865

Gong, Z., Wang, C., Liu, B., Li, B., Tu, W., Chen, Y., Deng, Z., & Zhao, P. (2025). Multi-spatial urban function modeling: A multi-modal deep network approach for transfer and multi-task learning. *International Journal of Applied Earth Observation and Geoinformation*, *136*, 104397. https://doi.org/10.1016/j.jag.2025.104397

Healey, M., & Clark, D. (1984). Industrial decline and government response in the West Midlands: The case of Coventry. *Regional Studies*, *18*(4), 303–318. https://doi.org/10.1080/09595238400185301

Lopane, F. D., Kalantzi, E., Milton, R., & Batty, M. (2023). A land-use transport-interaction framework for large scale strategic urban modeling. *Computers, Environment and Urban Systems*, *104*, 102007. https://doi.org/10.1016/j.compenvurbsys.2023.102007

Wilson, A. (2018). The Future of Urban Modelling. *Applied Spatial Analysis and Policy*, *11*(4), 647–655. https://doi.org/10.1007/s12061-018-9258-6

Characterizing User Exposure in Mobility Data: Insights from Varying User Behavioral Profiles

LUCAS FÉLIX[1,2], ANNE J. KOUAM[1], ALINE C. VIANA[1], NADJIB ACHIR[1], AND JUSSARA M. ALMEIDA[2]

[1] INRIA, France
[2] Universidade Federal de Minas Gerais, Brazil
[*] lucasgsfelix@ufmg.com

1. INTRODUCTION

The increasing availability of large-scale mobility datasets presents new opportunities to analyze and model human movement at unprecedented scales. However, such data can also compromise user privacy by revealing sensitive personal information (e.g., home location). Studies have shown that just a few spatiotemporal points in a user's trajectory can be enough to uniquely identify them within large datasets [1]. This has spurred growing research on quantifying user exposure—the risk of distinguishing an individual from others—in mobility data [2]. A widely used approach is the *Uniqueness* measure [1], which relates exposure to the unicity in the subsequences of locations visited by the user (i.e., points in the user's trajectory)

Recently, we proposed MoBES—Mobility Behavior-based user Exposure Score—which evaluates exposure based on how distinct a user's *mobility behavior* is from that of others in the dataset [2]. Unlike Uniqueness, which relies on raw trajectory data, MoBES employs a set of mobility metrics to capture multiple complementary facets of mobility behavior, enabling the identification of users with highly distinctive patterns and, consequently, greater exposure risk.

Previous analyses of Uniqueness and MoBES have focused solely on population-level results. However, user mobility behavior is highly heterogeneous and can be categorized into distinct profiles [3]. This raises an open research question: *Which types of users, based on their mobility profiles, are more prone to exposure in mobility datasets?* To our knowledge, neither Uniqueness nor MoBES has been examined through the lens of mobility profiles, a research gap that we address next.

2. METHODOLOGY

We discuss next how we cluster users into mobility profiles, and present the two techniques employed to quantify user exposure.

User mobility profiling: We adopt the proposal by Amichi *et al.* [3] to categorize users into three commonly observed mobility profiles, namely: (i) *Routiners*, or users who mostly follow routine movements (e.g., home-work) with occasional novelty; (ii) *Scouters*, who frequently explore new areas, showing high diversity and low regularity; and (iii) *Regulars*, who exhibit a balanced mix of routine and exploration.

To assign a profile to a user, we follow a two-step methodology. First, each visited location in the user's trajectory is classified based on visit frequency as either an exploration (new or rarely visited place) or a return (frequently visited place). From these classifications, we compute two metrics: intermittency (how often users alternate between exploration and return) and degree of return (how frequently users revisit a return location after an exploration). Second, users are clustered based on these metrics using a Gaussian Mixture Model (GMM), with three clusters. We assess cluster quality using the Silhouette Score.

Quantifying user exposure: We here quantify the the exposure of a user using Uniqueness [1] and MoBES [2]. Uniqueness assesses user exposure by counting the number of unique time-ordered (sub)sequences of locations in a user's trajectory, given an aggregation time window t_w. A Uniqueness equal to 1 indicates the user has unique (sub)sequences, being thus exposed.

MoBES, in turn, quantifies a user's exposure by measuring the distance to their k-nearest neighbors in a multidimensional behavioral space, where each dimension represents a different mobility metric [2]. The underlying assumption is that users who are farther from their nearest neighbors exhibit more distinctive mobility behaviors and are thus more exposed.

Formally, given a set $\mathcal{U}$ of users and a set $\mathcal{M}$ of m mobility metrics, each user $i \in \mathcal{U}$ is represented by a vector $\mathcal{M}_i = \{m_1^i, m_2^i, \ldots, m_m^i\} \in \mathbb{R}^+$, where m_r^i corresponds to the value of metric m_r computed from user i's mobility trajectory. Thus vector $\mathcal{M}_i$ places user i as a point in an m-dimensional behavioral space. To compute the exposure of a user i, we first calculate the distance from i to each other user $j \in \mathcal{U}$ as an average relative absolute difference across all metrics, i.e., as $D_{i,j} = \frac{1}{m}\sum_{r=1}^{m}\left|\frac{m_r^i - m_r^j}{m_r^i}\right|$. Given the set NN_i^k of the k nearest neighbors of i according to the defined distance function, we then compute the MoBES value for user i as the average distance from i to its k-nearest neighbors, i.e., $MoBES_i = \frac{\sum_{j \in NN_i^k} D_{i,j}}{k}$.

MoBES provides each user with a continuous exposure score. Depending on the context, multiple thresholds can be applied to define different exposure categories. Here, we adopt a binary classification by using a single threshold τ: a user i is considered exposed if $MoBES_i > \tau$; otherwise, i is deemed protected due to behavioral similarity with their neighbors.

We emphasize that Uniqueness and MoBES capture user exposure from complementary perspectives. While the former explores only the raw trajectory data, the latter leverages mobility metrics to encapsulate processed insights about how users occupy space over time, thus capturing multiple facets of user mobility behavior. In this work, we set $t_w = 4$ for the time window used by Uniqueness, and instantiate MoBES with $k = 10$, $\tau = 0.05$ and a set $\mathcal{M}$ of $m = 13$ metrics, widely used in the mobility literature. The selected metrics are: Radius of Gyration (RoG), 2k Radius of Gyration (2KRoG), Maximum Distance between Consecutive Locations (MDIST), Number of Visits (NOV), Average and Standard Deviation of Jump Lengths (AJL/SJL), and Number of Unique Locations (NUL), Entropy (ENT), Stationarity (STAT) and Regularity (REG), Average and Standard Deviation of Waiting Time (AWT/SWT). We refer to [2] for a detailed description of these metrics as well as parameter choice.

3. EXPERIMENTAL EVALUATION

Datasets: We assess user exposure using two anonymized cellular datasets from the metropolitan areas of Shanghai and Shenzhen, China. The Shanghai dataset provides near-hourly user locations, while the Shenzhen dataset is event-based, recording locations only when users make calls or send SMS. These datasets differ significantly in properties and mobility patterns; full details on preprocessing and characterization can be found in [4].

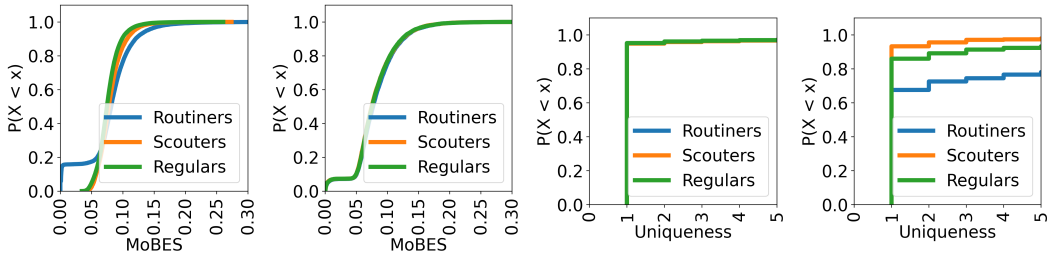

(a) MoBES - Shanghai **(b)** MoBES - Shenzhen **(c)** Uniqueness - Shanghai **(d)** Uniqueness - Shenzhen

Fig. 1. Distributions of MoBES and Uniqueness values in Shanghai and Shenzhen datasets

In brief, the Shanghai dataset includes 58,502 users tracked over 10 days, with an average of 180.37 records and 12.4 unique locations per user. The Shenzhen dataset covers 31,859 users over 15 days, averaging 136.01 records and 9.67 unique locations—indicating more limited mobility in both distinct places and total visitations. User clustering based on mobility profiles [3] shows that in Shanghai, 22.2% are routiners, 56.4% regulars, and 21.4% scouters. In Shenzhen, the distribution shifts to 31.4% routiners, 44.5% regulars, and 24.1% scouters. The higher proportion of routiners in Shenzhen reflects the more limited observation area, with mobility often confined to a few locations (e.g., home and work). The higher fraction of Scouters in Shenzhen, in turn, is due to its event-based collection, which increases data sparsity, recording users in more varied locations.

User exposure: Figure 1 shows the distributions of MoBES and Uniqueness scores across mobility profiles for both datasets.

In Shanghai, Uniqueness yields similar exposure levels across profiles, with most users fully exposed (Uniqueness equal to 1), likely due to: (i) the larger geographic coverage reduces overlap in users' visited areas, increasing the likelihood of capturing unique movements; (ii) frequent location updates (nearly hourly) enhance the detection of individual variations in movement and stationarity. In contrast, MoBES shows clear differences across profiles: nearly 20% of routiners have behaviors very similar to their neighbors, with MoBES values close to zero. Those are overly stationary users, who mostly stay on the same location (e.g., home), being thus "protected" by the behavioral similarity. Scouters and regulars, in turn, display higher MoBES scores, with scouters being slightly more exposed due to greater behavioral variability across multiple mobility metrics.

In Shenzhen, the profiles have distinct Uniqueness distributions: scouters tend to be the most unique, consistent with their higher mobility and greater diversity of visited locations, followed by regulars and routiners. MoBES, in turn, has the same distribution across profiles, despite their distinct mobility patterns. An analysis of the contribution of each metric for MoBES computation reveals also great similarities, with 2k Radius of Gyration, Entropy, and Number of Unique Locations being the most important ones for all profiles. These contrasting results between MoBES and Uniqueness likely reflect intrinsic differences in mobility patterns of both cities, as captured by each dataset.

As argued in [2], regardless of parameterization (e.g., τ value), the sets of users considered exposed according to Uniqueness and MoBES do exhibit important differences, which corroborates that both techniques capture complementary notions of exposure. Due to space constraints, we here focus on the users flagged as exposed by MoBES, using $\tau = 0.05$. In such setup, we find that roughly 82% routiners, 95% regulars, and 98% scouters in Shanghai as well as 90% routiners, 90% regulars, and 90% scouters in Shezhen have MoBES value greater than 0.05, being thus considered exposed.

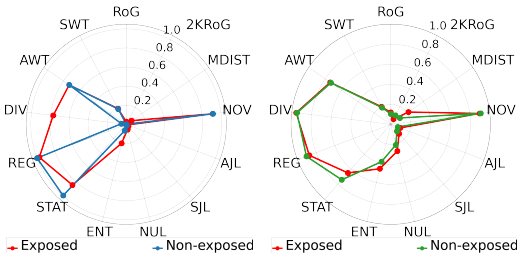

(a) Routiners **(b)** Scouters

Fig. 2. Average user behavior in Shanghai (user exposure according to MoBES with $\tau = 0.05$).

For illustration, Figure 2 presents radar charts showing the normalized average values of mobility metrics for exposed and non-exposed routiners and scouters (according to MoBES) in Shanghai. Clearly, routiners deemed as exposed tend to exhibit more exploratory behavior, with greater diversity/entropy and lower stationary. Conversely, exposed scouters tend to exhibit subtler differences but across more metrics, also related to more exploration, notably lower stationarity, higher entropy as well as longer maximum distance between consecutive locations.

4. CONCLUSION

In conclusion, user behavior may significantly affects exposure, whether considering spatial-temporal patterns or mobility behavior. Yet, this effect is strongly influenced by dataset characteristics. The differences observed for MoBES and Uniqueness in both Shanghai and Shenzhen highlight the complementarity of these measures and their potential for combined use in exposure evaluation. Additionally, mobility profiles can help guide personalized protection strategies, enabling more targeted privacy measures while preserving data utility.

Acknowledgments: This work is part of the Mob Sci-Dat Factory project (ANR-23-PEMO-0004) under the France 2030 program. It was partially supported by *Conselho Nacional de Desenvolvimento Científico e Tecnológico* (CNPq), by *Fundação de Amparo à Pesquisa do Estado de Minas Gerais* (FAPEMIG), by *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* (CAPES), and by the CAPES-STIC-AMSUD 22-STIC-07 LINT project.

REFERENCES

1. Y.-A. De Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel, Sci. Reports (2013).
2. L. Félix, A. J. Kouam, A. C. Viana, *et al.*, "The silent signature: Behavior-based user exposure in mobility data," in *Proc. IEEE MDM*, (2025).
3. L. Amichi, A. C. Viana, M. Crovella, and A. Loureiro, "Unders. individuals' proclivity for novelty seeking," in *Proc. ACM SIGSPATIAL*, (2020).
4. J. P. Esper, A. C. Viana, and J. M. Almeida, "Beauty or beast: Human behavioral insights and learning power of federated mobility prediction," in *Proc. ACM SIGSPATIAL*, (2024).

Understanding collective and individual phenomena in urban mobility using mobile phone datasets

ALBERTO AMADUZZI[1, 2, 3], FILIPPO DALLA[1, 2, *], DRAGOS DUMITRU IOAN[1, 2], GREGORIO BERSELLI[1, 2], DAVIDE MICHELI[4], ALDO VANNELLI[4], AND ARMANDO BAZZANI[1, 2]

[1] *Dipartimento di Fisica e Astronomia "Augusto Righi", Universita di Bologna, via Irnerio 46, Bologna, Italy*
[2] *INFN, Sezione di Bologna, via Irnerio 46, Bologna, Italy*
[3] *Fondazione Bruno Kessler, Via Santa Croce 77, Trento, Italy*
[4] *TIM S.p.A., Roma, Italy*
[*] *filippo.dalla3@unibo.it*

1. INTRODUCTION

The development of future smart cities has driven a new research field where different disciplines cope with the complexity of modern cities: Science of Cities. Complex Systems Physics has contributed significantly by applying Network Theory and Statistical Mechanics to uncover emergent properties through big data analytics.[1]. In particular, urban mobility has become one of the most intensely studied topics, thanks to datasets provided by new Information and Communication Technologies (ICT), which make it possible to relate microscopic individual behavior to the statistical laws of urban mobility through the ubiquitous use of mobile phones. The MDT Tim dataset gives georeferenced, anonymized information on the displacements of mobile devices, with a time frequency of a few seconds, when they're connected to the telecommunication network[2]. In this way, it is possible to study the mobility of individuals moving on the road network[3].
Specific algorithms have been developed to reconstruct the dynamics of mobility paths on the road network that allow one to classify trajectories into distinct typologies and to highlight individual behavior. We first process the MDT GPS dataset, collected during several days between 2022 and 2025 in Bologna, obtaining a sequence of points that can be georeferenced on the road network using OpenStreetMap as the cartography layer, and give dynamical information on a mobility path. The MDT data set is collected by Telecom Italia, one of the main telecommunication providers in Italy with a coverage of 30% of the market. Previous studies[3] estimated the penetration of the MDT datset $\simeq$ 8% of the mobile phone population. We explicitly studied mobility in the metropolitan area of Bologna (a city in the North Italy with $\simeq 900,000$ inhabitants) to understand how mobility depends on the structure of the urban road network and how the congestion effects emerge. This study is part of a city digital twin project. We classify the trajectories to reveal the empirical statistical laws of urban mobility and to propose a measure which define traffic state in the multi-layer road network.[4].
We also used an open dataset measuring traffic flows via magnetic coils[1] which allowed us to estimate $\simeq$ 3% for the penetration of MDT paths with respect the whole traffic flows. Our goal is to analyze the impact of large-scale tram line construction (started in April, 2023), focusing on some key districts.

2. METHODS AND RESULTS

Individual mobility behavior is influenced by the structure of the transport network. Route choices depend both on the road infrastructure and on users' knowledge of the traffic state. To verify these assumptions, we analyzed the dynamic features of MDT paths to highlight the hierarchical (multilayer) structure of Bologna's road network. We applied c-means clustering (FCM) algorithm[5] to define the road sub-network structure associated with different mobility patterns. The algorithm proves efficient to classify $\simeq$ 96% of the selected paths (more details on the application of fuzzy c-means clustering algorithm to the MDT paths are reported[3]). This analysis groups in a sub-network the roads with similar travel speeds and traffic flows. We identify four mobility classes: the slow mobility (class 0), the traffic in the historical center or on narrow streets (class 1), the traffic on main urban roads (class 2), and the highway traffic (class 3); see Fig. 1 for a map representation of the road sub-networks.

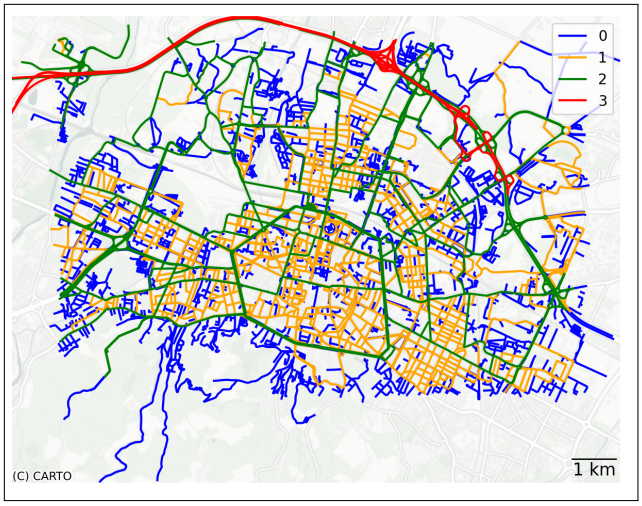

Fig. 1. Map of the hierarchical road sub-networks, each class is characterized by a different color.

In Fig. 2, we show as the velocity distribution of the each class can be interpolated by a Gaussian shape. Each individual performs an origin-destination mobility realizing paths that use the roads belonging to the different classes. This behavior explains the empirical distributions of the path length and the path duration that are shown in Fig. 3,4[4]. The empirical path length distribution suggests that the urban mobility demand in Bologna is well interpolated by an exponential distribution with an average length $\bar{l} = 4.0$ km. The exponential distribution suggests that the individuals performing the urban mobility realize a statistical maximum entropy equilibrium with a mobility energy (the path length). Conversely, the path duration distribution is better fitted by a power law that is an emergent property of the multilayer structure of the urban road network.

[1] https://opendata.comune.bologna.it/pages/home/

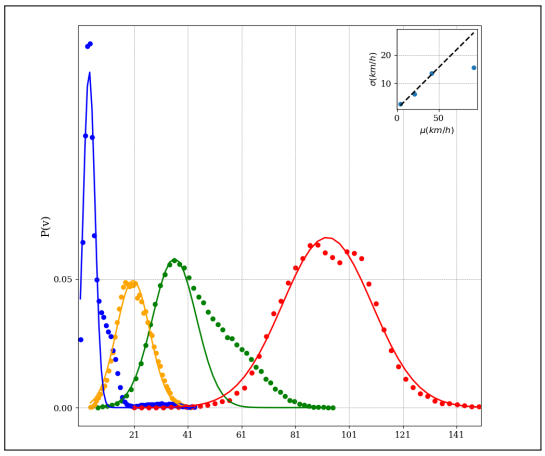

Fig. 2. The probability distribution of average travel speeds for the roads belonging to the 4 classes: the continuous lines refers to an interpolation with a Gaussian distribution.

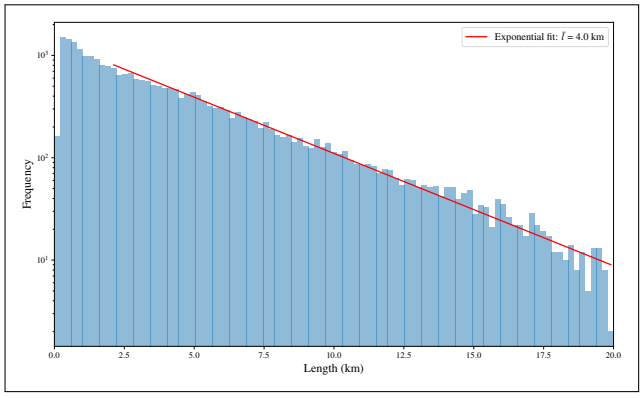

Fig. 3. MDT path lengths distribution for the classes 0,1,2: the histogram refers to frequency of the MDT path length and the continuous line to an exponential interpolation with a characteristic length $\bar{l} = 4.0$ km.

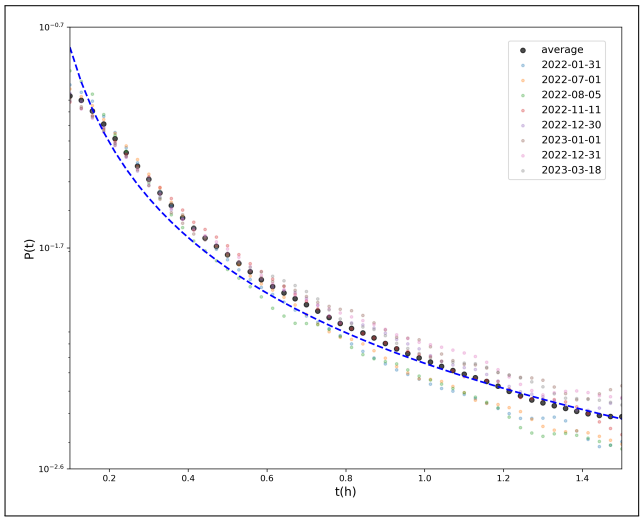

Fig. 4. MDT path duration daily distributions for the classes 0,1,2: the points refer to the different days in the MDT data set and the continuous line to a power law interpolation $p(t) \propto t^{-\alpha}$ with $\alpha = -1.36$.

3. DISCUSSION

We show that the Telecom MDT data set offers the possibility to study the individual behavior during urban mobility and to highlight the hierarchical structure of an urban road network, allowing to explain the macroscopic statistical laws of the urban mobility. The case study of Bologna metropolitan area is particularly interesting due to the ongoing transformations of the urban transport network with the construction of the tram. In Fig. 5 we show how the construction sites for the tram reduce the traffic flows in the city. Using the MTD paths, we can compare

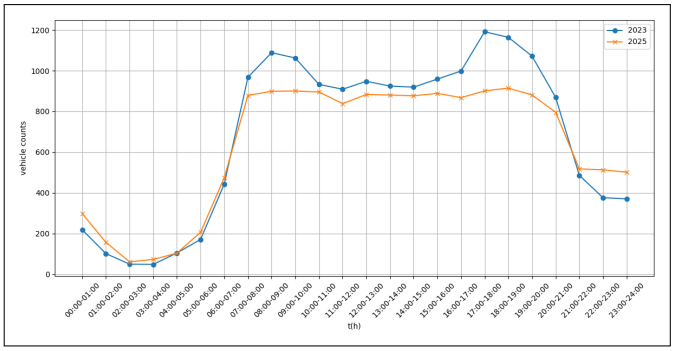

Fig. 5. Comparison of the hourly traffic load in via Stalingrado (heavily affected by tram line construction) between 2023 and 2025, averaged over a typical working week.

how individual changes their mobility habits according to the restrictions on the road networks. The understanding of this decision mechanism could be useful to build predictive models for the impact of new transport infrastructures on the urban mobility of a historical city like Bologna where the mobility demand has a well defined spatial scale. This study is part of a work in progress.

Funding. This study was carried out within the MOST – Sustainable Mobility National Research Center and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.4 – D.D. 1033 17/06/2022, CN00000023).

Acknowledgment. The authors are in debt with Dott. Riccardo Gallotti of Fondazione Bruno Kessler for several useful discussions.

Disclosures. The authors declare no conflicts of interest.

Data Availability Statement. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

REFERENCES

1. M. Batty, *Cities and Complexity* (The MIT Press, 2007).
2. D. Micheli and R. Diamanti, "Statistical analysis of interference in a real lte access network by massive collection of mdt radio measurement data from smartphones," in *2019 PhotonIcs Electromagnetics Research Symposium - Spring (PIERS-Spring)*, (2019), pp. 1906–1916.
3. C. Mizzi, A. Baroncini, A. Fabbri, *et al.*, EPJ Data Sci. **12** (2023).
4. A. Amaduzzi, "Emergent patterns in human mobility: A statistical physics approach," Ph.D. thesis, Dipartimento di Fisica e Astronomia "Augusto Righi", Università di Bologna (2025).
5. F. W. Bezdek JC, Ehrlich R, Comput. Geosci **10 (2-3)**, 191 (1984).

Data Challenge

talks

Protecting participants or population? Comparison of k-anonymous Origin-Destination matrices

PIETRO ARMENANTE[1], KAI HUANG[1], NIKHIL JHA[1], AND LUCA VASSIO[1,*]

[1] *Politecnico di Torino, Italy*
[*] *luca.vassio@polito.it*

1. INTRODUCTION

Origin-Destination (OD) matrices are a core component of research on users' mobility and summarize how individuals move between geographical regions. These regions should be small enough to be representative of user mobility, without incurring substantial privacy risks.

There are two added values of the NetMob2025 challenge dataset. Firstly, the data is extensive and contains a lot of socio-demographic information that can be used to create multiple OD matrices, based on the segments of the population. Secondly, a participant is not merely a record in the data, but a statistically weighted proxy for a segment of the real population. This opens the door to a fundamental shift in the anonymization paradigm. A population-based view of privacy is central to our contribution. By adjusting our anonymization framework to account for representativeness, we are also protecting the inferred identity of the actual population, rather than survey participants alone.

The challenge addressed in this work is to produce and compare OD matrices that are k-anonymous for survey participants and for the whole population. We compare several traditional methods of anonymization to k-anonymity by generalizing geographical areas. These include generalization over a hierarchy (ATG and OIGH) and the classical Mondrian. To this established toolkit, we add a novel method, i.e., ODkAnon, a greedy algorithm aiming at balancing speed and quality.

Unlike previous approaches, which primarily address the privacy aspects of the given datasets, we aim to contribute to the generation of privacy-preserving OD matrices enriched with socio-demographic segmentation that achieves k-anonymity on the actual population.

The full report for this work is available on arXiv[1], with open-source reproducible code available on GitHub[2].

2. METHODOLOGY

A. Geographical area hierarchy definition

We assume a fixed geographical generalization hierarchy of zones. A generalization hierarchy describes how smaller zones should be merged together. We can think of this hierarchy as a tree where the leaves are the initial small zones, and a parent node represents the union of the areas of its child nodes.

In this work, we use the well-known, Uber-developed H3 hierarchy.[3] In H3, the space is organized in hexagonal cells: every cell, up to the maximum resolution supported by H3, has seven child cells below it in this hierarchy—see Figure 1. Hexagons do not cleanly subdivide into finer ones; however, a subdivision into seven cells can be approximated.

B. Algorithms for k-anonymity

We want to find an aggregation of origins and destinations to ensure that each entry corresponds to a group of at least k indistinguishable individuals. We compare algorithms that perform an adaptive generalization such that, for every trip in our anonymized dataset, at least k-1 other trips exist that start in the same origin and finish in the same destination (possibly different from the origin). We consider the state-of-the-art ATG algorithm [1], OIGH uniform tree generalisation algorithm [2], and Mondrian [3]

We propose a new OD-matrix anonymization algorithm, ODkAnon, that is greedy, iterative, and aims at generalizing OD locations in a hierarchy to satisfy k-anonymity [4] constraints. The algorithm considers the fixed, H3-based generalization hierarchy, and iteratively selects and aggregates cells related to lower-density areas (considering both origins and destinations) until the k-anonymity condition is met. This algorithm (as OIGH) creates *homogeneous* geographical areas, i.e., origin areas choice does not depend on the given destination and vice-versa. This means that the final OD matrix will comprise non-overlapping areas. A detailed description of ODkAnon is available in our report, and its implementation is in our GitHub repository.

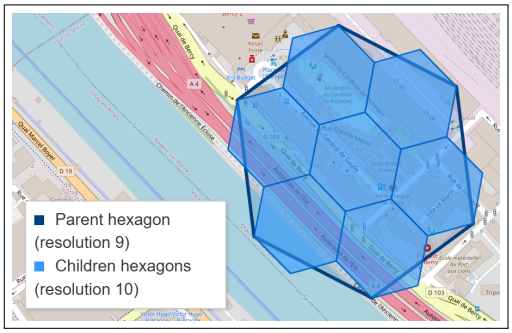

Fig. 1. Example of a H3 hierarchy of hexagons in Paris.

C. Protecting the participants vs. protecting the population

Common algorithms aim at obtaining a k-anonymous dataset, where every item—in this case, a trip—cannot be distinguished from at least k-1 other ones. However, each dataset only represents a subsample of the population, some segments of which may be over- or under-represented by the records in the dataset. The awareness of this data skewness could help re-identification attacks by an adversary. Therefore, risks are not only towards the participant from whom the data has been collected, but also towards the original population that has been sampled.

The weight attribute that the NetMob2025 challenge dataset offers hints at a solution in this direction, indicating the number of people in the Île-de-France area that each participant represents. By using this number to weight each entry in the dataset, we can tweak the algorithms to account for the population rather than the dataset, strengthening the privacy guarantees of the algorithms. Naively, the algorithms will see the same trip repeated as many times as the survey participant's population representativeness.

[1] https://arxiv.org/abs/2509.12950
[2] https://github.com/SmartData-Polito/ODkAnon
[3] https://h3geo.org/docs/

Table 1. Challenge dataset relevant characteristics.

Participants with GNSS records	3,320
Represented population	9,001,164
Trips with GNSS points	81,291
H3 hexagons at resolution 10	44,011
Relevant socio-demographic features	sex, age, profession

D. Performance metrics

The anonymization alters the values of the OD matrix, as well as the zone sizes, depending on the aggregation of the geographical areas in the hierarchy.

Privacy metrics validate the strength of the anonymization. We measure the minimum k-anonymity obtained for participants when protecting the population, and the other way around.

Data utility metrics measure how well the anonymized OD matrix preserves the original data characteristics, comparing the pre- and post-anonymization versions. Among the ones we used, we report here: a generalisation error ($\bar{G}$) that counts the number of merges in the hierarchy [1], a discernibility metric (C_{DM}) that accounts for how many trips are aggregated [3], and a reconstruction loss (E) [1].

3. DATASET CHARACTERISTICS

The NetMob2025 challenge dataset offers a multi-dimensional view of mobility behavior in the Île-de-France region over a continuous 7-day period. We select only participants with complete GNSS records. Then, we consider a trip completed after the absence of GNSS records for 3 minutes. Finally, we extract the origin and destination coordinates of each trip. These are anonymized, generalized to the center of the closest H3 hexagon of resolution 10.

Together with origins and destinations, we employ the weight feature assigned to every trip, describing how many people the record is representative of, and three socio-demographic features (sex, age, and profession) to build segment-specific OD matrices. The main characteristics relevant to our work are reported in Table 1.

4. MAIN RESULTS AND DISCUSSION

The complete results are available in our report. We run the OD-kAnon algorithm on the trips within the Île-de-France, obtaining a single OD matrix aggregating all data. We first consider protecting the participants in the survey, and we set $k = 10$ for obtaining a k-anonymous OD matrix. We allow for suppression of up to 10% of the trips. Merging over the hierarchy, the original thousands of resolution 10 hexagons, we obtain the same 29 zones both for the origins and destinations. Figure 2 reports a zoom over Paris of the obtained generalized areas (blue hexagons).

When protecting the population, k should be adapted, accounting for the representativeness of each participant. Given that a participant on average accounts for 2,674 people, we use $k = 10 \times 2,674$ in order to keep a fair comparison of the two approaches. We now obtain the same 29 destination zones, but more fine-grained 35 origin zones. We show the origin zones in Figure 2 as red hexagons. Protecting the population produces a different anonymization. In particular, 7 smaller hexagons are now generalized to their parent node (the large blue hexagon). Likely, these zones contain few trips from the participants, but are related to a large portion of the population. Even more interesting, when the protection is applied to the population, the participants' OD matrix loses k-anonymity: 13 hexagons fall below the threshold (k = 10), with the smallest value equal to 4.

Table 2. Data utility metrics, after anonymization with OD-kAnon. Metrics are computed considering participant trips. A lower value means that the OD matrix retains more utility.

Protection	Data	OD dimension	$\bar{G}$	C_{DM}	E
Participants	All	29x29	601	5.4×10^7	1.91
Participants	Men	5x2	4,388	4.6×10^7	1.43
Participants	Women	29x29	447	1.8×10^7	1.51
Population	All	35x29	539	5.3×10^7	1.91
Population	Men	2x2	6,727	1.3×10^8	1.86
Population	Women	29x29	450	1.8×10^7	1.87

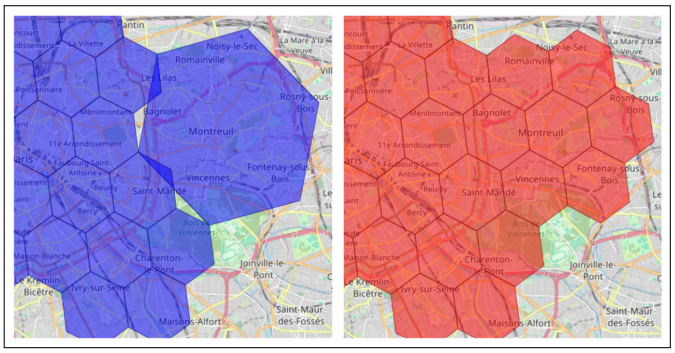

Fig. 2. Detail over Paris of the origin generalization hexagons for the participant-protecting (left, blue hexagons) and the population-protecting (right, red hexagons) definitions.

Moreover, we observed amplified differences when segmenting the population. Consider the example of sex (men and women), with the same k thresholds. The two segments have a similar total number of trips. When protecting the participants, anonymizing the men dataset produces a coarse 5x2 OD matrix, whereas anonymizing the women dataset results in a fine-grained 29×29 matrix. This indicates that protecting men is more challenging, as it requires very coarse hexagons, while for women the resulting hexagons remain much finer. Furthermore, when applying protection to the population, the difference becomes even more pronounced: the anonymized men dataset reduces to a 2×2 matrix.

We report the data utility metrics for the obtained anonymized datasets in Table 2. We computed these metrics considering the survey's participant trips. All these metrics indicate more utility when lower. Notice how the men segment loses utility when anonymized. Full comparisons, including (i) different anonymization algorithms, (ii) multiple socio-demographic segmentations, and (iii) metrics computation for population, are available in our arXiv report.

Our results showed that significant differences may exist when considering population-protecting for OD matrices anonymization, rather than survey participant-protecting ones. Moreover, we showed how these differences can even be amplified across socio-demographic segments.

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. B. Matet, A. Furno, M. Fiore, *et al.*, Transp. Res. Part C: Emerg. Technol. **154**, 104236 (2023).
2. W. Mahanan, W. A. Chaovalitwongse, and J. Natwichai, World Wide Web **24**, 1551 (2021).
3. K. LeFevre, D. J. DeWitt, and R. Ramakrishnan, "Mondrian multidimensional k-anonymity," in *ICDE'06*, (IEEE, 2006), pp. 25–25.
4. P. Samarati, IEEE Trans. on Knowl. Data Eng. **13**, 1010 (2002).

On commuting duration and the daily travel time budget

HAO GUO[1], JUNLONG ZHANG[1], JUNJIE YANG[1], AND LEI DONG[1*]

[1] Institute of Remote Sensing and Geographical Information Systems, School of Earth and Space Sciences, Peking University, Beijing, China
[*] leidong@pku.edu.cn

BACKGROUND

There has been accumulated evidence that individuals' average daily travel time is relatively constant at a regional or national scale[1]. This constant, around or slightly more than an hour, is referred to as the travel time budget or Marchetti's constant, assumed to reflect a behavioral regularity of humans[2, 3].

Among the everyday travel purposes, commuting trips (i.e., individuals' travels between their places of residence and work) play a crucial role in urban mobility[4]. Considering the diverse distribution of commuting time within a city[5], we may ask the following question: Would residents with longer commuting times spend less time traveling for other purposes? The answer seems to be yes, as a consequence of the daily travel time budget. Here, we present contrary evidence from the individual-level travel diary data in the Île-de-France (IDF) region (referred to as the NetMob25 dataset hereinafter): residents' daily non-commuting travel time, averaged over a week, remains relatively stable to the change of commuting time.

DATA AND METHODS

The NetMob25 data[6] is collected from a GPS-based mobility survey from October 2022 to May 2023, covering 3,337 volunteer residents aged 16 to 80 in the IDF region. For each resident, the travel diary dataset contains the start/end time and the purpose (activity type at the origin and the destination) of each trip during a week. Among the seven recorded days for each individual, we first exclude days with no traces or when the resident is outside IDF, and label the remaining days as *valid*. Furthermore, a *valid* day is labeled as a workday if and only if it involves any work or study activity. For each resident, we calculate T_{tot} as the average daily travel time over *valid* days, and T^w_{tot} as the average travel time over workdays.

To estimate the average commuting time of each individual, a common approach is to classify the trips into commuting and non-commuting trips. However, one may stop at other locations on the way of commuting, resulting in a chain of multiple trips. The commuting time would be overestimated if all these trips were considered commuting. To avoid such bias, we only consider direct trips between home and work (or school) which are not separated by other stops, and define their average duration as the individual commuting time T_{comm}.

Next, the average daily non-commuting travel time T_{ncom} for each individual is calculated as

$$T_{ncom} = \frac{T_{tot} \times \#\text{valid_days} - T_{comm} \times \#\text{comm_trips}}{\#\text{valid_days}} \quad (1)$$

where #comm_trips is the number of commuting trips (possibly with intermediate stops) over the week. In other words, we subtract T_{comm} from the total travel time for each commuting trip, then average over *valid* days. Similarly, we compute the average non-commuting travel time for workdays as T^w_{ncom}. Our analysis considers 2,240 individuals with at least 3 *valid* days and 3 direct commuting trips to mitigate the effect of data uncertainty. For analysis over workdays, we further exclude residents with fewer than 3 workday records, resulting in 1,917 residents.

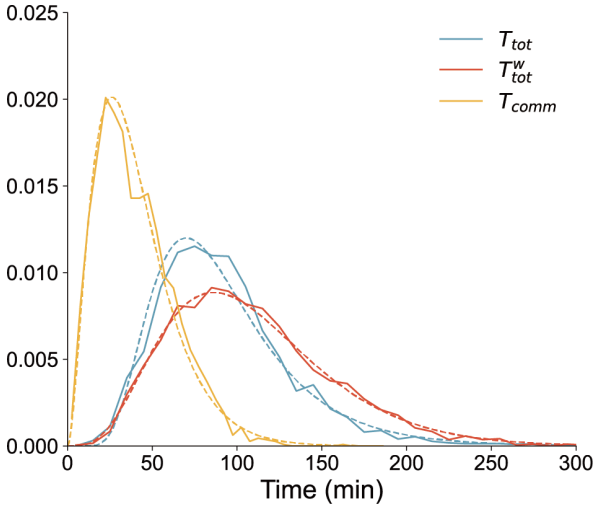

Fig. 1. Distribution of the average daily travel time T_{tot}, the average daily travel time on workdays T^w_{tot}, and average commuting time T_{comm} over the volunteer residents. T_{tot} is fit with a lognormal distribution, while the other two curves are fit with gamma distributions.

RESULTS

Figure 1 shows the distribution of T_{tot}, T^w_{tot}, and T_{comm}, where each point represents a resident. The daily travel time T_{tot} has a mean of 92.5 min and a median of 86.5 min. Generally, residents spent more time traveling on workdays, yielding a mean of 108.8 min and a median of 101.3 min for T^w_{tot}. The average commuting time over 2,240 residents is 40.0 min, and the quartiles are 23 min, 36 min, and 54 min. The distribution of T_{comm} is best modeled with a gamma distribution according to the Akaike information criterion (candidates include lognormal, gamma, exponential, and Pareto), with the probability density function

$$f(t) \propto x^{\alpha-1} e^{-\frac{x}{\theta}} \quad (2)$$

where $\alpha = 5.67$ and $\theta = 16.32$. The gamma property still holds when we consider the duration distribution of the 14,627 direct commuting trips, and when the trips are grouped by the department of residence or the professional group of the individuals. These results suggest the gamma distribution as a stable temporal property of urban commuting.

To explore the association between commuting time and non-commuting travel time, the simple linear regression yields

$$T^w_{ncom} = 38.612 - 0.199\ T_{comm},\ R^2 = 0.015 \quad (3)$$
$$(1.646) \quad (0.037)$$
$$T_{ncom} = 48.795 - 0.053\ T_{comm},\ R^2 = 0.001 \quad (4)$$
$$(1.416) \quad (0.031)$$

where standard errors are given in parentheses. The slope estimate suggest that non-commuting travel time is only slightly affected by the change of commuting time. In contrast, assuming T^w_{tot} is constant, the slope parameter in Eq. 3 should be

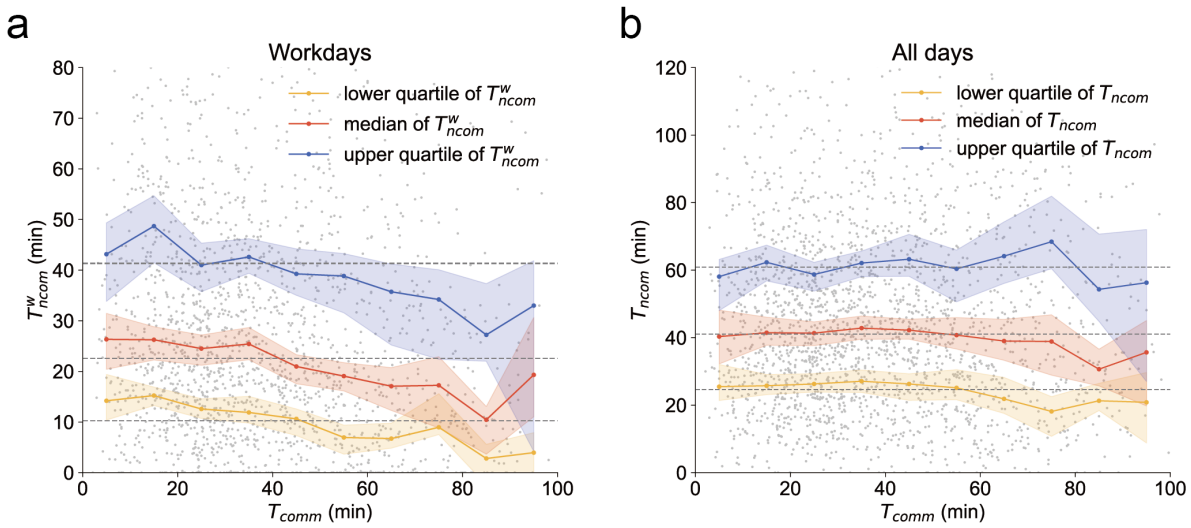

Fig. 2. The relationship between daily non-commuting travel time and commuting time over (a) workdays (T^w_{ncom} against T_{comm}) and (b) all days (T_{ncom} against T_{comm}). Each gray dot represents a resident in the dataset. Yellow, red, and blue lines show the quartiles of non-commuting travel time for each 10 min interval of commuting time with 95% confidence interval from the bootstrap method. Gray dashed lines are the overall quartiles.

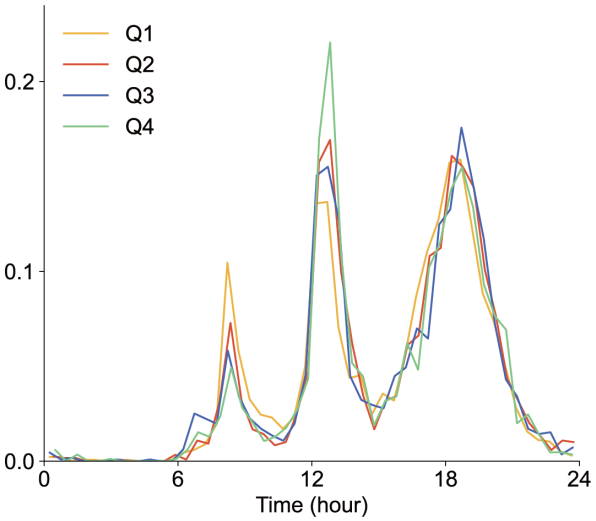

Fig. 3. Start time distribution of non-work activities throughout the day for quartile groups of commuting time.

around -2.0 given two commuting trips per day. Furthermore, if weekends are included, the association is not even significant ($p = 0.087$ in Eq. 4). This finding is further supported by Fig. 2, where we group the residents using commuting time intervals of 10 min and consider non-commuting travel time of each group. For the medians, the non-commuting travel time on workdays of the 50-60 min group is only 7.3 min less than the 0-10 min group (Fig. 2a). If weekends are included, non-commuting travel time remains remarkably stable around the overall median of 41.0 min (Fig. 2b). Instead of the daily travel time budget, these empirical results seem to suggest a hypothesis of "non-commuting travel time budget": the daily non-commuting travel time is relatively stable regardless of the time spent on commuting.

The association between non-commuting travel time and commuting time exhibits heterogeneity among professional groups. A linear regression model of $\ln T_{ncom}$ against $\ln T_{comm}$ (with gender, age, car ownership, and professional groups as dummy control variables) gives an elasticity of -0.076 ± 0.028. When adding interaction terms of T_{comm} and dummy variables indicating professional groups, workers exhibit the strongest elasticity (-0.385 ± 0.126) of all groups. This may indicate potential inequality in access to spare-time activities faced by workers with lower socioeconomic status.

DISCUSSION

We find that the non-commuting travel time of individuals is not necessarily affected by commuting burden, indicating a basic need for spare-time activities such as shopping, recreation, and accompanying. We assume residents accommodate longer commuting time in two possible ways: (i) arranging more non-commuting travels on weekends to compensate for the loss of spare time on workdays, as evidenced in Fig. 2; (ii) rescheduling some non-commuting travels from the morning before work to the lunch break, as suggested in Fig. 3.

Our analysis is currently confined to the IDF region, limiting the generality of results. In the future, we will incorporate individual-level mobility data from other countries to validate our findings.

Funding. This work is supported by the National Natural Science Foundation of China (42422110).

Data availability. This work was performed using data from the Netmob 2025 Data Challenge [6].

REFERENCES

1. P. L. Mokhtarian and C. Chen, Transp. Res. Part A: Policy Pract. **38**, 643 (2004).
2. Y. Zahavi, "The UMOT Project," Tech. rep., U.S. Department of Transportation (1979).
3. C. Marchetti, Technol. Forecast. Soc. Chang. **47**, 75 (1994).
4. T. Louail, M. Lenormand, M. Picornell, *et al.*, Nat. Commun. **6**, 6007 (2015).
5. K. S. Kung, K. Greco, S. Sobolevsky, and C. Ratti, PLOS ONE **9**, 1 (2014).
6. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The NetMob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," (2025).

On the Relationship between Space-Time Accessibility and Leisure Activity Participation

YUAN LIAO[1,2,*], ANTONIO DESIDERIO[2], RAFAEL H. M. PEREIRA[3], JORGE GIL[4], SILVIA DE SOJO CASO[2], AND LAURA ALESSANDRETTI[2]

[1] Department of Space, Earth and Environment, Chalmers University of Technology, Gothenburg, Sweden
[2] Department of Applied Mathematics and Computer Science, Technical University of Denmark, Lyngby, Denmark
[3] Institute for Applied Economic Research (Ipea) - Brazil, Brasília, Brazil
[4] Department of Architecture and Civil Engineering, Chalmers University of Technology, Gothenburg, Sweden
[*] yuan.liao@chalmers.se

Ensuring equitable access to opportunities through urban transport systems is a cornerstone of inclusive and sustainable cities. While transport accessibility is commonly evaluated through network-based models of potential access to opportunities, such measures do not fully capture how different population groups actually experience and navigate urban mobility systems, particularly to those third-place activities that foster meaningful social interactions. Studies have shown that greater proximity to certain destinations is associated with higher participation rates in those activities. To date, our understanding of how transport accessibility shapes leisure activity participation remains largely static. It remains unclear how transport accessibility, when accounting for individual constraints, links to participation in third-place activities and how this varies across population groups. In this study (GitHub Repository), we i) quantify individual space–time accessibility to leisure opportunities, ii) examine its explanatory power for total travel time and leisure activity participation (location diversity), and iii) assess how these relationships vary across population groups, using high-frequency GNSS traces from 3,337 individuals in Île-de-France, enriched with socio-demographic data and detailed travel diaries [1].

First, we identify each participant's home and work zones and visited destination zones for leisure activities (H3 resolution 10 grid centroids), based on trip records and GPS trajectories within Paris' 5,142 IRIS zones. We focus on *Social & Leisure* (67 k) Points of Interest (POIs) derived from Overture's primary labels (confidence level > 0.7), consolidated through a combination of GPT-4o-assisted and manual labeling. Next, we operationalize individual-based space–time accessibility (SPA) as a capability metric that quantifies the number of feasible leisure opportunities reachable within a fixed daily time budget (90 min), considering actual home and work locations, transport mode, and multimodal infrastructure (Figure 1). SPA is computed based on trip chaining feasibility—assuming individuals engage in one-stop leisure activities between work and home—using multimodal travel time calculations from GTFS schedules and OpenStreetMap data, implemented through the `r5r` package. Our behavioral alignment analysis compares individuals' observed leisure activity locations against their modeled opportunity landscapes (SPA-based capability set), demonstrating that SPA effectively captures not only the chosen activity locations but also their relative rankings in terms of accessibility. To further capture realized mobility, we compute daily total travel time and quantify leisure participation through the Hill number of order 1, reflecting the diversity of visited leisure locations. Finally, we estimate a Structural Equation Model (SEM) to examine the pathways through which SPA, transport mode, individual attributes, and travel behavior influence leisure activity participation. The SEM structure is informed by a Directed Acyclic Graph (DAG) and estimated using the diagonally weighted least squares (DWLS) method, incorporating exogenous variables (e.g., household structure, education), mediators (e.g., transport mode), and outcomes (e.g., travel time, leisure diversity), enabling decomposition of direct and indirect effects.

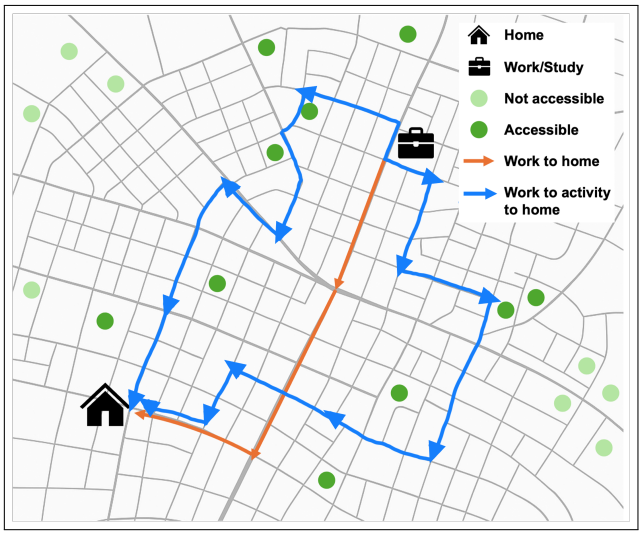

Fig. 1. Space–time accessibility (SPA). Adapted from [2].

Figure 2a shows the share of individuals with SPA > 0, with higher values concentrated in central urban areas—likely due to abundant leisure opportunities—and in select suburban zones, reflecting greater flexibility in movement, primarily by car. We further examine how SPA relates to total travel time and leisure activity participation, stratified by main transport mode (car vs. public transit) in Figures 2b-d. Figure 2b shows that individuals with SPA $\neq 0$ have substantially lower travel times, suggesting that feasible opportunities beyond home and work enable more efficient activity engagement. Figure 2c illustrates a clear negative association between SPA and travel time for both car and public transit users, with the effect more pronounced among car users. Figure 2d shows that leisure activity diversity increases with accessibility, particularly for public transit users. Taken together, these results indicate that higher accessibility not only facilitates efficiency but also broadens the diversity of leisure participation, though the magnitude and pathways of these effects differ between car and public transit users. The correlational results (Figures 2b-d) support our assumptions regarding the interdependence of individual attributes, transport mode choice, and space–time accessibility (Figure 2e). They further suggest that an individual's space–time accessibility shapes both trip-making behavior and actual leisure activity participation. SEM

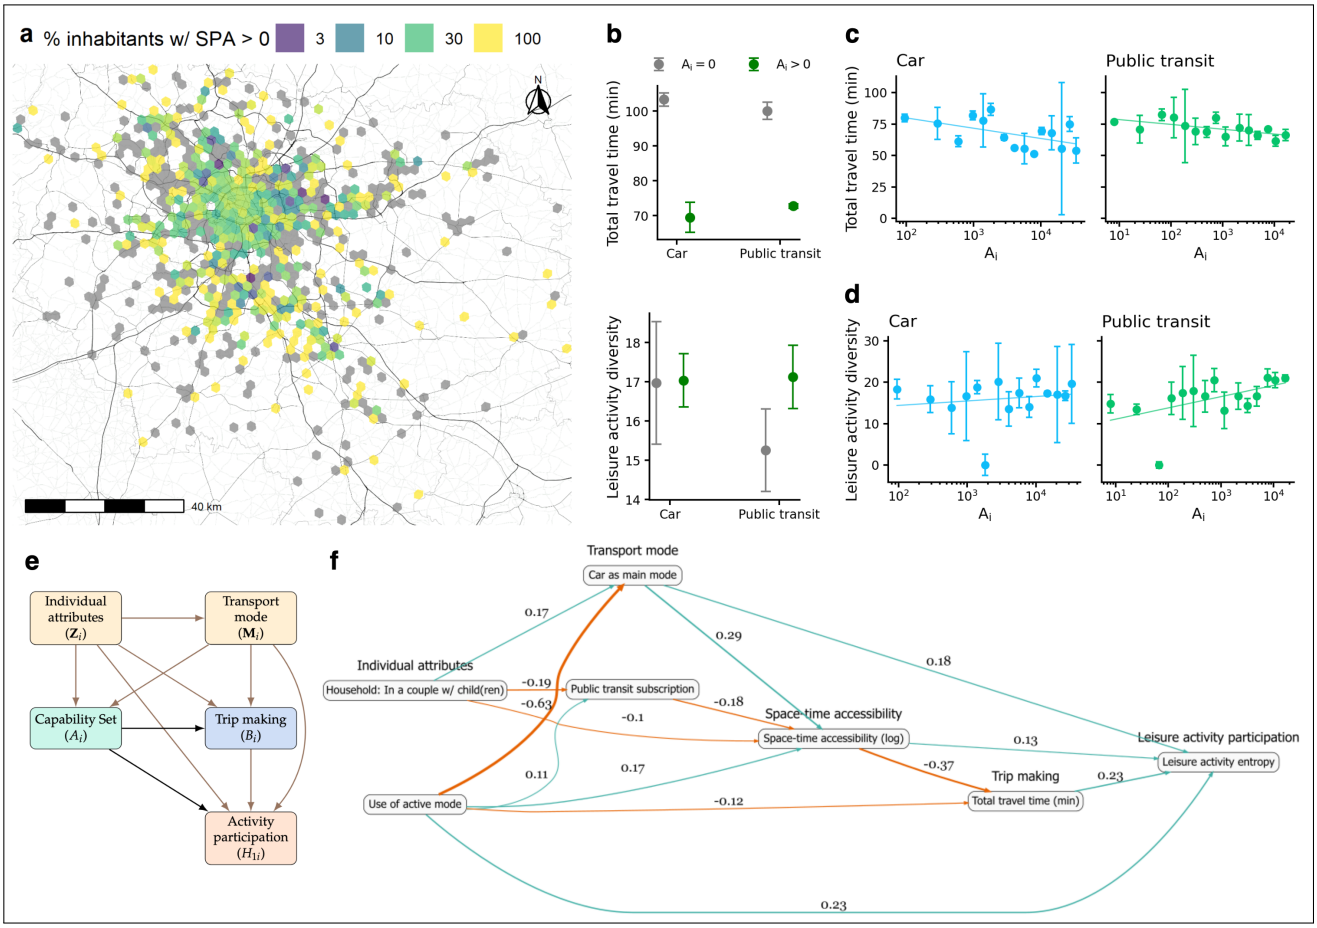

Fig. 2. Space–time accessibility (SPA, A_i) effect on the leisure activity participation. **a**, Share of inhabitants with SPA > 0. **b**, Total travel time and leisure activity diversity, comparing those with $A_i = 0$ (grey) and $A_i > 0$ (green). **c**, Total travel time as a function of A_i for car and public transit users ($A_i > 0$). **d**, Leisure activity diversity ($A_i > 0$) by mode. Error bars indicate bootstrap median estimation errors. **e**, Pathway structure. This directed acyclic graph shows the hypothesized and empirically validated pathways. Exposure=Capability set, Outcome=Activity participation, Black arrows=Causal paths, Brown arrows=Biasing paths. **f**, Structural equation model of factors shaping leisure activity participation, with a Comparative Fit Index (CFI) of 0.998 and Tucker–Lewis Index (TLI) of 0.963. Standardized path coefficients ($p < 0.05$, $\beta > 0.1$) are shown along the arrows, with positive effects in teal and negative effects in orange. Individual attributes (left) influence transport mode choices and space–time accessibility (center), which in turn shape travel behavior (trip making) and leisure activity participation (right).

outcomes further disentangle their relationships.

Space–time accessibility (SPA) is significantly shaped by sociodemographic factors and transport mode (Figure 2f). Individuals using active modes (walking or cycling) or cars have higher SPA, while those with public transport subscriptions or caregiving responsibilities (e.g., couples with children) face lower accessibility due to constrained flexibility. SPA, in turn, is associated with shorter total travel time, suggesting that accessible environments reduce the need for long-distance travel. Leisure activity diversity increases with both higher SPA and longer travel durations, indicating that time spent traveling can support meaningful leisure engagement. Active mode users, in particular, exhibit greater diversity in leisure locations, highlighting the value of flexible, proximity-based movement. Decomposition analyses reveal a nuanced role for SPA: while it directly promotes leisure diversity, its effect is partially offset by travel time, resulting in a marginal total effect—emphasizing the importance of considering both spatial opportunity and travel behavior in shaping activity participation.

This study demonstrates that individual-based space–time accessibility (SPA) significantly shapes both travel efficiency and leisure activity diversity in the Paris region. Higher SPA allows individuals to meet activity demands with shorter travel times while enabling richer leisure participation, though these effects can partially offset each other. Structural constraints—such as caregiving responsibilities—limit accessibility and participation for certain groups, emphasizing the need for inclusive mobility policies. By framing accessibility as a capability, this research highlights the importance of not just infrastructure, but how individuals convert mobility into meaningful participation—offering a useful lens for future policy design and evaluation. Future research should explore varying levels of travel time budgets to examine how space–time accessibility (SPA) outcomes shift under different temporal constraints.

REFERENCES

1. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," (2025).
2. M. Saraiva and J. Barros, Appl. geography **144**, 102731 (2022).

A Simulation Study on Equitable Mobility During City Emergencies, Focusing on Vulnerable Groups

YOUSSEF M. ABD EL HAMEID[1] AND NOHA GAMAL EL DIN[2]

[1] School of Computational Sciences & Artificial Intelligence, University of Science & Technology, Zewail City, **s-youssef.hameid@zewailcity.edu.eg**
[2] Computer Science Program, School of Information Technology and Computer Science, Nile University, **ngamal@nu.edu.eg**

1. INTRODUCTION

Modern cities are systems of profound fragility. The intricate web of infrastructure supporting daily life is predicated on a state of normalcy that can be shattered by an emergent, large-scale shock. Events such as the 2015 Paris terror attacks and the 2012 Hurricane Sandy flooding in New York have shown clearly that when city systems fail, the primary challenge is not merely logistical but human: how to safely move millions of people from a zone of escalating danger.

Traditional emergency management has long approached this problem by modeling populations as homogeneous flows routed through a network. As argued by scholars in disaster studies, this perspective is dangerously incomplete. A city's population is not a uniform mass. As foundational work by Cutter et al. and Flanagan et al. established, social vulnerability is a pre-existing condition: a person's age, physical ability, household structure, and socioeconomic status fundamentally dictate their resilience before the disaster even strikes.

This research addresses this critical gap by moving beyond traditional simulations that model only infrastructure and instead focusing on evacuation equity. To do this, we implement a two-stage methodology. First, we leverage the multi-layered NetMob25 dataset [1] to construct a high-resolution Social Vulnerability Index (SVI). Second, we use this SVI as the core driver for a flexible, high-fidelity agent-based model (ABM). This virtual laboratory allows us to explore a central question: how do pre-existing social inequalities manifest as physical, life-or-death outcomes during an evacuation? Our primary objective is to move beyond identifying where network bottlenecks form and, instead, to analyze who becomes trapped, thereby providing a quantitative assessment of how emergency outcomes are shaped by the interplay among social vulnerability, transport-mode choice, and dynamic network congestion.

2. METHODOLOGY

Our research is underpinned by a two-stage methodological framework designed to first quantify the latent condition of social vulnerability and then simulate its emergent consequences.

A. Social Vulnerability Index (SVI) construction

Our SVI construction follows established frameworks [2, 3]. The goal is a multi-dimensional, individual-level vulnerability score in which higher values consistently indicate greater vulnerability.

A.1. Vulnerability-aligned feature engineering

We translate raw NetMob25 features into vulnerability-aligned components. Categorical variables are encoded (binary or ordinal) and scales are inverted where necessary so that higher numerical values always correspond to higher vulnerability. Ordinal variables (e.g., education) are mapped to an ordered numeric scale and then inverted when higher original values represented greater capacity.

Nonlinear resource modeling Resource indicators (e.g., number of cars) have diminishing marginal effects on resilience. To capture this, we apply nonlinear transforms that compress high values while preserving variation among the resource-poor. We use:

$$\boxed{f_1(x) = \frac{1}{1+\log(1+x)}, \qquad f_2(x) = \log(1+x)}$$

The addition of $+1$ inside the logarithm avoids $\log(0)$ and ensures numerical stability.

Age-related vulnerability Vulnerability varies nonlinearly with age (a U-shaped relationship). We model the age component centered at a resilience age μ:

$$\boxed{v(\text{AGE}) \;=\; \log\bigl(1+|\text{AGE}-\mu|\bigr)}$$

A.2. Data-driven weighting via PCA

To avoid subjective weighting, all features are standardized (zero mean, unit variance) and Principal Component Analysis (PCA) is applied. The loadings from the first principal component (PC1) provide data-driven weights v_j. We perform a sign-check to ensure the index orientation is correct (higher SVI = higher vulnerability). The final SVI for individual i is:

$$\boxed{\text{SVI}_i \;=\; \sum_{j=1}^{p} v_j\, z_{ij}}$$

where z_{ij} is the standardized value of feature j for individual i.

B. Agent-Based Model (ABM) for evacuation simulation

The second stage uses the SVI to parameterize a high-fidelity agent-based model implemented in the AgentPy framework. The ABM integrates a multi-modal transport network, realistic agent initialization from NetMob25 traces, and SVI-driven behavioral heuristics.

B.1. Simulation environment construction

We build separate, topologically simplified network graphs for driving, walking, and cycling using OpenStreetMap via OSMnx, and incorporate public-transport schedules and stop locations from the official Île-de-France GTFS feed [4]. Safe havens (e.g., hospitals, schools) are selected from OSM and designated as valid shelters only if they lie outside the evacuation polygon.

B.2. Empirically-grounded agent population

Agents are instantiated from NetMob25 GPS traces using the `AgentsGatherer` module. A participant is included in a scenario only if they have one or more nighttime/temporal samples inside the evacuation polygon at the scenario start (temporal filtering uses a primary window of ± 4 hours, with a deterministic fallback hierarchy). For participants with multiple candidate samples we select the sample nearest the evacuation-center by great-circle distance (Haversine formula).

$$\boxed{\begin{aligned} &\Delta\varphi = \varphi_2 - \varphi_1, && \Delta\lambda = \lambda_2 - \lambda_1, \\ &a = \sin^2\!\left(\tfrac{\Delta\varphi}{2}\right) + \cos\varphi_1 \cos\varphi_2 \sin^2\!\left(\tfrac{\Delta\lambda}{2}\right), && \\ &c = 2\,\operatorname{atan2}\!\left(\sqrt{a}, \sqrt{1-a}\right), && d = R\,c. \end{aligned}}$$

Per-agent speed outliers are removed using the interquartile-range (IQR) rule before aggregation; missing per-agent medians are imputed with population-level medians. Each agent's home is

estimated from nighttime samples (22:00–05:00) using the spherical centroid method (averaging unit vectors and converting back to latitude/longitude) or the arithmetic mean of coordinates, depending on dispersion.

B.3. SVI-driven agent behavior

The ABM links each agent's precomputed SVI to three behavioral mechanisms:

Reaction delay We model SVI-induced activation latency with a simple linear mapping:

$$\Delta t_{\text{start}} = s \cdot D_{\max}, \qquad t_{\text{eff}} = t_{\text{base}} + \Delta t_{\text{start}},$$

where s is the normalized SVI (recommended $s \in [0, 1]$) and $D_{\max}$ is the maximum SVI-based start delay.

Speed penalty An agent's effective travel speed is penalized as a function of their SVI:

$$v = v_0 \cdot \max(\text{min_mult}, 1 - s\,\kappa)$$

where v_0 is the baseline speed, κ is the SVI speed-sensitivity parameter, and min_mult enforces a lower bound (e.g., 0.1).

Patience and rerouting Patience is modeled as an inverse function of SVI:

$$m = \max(\varepsilon, 1 - s), \qquad T_{\text{pat}} = P_{\text{base}} \cdot m,$$

where P_{base} is base patience and ε a minimum multiplier (e.g., 0.1).

B.4. Simulation dynamics

The simulation proceeds in discrete 60-second time steps. Agents plan routes using A* on the appropriate modal graph (road, walking, or bike). A dynamic congestion feedback loop aggregates agent reports per edge to compute a congestion index:

$$\text{Congestion Index} = \frac{\text{Load}}{\text{Capacity}}$$

This index increases the traversal cost for the next planning round; we use an exponential penalty:

$$\text{Cost} = \text{BaseCost} \times 2^{\text{Congestion Index}}$$

These mechanisms allow bottlenecks to emerge organically from agent interactions.

3. PRELIMINARY RESULTS

We report results from a single simulation run for an evacuation area with a 50 km radius centered on Paris, executed over a 180-minute horizon. The simulation produced an overall evacuation success rate of 18.3%.

A. Mode choice and vulnerability

Mode-choice correlates non-monotonically with vulnerability. The 'Low' SVI group is predominantly pedestrian and transit-oriented; 'Moderate' and 'High' groups rely more on private cars. Figure 1 illustrates the mode distribution by SVI group.

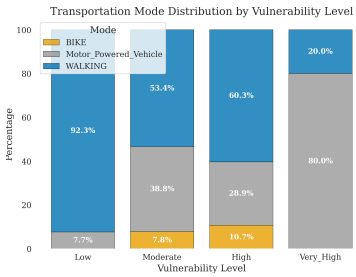

Fig. 1. Transport-mode distribution by SVI group.

B. Evacuation outcomes by vulnerability

The simulation quantifies an "Evacuation Equity Gap": evacuation success is non-monotonically related to SVI. Middle-vulnerability, car-dependent cohorts experience severe delays from congestion; the highest-vulnerability cohorts exhibit low completion rates due to limited access and isolation; low-vulnerability pedestrians often travel slowly but are less impacted by car congestion. Figure 2 presents evacuation success rates by SVI group.

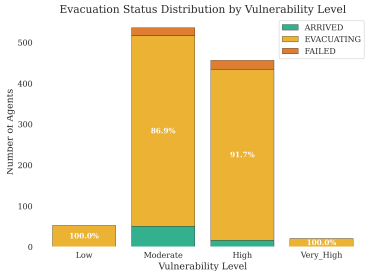

Fig. 2. Evacuation success rates by SVI group.

C. Geospatial bottlenecks

Figure 3 maps emergent bottlenecks on the road network and colors edges by the mean SVI of agents present. The map demonstrates that severe network failures disproportionately trap high-SVI individuals.

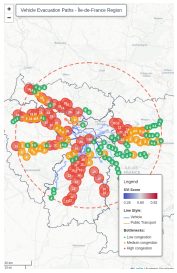

Fig. 3. Emergent road-network bottlenecks colored by mean agent SVI. Zenodo repository (data): (zenodo.org/records/17074045).

4. DISCUSSION

Our results expose a "Tale of Two Failures": (1) *Failure by Gridlock* for car-dependent, middle-vulnerability cohorts; and (2) *Failure by Isolation* for resource-poor extremes. The non-monotonic relationship between SVI and evacuation success shows that vehicle access alone does not guarantee advantage—systemic social vulnerabilities can amplify congestion effects and reduce effective mobility for certain groups.

Limitations include first-order behavioral heuristics and GTFS-based transit modeling. Future work will enrich agent decision models, incorporate dynamic transit interruptions, and perform sensitivity analyses on key parameters.

In conclusion, by quantifying the Evacuation Equity Gap and demonstrating physical consequences of social vulnerability, we support policy interventions that prioritize equitable emergency planning and address differential capacities of urban residents.

REFERENCES

1. A. Chasse *et al.*, "The NetMob25 dataset: A high-resolution multi-layered view of individual mobility in the Greater Paris region," Dataset paper (2025).
2. S. L. Cutter, B. J. Boruff, and W. L. Shirley, "Social vulnerability to environmental hazards," *Social Science Quarterly*, vol. 84, no. 2, pp. 242–261, 2003.
3. B. E. Flanagan *et al.*, "A Social Vulnerability Index for disaster management," *Journal of Homeland Security and Emergency Management*, 2011.
4. Ile-de-France Mobilités GTFS feed, official schedule dataset.
5. T.-R. Roan, "Developing an agent-based evacuation simulation model...," PhD thesis, UCL, 2013.
6. S. Vyas and L. Kumaranayake, "Constructing socio-economic status indices: How to use principal components analysis," *Health Policy and Planning*, 2006.

Summarizing Trajectories to Uncover Mobility Behaviors and Urban Contexts

CHIARA PUGLIESE[1,*], GIULIO LODDI[2], FRANCESCO LETTICH[3], CHIARA RENSO[3], AND FABIO PINELLI[2,3]

[1] Institute of Informatics and Telematics (IIT), CNR, Pisa, Italy
[2] IMT School for Advanced Studies, Lucca, Italy
[3] Institute of Information Science and Technology (ISTI), CNR, Pisa, Italy
[*] chiara.pugliese@iit.cnr.it

1. INTRODUCTION

The rapid proliferation of location-aware devices has generated massive volumes of GPS trajectories, opening new possibilities for analyzing urban mobility. However, linking raw GPS trajectory data with socio-demographic profiles and trip purposes remains challenging due to privacy and confidentiality issues concerning the movements of individuals, thus leading to a scarcity of datasets with such mixture of information. This gap limits our ability to understand why people move, how demographic factors shape mobility patterns, and whether certain groups of individuals face inequalities. We thus argue that datasets integrating movement traces with user attributes might enable the development of novel analytical methods that target the above-mentioned research areas, which in turn have the potential to greatly enhance urban planning and transport policy by enabling more accurate models of accessibility, equity, and urban infrastructure use.

Some of our recent research work has focused on the construction, management, and analysis of multidimensional semantic trajectories [1, 2], that is, trajectories augmented with different semantic dimensions such as points of interest (POI), transportation modes, weather data, social networks, and more. Given the potential size, scale, and complexity of this type of trajectory, in a recent past work, we proposed MAT-SUM [3], our method for summarizing semantic trajectories. This approach aims at preserving essential semantic information while significantly reducing the overall size of the original data. It employs a uniform grid in which the cells are enriched with contextual geographic information, such as land-use categories, POI data, and more. The grid is then used to transform the original trajectories, which can already be semantically enriched or not, into sequences of said cells, yielding summarized semantic trajectories. Overall, this new type of trajectory preserves essential semantic information while dramatically reducing data volume, outperforming traditional compression techniques in both efficiency and interpretability.

With this context in mind, we see the NetMob 2025 dataset as an essential resource to validate MAT-SUM, and investigate further its potential as well as that of semantic trajectories in general. More precisely, we believe the dataset's integration of demographic, semantic, and movement data offers a unique opportunity to analyze the mobility behaviors of individuals within urban environments using summarized semantic trajectories. In particular, we aim to demonstrate how MAT-SUM produces summarized semantic trajectory datasets that retain semantic information essential for the analysis of key equity-related metrics such as accessibility, spatial fairness, and reachability equity, which are all relevant to the concept of 15-minute cities [4–6]. Furthermore, we would like to highlight that we have already participated in the NetMob 2023 data challenge [7, 8], which allows us to access mobile traffic data of several French cities: we believe that combining the two datasets has the potential to provide a more comprehensive view of the mobility behaviors and urban structures within French cities.

2. RESEARCH CHALLENGE

Since the NetMob 2025 dataset comes with an additional layer of semantic information, we aim to investigate if it can further improve and empower context-aware summarization. For example, by conditioning the summarization on trip purpose, we may achieve finer control over the summarization process while preserving relevant geographic and functional information: work or education trips might follow repetitive paths, and might be summarized more aggressively – at least spatially – than irregular leisure trips. Furthermore, incorporating socio-demographic features permits to investigate the inter-dependencies between mobility and urban contexts – for example, study the access to essential services in the framework of 15-minute cities, explore whether Paris exhibits monocentric or polycentric structures, and evaluate spatial equity among different population groups. We argue that this multi-dimensional representation of mobility enables a richer, more inclusive understanding of city dynamics.

We conceptualize a trajectory not simply as a sequence of geographic coordinates, but as a multi-dimensional object that can embed movement, intent, and other interesting information. A trajectory thus becomes a semantic trace of individual behavior, situated within urban space and conditioned by both the purpose of the movement and the characteristics of the individual. Let $\mathcal{T} = (U, \Pi, \Sigma)$ denote a semantic trajectory, where U is a vector of socio-demographic descriptors, such as age, gender, occupation, or residence zone; $\Pi \in \mathcal{P}$ represents the purpose of the trip, with $\mathcal{P}$ being a finite set of annotated purposes (e.g., work, leisure, education); $\Sigma = \langle s_1, s_2, \ldots, s_n \rangle$ is the raw trajectory underlying the semantic one, where each element $s_i = (p_i, t_i)$ represents the geographical location (latitude and longitude) recorded at timestamp t_i. The summarization process is a transformation:

$$\mathcal{T} \to \mathcal{T}' = (U, \Pi, \Sigma') \quad \text{with} \quad \Sigma' = \langle \tilde{p}_1, \ldots, \tilde{p}_m \rangle, \; m \ll n,$$

where each $\tilde{p}_j$ is a representative semantic unit, possibly weighted by temporal exposure, and selected or aggregated from the original samples in Σ based on both the purpose Π, the user profile U, and the semantic context associated with the spatial unit. This formulation makes explicit the dual conditioning of summarization: we do not reduce spatial paths uniformly, but instead encode the contextual importance of segments relative to why the trip was made. This allows us to reason about mobility not just as displacement, but as structured behavior situated within a personal and functional narrative. Moreover, embedding trajectories in this high-level semantic space enables comparisons across individuals, purposes, and spatial contexts, supporting inquiries about several issues that concern urban environments, such as urban equity, service accessibility-fairness, systemic spatial organization, and reachability equity [4–6].

3. PRELIMINARY RESULTS

The application of the MAT-Sum method to GPS trajectories allows us to analyze behavioral patterns by combining spatial tessellation with semantic enrichment. The first step of the method consists of tessellating the study area. We relied on the INSEE 200×200 grid of the Île-de-France region, a partitioning commonly used for statistical analysis. Each tile was then enriched with POIs downloaded from OpenStreetMap via Geofabrik. For consistency and comparability, we focused on POI categories that match with trip purposes, including *health*, *leisure*, *sports*, *catering*, *accommodation*, and *shopping*. This mapping allowed us to filter the semantic contexts of traversed tiles according to the annotated purpose of each trip, and thus to measure the semantic diversity of the trajectories, i.e., the variety of facilities encountered during displacements.

Our first set of results shows that, in most cases, individuals move through a limited number of semantic regions, and that many of these are "empty" tiles without POIs. This indicates that movements often take place in poorly equipped areas, while only a subset of trips intersect with semantically rich urban zones. Nevertheless, clear differences emerge between purposes. Movements for `HEALTH`, `ACCOM`, or `PURCHASE` purposes usually involve very few semantic regions, confirming that they are short, direct, and goal-oriented. On the other hand, `RETURN_HOME` and `STUDIES` trips are much more diverse: returning home frequently implies crossing heterogeneous urban areas, while study-related trips display higher averages, probably because younger people combine school or university attendance with activities such as eating out or visiting leisure facilities.

The distribution of distinct regions per trajectory confirms this contrast. While health and accommodation trips are concentrated on very few tiles, home-related, study-related, and business trips cover a wider range of contexts and display significant variability. This heterogeneity is also evident in the top three regions where individuals spend the most time (Fig. 1). For most purposes, these top regions are empty, suggesting that people spend a considerable share of their trajectories simply passing through. Exceptions, however, reveal meaningful patterns: health-related trips often occur near `fast-food outlets`, `parks`, or `swimming pools`, while study and business trips intersect with central facilities such as `hotels`, `sports centres`, and `restaurants`.

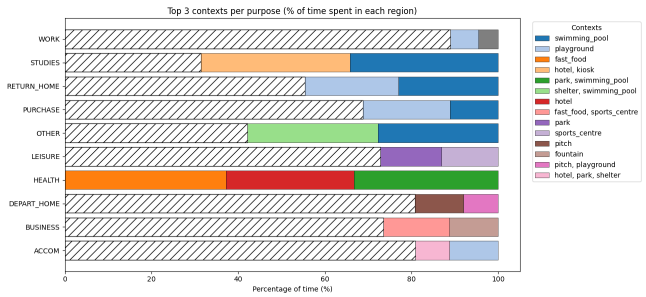

Fig. 1. Top three semantic contexts per trip purpose, expressed as the percentage of time spent in each region.

When focusing on demographic differences, the role of age becomes evident. Older individuals (over 60) tend to display more constrained and goal-oriented mobility: their trips are dominated by empty regions, with only limited exposure to POIs, except for purchase and leisure activities. By contrast, younger individuals (under 30) show richer semantic diversity. Study trips intersect not only with schools or universities but also with hotels, playgrounds, and sports facilities; leisure trips are linked to contexts such as food courts, memorials, and sports centres; purchase trips involve smaller-scale commercial facilities like kiosks, greengrocers, and convenience stores. This suggests that younger participants navigate more multifunctional and socially active environments, whereas older participants restrict their movements to essential and direct routes.

Overall, our analysis highlights that semantic diversity in mobility is not evenly distributed but depends strongly on the purpose of the trip and the demographic profile of the traveler. Short, utilitarian trips concentrate on a narrow set of contexts, while home-related and study-related trips expose individuals to more heterogeneous urban regions. Age amplifies these differences, with younger individuals engaging more with semantically rich areas and older individuals relying on simpler, more direct trajectories.

4. CONCLUSION

In summary, our analysis highlights how the semantic diversity of trajectories varies across purposes and age groups. While older individuals tend to follow more direct and goal-oriented paths, largely traversing semantically "empty" regions, younger participants show a richer exposure to urban facilities, particularly for study, leisure, and purchase-related trips. This contrast underlines the role of demographic factors in shaping mobility behaviors.

For future work, we plan to extend this investigation by incorporating additional demographic features (e.g., gender, socioeconomic background) and by analyzing the spatial distribution of semantic regions in greater detail, in order to better understand how urban form and demographics jointly influence human mobility. Another important direction will be to improve the mapping between POIs. In addition, we plan to introduce stop detection techniques, with the dual goal of (i) distinguishing between regions where people only move through and those where they actually stop, thus refining the analysis presented in this report, and (ii) identifying the POIs that individuals effectively visited.

REFERENCES

1. R. d. S. Mello, V. Bogorny, L. O. Alvares, *et al.*, "MASTER: A multiple aspect view on trajectories," Trans. GIS **23**, 805–822 (2019).
2. F. Lettich, C. Pugliese, C. Renso, and F. Pinelli, "Semantic enrichment of mobility data: A comprehensive methodology and the mat-builder system," IEEE Access **11**, 90857–90875 (2023).
3. C. Pugliese, F. Lettich, F. Pinelli, and C. Renso, "Summarizing trajectories using semantically enriched geographical context," in *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems*, (Association for Computing Machinery, New York, NY, USA, 2023), SIGSPATIAL '23.
4. T. Jin, K. Wang, Y. Xin, *et al.*, "Is a 15-minute city within reach? measuring multimodal accessibility and carbon footprint in 12 major american cities," Land Use Policy **142**, 107180 (2024).
5. J. Zhang, B. Y. Chen, C. Fu, *et al.*, "Uncovering horizontal and vertical inequities of individual accessibility using mobile phone data," Transp. Res. Part D: Transp. Environ. **143**, 104755 (2025).
6. T. Abbiasov, C. Heine, E. L. Glaeser, *et al.*, "The 15-minute city quantified using mobility data," Tech. rep., National Bureau of Economic Research (2022).
7. "Netmob 2023 book of abstracts," https://netmob.org/www23/book_of_abstract/book_of_abstract.pdf.
8. G. Loddi, C. Pugliese, F. Lettich, *et al.*, "Urban region embeddings from service-specific mobile traffic data," in *IEEE MDM 2025*, pp. 72–78.

Persons with reduced mobility facing invisible obstacles: A contrastive analysis on mobility patterns and behaviours based on GNSS data

COLIN CHAIGNEAU[1,*] AND ARNAUD LE GUILCHER[1]

[1] *Univ Gustave Eiffel, ENSG, IGN, LASTIG, F-77420 Champs-sur-Marne, France*
[*] *Colin.Chaigneau@ign.fr*

1. INTRODUCTION

Access to large amounts of data in recent years has enabled new approaches in urban planning and assessment of mobility issues. By analysing the mobility patterns and behaviours of individuals directly from data, local players gain access to insightful information that can help their decision making.

Geolocation data for instance play an important role for public transport. Once processed, they help devise better transportation planning, accident alerts systems, mobility dashboards for decision taking and so on. On the topic of urban planning, this information provides detailed and accurate insights about the daily mobility patterns and behaviours of inhabitants, which can then be leveraged to improve quality of life when using public transportation or private means of transport.

Among the individuals that benefit from infrastructure improvements, persons with reduced mobility, or PMR, are a sensitive socio-demographic profile. Indeed, PMR individuals encounter many obstacles, often invisible for able-bodied individuals, in their daily mobility. A missing ramp in a stair, a lack of elevators or even a steep hill are all obstacles that can disturb their daily commutes. This motivates a need for analyses of PMR individuals mobility in contrast to able-bodied people.

2. DATASET OVERVIEW

This dataset was derived from a GNSS-based mobility survey conducted in the Île-de-France region (Greater Paris Area, France) between October 2022 and May 2023 [1].

Over 3,300 volunteers were given instructions to carry a GPS device with them alongside a diary they had to fill with details about their trips, like purpose or main means of transport. More than 80,000 trips were recorded, offering a rare opportunity to explore and analyse the mobility behaviours and patterns both at a large scale, with trips spanning over 12,000 km^2, and at a smaller scale, with 2-to-3 seconds time frequency GPS trajectories. Additionally a household survey was also conducted on the participants to provide socio-demographic information.

Focus on Persons with Reduced Mobility (PMR)

One socio-demographic profile we decided to focus on in our analysis are *Persons with Reduced Mobility (PMR)* and their mobility behaviours in contrast with able-bodied individuals. Note that PMR individuals are already labelled within the dataset. To ensure statistical significance we checked the representativeness of the PMR persons within the dataset.

Demography. Among the +3,300 participants of the survey, around 3.4% are described as PMR, amounting to 114 individuals, a ratio close to the national statistic[1] which illustrates the statistical representativeness of the dataset.

[1]The *DREES*, a French national statistic department, reports for between 3% and 6% of PMR individuals in France [2], depending on how mobility impairment is defined.

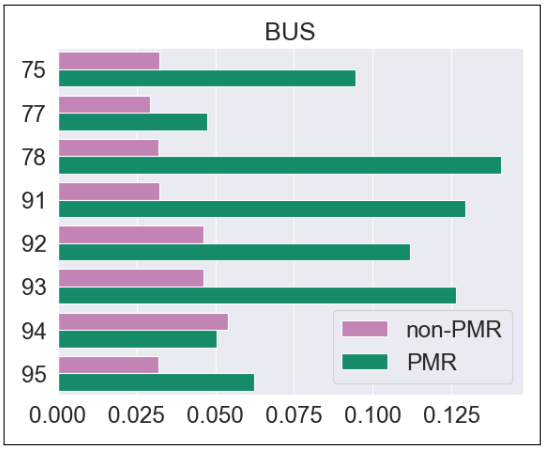

Fig. 1. Share of PMR individuals (green) and non-PMR (purple) individuals taking buses, computed by department. We can notice that, except from the department #94, PMR persons are more likely to use buses in their mobility. This could be explained by a strong effort to improve accessibility to the buses network [3]

Parity. Among the PMR persons we observed a very large gender bias with 90 individuals labelled as women over 112 participants. However, by scrutinizing the share of transport modes for PMR persons we did not observe discrepancies at the global scale between women and men. Hence it would suggest that while being unbalanced this gender gap would not interfere with our results.

Spatial bias. We performed a z-score on the expected number of PMR individual IDs in each department of *Île-de-France*. We observed an significant over-representation of PMR persons in the #95 department ($+3\sigma, p = 0.0014$), and some under-representation in the #75 and #92 departments (with $-1.6\sigma, p = 0.057$ for both).

3. ANALYSIS RESULTS

A. Mean of transport bias

We analysed the main transport modes used by PMR and non-PMR individuals in each department. All trajectories in the dataset are already labelled with their main transport mode.

The main observation we extracted from our analysis regarding the mobility behaviour of PMR persons is their higher proneness to use buses than non-PMR individuals as depicted in Figure 1.

The result in Paris, department 75, is interesting because this bias toward buses (PMR = 10.5% / non-PMR = 3.2%) is flipped when focusing on subway use (PMR = 12.9% / non-PMR = 22.3%). A lack of accessibility in the Parisian subway network,

reported in [3], could explain the observed difference.

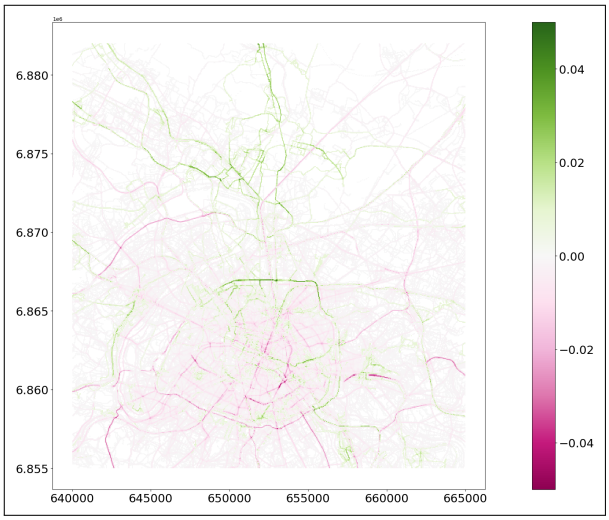

Fig. 2. Rasterized analysis of the number of PMR and non-PMR unique IDs. We computed the metric described in Equation 1 on Paris and its northern surroundings. Green lines represent higher sample rates of PMR individuals while purple lines indicate lower sample rates.

B. Spatial bias

Spatial biases are areas where the sampling rate of PMR and non-PMR persons would differ. To measure it we split the region *Île-de-France* in a grid of 50m×50m cells. Inside each cell we counted the number of unique PMR and non-PMR individuals ID. Then, we leverage our rasterized IDs counting by defining a metric which aims at identifying areas where PMR individuals would avoid while non-PMR individuals would explore. We decided to use the following **sampling rate difference** metric to measure the exploration bias :

$$\mu_{cell}\left(P_{cell}, V_{cell}\right) = \frac{P_{cell}}{P_{total}} - \frac{V_{cell}}{V_{total}} \quad \textbf{(1)}$$

where P_{cell} (resp. V_{cell}) is the number of PMR (resp. non-PMR) individuals unique IDs in *cell*, P_{total} (resp. V_{total}) is the number of PMR (resp. non-PMR) individuals unique IDs in the dataset. A higher sample rate of PMR individuals visiting an area the metric would yield positive values, while negative values indicate a higher sampling rate for non-PMR persons.

We report our computation in Figure 2. Through this map we can observe several green and purple axes, suggesting area with a discrepancy between the sampling rate of PMR or non-PMR persons:

- The northern part of Paris and its surrounding exhibit multiple green axes, an indicator of a higher sampling rate of PMR individuals. This is consistent with our previous observation that PMR persons are over-represented in department 95 which is located in the north of the region.
- Several purple areas can be observed within the inner Paris area. One of them is located in the north east part of Paris around the district of *Ménilmontant*. It is noteworthy that this area is located on a hill, hence suggesting that accessibility to this place is more complex to PMR people. This argument is also consistent with the lack of accessibility in the Paris subway network, especially for the oldest lines [3], which stop at this specific areas.

C. Mobility indicators discrepancy

In our analysis we identified some discrepancies between PMR and non-PMR persons based on the results from the average speed and radius of gyration mobility indicators.

Figure 3 depicts the average speed and radius of gyration of PMR and non-PMR trajectories starting from each department. We can derive the following observations from our results:

- no discrepancies can be found for Paris (department 75) regarding average speeds or radius of gyration;
- departments 77, 91, 95 exhibit small to inexistent differences between the average speeds of PMR and non-PMR individuals, however they do exhibit a discrepancy regarding the average radius of gyration;
- departments 78, 92, 93, 94 exhibit discrepancies for both average speed and radius of gyration, suggesting slower and more compact mobility from PMR persons compared to non-PMR persons.

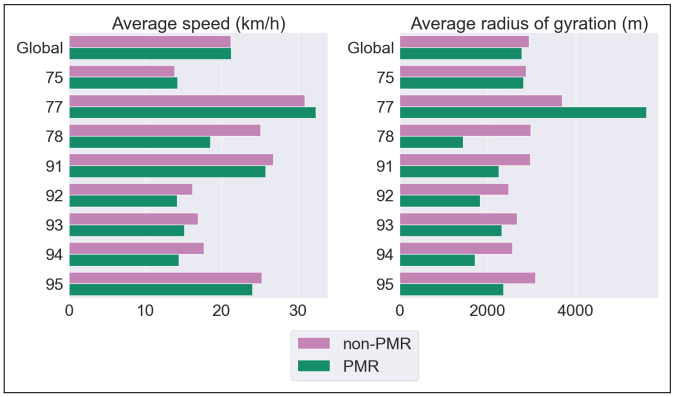

Fig. 3. Average speed and radius of gyration computed for the *Île-de-France* region (*global*) and by department for PMR (green) and non-PMR (purple) individuals. We can notice that the average speed and radius of gyration for PMR persons (green) are lower than for non-PMR persons in several departments.

4. CONCLUSION

Throughout this analysis we described three kinds of biases in the mobility patterns and behaviours of PMR persons in contrast to non-PMR individuals. While this analysis captures meaningful statistical differences for both groups, we could investigate this contrast deeper by exploring at a trajectory level their mobility patterns and behaviours.

ACKNOWLEDGMENT

This work is funded by the French National Research Agency as part of the Mob Sci-Dat Factory project (ANR-23-PEMO-0004) under the France 2030 program.

REFERENCES

1. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," https://arxiv.org/abs/2506.05903 (2025).
2. Direction de la recherche, des études, de l'évaluation et des statistiques, "Le handicap en chiffres," https://drees.solidarites-sante.gouv.fr/publications-communique-de-presse-documents-de-reference/panoramas-de-la-drees/241128_Panorama_Handicap2024 (2024).
3. "Ratp network accessibility report (2023)," https://www.ratp.fr/sites/default/files/inline-files/Mobility%20for%20all%202023.pdf.

Integrating Temporal Context into Mobility Embeddings: A Doc2Vec Approach to User Profiling

WENLAN ZHANG[1,2,*], ZHENGZI ZHOU[1], AND ANDREW RENNINGER[1]

[1] The Bartlett Centre for Advanced Spatial Analysis, University College London, Gower Street, London, WC1E 6BT
[2] Leverhulme Centre for Demographic Science, University in Oxford, Park End St, Oxford, OX1 1JD
[*] wenlan.zhang.21@ucl.ac.uk

1. INTRODUCTION

Understanding human mobility patterns is central to a wide range of applications, from urban planning and transportation modelling to public health and socio-economic analysis. As cities become increasingly data-rich, large-scale mobility datasets have opened new possibilities for uncovering how individuals move through urban space and how these patterns are shaped by social and demographic factors.

Recent advances in machine learning, particularly in representation learning and natural language processing, have introduced new tools for modelling sequential data like human language and human trajectories. By treating sequences of visited locations as analogous to sentences, and locations as tokens, these methods allow researchers to learn dense embeddings that capture semantic relationships between places and behavioural similarities between individuals [1]. This shift enables a move away from hand-crafted features and towards more flexible, data-driven approaches to understanding urban behaviour.

However, despite the promise of these techniques, most prior work has mainly focused on the geographic dimensions of mobility in isolation, often overlooking the role of temporal information in shaping movement [2]. Furthermore, word embedding models have demonstrated their ability to capture deeper cultural and social meanings, such as dimensions of class and status inferred from the semantic relationships between words over time [3]. However, this potential remains under-explored in the context of mobility data, where similar techniques could offer insights into latent socio-economic structures underlying human movement. In particular, few studies have attempted to infer user-level demographic or socio-economic attributes directly from mobility traces, especially in ways that are both scalable and privacy-preserving.

In this study, we aim to integrate both spatial and temporal information from individuals' trips to build comprehensive vector representations of mobility behaviour, ultimately supporting the inference of user profiles. Our broader objective is to capture not only where people go but also when and how they move through urban space, recognising that the meaning and function of locations often depend on time. While this is the long term goal, the current stage of the project focuses on trip level data, using sequences of named origin and destination locations to learn location embeddings based on co-occurrence patterns. These preliminary experiments serve as a foundation, allowing us to assess whether such sequence based representations already reflect latent geographic structure. In the next phase, we will incorporate the raw GPS trajectories to capture finer grained spatial detail and movement geometry, enabling more nuanced and temporally aware behavioural modelling.

2. DATA AND METHODS

Mobility traces: We use NetMob25 GPS data for 3,000 participants in Paris (Oct 2022–May 2023). Stops were detected using `InfoStop`, then matched to IRIS census units, producing sequences of visited locations per user.

Socio-demographics: Metadata include age, sex, profession, education, and car ownership. These were not used in training, but serve as external labels for evaluation.

Embedding technique: We train Doc2Vec models with two tokenisation: (i) *location* (IRIS only), and (ii) *location_time* (location with time-of-day category: daytime, evening, night). Doc2Vec outputs word embeddings (stops or stop-time tokens) and document embeddings (users). Evaluation uses silhouette, nearest-neighbour purity, compactness, and pairwise cosine similarity between groups.

Evaluation metrics. To assess structure in the embeddings, we apply four complementary metrics, grouped into within-group and between-group measures:

- *Within-group metrics:*
  - **Silhouette score:** quantifies global cluster separation by comparing cohesion of embeddings within the same group to distances from other groups.
  - **Nearest-neighbour (NN) purity:** measures the proportion of an embedding's nearest neighbours that share its label, capturing local consistency.
  - **Compactness:** average pairwise cosine similarity among embeddings of the same group, reflecting how tightly they cluster.
- *Between-group metrics:*
  - **Pairwise cosine similarity:** average cosine similarity between embeddings from two different groups, revealing overlap or separation between categories.

Together, these provide complementary perspectives: within-group measures assess cohesion and purity of subgroups, while between-group similarity captures relative distances across categories.

3. RESULTS

Word-level embeddings. Location embeddings preserve geography, with PaCMAP projections showing smooth gradients by latitude and longitude (Figure 1a). Incorporating time-of-day sharpens local structure: night-time tokens form compact clusters near the centre of the embedding space, while daytime and evening tokens remain more dispersed. Metrics confirm this: silhouette scores are close to zero (daytime 0.012, evening –0.002, night –0.006), indicating weak global separation. Compactness is highest at night (0.903), while NN purity is greater for daytime (0.636) and evening (0.600) than for night (0.378). Between-time similarities further highlight divergence, with daytime–evening highly similar (0.77), but day–night (0.54) and evening–night (0.60) substantially lower.

User-level embeddings

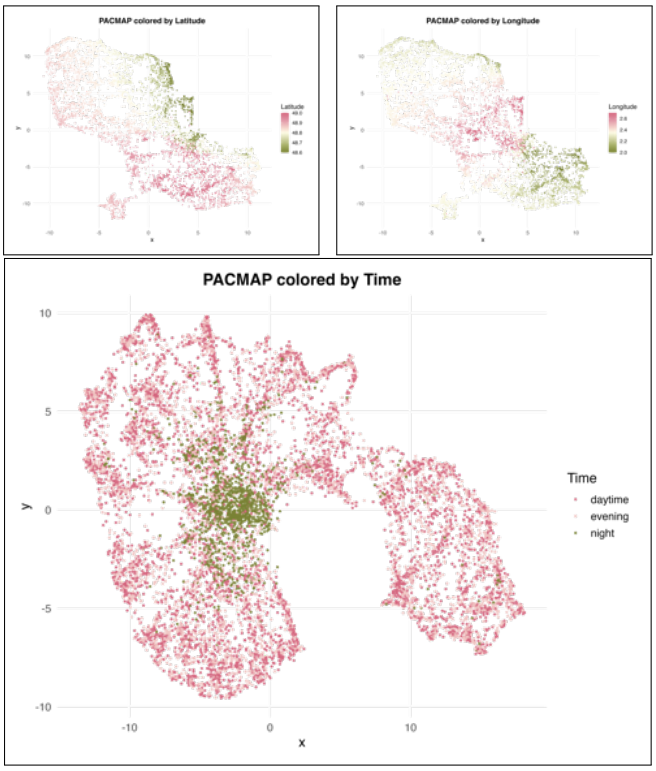

Fig. 1. PaCMAP projection of visualisations. **Top:** `location` embeddings coloured by latitude and longitude. **Bottom:** `location_time` embeddings coloured by time of day.

For the *location*-based model, document embeddings show limited global separation across demographic groups, with silhouette values near zero or negative. Still, several subgroups exhibit stronger local cohesion. Women achieve higher NN purity (0.60) compared to men (0.42). Retired individuals are the only professional group with a positive silhouette (0.015) and show high compactness (0.421). Adolescents (<18) stand out among age groups, with the highest silhouette (0.098) and compactness (0.511). Car ownership shows the clearest contrast: although silhouettes remain near zero, participants with one or more cars achieve much higher NN purity (0.78) than non-owners (0.42).

Pairwise cosine similarities illustrate these patterns. Sex differences are modest (Man vs Woman: 0.36). Among professions, *Employee vs Retired* reaches 0.40, while *Tradesman vs Senior* is lower (0.31). Across age groups, values range from 0.34 (18–30 vs 31–45) to 0.45 (60+ vs <18). Car ownership categories are moderately similar (0.35). Visualisations (Figure 2b) confirm car ownership as the most visibly differentiating attribute, though global metrics remain modest.

Results from the *location_time* model are highly consistent with these findings: the same subgroups (women, retirees, adolescents, and car owners) show relatively stronger local cohesion, and overall demographic separability remains weak.

Summary In sum, embeddings capture strong spatial and temporal regularities and reveal localised cohesion for certain subgroups, but demographic separation at the global level remains weak.

4. DISCUSSION

Our results show that distributional embeddings of mobility sequences capture clear spatial and temporal regularities

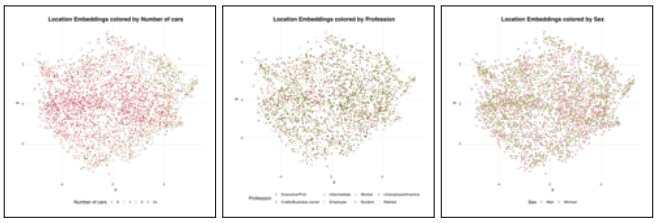

Fig. 2. Two-dimensional visualisation of user embeddings from the *location*-level Doc2Vec model using PaCMAP. Points represent individuals, coloured by age group, education (diploma), number of cars, profession, and sex. Car ownership shows the clearest visual separation among groups.

but only weak demographic separation. At the location level, embeddings align strongly with geography: PaCMAP projections reveal smooth gradients by latitude and longitude, while temporal augmentation highlights distinct night-time clusters. Metrics confirm that daytime and evening remain highly similar, whereas night-time tokens are more compact but globally less separable.

At the user level, document embeddings display limited global clustering across socio-demographic groups, with silhouette values near zero or negative. Nevertheless, local cohesion is evident for certain subgroups: women, retirees, adolescents, and car owners with one or more vehicles show higher purity or compactness. Visualisations emphasise car ownership as the most visibly distinct factor, though quantitative measures remain modest. Results from the *location_time* model are highly consistent with these patterns, indicating that temporal information sharpens token-level distinctions but does not materially improve user-level profiling.

Taken together, these findings suggest that embeddings are best suited to capture local behavioural structure and temporal routines, rather than to produce globally separable demographic clusters. Future work will extend the framework with raw GPS trajectories, richer temporal encodings, and evaluation on predictive tasks. Given the individual distinctiveness of mobility embeddings, privacy safeguards such as aggregation or differential privacy remain essential. Overall, embedding approaches provide a scalable, flexible foundation for mobility analysis, linking behavioural regularities with urban context while demanding careful methodological and ethical consideration.

5. CONCLUSION

Embedding methods applied to mobility sequences recover strong spatial and temporal regularities but only weak demographic separability. These results highlight both the promise and limits of embedding-based mobility analysis: they capture fine-grained behavioural structure and time-specific patterns, but group-level inference is limited. Future extensions include raw trajectory geometry, richer temporal encoding, and extrinsic predictive tasks. Given the distinctiveness of individual embeddings, privacy safeguards remain essential.

REFERENCES

1. D. Murray, J. Yoon, S. Kojaku, *et al.*, Proc. National Acad. Sci. **120**, e2305414120 (2023).
2. Y. Yao, Z. Guo, C. Dou, *et al.*, Comput. Environ. Urban Syst. **104**, 102009 (2023).
3. A. C. Kozlowski, M. Taddy, and J. A. Evans, Am. Sociol. Rev. **84**, 905 (2019).

Who Breathes More Smog? Modeling Individual PM2.5 Exposure and Inequality with GPS Mobility Data

FANGYI DING[1,4], YAMIN WANG[1], AND ZHAN ZHAO[1,2,3,*]

[1] Department of Urban Planning and Design, The University of Hong Kong, Hong Kong SAR, China
[2] Urban Systems Institute, The University of Hong Kong, Hong Kong SAR, China
[3] Musketeers Foundation Institute of Data Science, The University of Hong Kong, Hong Kong SAR, China
[4] Department of Civil and Environmental Engineering, National University of Singapore, Singapore, 117576, Singapore
[*] zhanzhao@hku.hk

1. INTRODUCTION

Air pollution remains a major public health challenge in urban areas, with individual exposure levels shaped by both daily behaviors and structural factors[1]. In Paris, despite recent improvements from measures like low-emission zones[2], PM2.5 concentrations still exceed WHO-recommended limits in many neighborhoods [3]. Understanding who is most exposed—and why—is critical. This study integrates large-scale GPS mobility data from the NetMob challenge with high-resolution pollution maps from Airparif[3] to estimate personalized PM2.5 exposure across Greater Paris. By linking individual mobility, activity patterns, and socio-demographics with spatial pollution data, we develop interpretable models that reveal the behavioral and structural drivers of exposure inequality. Our framework enables both a fine-grained behavioral diagnosis of exposure and a group-based assessment of environmental injustice.

2. METHODOLOGY

A. Mobility Trajectory Reconstruction

To compute exposure, we reconstruct each individual's daily trajectory using three steps: **1) Missing GPS Imputation:** For trips without GPS records, we estimate the trajectory using the shortest path on the street network between the origin and destination. **2) Trip-to-Activity Conversion:** Trips are transformed into activity episodes based on inter-trip intervals, as illustrated in Fig. 1. **3) Grid Projection:** All locations from trips and activities are mapped onto Airparif pollution grid, enabling exposure computation based on grid-level PM2.5 values.

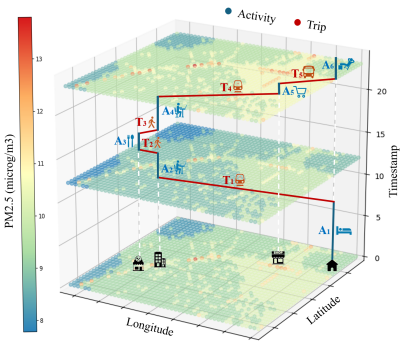

Fig. 1. Schematic Diagram of Daily Spatiotemporal Patterns of Travel and Activity.

B. PM2.5 Exposure Estimation

We estimate individual-level PM2.5 exposure by computing trip- and activity-related exposure separately, adjusted for time-of-day and inhalation rate.

1. **Mobility-based Exposure.** For each trip T_i with GPS points $\{(g_k, t_k^{start}, t_k^{end}, m_k)\}$ denoting grid location, start time, end time, dwell duration, and travel mode, the exposure is:

$$E_{T_i}^{\text{mob}} = \sum_{k=1}^{K_i} \text{PM}_{g_k} \cdot \Delta t_k \cdot \alpha_{t_k} \cdot R_{m_k}^{I/O} \cdot \text{IhR}_{m_k} \quad (1)$$

where PM_{g_k} is the annual average PM2.5 concentration at grid cell g_k, α_{t_k} is the time-of-day adjustment factor, and $R_{m_k}^{I/O}$ and IhR_{m_k} are the mode-specific concentration ratio and inhalation rate, respectively.

2. **Activity-based Exposure.** For each stationary activity A_i located in grid g_{A_i} with duration Δt_{A_i} and activity type P_{A_i}, the exposure is:

$$E_{A_i}^{\text{act}} = \text{PM}_{g_{A_i}} \cdot \Delta t_{A_i} \cdot \alpha_t \cdot R_{A_i}^{I/O} \cdot \text{IhR}_{A_i} \quad (2)$$

3. **Exposure Intensity.** Exposure intensity quantifies PM2.5 burden per hour. Let an individual's total exposure be $E^{\text{mob}}, E^{\text{act}}$and durations be $\Delta T^{\text{mob}}, \Delta T^{\text{act}}$, then:

$$EI^{\text{mob}} = \frac{E^{\text{mob}}}{\Delta T^{\text{mob}}}, \qquad EI^{\text{act}} = \frac{E^{\text{act}}}{\Delta T^{\text{act}}} \quad (3)$$

C. Exposure Disparity Measurement

To quantify inequality in PM2.5 exposure intensity across population groups, we adopt the Theil index, a widely used entropy-based measure of distributional disparity. It captures both within-group and between-group variations.

Given a population of N individuals with exposure intensities $\{EI_i\}_{i=1}^N$ and mean $\bar{EI}$, the Theil index is defined as:

$$TI = \frac{1}{N} \sum_{i=1}^{N} \frac{EI_i}{\bar{EI}} \ln\left(\frac{EI_i}{\bar{EI}}\right) \quad (4)$$

D. Modeling Strategy

Fig. 2 presents the modeling framework. We extract activity attributes (e.g., daily time allocation across activity types), mobility attributes (e.g., daily trip distance, duration, mode shares), and individual socio-demographics from the mobility dataset. With PM2.5 exposure intensity estimated, two Explainable Boosting Machine (EBM) models are built separately for activity and mobility exposure to identify key behavioral drivers and test group disparities. Theil index is then used to quantify inequality and inform policy implications.

3. KEY FINDINGS

A. Spatial Gradient Patterns in Exposure Intensity

The distribution of PM2.5 exposure intensity across participants' residential IRIS units is illustrated in Figure 3y. The results reveal a clear spatial gradient in exposure intensity, with central districts exhibiting the highest levels of PM2.5 exposure in both activity and mobility dimensions.

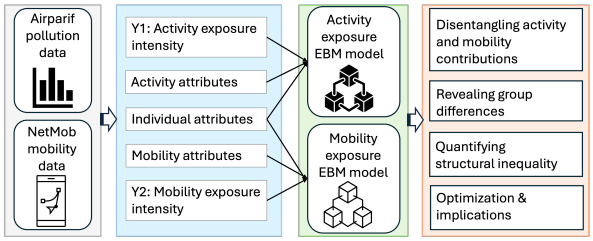

Fig. 2. The Framework of Modeling.

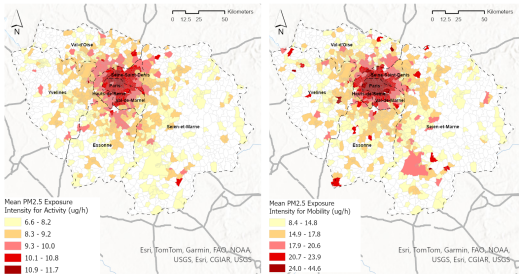

Fig. 3. Spatial Variation in PM2.5 Exposure Intensity.

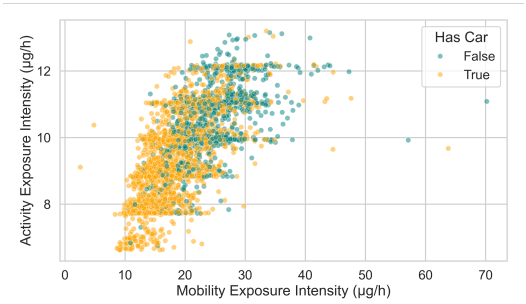

(a) Mobility vs Activity PM2.5 Exposure Intensity.

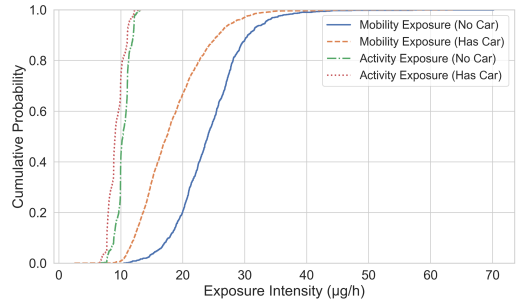

(b) CDF of PM2.5 Exposure Intensity by Car Ownership.

Fig. 4. Exposure Intensity Disparities by Car Ownership.

B. Activity and Mobility Patterns Drive Differences in Exposure Intensity

Individual activity and mobility behaviors significantly explain variation in PM2.5 exposure intensity. Using EBM models controlling for socio-demographics, we isolate the marginal effects of behavioral attributes:

- **Activity exposure intensity** decreases with greater time spent on accompanying activities, but increases as more time is allocated to business-related activities.
- **Mobility exposure intensity** is influenced by both travel mode and trip structure.
  - *Travel mode composition:* Higher private car mode share is associated with lower exposure intensity, while walking and biking lead to increased exposure likely due to elevated inhalation rates in these active travel modes.
  - *Trip structure:* Longer average trip distances are associated with lower exposure intensity, suggesting higher mobility capacity. Conversely, longer trip durations increase exposure, indicating that inefficient travel modes may lead to higher pollution burden.

C. Car Ownership Is a Key Divider in Exposure Burden

Fig. 4(a) visualizes individual-level PM2.5 exposure intensity (activity vs. mobility), colored by car ownership. Individuals without a car tend to cluster in the upper right—experiencing high exposure in both domains.

Fig. 4(b) further confirms this disparity, showing rightward-shifted CDFs for non-car owners in both activity and mobility exposure. Quantitatively, the EBM reveals strong marginal effects of car ownership: (1) **Non-car owners:** Δ exposure intensity $\approx +0.437$ (activity), $+0.766$ (mobility) (2) **Car owners:** Δ exposure intensity ≈ -0.197 (activity), -0.333 (mobility)

These results confirm that access to private vehicles reduces environmental exposure risks, while lacking it amplifies them during both daily activities and travel.

D. Transport Poverty as a Structural Source of Inequality

To assess whether the disparity reflects systemic inequality, we decomposed the Theil index for exposure intensity into within-group and between-group components (based on car ownership). (1) **Mobility exposure:** 20% of total inequality is attributable to between-group differences. (2) **Activity exposure:** 15% is attributable to between-group differences.

Although most variation remains within groups, the non-negligible between-group share reveals that pollution exposure inequality is not purely behavioral, but partially structural—linked to disparities in transport access.

4. DISCUSSION & IMPLICATIONS

Our findings underscore the need for differentiated policy interventions to mitigate disparities in air pollution exposure. Individuals with high exposure intensity—particularly those who walk, cycle, rely on public transport, or frequently engage in business-related activities—require targeted protective measures. Potential strategies include expanding access to low-exposure public spaces (e.g., green rest areas and low-pollution business districts) and enhancing route design for micromobility users through alternative pathways or physical buffering from traffic emissions.

Funding. None.

Acknowledgment. The authors would like to thank NetMob for providing the data that makes this study possible.

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. I. Testi, A. Wang, S. Paul, *et al.*, Nat. Cities **1**, 512 (2024).
2. A. Poulhès and L. Proulhac, Transp. Res. Part D: Transp. Environ. **98**, 102977 (2021).
3. Airparif, "Bilans et cartes annuels de pollution," (2024).

Day and Night in Paris: Vibrancy and Perceived Security Analysis from Mobility Data

OĞUZ YÜCEL[1, 2, *], İREM BETÜL KOÇAK[3, 4], AND ASLIGÜL AKSAN ERAKIN[1, 3]

[1] Management Engineering Department, Istanbul Technical University, Reşitpaşa, 34467 Sarıyer/İstanbul
[2] Management Department, Bogazici University, Bebek, 34342 Beşiktaş/İstanbul
[3] Industrial Engineering Department, Bogazici University, Bebek, 34342 Beşiktaş/İstanbul
[4] Value Analytics Labs, 100 Cambridge Street, Suite 1400 Boston, MA, 02114, USA
[*] yucelog@itu.edu.tr

1. INTRODUCTION

Mobile phone data has become an increasingly valuable resource in urban analysis. This study explores the dynamics of the city of Paris using the NetMob Data Challenge 2025 dataset [1], with a focus on identifying areas of high vibrancy, frequent visitation, and population density. Furthermore, the study aims to detect potentially less secure areas based on observed mobility patterns and the demographics of individuals present in those locations.

The Global Positioning System (GPS) enables precise geolocation by connecting to at least four satellites, allowing smartphones and GPS-equipped vehicles to capture high-resolution spatiotemporal movement data [2].This makes GPS a valuable tool for analyzing human mobility in smart cities. Alessandretti et al. (2018) [3] highlight the balance between exploratory and preferential behaviors in human movement and reveal a link between preferred location diversity and social interactions. After accounting for trajectory-specific distance variations and anisotropy, Gonzalez et al. (2008) [4] show that individual mobility patterns converge into a universal spatial distribution, suggesting that human movement follows reproducible and simple structures despite varied travel histories.

2. METHODOLOGY AND DISCUSSION

We merge a high-frequency GPS dataset (collected every 2–3 seconds from 3,300 individuals) with a GeoJSON file of Paris administrative boundaries. Using individual-level latitude and longitude coordinates, we analyze mobility patterns and displacement across the day, drawing demographic data from the individuals dataset and movement details from the trips dataset. Our study explores city vibrancy and perceived security: for vibrancy, we map population distribution, measure spatial displacement with the Haversine metric, and construct a network of the most visited areas, applying network analysis to interpret temporal density patterns and regional transitions.

To examine urban security, we analyze the aggregated presence of women and men separately. Prior research shows that the visible presence of women in public spaces often correlates with higher perceived and actual safety [5], though perceived and actual security may diverge in large cities [6]. Factors such as street lighting, crowd density, and time of day also shape perceptions of safety [7].

Using network analysis, we detect the most visited places in Paris for both genders. Each polygon from the Paris GeoJSON file is treated as a node, with node size reflecting total dwell time, and edges representing transitions between areas. Figure 1 highlights the top 10 areas by dwell time overall and by gender, with Centre-Est consistently among the most visited. While overall hotspots are similar, transition patterns differ between men and women, offering insights into gendered mobility and space use in the city.

The Haversine formula is used to calculate the distance between two points based on their latitude and longitude coordinates. In Figure 2, the total displacement of all users is aggregated and stratified by gender on an hourly basis. The x-axis represents the hours of the day and offers insight into whether men and women exhibit different levels of mobility. The results show that men and women follow similar overall mobility patterns. Therefore, the lower presence of women in certain areas during late evenings and nights cannot be explained by reduced displacement alone; rather, it suggests that women actively choose not to spend time in those areas during those hours.

Fig. 1. Top 10 Nodes Representation of Paris Areas

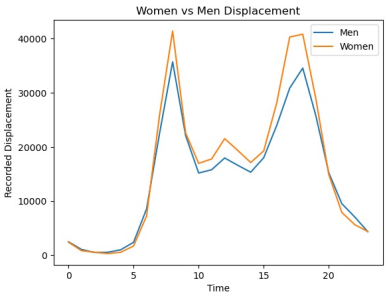

Fig. 2. Displacement Summations of Both Genders in Different Hours of a Day

To study mobility, we focused on walking trips, as pedestrian behavior reflects both vibrancy and perceived security [8]. For each trip, we calculated the shortest path using the OpenStreetMap API and compared it with the actual route using the standardized Fréchet distance. Data were stratified by age, Young (15–40), Middle-Aged (40–65), and Old (65+), and by gender. The metrics show that men display greater variability and more outliers, indicating longer, wandering routes, while women generally follow the shortest paths. One-way ANOVA across six age-sex groups confirmed a significant effect on Fréchet distance, $F(5, 2477) = 9.30$, $p < 0.001$, supporting gender differences in route choice. Deviations from the shortest path reflect safety concerns: women may take longer routes to avoid unsafe areas or stick to the shortest path to minimize exposure. They also walk less than men at night and before 6 a.m., with longer morn-

ing routes after sunrise suggesting cautious behavior. Overall, path deviations reveal gendered mobility patterns and varying security awareness across different times of day.

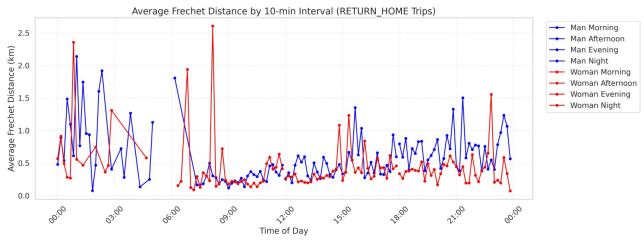

Fig. 3. Average Fréchet Distance in Different Hours of the Day

Origin and destination points for each trip were derived from the GPS data, and their similarity was assessed using DBSCAN clustering. The k-distance method helped select DBSCAN parameters, with k as the minimum points per cluster and eps as the maximum distance between points, calculated using the Haversine formula. Using k = 45 and eps is around 400 meters, clusters were identified, though some trips remained unclustered, reflecting trajectory variability in Figure 4. The spatial distribution of clustered trips is shown in the figure, and cluster similarity can be further evaluated using Jaccard similarity and frequency counts.

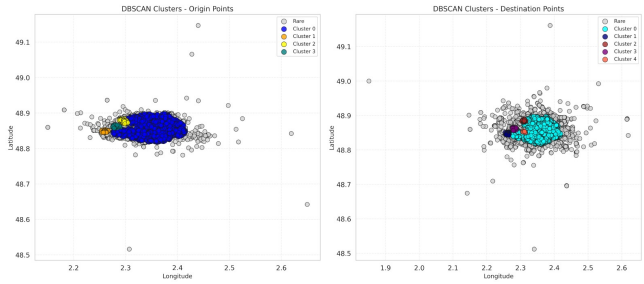

Fig. 4. DBSCAN Clustering

Finally, we found the crime rates of Paris in 2023 and the density of street lamps in September 2025[9] in Paris to compare them with deviations from the shortest paths. Although the analysis is limited by sample size, it still provides part of the picture. Street lighting, in particular, plays an important role in route choice, especially for women, when deciding which path to take. We compare these factors with the mean deviations from the shortest paths. Figure 5 presents choropleth maps of crime rates, the number of street lamps, and deviations from the shortest paths across the 20 Paris arrondissements.

The data suggest a complex relationship between route deviations, crime, and street lighting. Figure 5 shows the highest mean deviations from the shortest path in Arrondissements 17 and 1, while the highest crime numbers are in Arrondissements 18 and 15, indicating little direct correlation between crime and path choice. In contrast, a stronger relationship emerges with street lamps: Arrondissement 17, with many lamps, also has high mean deviations, suggesting that route choices are influenced more by perceived safety, such as lighting, than by crime. This is particularly relevant for women, whose decisions appear sensitive to lighting. Although limited by arrondissement-level aggregation, these results indicate that people may take longer, better-lit routes for safety rather than avoid areas with high crime, highlighting the importance of environmental factors in urban mobility.

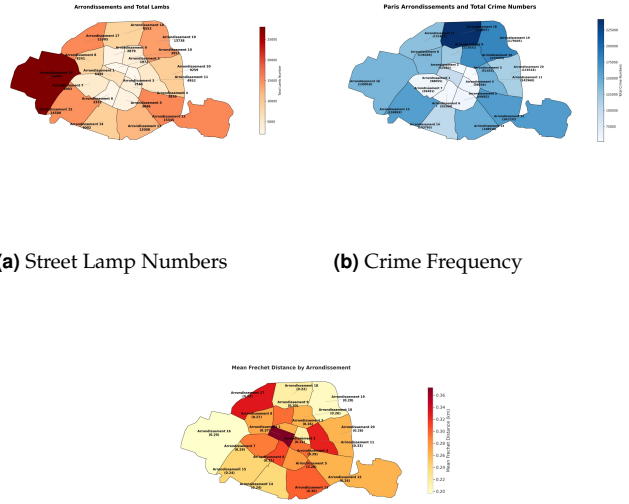

(a) Street Lamp Numbers **(b)** Crime Frequency

(c) Average Fréchet Distance

Fig. 5. Choropleth Maps

3. CONCLUSION & FUTURE WORK

This study integrates data and methods to explore urban vibrancy, providing preliminary insights into Paris's mobility and activity patterns across socio-economic groups. Using limited variables and simplifying assumptions, such as dividing the day into four intervals and tracking regional presence, we analyzed route deviations (Fréchet distance) alongside crime and street lighting, finding perceived safety, influenced by lighting, shapes behavior more than actual crime rates.

Analyses at the arrondissement level may be too coarse, so future work will incorporate localized data—high-crime areas, parks, open stores, and demographics—using buffer zones to assess route safety. We also plan to develop a *Dynamic Security Index (DSI)* combining trajectory patterns, temporal factors, and demographic heterogeneity to capture evolving urban safety.

REFERENCES

1. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, arXiv preprint arXiv:2506.05903 (2025).
2. J. Wang, X. Kong, F. Xia, and L. Sun, ACM SIGKDD explorations newsletter **21**, 1 (2019).
3. L. Alessandretti, P. Sapiezynski, V. Sekara, *et al.*, Nat. human behaviour **2**, 485 (2018).
4. M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, Nature **453**, 779 (2008).
5. N. Yavuz and E. W. Welch, Urban studies **47**, 2491 (2010).
6. M. Paydar, A. Kamani-Fard, and R. Etminani-Ghasrodashti, Cities **60**, 289 (2017).
7. L. J. Loewen, G. D. Steel, and P. Suedfeld, J. environmental psychology **13**, 323 (1993).
8. Y. Wang, C. K. Chau, W. Ng, and T. Leung, Cities **50**, 1 (2016).
9. City of Paris, "Paris Data: Open Data Portal of the City of Paris," https://opendata.paris.fr/pages/home/ (2025). Accessed: September 10, 2025.

Machine Learning-Based Netmob25 Mobility Model for Network Simulators

RAPHAEL COUSIN[1] AND THI MAI TRANG NGUYEN[2,*]

[1] SCAI, Sorbonne Université, 4 Place Jussieu, 75005 Paris, France
[2] L2TI, Universite Sorbonne Paris Nord, 99 avenue Jean-Baptiste Clement, 93430 Villetaneuse, France & LIP6, Sorbonne Université, Paris, France
[1] Raphael.Cousin@sorbonne-universite.fr
[*] ThiMaiTrang.Nguyen@univ-paris13.fr

1. INTRODUCTION

Human mobility plays an important role in mobile network dimensioning and deployment. User Equipment (UE) generates network traffic which must be supported by the network infrastructure. When users move, handovers happen and network traffic also moves from one base station to another. By taking into account user mobility, mobile operators can decide the number of base stations needed, their location and configuration to satisfy all traffic demands.

Network simulators such as NS3 [1] or OMNeT++ [2] are the main tool to carry research experiments on mobile networks. In network simulation, one of the most important steps is simulating the trajectory of UEs. Current mobility models inherently integrated in these network simulators are quite simple such as random walk or constant velocity.

In this paper, we use machine learning techniques to create a new mobility model from the Netmob25 dataset. Researchers can generate directly from NS3 a number of UE trajectories of different types such as car, metro or pedestrian which have mobility patterns matching the properties of the Netmob25 dataset. Therefore, researchers can simulate a network dimensioning or deployment problem with more realistic UE trajectories.

2. DATA PREPROCESSING

The NetMob25 dataset contains GPS trajectories from 3,337 individuals in the Île-de-France (IDF) region collected over one week. To prepare this data for machine learning models, we implemented a comprehensive preprocessing pipeline with the following steps:

Filtering: We applied sequential filters to ensure data quality: (1) removal of trips with missing date/time values (5.0% removed), (2) exclusion of trips outside IDF bounds (0.7% removed), (3) exclusion of trips longer than 3 hours and shorter than 30 seconds (0.3% removed), (4) filtering of excluded transport modes (OTHER, PLANE, ON-DEMAND) (0.05% removed).

Mode Categorization: We simplified transport modes into five predominant categories sufficient to capture the essential speed patterns needed for realistic network simulation experiments: CAR (including private cars, passengers, and taxis), PUBLIC_TRANSPORT (bus, train, subway, tramway), WALKING, BIKE (including e-bikes and e-scooters), and MIXED (multimodal trips).

Speed Outlier Management: Outliers exceeding the 99th percentile speed threshold were clipped to this maximum value by transport mode.

Trajectory Processing: GPS traces were interpolated to 1-second intervals and resampled to 2-second intervals. Sequences were cut into fixed-length chunks of 2000 points (approximately 66 minutes, which represents more than 97 % of the trips) with padding for shorter sequences. This resulted in 71,655 sequences from 68,070 validated trips.

Feature Scaling: We used MinMaxScaler to normalize latitude, longitude, and speed features, creating a final dataset with shape (71655, 2000, 3).

Table 1. Transport Mode Statistics in Preprocessed Dataset

Mode	Trips	Duration (min)
BIKE	6,115	19.6 ± 15.0
CAR	25,876	19.5 ± 15.7
MIXED	9,528	39.4 ± 21.8
PUBLIC_TRANSPORT	10,012	27.7 ± 18.1
WALKING	20,135	12.1 ± 11.5

3. MODEL ARCHITECTURE

To generate realistic trajectories from the NetMob25 dataset, we explored Variational Autoencoders (VAEs) [3]. VAEs are generative models that learn to encode data into a continuous latent space and decode it back to the original space. Unlike standard autoencoders, VAEs impose a probabilistic structure on the latent space, enabling generation of new, diverse trajectories by sampling from the learned distribution.

Our Conditional VAE (CVAE) conditions on transport mode and trajectory length to capture mode-specific movement patterns and generate trajectories with desired characteristics.

$\mathcal{L} = \|M \odot (x - \hat{x})\|_2^2 + \beta \cdot D_{KL}(q_\phi(z|x,c,l)||p(z))$ where $x \in \mathbb{R}^{T\times 3}$ represents trajectory sequences (latitude, longitude, speed), c denotes transport mode, l is trajectory length, $z \in \mathbb{R}^d$ is the latent representation, $\hat{x} = \text{Decoder}\theta(z,c,l)$, M masks padded sequences, and β balances reconstruction and regularization. We evaluated four encoder-decoder architectures: Multi Layer Perceptron (MLP), Long Short-Term Memory (LSTM) [4], Attention [5], and Convolutional Neural network (CNN) [6].

4. TRAINING METHODOLOGY

Data Split & Computational Setup: We employed an 85/15 split for training and validation. Training required 2-15 hours on NVIDIA A100 GPU with 40GB RAM depending on model complexity.

Hyperparameter Optimization: We employed the Hydra framework for systematic hyperparameter exploration across different architectural backbones. Key search parameters included hidden dimensions (16-512), latent dimensions (8-256), and batch sizes (16-512). To mitigate posterior collapse, we implemented a beta annealing strategy, gradually increasing β from 0 to 0.5 over the first 100 epochs. This approach encourages the model to initially focus on reconstruction quality before progressively enforcing the latent space regularization. Model selection

was based on validation performance using the reconstruction loss (we relied on validation set due to limited data quantity). We retained only models achieving a KL divergence below 0.5, ensuring meaningful latent representations.

Training Configuration: Adam optimizer ($\alpha = 10^{-4}$), ReduceLROnPlateau scheduling (patience=15, factor=0.5), gradient clipping (max norm=1.0), early stopping (patience=30, $\delta = 0.0001$).

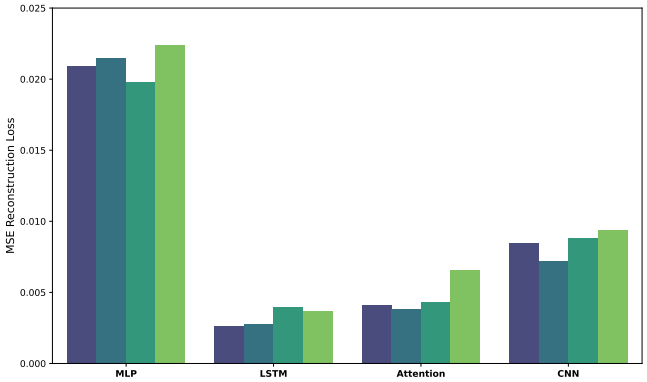

Fig. 1. Best Reconstruction loss comparison across model architectures.

A. Key Findings

Architecture Performance: The LSTM architecture achieved the best overall performance with a Mean Square Error (MSE) reconstruction loss of 0.0026, followed by Attention (0.0045), CNN (0.0085), and MLP networks (0.0209). The best model demonstrates trajectory reconstruction capabilities across all evaluation metrics, Speed Mean Absolute Error (MAE) of 6.749 km/h, Point Distance MAE of 1.985 km, Total Distance MAE of 4.164 km, and straight-line Distance MAE of 2.347 km.

5. GENERATION QUALITY EVALUATION

Methodology: To ensure statistical comparison, we generated trajectories with identical conditioning parameters (duration, transport_mode) as the real data and computed global metrics across all generated samples. Key evaluation criteria include speed distribution alignment, straight-line distance preservation, and sinuosity analysis ($\sigma = \frac{d_{straight}}{d_{total}}$, where $d_{straight}$ represents Euclidean distance between trajectory endpoints and d_{total} denotes cumulative path length).

Results Analysis: Generated trajectories exhibit significant deviations from real data in distance metrics across all transport modes. For most modes, the model underestimates distances: CAR shows 38.6% reduction in straight-line distance and 27.1% in total distance; BIKE exhibits 35.5% and 29.7% reductions respectively; PUBLIC_TRANSPORT displays 35.5% and 20.5% reductions; and MIXED mode shows the largest underestimation with 52.5% and 42.7% reductions.

Notably, WALKING mode presents an opposite pattern with substantial overestimation: straight-line distance increases by 25.2% and total distance by 37.7%. This divergent behavior suggests the model struggles to capture mode-specific characteristics. The consistent reduction in variance across all generated trajectories (evident from the narrow standard deviations) confirms mode collapse, where the VAE converges to average behaviors rather than reproducing the full distribution of real-world trajectory variations.

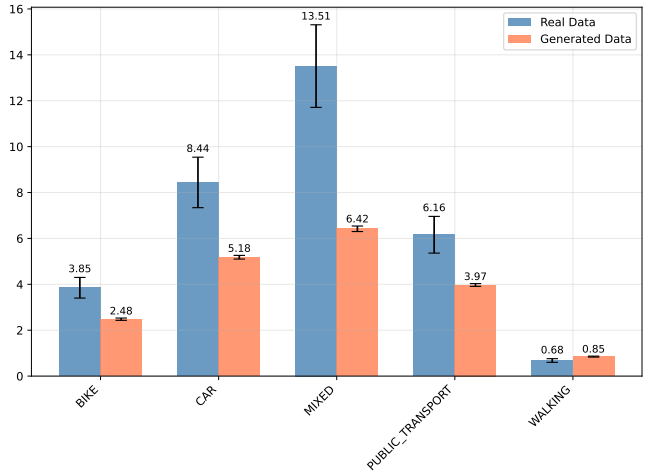

Fig. 2. Straight-line distance comparison.

6. NETMOB25 MOBILITY MODEL IN NS3

In order to use the trajectory generator in NS3, we define a new mobility model called `Netmob25MobilityModel` [7] integrated in the `src/mobility` module of NS3. Libtorch is used to run the VAE model from a C++ program.

This mobility model provides developers with APIs to choose the transport mode and generate a number of trajectories in IDF. These trajectories can be installed on UE nodes using conventional methods in NS3. Even though the generated trajectories don't really fit into all the streets in the IDF map, they still reflect the realistic UE density and movement tendency in different mobility areas. These factors are essential elements to be considered in cellular network capacity dimensioning and design. Study on handover needs further improvement on the model.

7. CONCLUSION

We presented a ML-based mobility model bridging real-world GPS data and network simulation. Despite reconstruction limitations, our CVAE framework captures aggregate mobility patterns essential for network research. Our framework provides a complete package for training trajectory models with seamless NS3 integration, enabling realistic mobility simulation for network experiments.

Future work will explore vector-quantized variational autoencoders (VQ-VAEs) and diffusion models, and incorporate temporal conditioning for improved trajectory diversity and realism. With the growing trend of using ML in network simulators, it would be valuable to develop direct Python bindings for NS3, enabling more natural integration of ML-based trajectory experiments.

Availability: The complete implementation, including training scripts and NS3 integration, is available at [7].

REFERENCES

1. https://www.nsnam.org/.
2. https://omnetpp.org/.
3. D. P. Kingma and M. Welling, Int. Conf. on Learn. Represent. (ICLR) (2013).
4. S. Hochreiter and J. Schmidhuber, Neural Comput. **9**, 1735 (1997).
5. A. Vaswani, N. Shazeer, N. Parmar, *et al.*, Adv. Neural Inf. Process. Syst. **30** (2017).
6. Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, Proc. IEEE **86**, 2278 (1998).
7. https://gitlab.lip6.fr/trnguyen/netmob25.

Mobility Under Pressure: Socioeconomic Inequities in Urban Travel Behavior During Urban Strikes

SAUMITRA KULKARNI[1], BIJIN JOSEPH[1,2], AND ESTEBAN MORO[1,2,*]

[1] Network Science Institute, Northeastern University, Boston, MA, USA
[2] Department of Physics, Northeastern University, Boston, MA, USA
[*] e.moroegido@northeastern.edu, bi.joseph@northeastern.edu

1. INTRODUCTION

Mobility facilitates life in cities, enabling access to employment, education, healthcare, and social participation. However, this everyday function is increasingly subject to disruption, from climate events to labor strikes to infrastructure failures. Understanding how individuals are affected and adapt their mobility behavior in response to such disruptions is critical not only for transportation planning or urban resilience and sustainability, but also to ensure equal accessibility across urban populations [1].

At the same time, disruption creates natural or semi-natural experiments that can modify individual choices to new mobility options. For example, during strikes, individuals may be forced to explore alternatives like biking, walking, or staying home [2]. These temporary changes can have lasting effects, with some users adopting new mobility routines permanently. In this way, disruptions not only reveal vulnerabilities in urban mobility but might also catalyze long-term behavioral changes [3, 4].

However, the analysis of those changes is limited by the aggregated, coarse-grained nature of mobility surveys, which often lack the temporal resolution and behavioral context needed to capture short-term adaptations and their long-term consequences. As a result, our understanding of disruption-induced behavioral change and its implications for equity, sustainability, and system design remains incomplete.

2. METHODS

In this study, we leverage the NetMob25 dataset to quantify behavioral responses to public transportation strikes in the Île-de-France region [5]. The dataset originates from a high-resolution GPS-based mobility survey involving 3,337 individuals, who were tracked over seven consecutive days between October 2022 and May 2023. In addition to raw GPS traces collected every 2–3 seconds, the dataset includes rich individual-level information on demographics, socioeconomic status, and detailed records of each trip, covering origin, destination, timing, duration, transport mode, and purpose.

Crucially, the observation period coincides with several major labor strikes and public demonstrations in France, including the mass mobilizations against pension reform [6] and cost of living protests [7]. These events drew hundreds of thousands of participants and caused widespread disruptions to public transportation and road traffic across the region. Since 2007, French law has required a minimum level of public transport service during strikes, which ensures over 50% service in peak hours in Île-de-France but often leaves rural and less dense areas with far fewer options, as the law doesn't mandate fixed service levels and staffing shortages limit effectiveness. The presence of multiple strike days creates a unique opportunity for quasi-experimental analysis. By comparing mobility patterns on strike and non-strike days within the same individuals, we can isolate the behavioral impact of transportation disruptions and examine how different population groups respond to these shocks.

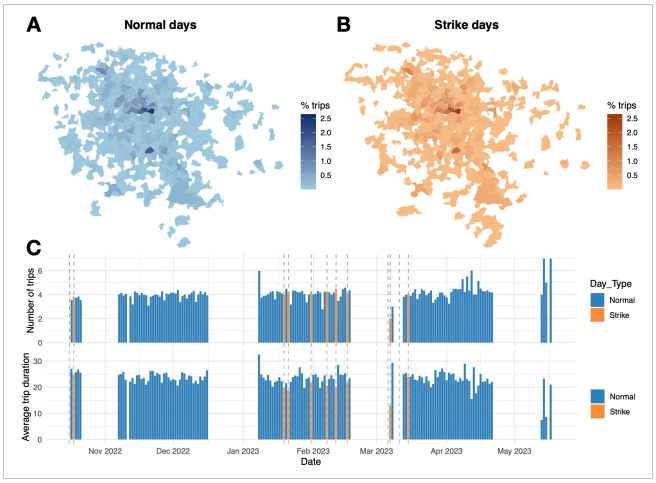

Fig. 1. Comparison of visitation patterns on normal days versus strike days.(A-B) Total number of visits to different IRIS in Paris on normal and strike days.(C) Number of trips (top), and average duration (bottom) for each day in our dataset. Days are colored according to strike or not.

3. RESULTS

As a first approximation to evaluate the effect of disruptions on mobility behavior, we compute differences in the total number of trips, their duration, and their area of destination. As shown in Figure 1A, there is a noticeable decrease in the number of visits to certain areas. Interestingly, in Figure 1C, we find that while the total number of trips does not change significantly during strike days, their average duration clearly decreases. This observation suggests that individuals adapt their behavior by maintaining similar levels of travel frequency, but shortening or abandoning longer trips. Taken together, these patterns indicate that mobility during strikes is reshaped rather than suppressed. To move beyond descriptive comparisons, we implement causal inference techniques to estimate the strike-day effect on mobility outcomes. Specifically, we use a difference-in-differences (DiD) design that compares strike versus non-strike days within individuals, while controlling for individual fixed effects and clustering standard errors at the person level. This approach isolates the causal impact of disruptions from unobserved heterogeneity in baseline travel behavior. Moreover, because the DiD framework can be extended to include interactions with demographic characteristics, it also allows us to investigate whether the effects of strikes differ systematically across population groups.

Formally, we estimate:

$$Y_{it} = \alpha_t + \beta\,\text{Strike}_t + \gamma_i + \epsilon_{it} \quad (1)$$

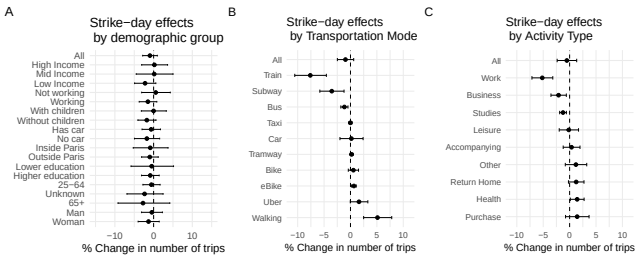

Fig. 2. Effect of strikes on the daily trips by demographic group, transportation mode, and trip purpose. (A) Conditional Average treatment effect (in percentage change) in the number of trips by demographic group. (B) Average treatment effect (in percentage change) in the number of trips due to strike days by main mode of transportation by trip.(C) Average treatment effect (in percentage change) in the number of trips due to strike days by trip purpose.

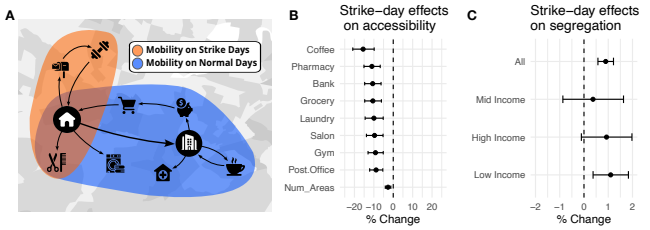

Fig. 3. Effect of strikes on the daily accessibility and experienced income segregation. (A) Illustration of changes in the mobility of individuals from long-distance or work-related trips of normal days to more localized activities around home during the strike days. (B) Average treatment effect (in percentage change) in the daily accessibility by different types of essential services. (C) Average treatment effect (in percentage change) in the experienced segregation due to strike days by income groups.

where Y_{it} denotes a mobility outcome (e.g., number of trips, total duration) for individual i on day t, Strike_t is an indicator for strike days, γ_i and α_t are individual and monthly fixed effects, and ϵ_{it} is the error term. Because participants are observed only for short survey periods, including high-frequency time fixed effects (e.g., day-of-week) would lead to under-specification. The coefficient β thus captures the Average Treatment Effect (ATE), i.e, the causal impact of strikes on mobility outcomes.

Building on the DiD model, the main results can be summarized as follows:

- Strikes have a negligible impact on the number of trips (ATE = -0.95%, $p = 0.36$) but significantly shorten their duration (ATE = -13.4%, $p = 0.00512$), showing that people maintain travel frequency while abandoning longer journeys.
- Subgroup analysis (by age, gender, education, car ownership, household type, residence, and professional status) reveals uniformly small and statistically insignificant effects, indicating a broadly shared population response.
- Public transport trips (train, subway, bus) decline sharply, while walking and ride-hailing increase, keeping trip counts stable but reducing average trip length.
- Activity types shift: work and study trips drop most (-5% to -1%), while health and shopping trips slightly rise (+1%), reflecting substitution toward essential and local activities.
- Accessibility contracts: while geographical coverage decreases modestly (-2.7%, $p = 0.003$), exposure to essential services falls much more sharply, including pharmacies (-11.0%), banks (-10.6%), groceries (-10.5%), and coffee shops (-15.5%), among others.
- Income segregation intensifies: by reducing long-distance trips, strikes increase the likelihood that individuals visit areas with similar income levels to their own, amplifying experienced daily segregation by 1% daily.

These results show that while the contraction in the spatial extent of daily mobility is modest, the loss of exposure to essential amenities is much larger, underscoring how strikes reshape not only the quantity of mobility but also its functional value for urban life. During strikes, individuals lose access to critical places in their daily lives.

4. CONCLUSION AND FUTURE WORK

Previous studies have examined the effects of transport strikes on ridership, congestion, and travel times, but less on their broader social and equity consequences [2, 8–10]. Strikes are a recurring feature of urban life in France, affecting millions of commuters and shaping access to opportunities. Using detailed GPS data and a causal inference approach, we find that strike days barely reduce the *number* of trips but significantly reduce travel *time*. This shift is consistent across demographic and socioeconomic groups, reflecting common adaptation strategies: declines in rail, metro, and bus travel are offset by increases in walking and driving, with fewer trips to work or school. Mobility becomes more local and less diverse, with reduced access to essential services and greater income segregation. Overall, strikes do not suppress mobility but shorten, localize, and narrow it, with important welfare and equity implications.

Future work will combine GPS traces with Foursquare POIs to examine venue choices, activity durations, and trip-chaining during disruptions. We will strengthen causal validity with placebo tests and alternative models, and assess external validity by replicating in other cities and extending to non-strike disruptions (e.g., weather or service incidents). Finally, integrating route-level service data will allow us to estimate dose–response functions of strikes and spatial spillovers, linking mobility outcomes to transport supply shocks.

REFERENCES

1. T. Yabe, N. K. W. Jones, P. S. C. Rao, *et al.*, Comput. Environ. Urban Syst. **94**, 101777 (2022).
2. M. Saberi, M. Ghamami, Y. Gu, *et al.*, J. Transp. Geogr. **66**, 154 (2018).
3. E. A. Martinez, M. H. Rubio, R. M. Martinez, *et al.*, "Measuring Economic Resilience to Natural Disasters with Big Economic Transaction Data," (2016). ArXiv:1609.09340 [cs].
4. S. Larcom, F. Rauch, and T. Willems, The Q. J. Econ. **132**, 2019 (2017).
5. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The NetMob25 Dataset: A High-resolution Multi-layered View of Individual Mobility in Greater Paris Region," (2025). ArXiv:2506.05903 [cs].
6. Wikipedia contributors, "2023 french pension reform strikes — wikipedia, the free encyclopedia," https://en.wikipedia.org/w/index.php?title=2023_French_pension_reform_strikes (2025). Accessed July 17, 2025. Page Version ID: 1294984378.
7. Wikipedia contributors, "2022 french protests," https://en.wikipedia.org/w/index.php?title=2022_French_protests (2025). Accessed July 17, 2025. Page Version ID: 1270077945.
8. I. Spyropoulou, Transp. Res. Part A: Policy Pract. **132**, 651 (2020).
9. S. Bauernschuster, T. Hener, and H. Rainer, Am. Econ. Journal: Econ. Policy **9**, 1 (2017).
10. M. L. Anderson, Am. Econ. Rev. **104**, 2763 (2014).

Mobility Patterns, Segregation, and Pollution Exposure in the Context of the Parisian Low-Emission Zone

FRANZISKA ZOLLNER[1,*]

[1] Chair of Policy and Sustainability, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
[*] franziska.zollner@epfl.ch

1. INTRODUCTION

Low-emission zones (LEZs) are increasingly used across Europe to reduce air pollution and promote sustainable mobility by restricting access to urban areas for high-emission vehicles. Between 2019 and 2022, the number of active LEZs increased from 228 to 320, with projections of over 500 by 2025 [1]. While a growing body of research has examined the environmental and health impacts of LEZs (e.g., [2–7]), there is limited evidence on how LEZs affect mobility patterns across socioeconomic groups.

This study examines the heterogeneous implications of the Parisian LEZ for mobility, spatial segregation, and air pollution exposure. Paris first introduced its LEZ within the inner-city ring road in 2017, expanding it to the Greater Paris Area in July 2019. A 2024 Airparif report attributes 3-6% reductions in air pollution to the LEZ, depending on the pollutant [8]. Despite this success in terms of pollution reduction, the policy remains controversial. It has been criticized for creating a form of "social segregation by excluding millions of French citizens from city centers" [9]. Concerns over the disproportionate burden placed on low-income residents have lead up to the National Assembly's vote to abolish LEZs in May 2025 [10].

In this project, I explore how mobility patterns, spatial segregation, and air pollution exposure vary across socio-professional (SP) category within the context of the Parisian LEZ. I document heterogeneous travel behaviors: some groups undertake longer and more frequent trips with higher car reliance, whereas others tend to rely more on public or active transport. While no group appears entirely excluded from central areas, differences in mobility characteristics suggest that certain individuals — particularly those who are car-dependent and travel longer distances — may face a greater burden from the regulation. Air pollution exposure does not show strong systematic differences by SP category.

2. METHOD

Mobility Patterns and Exposure to LEZ: I provide a descriptive analysis of mobility patterns across SP categories, focusing on differences in trip frequency, duration, length, and mode choice. Trips are classified based on LEZ exposure using two approaches: (i) trips starting or ending within the LEZ during active hours (weekdays, 8 AM–8 PM) and (ii) trips that traverse the LEZ during these hours, using GPS traces for a more granular measure. To capture potential behavioral responses to the LEZ, I analyze the timing of LEZ entry points to detect bunching around regulation hours and compute detour ratios (actual trip length over straight-line distance) to identify potential route deviations around the LEZ. I aggregate origin–destination (OD) flows by SP category and transport mode, with car trips further disaggregated by travel times.

Experienced Segregation and Spatial Inequality: To assess intergroup mixing in daily activity spaces, I calculate the SP composition of visited IRIS using GPS pings. The spatial distribution of experienced segregation is captured using an entropy index at the IRIS level:

$$I_z = -\sum_{g=1}^{G} p_{zg} \times \ln(p_{zg})$$

where p_{zg} is the proportion of SP category g in IRIS z. This index is normalized by dividing it through $ln(G)$, where G is the number of SP categories. An index of 0 represents no mixing of groups, whereas an index of 1 means all groups are represented equally. I then derive individual exposure to other SP categories as a weighted average of encountered categories in the visited IRIS units.

Environmental Justice and Pollution Exposure: A common criticism of LEZs is that lower-income individuals may bear disproportionate costs without corresponding environmental benefits. To examine differential exposure to air pollution, I link hourly GPS-based locations to NO_x concentrations from Airparif [11] to examine differential pollution exposure across SP groups. Hourly mean locations are estimated from GPS pings. If the individual is in motion, this is calculated as the mean latitude and longitude of GPS pings within the hour. If stationary, I assign the last recorded GPS location. Each hourly location is then matched to the nearest Airparif measurement station, and the NO_x value at that time is assigned.

3. RESULTS

Mobility Patterns and Exposure to LEZ: Table 1 reports the weighted mean number of trips per week, trip duration, trip length, and likelihood to use a car at any point of the trip by SP category. Trip frequency and durations are fairly homogeneous across groups. Workers stand out by recording more frequent, longer, and further trips compared to other categories. Retired and unemployed individual do shorter trips, suggesting more localized mobility patterns. In terms of main transport mode, car use dominates among workers who use a car for about half of their trips. Students show the lowest car-reliance.

SP Category	# Trips	Duration (min)	Length (km)	% Car Use
Craftsman	29.1	23.2	8.3	41.0
Employee	27.8	24.0	8.0	43.7
Executive	27.9	24.0	7.3	30.5
Intermediate	27.5	24.4	8.4	41.1
Retired	26.2	21.6	5.6	40.0
Student	23.4	28.7	7.7	17.4
Unemployed	27.2	21.4	5.6	29.4
Worker	34.1	27.9	10.4	51.1
Total	27.5	24.4	8.9	35.6

Table 1. Weighted mean of mobility indicators by SP category

The left panel of Figure 1 shows the distribution of trip categories by SP category. Patterns are relatively homogeneous:

across all groups, a large share of trips start or stop within the active LEZ. Traversing trips are rare, with slightly higher shares for students, unemployed, and retirees. Trips within the LEZ outside active hours account for 15–25% of trips. Employees, workers, unemployed, and retirees record the highest share of trips entirely outside the LEZ. The right panel plots the main transport mode used by SP category for trips in the LEZ during active hours. Workers and retirees are the most car-reliant categories. Students mainly use public transport, while executives and unemployed rely on active transport modes.

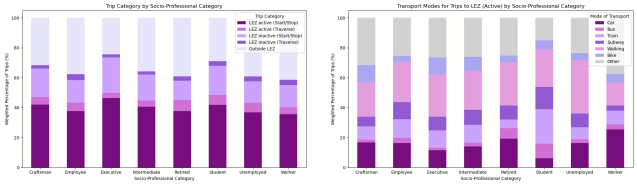

Fig. 1. Trip categories & transport mode by SP category

Figure 2 illustrates OD flows for workers as an exemplary SP category, disaggregated by transport mode (upper panel) and for car trips only, distinguishing between LEZ active and inactive hours (lower panel). A pronounced spatial divide emerges between car and non-car flows, a pattern consistent across all SP categories. Car trips predominantly originate from and extend toward the outer suburbs and peri-urban areas, while non-car trips are more concentrated close to the LEZ. Restricting the analysis to car trips, flows during and outside of LEZ active hours display broadly similar spatial patterns, though overall volumes are lower outside active hours. I do not find evidence that any SP category systematically shifts car travel into the LEZ to periods outside regulation hours.

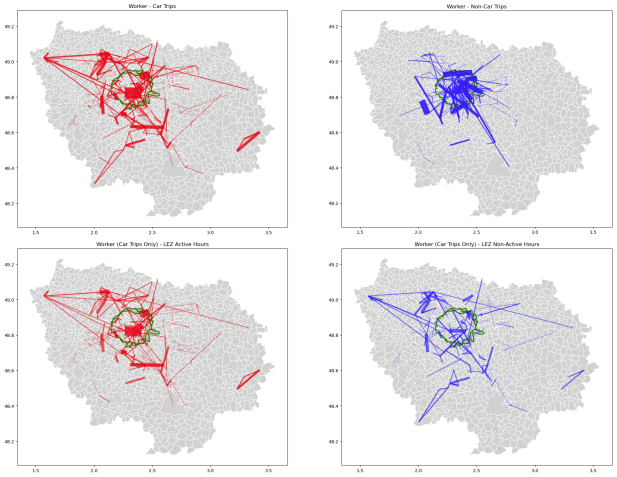

Fig. 2. Worker's OD flows by transport mode & LEZ hours

No systematic bunching of trips around the LEZ regulation hours is observed. Similarly, trips with unusually high detour ratios or long distances do not cluster near the LEZ boundary.

Experienced Segregation: Figure 3 plots the normalized entropy index across Île-de-France at the IRIS level. The peripheral regions are characterized by lower levels of population mixing.[1] Population mixing within the LEZ, outlined in green, is high, consistent with the concentration of economic activity in Paris.

[1]Lower mobility in peripheral zones in the data can inflate experienced segregation values.

Some parts of lower population mixing can be observed along the LEZ boundary, particularly in the northeast.

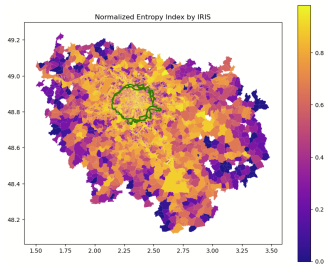

Fig. 3. Distribution of entropy index by IRIS

Individual exposure to other SP categories is highly homogeneous across SP categories. Overall, people are mostly exposed to executives, intermediate professionals, and employees. There are no significant differences of exposure to other SP categories between LEZ active and inactive hours.

Air Pollution Exposure: The left panel of Figure 4 shows mean weekly NO_x exposure by residential IRIS, accounting for individuals' movements throughout the day. Exposure is highest for residents in the eastern part of Paris, within and near the LEZ boundary. Peripheral areas do not exhibit systematically high pollution levels, suggesting that residents farther from the city center are not disproportionately affected. The right panel presents the distribution of weekly NO_x exposure by SP category. No clear relationship is observed: extreme values occur across groups, and mean exposure levels are broadly similar.

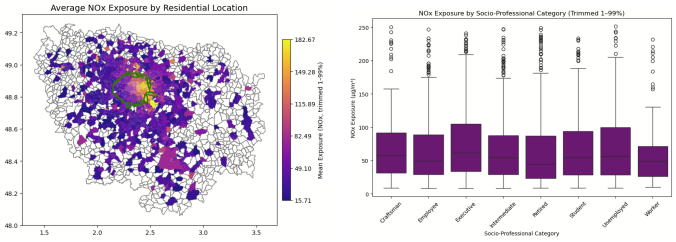

Fig. 4. Weekly NO_X exposure by residence & SP category

REFERENCES

1. Z. Azdad, B. Stoll, and J. Müller, "CLEAN CITIES The development trends of low and zero-emission zones in Europe," (2022).
2. H. Wolff, The Econ. J. **124**, F481 (2014).
3. M. Gehrsitz, J. Environ. Econ. Manag. **83**, 121 (2017).
4. S. Margaryan, J. Health Econ. **76**, 102402 (2021).
5. H. Klauber, F. Holub, N. Koch, *et al.*, Am. Econ. Journal: Econ. Policy **16**, 220 (2024).
6. A. Rohlf, F. Holub, N. Koch, and N. Ritter, Econ. Lett. **192**, 109221 (2020).
7. L. Sarmiento, N. Wägner, and A. Zaklan, J. Public Econ. **227**, 105014 (2023).
8. Airparif, "Zone à faibles émissions mobilité dans la métropole du Grand Paris– Évaluation de l'impact de la ZFE-m," Tech. rep., Airparif (2024).
9. Fleet Europe, "Why does france want to abolish its low-emission zones?" https://www.fleeteurope.com/en/new-energies/france/analysis/why-does-france-want-abolish-its-low-emission-zones?a=FJA05&t%5B0%5D=France&t%5B1%5D=Emissions&curl=1 (2025). Accessed: 2025-08-18.
10. BBC, "French MPs vote to scrap low-emission zones in blow to Macron," https://www.bbc.com/news/articles/c0mrpl2208no (2025). Accessed: 2025-08-18.
11. Airparif, "mes_idf_horaire_nox," https://data-airparif-asso.opendata.arcgis.com/datasets/dd6773adac1c47848d5d341c7184a5ee/about (2023). Dataset, accessed 2025-09-16.

Impact of income-stratified mobility mixing and vaccination uptake on epidemic progression: a modeling approach on emerging pandemic flu

Lorenzo Lucchini,[1,2] Simone Centellegher,[1] Paolo Bosetti[3,*]

[1]*Fondazione Bruno Kessler, Via Sommarive 18, Trento, 38123, ITALY*
[2]*Bocconi University, Via Röntgen 1, Milan, 20136, ITALY*
[3]*Mathematical Modelling of Infectious Diseases Unit, Institut Pasteur, Université Paris Cité, INSERM U1332, CNRS UMR2000, Paris*
*paolo.bosetti@pasteur.fr

1. INTRODUCTION

In recent years, the world has witnessed an increasing emergence of epidemics and pandemics caused by infectious pathogens. These events have posed serious challenges to global health systems, social structures, and economies [1]. Of growing concern is the evolving threat posed by highly pathogenic avian influenza viruses, particularly H5N1. The virus has become widely endemic in avian populations and has recently been reported in an expanding range of species [2]. Alarming developments in the United States have highlighted a rapid spread of H5N1 among cattle, with confirmed cases of human infections, particularly among livestock workers, raising fears of potential zoonotic spillover and sustained human-to-human transmission [3, 4].

These trends emphasize the urgent need for proactive surveillance systems and the planning of timely public health responses to mitigate the threat of future pandemics [5, 6]. The experience of the COVID-19 pandemic has shown both the vulnerabilities and strengths of current health infrastructures. In many Western countries, a combination of nonpharmaceutical interventions followed by large-scale vaccination campaigns succeeded in significantly reducing the impact of the initial waves of SARS-CoV-2 and transitioning the pandemic into a more manageable phase. However, the deployment of vaccines, despite unprecedented speed and innovation, revealed considerable heterogeneities in vaccine uptake [7]. These disparities often reflected underlying socioeconomic inequalities. For instance, in France, geographic areas characterized by lower average household income exhibited systematically lower vaccination rates compared to wealthier regions when comparing the same age categories [8]. Such observations underscore the need to account for structural and behavioral heterogeneities when modeling epidemic dynamics and designing effective intervention strategies. In this context, understanding spatial mixing patterns of human mobility and the integration of such information into epidemic models for infectious disease spread has become more and more important. Being able to efficiently leverage GPS sourced mobility data to reliably inform such models is thus a pressing challenge for the infectious disease community [9].

Here, we develop a metapopulation model to investigate how differences in vaccination coverage, shaped by socioeconomic disparities, can influence the spread of a flu pandemic across the various geographic regions of the Paris metropolitan area. The model incorporates income stratified mobility matrices derived from NetMob25 challenge dataset and heterogeneous vaccination uptake. By capturing geographical aspects of population structure and by leveraging mobility mixing patterns at the arrondissement level, our approach offers insights into how targeted vaccination strategies could reduce health inequities and enhance epidemic control.

2. RESULTS

We use the NetMob25 dataset to extract income-stratified mobility patterns across the Paris metropolitan area. The NetMob25 dataset contains individual-level GPS mobility data combined with several sociodemographic attributes. However, it does not include explicit information on household income. To address this limitation, we infer the household income category for each user by training an XGBoost classifier using data from the external IMMUNE Behavioural survery [10], also conducted in the Île-de-France region (see Methods for full details). Each user is thus assigned to one of three monthly household income categories: low (less than €2,000), middle (between €2,000 and €3,300), or high (more than €3,300).

Using this income classification, we analyze patterns of mobility through the Trips dataset, which records detailed travel flows. We aggregate individual-level trips to construct origin-destination (O/D) matrices stratified by income group. These matrices are generated by grouping mobility data at the level of the nine arrondissements of the Métropole du Grand Paris and the 20 arrondissements of Paris. The resulting matrices (see Fig. 1) allow us to capture differences in movement behavior across socioeconomic segments. Fig. 2 shows the trip distribution of the users stratified by income group showing differences in mobility patterns.

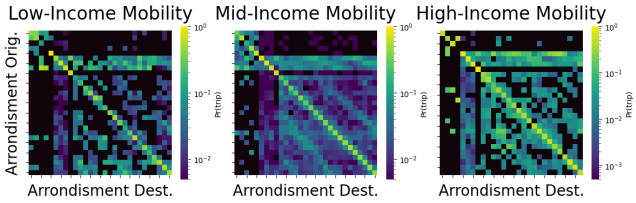

Fig. 1. O/D matrices of users classified as low, middle and high household income. Origins and destinations include the arrondissement within the greater Paris metropolitan area and the municipal arrondissement of Paris.

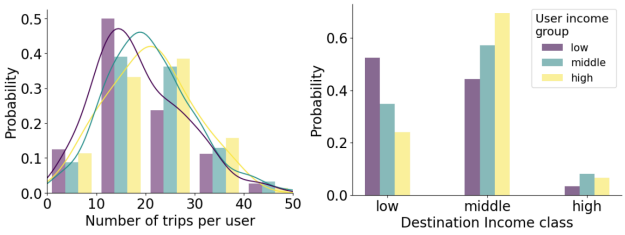

Fig. 2. (left) Distribution of number of trips for each income group. (right) Trip income distribution of users assigned to one of three monthly household income categories we defined (low,middle, high), *looking at the household income class of the commune destinations of the trips.*

Building on this stratified mobility data, we are developing a metapopulation epidemic model that explores how disparities in vaccination coverage, shaped by socioeconomic inequality, influence the

spread and attack rate of a respiratory infectious disease across the metropolitan area. The model incorporates the income-stratified mobility matrices, enabling us to simulate transmission dynamics more accurately across population subgroups and locations.

To inform the vaccination component of the model, we use real-world COVID-19 vaccination data as a proxy for the uptake of a hypothetical vaccine against an airborne pathogen. For each commune, we fit an innovation diffusion model, where vaccine adoption is treated as a form of innovation diffusion (see Methods). From this model, we estimate the carrying capacity parameter K, which represents the maximum fraction of the population expected vaccine uptake in the absence of vaccine mandates (see Fig. 3). By integrating heterogeneous vaccination coverage with income-based mobility patterns, our epidemic model can capture the interaction between population structure, mobility behavior, and socioeconomic inequalities. This approach provides insights into how targeted vaccination strategies, tailored to the mobility and uptake patterns of different income groups, could mitigate health disparities and improve epidemic control outcomes.

3. METHODS

A. Household income inference

To infer the household income for the users in the Netmob25 dataset, we integrate data with the IMMUNE behavioural study [10], collected also in the Île-de-France region. This dataset includes, among other variables, ZIP code–level geographic information and self-reported household income.

For each survey respondent, we used the following variables enriched with census data from INSEE[1234], incorporating additional information such as: *Demographic and Socioeconomic Variables* (Age, Sex, Education, Occupation, Main mode of transport, Household size, Number of household members aged 0–10, Number of household members aged 11–17, Number of household members aged 18–24, Number of household members aged 25–64, Number of household members aged 65 and over, Household includes a minor, Household includes an elderly person); *Community Characteristics* (Commune population density, Commune altitude); *Availability of Facilities (*Number of facilities for Particular services, Shops, Education, Health and social services, Transport and travel, Sports, leisure and culture, and Tourism); *Economic Indicators* (Median household net income, Interdecile income ratio, Poverty rate, Deprivation level).

Using this enriched dataset, we train an XGBoost classifier that predicts monthly household income categories defined as (i) *low*: less than €2000; (ii) *middle*: between €2000 and €3300; (iii) *high*: more than €3300. We achieve an accuracy of 73%. Each user in the NetMob25 dataset is then assigned a predicted household income category based on the trained model.

B. Estimation of vaccine uptake

To simulate the uptake of a hypothetical airborne disease vaccine, we build on the framework proposed in [11]. The impact of COVID certificates on vaccination uptake is estimated using a model rooted in innovation diffusion theory, where vaccination is treated as an innovation, and each eligible individual may choose to adopt it over time.

Let t_0 denote the date of vaccine introduction, and let $x(t) \in [0, 1]$ represent the cumulative fraction of the population that has received at least one dose by day t. By construction, $x(t)$ is non-decreasing, with $x(t) = 0$ for all $t <= t_0$, and $x(t) <= 1$ for all t. The innovation diffusion model [11] is described by the ordinary differential equation with the unique solution:

$$x(t) = K\frac{1-e^{-(p+q)(t-t_0)}}{1+\frac{q}{p}e^{-(p+q)(t-t_0)}}, t \geq t_0 \; and \; x(t) = 0 \; elsewhere.$$

The innovation diffusion model depends on three key parameters:

- $p > 0$: the adoption rate, i.e., the probability that an unvaccinated individual gets vaccinated independently of others.
- $q > 0$: the coefficient of imitation, the extent to which the probability of getting vaccinated increases with the proportion of already vaccinated individuals.
- $K \in (0, 1]$: the carrying capacity, representing the maximum share of the population expected to eventually adopt the vaccine.

The parameters p, q, t_0, and K are estimated via least-squares fitting to empirical data on COVID-19 vaccine uptake[5].

The fitting is performed for each commune and arrondissement in the Métropole du Grand Paris, using the 28 weeks leading up to the announcement of the French COVID-19 health pass (Week 28 of 2021, i.e., July 12, 2021).

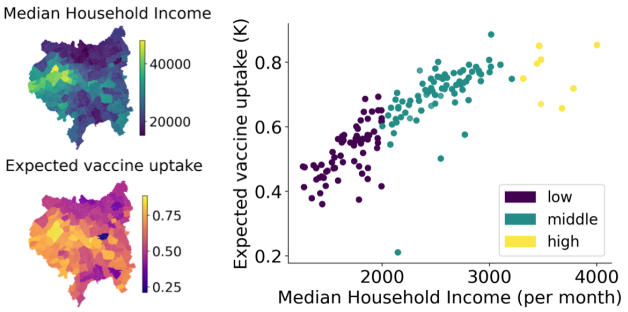

Fig. 3. (left) Spatial distribution of median household income and K, the vaccine uptake (i.e., maximum share of the population expected to eventually adopt the COVID-19 vaccine) in the Métropole du Grand Paris. (right) is the correlation between the median household per month and K in each commune (Pearson correlation = 0.817, p-value of 0.0)

References

1. Baker, Rachel E., et al. "Infectious disease in an era of global change". *Nature Reviews Microbiology* 20.4 (2022): 193-205.
2. Lewis, Nicola, and Martin Beer. "Stop H5N1 influenza in US cattle now." *Science* 385.6705 (2024): 123-123.
3. Caserta, Leonardo C., et al. "Spillover of highly pathogenic avian influenza H5N1 virus to dairy cattle." *Nature* 634.8034 (2024): 669-676.
4. Garg, Shikha, et al. "Highly pathogenic avian influenza A (H5N1) virus infections in humans." *New England Journal of Medicine* 392.9 (2025): 843-854.
5. Goodman, Jesse L., et al. "Prepare now for a potential H5N1 pandemic." *Science* 387.6738 (2025): 1047-1047.
6. Bartlett, Maggie L., et al. "Enhancing the response to avian influenza in the US and globally." *The Lancet Regional Health–Americas* 46 (2025).
7. Larsen, Sophie L., et al. "Quantifying the impact of SARS-CoV-2 temporal vaccination trends and disparities on disease control." *Science Advances* 9.31 (2023): eadh9920.
8. Vandentorren, Stéphanie, et al. "The effect of social deprivation on the dynamic of SARS-CoV-2 infection in France: a population-based analysis." *The Lancet Public Health* 7.3 (2022): e240-e249.
9. Kostandova, Natalya, et al. "Improving mobility data for infectious disease research." *Nature Human Behaviour* (2025): 1-4.
10. Melegaro, Alessia et al. "IMMUNE Behavioural Study." OSF, 9 July 2025. Web.
11. Oliu-Barton, Miquel, et al. "The effect of COVID certificates on vaccine uptake, health outcomes, and the economy," Nature Communications, 13.1 (2022): 3942.

[1] https://www.data.gouv.fr/fr/datasets/communes-et-villes-de-france-en-csv-excel-json-parquet-et-feather/

[2] https://catalogue-donnees.insee.fr/en/catalogue/recherche/DS_BPE

[3] www.insee.fr/fr/statistiques/5359146

[4] https://www.data.gouv.fr/fr/datasets/caracteristiques-socio-economiques-des-communes-dile-de-france/reuses_and_dataservices/

[5] https://datavaccin-covid.ameli.fr/explore/dataset/donnees-de-vaccination-par-indice-de-defavorisation-par-tranche-dage/table/

Commuting Patterns Across Household Structures: An Empirical Analysis of Time and Mode Choices in Île-de-France

XUECHEN WANG[1]

[1] Institute of Remote Sensing and Geographical Information Systems, Peking University, Beijing, China
[*] cherylwang0519@gmail.com

1. INTRODUCTION

Commuting is a vital aspect of mobility, closely linked to the labor market, urban structure, and phenomena like epidemic spread[1]. It is influenced by individual factors such as income, housing prices, and travel costs, with models like the gravity [2] and radiation models [3] providing quantitative insights. Additionally, interpersonal and household dynamics [4, 5], including evolving gender roles, increased female workforce participation, and family-related policies, play a significant role [5–7]. These dynamics shape how individuals balance work and domestic responsibilities, affecting commute lengths and job flexibility. While research has examined commuting in dual-earner households and multigenerational families, further empirical studies are needed to better understand how household structure impacts residence and workplace choices.

In this study, we analyzed the commuting patterns of 2,291 Île-de-France residents to examine how household composition influences travel behavior. Key findings include: (1) individuals living with parents have the longest commutes, followed by those who are married or have children; (2) parents favor private cars, while childless individuals rely more on public transportation; and (3) couples with adult children have longer commutes, despite minor differences between households with children under 18 and those without children. These results imply the significant role of childcare responsibilities and life-stage transitions in shaping commuting behavior.

2. DATA AND PREPROCESSING

We use the NetMob25 dataset [8], which tracks the daily mobility of 3,337 individuals aged 16–80 over one week (mid-October 2022 to mid-May), comprising: (i) demographic and household data, (ii) over 80,000 annotated trips with metadata (times, modes, purposes), and (iii) 500 million GPS location points.

For our analysis, we used records from October 17, 2022, to May 31, 2023, filtering trips where the origin was *DEPART_HOME* and destination *WORK*. Trips were included only if the home location matched the Individuals database. For each individual, we calculated average commuting time and identified their primary travel mode. This preprocessing provided commuting data for 2,291 individuals (1,257 women, 1,034 men).

3. RESULTS

A. Commuting time and mode choice for different house types

We begin by examining the characteristics of commuting time and mode. Figure 1 displays the distribution of commuting times across various household types. Individuals *Living with one or both parents* exhibit a wider distribution of longer commuting times, while the other four household types generally experience shorter and more similar commuting durations. In other words, individuals living with their parents tend to have the longest commutes, followed by those who are married, have children, or live alone. Across all household types, the most common commuting modes are *TRAIN* and *PRIV_CAR_DRIVER*, though preferences vary (Figure 2). Households without children favor train travel over private cars, with significantly less car use overall. Individuals *Living alone* prefer *SUBWAY*, while those *Living with parents* rely more on *TRAIN*. In contrast, families with children show reduced public transport use, with private cars becoming the dominant mode.

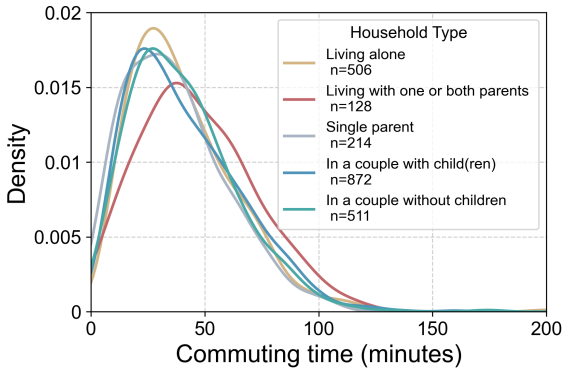

Fig. 1. Commuting time distribution for each household type.

B. Influence of children under 18 on commuting times

Figure 1 reveals minor differences in commuting times between couples with and without children, with childless households having slightly longer commutes. This raises the question: if children have little impact, what other factors influence commuting duration?

To account for varying childcare needs by age, we divided the group *In a couple with child(ren)* into two subgroups: those with at least one child under 18 and those with only adult children. For both men and women, commuting times are similar for couples with children under 18 and without children (Figure 3). However, couples with only adult children show longer commutes, with a slightly bimodal distribution (dashed lines). For men, the peak is around 60 minutes (Figure 3(a)), while for women, it is around 25 minutes (Figure 3(b)).

A possible explanation for the longer commuting times observed among couples with only adult children is the completion of childcare responsibilities. The presence of children often encourages parents to minimize their commuting times, a factor that may even influence decisions made at the stage of forming a partnership. Once childcare responsibilities are no longer a concern, individuals may become more flexible with longer commuting times due to the reduced constraints on their schedules. However, these empirical findings require further validation with larger sample sizes to ensure robustness.

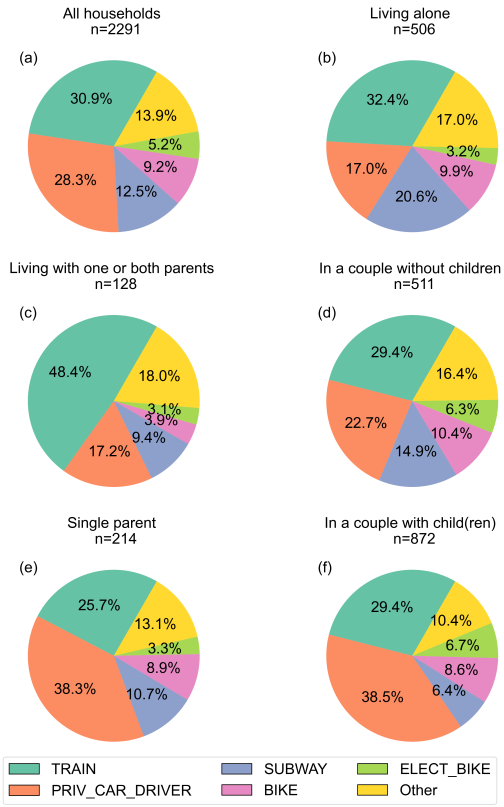

Fig. 2. Main mode composition by household type.

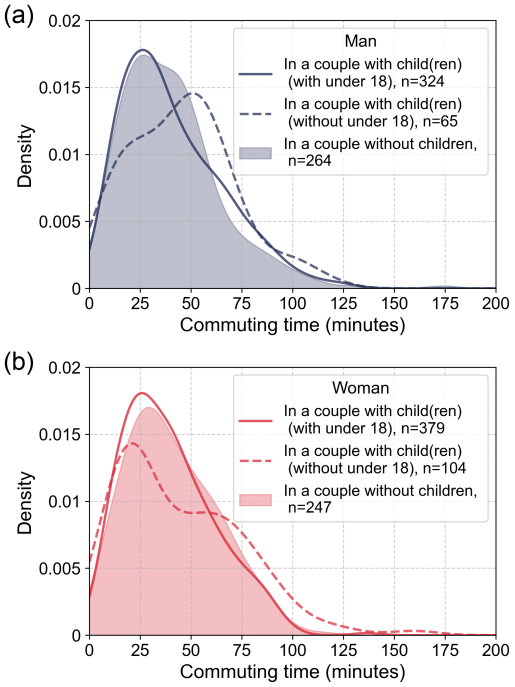

Fig. 3. Comparison of commuting time by sex for individuals in the following groups: *In a couple with child(ren) (with under 18)*, *In a couple with child(ren) (without under 18)* (adult children only) and *In a couple without children*. (a) Man; (b) Woman.

We also examine the impact of childcare responsibilities on individuals across three socio-professional categories with the largest sample sizes. Figure 4 shows that individuals in the *Employee* and *Executives and Senior Intellectual Professionals* categories tend to have longer commuting times. Notably, *Employees* display a bimodal commuting time distribution for couples with only adult children (dashed lines). However, further data collection is needed due to sample size limitations.

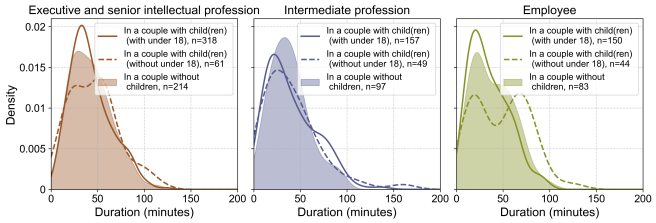

Fig. 4. Comparison of commuting time by socio-professional statuses for household members in the following groups: *In a couple with child(ren) (with under 18)*, *In a couple with child(ren) (without under 18)* and *In a couple without children*. Results are shown for *Executive and senior intellectual profession*, *Intermediate profession* and *Employee*.

4. DISCUSSION AND CONCLUSIONS

We analyzed the commuting patterns of 2,291 Île-de-France residents to understand how household composition affects travel behavior. Individuals living with their parents have the longest commutes and rely heavily on trains, likely due to limited housing flexibility. In contrast, married individuals and parents generally have shorter commutes, with parents favoring private cars while childless individuals rely more on public transportation. Notably, households with adult children experience longer commutes compared to those with children under 18 or no children.

These findings suggest that residential and workplace choices evolve with life stages: family formation encourages shorter commutes, while post-childrearing allows for greater location flexibility. Future research should confirm these trends using larger longitudinal data, focusing on how households balance work-family needs and how job type and income influence location choices across life transitions.

REFERENCES

1. A. S. Fotheringham and M. E. O'Kelly, *Spatial Interaction Models: Formulations and Applications* (Kluwer Academic Publishers, Dordrecht, Netherlands, 1989).
2. A. Wilson, "Entropy in urban and regional modelling: Retrospect and prospect," Geogr. Anal. **42**, 364 (2010).
3. F. Simini, M. C. González, A. Maritan, and A.-L. Barabási, "A universal model for mobility and migration patterns," Nature **484**, 96 (2012).
4. M. J. White, "A model of residential location choice and commuting by men and women workers," J. Reg. Sci. **17**, 41 (1977).
5. J. Feng, M. Dijst, B. Wissink, and J. Prillwitz, "The impacts of household structure on the travel behaviour of seniors and young parents in china," J. Transp. Geogr. **30**, 117 (2013).
6. S. Williams, W. Qiu, Z. Al-awwad, and A. Alfayez, "Commuting for women in saudi arabia: Metro to driving - options to support women employment," J. Transp. Geogr. **77**, 126 (2019).
7. S. H. Hong, B. S. Lee, and J. F. McDonald, "Commuting time decisions for two-worker households in korea," Reg. Sci. Urban Econ. **69**, 122 (2018).
8. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," (2025).

Post-Disaster Route Reliability Assessment Using Drone-Aided Mobility Data from NetMob25

MAHMOUD ERFAN[1] AND NOHA GAMAL[5]

[1] Mechatronics Engineering Program , Higher Technological Institute,10th of Ramadan,Egypt.
[2] School of Information Technology and Computer Science, Nile University, Giza, Egypt.
[*] mahmouderfan46@gmail.com, ngamal@nu.edu.eg

1. INTRODUCTION

Urban transport systems are vital for daily life, but they are also highly vulnerable to disruptions. In regions like Île-de-France, even moderate events such as transport strikes or floods can block commuting, slow logistics, and reduce access to essential services. Previous research has shown that strikes often push people to shift from public transport to walking or private cars, which increases congestion [1]. Other studies in network science highlight that Paris mobility depends heavily on a few central hubs, making the system fragile if those hubs are disrupted [2, 3].

Most existing studies, however, look at either behavior or network structure separately. What is missing is an integrated approach that connects mobility patterns, physical hazards, and practical response strategies.

This study fills that gap by developing a multi-phase framework for resilience assessment. We combine high-resolution GNSS mobility data from the NetMob25 dataset [4] with official flood hazard maps (DRIEAT Île-de-France), and test response strategies in a 3D simulation environment (Unreal Engine + AirSim + PX4). The workflow allows us to: (i) detect critical but fragile mobility zones, (ii) assess how floods affect these areas, and (iii) evaluate UAV-based routing strategies for emergency logistics.

By combining behavioral, structural, and hazard-aware perspectives, this approach gives a more complete view of mobility resilience and offers practical tools for improving preparedness and response.

2. METHODOLOGY

We developed a four-phase framework that combines mobility data, hazard information, and simulation tools to assess resilience in Île-de-France (Figure 1).

1. Data Sources

We used the *NetMob25 mobility dataset* [4], which includes over 80,000 trips and 500M GNSS points from 3,337 residents (2022–2023). It provides detailed trip records (origins, destinations, modes, and day types) and anonymized user profiles, enabling both behavioral and structural analysis of urban mobility.

2. Data Preprocessing

Invalid or incomplete trips were removed, and flows were aggregated into an origin–destination (OD) matrix. From this, we built a directed weighted graph and computed two key metrics:

- **Betweenness centrality:** to identify critical nodes that act as network bottlenecks.
- **Redundancy scores:** to detect fragile links with few or no alternatives [2, 5].

3. Hazard Integration

Official flood hazard polygons were intersected with OD nodes. Each zone was classified as *affected* (inside/near flood areas) or *non-affected*. KMeans clustering of exposed trips revealed disruption hotspots, which served as candidate waypoints for UAV-based interventions.

4. Simulation Testing

We implemented a 3D testbed using **Unreal Engine**, **AirSim** [6], and **PX4** [7]. Drone routes were simulated using three strategies:

- **TSP optimization** (shortest total route),
- **Greedy heuristic** (nearest neighbor),
- **Random routing** (baseline).

This setup tested how UAVs could support emergency logistics in flood-affected areas.

Fig. 1. Workflow combining mobility data, hazard layers, and UAV simulation.

3. RESULTS

A. Behavioral Adaptation

The NetMob25 data showed strong behavioral shifts during disruptions. On strike days, average trip durations dropped by nearly 30%, and suburban travel demand collapsed. As Figure 2 shows, *public transport* use (metro, train, bus) almost disappeared, while *private cars* and *walking* became dominant. This substitution reflects short-term resilience but also highlights systemic fragility, since increased car use leads to congestion.

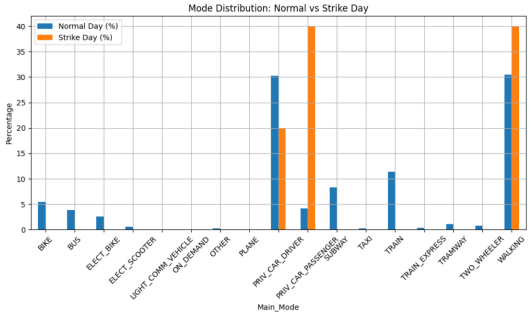

Fig. 2. Modal shifts during strike days. Public transport collapses, while car use and walking increase sharply.

B. Structural Vulnerabilities

Origin–destination graph analysis revealed strong hub dependency. PARIS-12E, PARIS-15E, and PARIS-17E concentrated most shortest paths, acting as bottlenecks. Redundancy analysis further showed that many OD pairs lacked alternative routes, making the network fragile to disruptions (Figure 3).

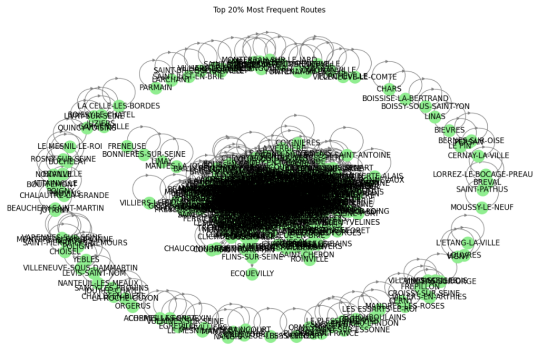

Fig. 3. Top 20% most frequent OD routes in Île-de-France. Connectivity is concentrated in central Paris, creating bottlenecks.

C. Compound Hazard Exposure and Simulation

Overlaying flood hazard maps revealed that several high-centrality nodes lie inside flood-prone areas, amplifying vulnerability (Figure 4). KMeans clustering identified disruption hotspots, which were used as drone waypoints. In simulation tests, TSP-based routing reduced UAV travel distance by up to 37% compared to random baselines, showing the value of optimization for emergency logistics.

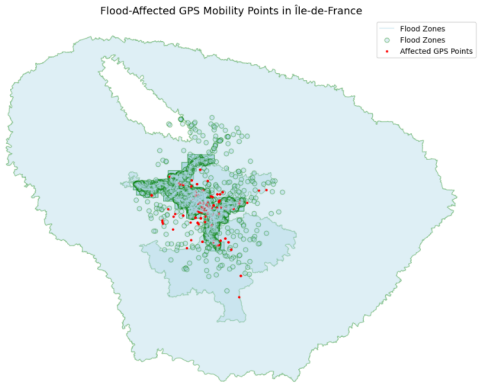

Fig. 4. Flood polygons overlaid with mobility nodes. Red markers show flood-exposed hubs.

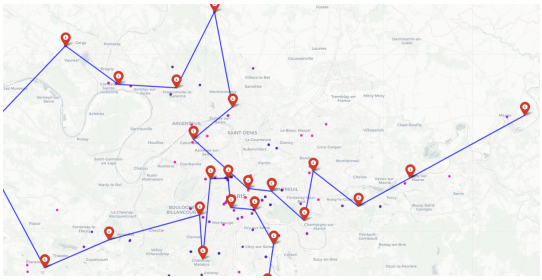

Fig. 5. Optimized UAV routing across flood-affected areas. Red points = waypoints, blue line = TSP path.

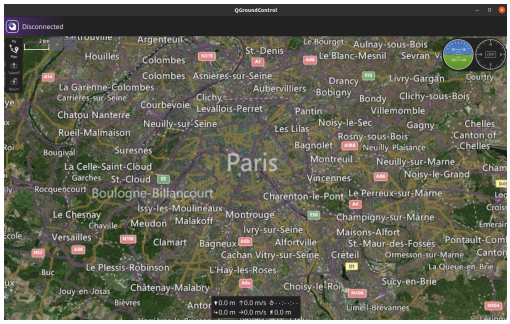

Fig. 6. QGroundControl interface showing UAV mission planning during simulation.

Key Takeaway

Behavioral adaptation helps during crises, but structural fragility and hazard exposure create systemic risks. Integrating hazard data with simulations provides actionable insights for improving urban resilience.

4. DISCUSSION

Our results highlight three key lessons for mobility resilience in Île-de-France.

First, individuals adapt during disruptions—mainly by shifting to cars and walking—but this comes at the cost of congestion and reduced accessibility. Behavioral flexibility alone is not enough to sustain system performance [1].

Second, the mobility network is highly centralized. A few arrondissements (e.g., PARIS-12E, PARIS-15E) dominate connectivity, while many OD pairs lack redundancy. This hub dependency makes the system fragile, as failures in central nodes can trigger region-wide impacts [2, 8].

Third, hazard overlays showed that some of these same hubs are also flood-prone, compounding structural and environmental risks [9, 10].

Finally, UAV simulations demonstrated that optimized routing strategies (TSP-based) can reduce response distances and improve efficiency, offering practical tools for emergency logistics [6, 11].

In sum, resilience requires more than individual adaptation. It calls for diversifying transport investments, strengthening redundancy, and embedding hazard-aware and simulation-based planning into mobility policy.

Funding. This research was supported by Nile University.

Acknowledgments. The authors thank the NetMob community for providing access to high-resolution mobility datasets and methodological guidance.

Disclosures. The authors declare no conflicts of interest.

Data Availability Statement. Data underlying the results presented in this paper are available through the NetMob25 challenge dataset, subject to the terms and conditions of the challenge organizers.

REFERENCES

1. O. Cats, Q. Wang, and Y. Zhao, "Impacts of recurrent disruptions on public transport demand and urban resilience," Transp. Res. Part A: Policy Pract. **96**, 35–47 (2016).
2. M. Barthélemy, "Spatial networks," Phys. Reports **499**, 1–101 (2011).
3. S. Porta, P. Crucitti, and V. Latora, "The network analysis of urban streets: A primal approach," in *Environment and Planning B: Planning and Design*, vol. 33 (SAGE, 2006), pp. 705–725.
4. L. Chassé and et al., "The netmob25 dataset: High-resolution gnss-based mobility survey for Île-de-france," in *NetMob 2025: International Conference on Computational Social Science*, (2025).
5. M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, "Understanding individual human mobility patterns," Nature **453**, 779–782 (2008).
6. S. Shah, D. Dey, C. Lovett, and A. Kapoor, "Airsim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and Service Robotics (FSR)*, (Springer, 2018), pp. 621–635.
7. L. Meier, D. Honegger, and M. Pollefeys, "Px4: An open source autopilot for research and industry," in *IEEE International Conference on Robotics and Automation (ICRA) Workshop*, (2021). Available at: https://px4.io.
8. T. Ji and et al., "Urban transport resilience to extreme events: A review," Sustainability **14**, 3880 (2022).
9. B. Barroca, P. Bernardara, J.-M. Mouchel, and G. Hubert, "Floods as a threat for urban resilience: lessons from the seine river in paris," Nat. Hazards Earth Syst. Sci. **13**, 3235–3249 (2013).
10. P. Cui, Z. Wang, H. E. Stanley, and S. Havlin, "Interdependent network resilience: Modeling, analysis, and applications," Proc. National Acad. Sci. **117**, 29569–29579 (2020).
11. A. Otto, N. Agatz, J. Campbell, *et al.*, "Optimization approaches for civil drone operations in disaster scenarios," in *Computers & Operations Research*, vol. 98 (Elsevier, 2018), pp. 134–147.

Main Conference

posters

Modeling Base Station Metadata Geolocation

ORLANDO E. MARTÍNEZ-DURIVE[1,2,*], STEFANOS BAKIRTZIS[3], CEZARY ZIEMLICKI[4], AND MARCO FIORE[1]

[1] *IMDEA Networks Institute, Madrid, Spain*
[2] *Universidad Carlos III de Madrid, Madrid, Spain*
[3] *Cambridge University, Cambridge, United Kingdom*
[4] *SENSE / Orange Research, Paris, France*
[*] *orlando.martinez@imdea.org*

1. BACKGROUND AND MOTIVATION

Mobile networks, together with smartphones, are among the most ubiquitous technologies in the modern world. This ubiquity allows mobile networks to act as large-scale passive sensors, continuously capturing vast amounts of data. Specifically, the metadata collected provides rich information about the movement, communication patterns, activities, and interests of large populations of subscribers, with high spatiotemporal resolution and at broad geographical scales. Therefore, such metadata has been used in various scientific fields, including networking, demography, geography, sociology, human mobility, and epidemiology, enabling dependable quantitative analyses of phenomena that can only be studied qualitatively otherwise. In all these studies, the metadata is usually geo-referenced at the level of the individual Base Station (BS). In other words, all metadata records (*e.g.*, the traffic generated by a device or the total demand of all devices for a given service) are associated with the geographical site of BS that serves the device at each time instant. Yet, network analyses building on such metadata need to be carried out over the continuous space of the target region and not at the discrete BS locations only. A clear example is when mobile metadata is related to other sources of data, such as socioeconomic indicators, where administrative areas are used as a spatial dimension; therefore, some transformation is needed to relate both spatial units. This makes the mapping of BS-referenced data to the underlying geographical space an essential step in any mobile network metadata processing pipeline and calls for *spatial diffusion* models that describe *the probability that a device attached to a BS is at a specific geographical location served by that BS*.

Accurate spatial diffusion models can be directly obtained from client-side measurements or estimated by post-processing the results of radio propagation solvers. However, these approaches are not viable in academic research, as they require massive measurement campaigns or access to confidential information about the network infrastructure that is typically not available. As a result, the vast majority of the academic literature has relied on simplistic approximations of spatial diffusion based on Voronoi tessellations of the space, assuming that the metadata associated with one BS is uniformly diffused over the surface of its Voronoi cell, as illustrated in Figure 1b. Evidently, this approximation poses significant performance limitations and yields substantial discrepancies when compared to ground truth data, as shown in Fig. 1a; yet, many works in the literature rely on this model due to its simplicity. Indeed, computing a Voronoi tessellation is a straightforward operation that only requires the geographical coordinates of the BSs, which are typically provided to academic researchers along with the metadata.

An alternative is a probabilistic density model that estimates spatial diffusion via probability density inference [2]. That allows partially overcoming the limitations of Voronoi cells by generating a statistical representation of the location of an User Equipment (UE) given solely its distance from the BS of attachment, computing an approximate probabilistic spatial diffusion of the BS. This model results in a smoothed Voronoi cell, controlled by a power attenuation parameter that needs to be properly configured through field measurements. However, such models need fine-tuning via measurement data collected in the target network and pre-request access to confidential information, such as the BSs transmitting power, the signal-to-noise and interference ratio (SINR) thresholds, interference from non-neighboring base stations. An in-depth discussion on alternative approaches can be found in [1].

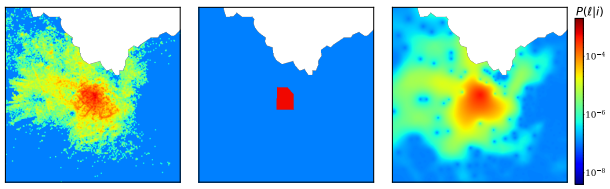

(a) Ground truth **(b)** Voronoi **(c)** DEEPMEND

Fig. 1. Spatial diffusion of one real-world BS, obtained from (a) elaborate post-processing of radio propagation simulations by a network operator based on complete RAN information, (b) Voronoi tessellation, (c) our proposed DEEPMEND. Colors denote diffusion probabilities over space. From [1].

2. PROPOSED SOLUTION

Our proposal, DEEPMEND [1] aspires to change the status quo and democratize dependable spatial diffusion by proposing a model that can combine the minimal input requirement of a Voronoi tessellation (*i.e.*, the sole locations of BSs, making the model accessible to any academic researcher) with the high accuracy granted by a full-fledged modeling process based on operator-proprietary information.

To this end, we use a rich set of spatial diffusion maps from almost 6000 outdoor BSs located in urban, suburban, and rural areas of France, making up the radio access network of Orange, a major local operator. The spatial diffusion dataset is provided directly by the network operator, and it is derived from realistic distributions of the *association probability*, $P(\ell|i)$, which describe the probability that a UE attached to the i-th BS is located at position ℓ. For each BS, the probability of association over a $3,600$ km^2 area is in the form of a 600×600 regular grid, with the BS located at the center of the grid and each grid element ℓ matching a 100×100 m^2 physical area.

These voluminous data are used to develop DEEPMEND, a deep learning-based spatial diffusion model [1] that exploits the concept of *knowledge distillation* [3]. Specifically, knowledge distillation involves a complex and computationally expensive model, the *teacher*, to derive soft labels that are then used to train a simpler model, the *student*. In our case, we consider as the teacher the model of [2], with its hyperparameter optimally selected via an exhaustive search, whilst the student is a deep convolutional encoder-decoder. The general concept of DEEPMEND is shown in Figure 2. As can be seen, the teacher model processes the

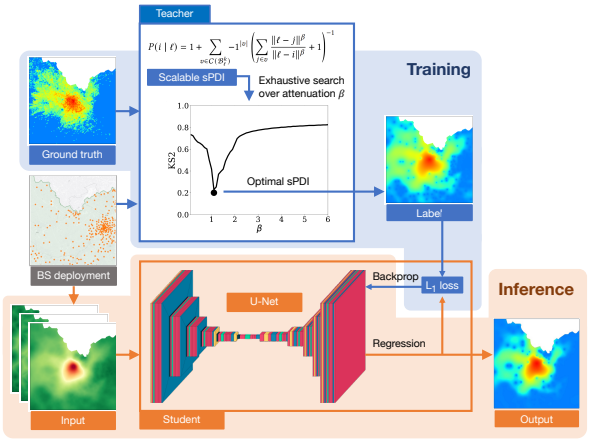

Fig. 2. Diagram of the teacher-student workflow of DEEPMEND, with inputs, outputs, and interactions. From [1].

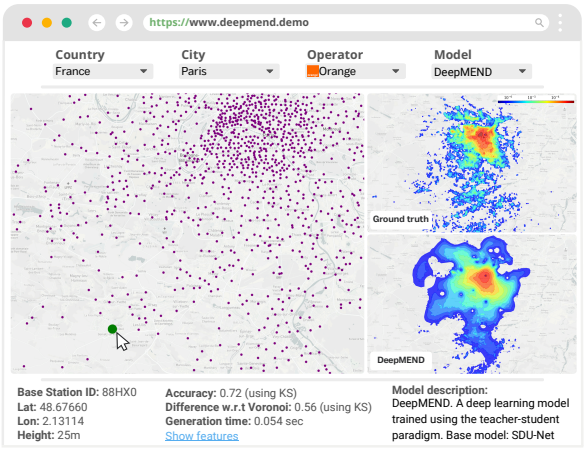

Fig. 3. Web app demonstrator of DEEPMEND.

BS locations and uses an exhaustive search to determine the hyperparameter of [2] which minimizes the Multidimensional Kolmogorov–Smirnov (KS2) test [4] error between the ground truth and the emerging model-based [2] spatial diffusion distribution. That corresponds to a *soft label*, *i.e.*, a high-fidelity representation of the ground truth as shown in Fig. 1a, which is thought to be easier to learned by the student than the ground truth itself. Finally, the student model uses these soft labels to understand how to transform simple BS location information into realistic spatial diffusion maps for each BS. Ultimately, the rationale for the design of DEEPMEND is that the weak correlations between the tangled ground-truth and the modest input represented by BS locations are difficult to learn directly. By introducing a teacher model, we generate soft labels that are close to the target ground truth yet more explicitly correlated to the BS positioning data, thus simplifying the learning process.

We assess the performance of DEEPMEND using a standard 5-fold cross-validation approach. In each fold, the dataset of BSs is split into training and test sets using an 80:20 ratio, a common practice in machine learning tasks. This setup enables a robust evaluation of our model's performance compared to existing geolocation methods. Specifically, DEEPMEND reduces the median error in modeling geolocation by 56.80%, 33.03%, and 39.27% with respect to Voronoi, data-driven scaling of Voronoi [5], and probabilistic density models [2], respectively. We refer the reader to the original work [1] for comprehensive results and discussions on performance, generalization, and scalability.

3. DEMONSTRATOR

DEEPMEND is implemented in Python and openly available on GitHub[1], allowing the research community to easily adopt, evaluate, and build upon our model for improved metadata geolocation tasks. However, to illustrate the improvements in metadata geolocation enabled by DEEPMEND and demonstrate its ease of integration and use, we have developed an interactive demo that we plan to present at the conference. Our demonstrator is a web app that allows the user to visually and quantitatively assess the performance of DEEPMEND among other geolocation map approaches. The dashboard of our demonstrator is shown in Fig. 3. The web app builds on *React.js* for the User Interface, *Mapbox* for map visualizations, and *Python* for data load and inference tasks.

Users can select the country, city, and mobile operator of interest using the menu at the top of the interface. Currently, the demo supports only France, offering a range of cities, including major urban areas such as Paris, Lyon, and Nantes, and Orange as the mobile operator. We are actively working to support additional countries, including the United Kingdom, Ireland, and Belgium, as well as other operators such as Proximus, Vodafone, Three, Free, and SFR. In the same menu, it is possible to select the model to infer geolocation maps: this list includes our proposed solution DEEPMEND together with other benchmarks.

After selecting the country, city, and operator, the app shows the corresponding base station deployment in the left panel. Then, the user can zoom in and out or move over the map; details will be displayed in a pop-up while the user hovers over a base station. When a base station is selected, the ground-truth geolocation map (if available) is shown in the top right panel. Based on the chosen model, the app requests the inferred geolocation map using the Python API, which is displayed in the bottom right panel. Both geolocation maps display the spatial probability density of the user equipment located over the space; a colorbar legend showing the probability density values is displayed on the top right panel.

The bottom panel displays relevant information about the selected base station, including latitude, longitude, and height. Furthermore, it also displays information about the selected model, including inference time, accuracy with respect to the ground truth (if available), and error metrics compared to the corresponding Voronoi cell and other benchmarks. There is also an option to show the features fed to the model.

[1] https://github.com/nds-group/DeepMEND

ACKNOWLEDGMENT

This work was supported by the ORIGAMI project (Grant no. 101139270) funded by SNS JU and the European Union; CoCo5G, grant no. ANR-22-CE25-0016, funded by the French National Research Agency, and 6G-IRONWARE funded by MICIU/AEI /10.13039/501100011033 and by the European Union NextGenerationEU/PRTR CNS2023-143870.

REFERENCES

1. O. E. Martínez-Durive *et al.*, "Deepmend: Reliable and scalable network metadata geolocation from base station positions," in *IEEE Conference on Sensing, Communication, and Networking (SECON)*, (2024).
2. V. A. Traag *et al.*, "Social event detection in massive mobile phone data using probabilistic location inference," in *Third International Conference on Privacy, Security, Risk and Trust*, (2011).
3. G. Hinton *et al.*, "Distilling the knowledge in a neural network," in *NIPS Deep Learning and Representation Learning Workshop*, (2015).
4. G. Fasano *et al.*, Mon. Notices Royal Astron. Soc. (1987).
5. O. E. Martínez-Durive *et al.*, "Voronoiboost: Data-driven probabilistic spatial mapping of mobile network metadata," in *IEEE Conference on Sensing, Communication, and Networking (SECON)*, (2022).

Cross-City Travel Mode Identification from Sparse Mobile Phone Data with Transformer Adapters

JUNYUAN LIU[1], JIAZHUANG FENG[1], AND TAO CHENG[1,*]

[1] SpaceTimeLab,University College London, Gower Street, London WC1E 6BT, United Kingdom
[*] tao.cheng@ucl.ac.uk

1. INTRODUCTION

Travel mode identification from mobile phone data is essential for urban mobility analysis and intelligent transportation systems. However, existing methods typically rely on densely sampled trajectories, which are not always available in practice. Mobile phone data is often sparse, irregular, and noisy, making conventional models difficult to apply without heavy preprocessing. In addition, models trained in one city often fail to generalize to new cities due to differences in geography, behavior, and data collection. Labeling new data is expensive, so it is crucial to develop models that can be transferred with minimal supervision.

In this paper, we propose a simple and effective framework that addresses both challenges. First, we introduce a daily input construction method, where data is structured into fixed-length sequences per user per day to help the Transformer model observe full-day mobility behavior and eliminate the need for trip segmentation or dense sampling. Second, we incorporate lightweight adapter modules into the Transformer encoder. During transfer, only the adapters and classifier are updated, keeping the rest of the model frozen. This design enables parameter-efficient adaptation to new cities with limited supervision.

We pretrain our model on the large-scale GeoLife dataset from Beijing [1] and transfer it to the smaller London dataset [2] via adapter tuning. Our approach achieves 76.0% accuracy on the London test set, significantly outperforming training from scratch (64.3%), demonstrating strong cross-city transferability.

In summary, our contributions are threefold: (i) a daily sequence structuring method enabling Transformer-based modeling of full-day mobility, (ii) a parameter-efficient adapter approach for rapid cross-city adaptation with limited supervision, and (iii) validation on Beijing and London datasets demonstrating strong few-shot performance.

2. METHODOLOGY

A. Framework Overview

Our method performs travel mode identification on sequence data using a Transformer-based encoder combined with adapter-based transfer learning. The framework is designed to handle sparse, real-world mobile phone data and to adapt across cities with minimal labeled data in the target domain. The pipeline includes: (1) structuring raw data into daily sequences, (2) extracting contextual and temporal features, (3) modeling sequences with a Transformer, and (4) applying adapters for efficient transfer.

B. Daily Sequence Construction

To mimic real-world sparse mobile data, we sample trajectories into daily sequences with 1440 time steps—one per minute. For each minute, a representative point is sampled if available; otherwise, the entry is left empty and masked during attention. This structure preserves full-day behavioral patterns while avoiding the need for trip segmentation or imputation.

C. Feature Extraction

At each time step, we extract a compact set of features capturing motion dynamics, spatial context, and temporal patterns following [3]:

- **Motion features:** distance difference (Δd), instantaneous speed (v), acceleration (a), and angular velocity (ω), defined as the rate of change in heading angle over time.
- **Public transit context:** N_i, the number of nearby bus or train stations within a fixed radius; and D_i, the distance to the nearest station.
- **Temporal features:** hour-of-day and minute-of-day are encoded as sine and cosine pairs to represent daily periodicity, helping the model distinguish between similar behaviors occurring at different times (e.g., commuting vs. leisure walking).

All features are standardized using statistics from the training set and concatenated into a feature vector at each time step.

D. Model Architecture

We use a Transformer-based multi-task model to predict travel modes and transition dynamics from constructed daily sequences. Input features are first projected into a latent space using a linear layer, followed by batch normalization. Positional encodings are then added to incorporate temporal order. The model consists of stacked Transformer encoder layers, each comprising a multi-head self-attention mechanism and a feedforward sublayer. The model includes two output heads:

- **Mode classification head:** a linear layer that predicts the travel mode at each time step.
- **Transition prediction head:** a linear layer that estimates the likelihood of mode transitions between consecutive time steps.

E. Adapter-Based Transfer Learning

To enable cross-city adaptation with minimal labeled data, after each encoder layer, a lightweight city-specific adapter module is inserted. Each adapter consists of a linear layer, a ReLU activation, and dropout, enabling efficient domain adaptation across cities. These adapters introduce a limited number of additional parameters—specifically, $L \times (d^2 + d)$ parameters, where L is the number of encoder layers and d is the model dimension. During transfer to the target city, only the adapters and classification head are updated, while the pretrained Transformer weights remain frozen.

3. EXPERIMENTS

A. Datasets

We evaluate our method on two datasets:

- **GeoLife (Beijing):** A widely-used public dataset collected by Microsoft Research Asia [1] in Beijing.
- **London Dataset:** A sparse real-world dataset collected in London [2].

Each dataset is preprocessed into daily sequences with up to 1440 one-minute steps per user. The London dataset is used for evaluating domain transfer performance.

B. Experimental Setup

We use a 3-layer Transformer encoder with 8 heads and a hidden size of 128. Input features are projected and combined with positional encodings. Dropout is set to 0.3. The model is trained using AdamW (learning rate 5e-4, weight decay 1e-3) with a batch size of 32.

C. Cross-City Transfer Results

We compare the following models on the same training set with 50 samples and the same test set:

- **Scratch:** Model trained from scratch on London data.
- **Zero-shot:** Model pretrained on Beijing, tested directly on London data.
- **Full:** Pretrained on Beijing, then fully fine-tuned on London.
- **Last-layer:** Only the last encoder layer is updated.
- **Add-layer:** An additional encoder layer is added and trained.
- **Adapter tuning:** Only lightweight adapters are updated.

Table 1 summarizes the model performances on the London test set. We observe that adapter-based fine-tuning outperforms all other methods, with over 10% absolute gain in accuracy compared to training from scratch in London. Full fine-tuning and last-layer tuning show marginal improvements over the source model, but do not reach the same level of generalization.

Table 1. Cross-city transfer performance on London test set

Model	Precision	Recall	F1
Scratch	0.736	0.643	0.637
Zero-shot	0.623	0.563	0.559
Full fine-tune	0.697	0.658	0.649
Last-layer tune	0.666	0.603	0.574
Add-layer tune	0.698	0.643	0.641
Adapter	**0.765**	**0.760**	**0.758**

D. Few-shot Adaptation with Adapters

We further test how adapter-based transfer performs under low-resource settings by varying the number of fine-tuning samples from the London dataset in Figure 1.

4. DISCUSSION

The experiments reveal several key insights regarding cross-city travel mode identification from sparse mobile phone data. First, the poor performance of the zero-shot model highlights the gap between cities, where direct transfer without adaptation leads to severe degradation. This validates the need for transfer learning.

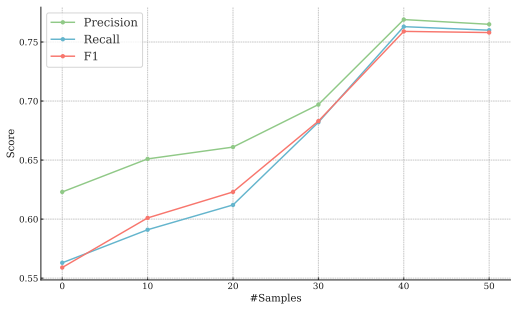

Fig. 1. Few-shot Adapter Fine-tuning with Different Samples

Second, Transformer models typically require large-scale datasets for effective training, especially when constructing daily data as input. In our setup, the Beijing dataset provides a larger pretraining source. However, the London dataset is small and sparse, conditions under which training from scratch performs poorly. Adapter-based fine-tuning enables effective transfer in this large-source, small-target regime by reusing pretrained knowledge and updating only a minimal number of parameters.

Third, while full fine-tuning and partial updates (e.g., last-layer tuning) provide moderate improvements, they remain sub-optimal in both accuracy and efficiency. In contrast, the adapter-based approach achieves the highest accuracy (76.0%) while being computationally light, learning well on the target area.

Finally, our few-shot experiments confirm the practical usability of the method. As the number of labeled daily sequences increases, the adapter-based model improves rapidly. In particular, once 40 or more labeled samples are used for fine-tuning, the model achieves a much better performance compared with training from scratch. This suggests that adapter tuning significantly reduces the annotation burden for real-world deployment, enabling transfer across cities. Interestingly, while performance generally improves with more labeled samples, we observe a slight dip in accuracy at 50 samples. This slight dip could result from overfitting or noise in the additional samples, highlighting that adapter tuning remains sensitive to data quality in low-resource settings.

5. CONCLUSION

We propose an effective framework for cross-city travel mode identification using Transformer encoders with adapter-based fine-tuning. By structuring data as daily sequences and learning over full-day patterns, our model achieves high performance on limited, sparse mobile trajectories in the target city.

Our results demonstrate that the adapter modules enable efficient and accurate domain transfer, outperforming other fine-tuning methods while requiring far fewer labeled samples. The ability to generalize with only a few training trajectories makes this approach practical for real-world deployment in new cities.

REFERENCES

1. Y. Zheng, X. Xie, W.-Y. Ma *et al.*, "Geolife: A collaborative social networking service among user, location and trajectory." IEEE Data Eng. Bull. **33**, 32 (2010).
2. A. Bolbol, T. Cheng, I. Tsapakis, and J. Haworth, "Inferring hybrid transportation modes from sparse gps data using a moving window svm classification," Comput. Environ. Urban Syst. **36**, 526 (2012).
3. J. Zeng, Y. Yu, Y. Chen, *et al.*, "Trajectory-as-a-sequence: A novel travel mode identification framework," Transp. Res. Part C: Emerg. Technol. **146**, 103957 (2023).

Estimating traffic flows from vehicle trajectories based on sparse mobile phone geolocation data

ROOS TEEUWEN[1,*] AND JORGE GIL[1]

[1] Department of Architecture and Civil Engineering, Chalmers University of Technology, Gothenburg, Sweden
[*] roos.teeuwen@chalmers.se

1. INTRODUCTION

Empirical traffic flow data are key in transport, mobility, and spatial planning. They can help understand congestion, traffic safety, environmental exposure and risks, among others. At scale, they can support data-driven decision making, helping to decide where interventions are most needed.

However, existing traffic flow data from sensors or traffic counts [1] lack spatio-temporal coverage and granularity. Other data, e.g. from navigation API's, are proprietary, commercial or limited-access, and unavailable to decision-makers. Large mobile phone traces data recently emerged as a promising source to capture dynamics at scale given their size, granularity, and coverage. They have been used to analyse travel demand (origin-destination), activity-locations, and individuals' activity spaces. Yet, despite their potential for exploring trajectories and traffic flows [1], dynamic applications other than understanding pedestrian routing behaviour [2] remain unexplored.

This study aims to explore how traffic flows with high spatial and temporal coverage and granularity can be estimated from vehicle trajectories based on sparse mobile phone geolocation data. We develop a methodology to create trajectories and flows from raw location data and test how various parameters affect the results. We contribute our methodology, code and data to allow for replication in other studies, and reflect on directions for future development.

2. DATA AND TOOLS

Raw mobile phone geolocation data from Sweden for 2024, containing high-resolution coordinates from GNSS (e.g. GPS) and other sources, timestamps and device ids, was acquired from the aggregator company *Pickwell*. We extract points in Stockholm (1292 million data points; from 3.2 million devices) and Gothenburg (670 million; 1.4 million), Sweden, including a 1 km buffer zone bounding box. We also collect three open ground truth data sets of average daily vehicle flows: (1) weekday flows in Stockholm, 2024, at 366 local and few major roads in Stockholm (obtained from *Stockholm Municipality*); (2) flows in Gothenburg, 2023, at 176 local and few major roads (*Gothenburg Municipality*); and (3) flows in and around Gothenburg, 2023, at 245 highway locations (Swedish Transport Administration *Trafikverket*). Lastly, we collect open data of roads accessible to vehicles and their speed limits (*Trafikverket*).

We process and analyse our data using open-source tools: *PostgreSQL* with *PostGIS* and *Python*. Data is collected and processed according to the European Union's General Data Protection Regulation and the data management plan has been approved by the Data Protection Office of the authors' institution. Our code and data are openly available on GitHub and Zenodo.

3. METHODOLOGY

A. From points to trajectories

We order all data points per device based on timestamp, and calculate the distance (d [m]), time difference (t [s]), and speed (s [km/h]) to the previous (i.e. incoming) and next data point (i.e. outgoing). To form trajectories, we exclude stationary points ($s \leq 3$ or $NULL$ and $d \leq 100$ for both incoming and outgoing points), isolated points ($s \leq 3$ and $d > 100$), singular points ($s = NULL$), and errors ($s > 250$). We then define trajectories as any set of points from the same device involving movement between stays (i.e. stationary), beginning with a *start* ($s_{incoming} \leq 3$ or $NULL$, and $d_{incoming} \leq 100$, and $s_{outgoing} > 3$ or $s_{outgoing} \leq 3$ but $d_{outgoing} > 100$), and assign those a common trajectory id.

These trajectories still require cleaning. We combine trajectories of the same device with a short stay in between ($t < 120$ and $d < 100$) (e.g. at a traffic light). We split trajectories in case of sudden jumps, i.e. when the next point is far away ($d > 1000$), the speed is too fast ($s > 250$), or too much time has passed ($t > 120$). In case a trajectory jumps far away and then back ($d \leq 100$) within reasonable time ($t < 120$) we only remove the jump. Lastly, we drop subsequent points with identical coordinates. We then create our final trajectories by linking points with the same unique trajectory id into linestrings, and calculate characteristics of the trajectory including the number of points, average distance between points and average speed.

B. Estimating flows per road segment

To estimate relative vehicle flows per road segment, we first match our trajectories to the road network using the *LCSS* map-matching algorithm [3]. The LCSS algorithm is a computationally efficient algorithm that iteratively matches (segments of) trajectories until an acceptable similarity fit has been achieved. We use the LCSS algorithm implementation from the *mappymatch* Python library, setting the similarity score parameter to 0.95 [3] and the other parameters to their default values. A limitation for our use case may be that the LCSS algorithm was developed for high-resolution data, while we use sparse data. However, since the algorithm recursively reduces the resolution of trajectories for computational efficiency as long as it yields an acceptable fit, in practice it handles trajectories of varying resolutions.

We match our trajectories with the collected road network, that we preprocess to resemble the data structure of *osmnx* networks [4] to allow integration with the LCSS algorithm as implemented in the mappymatch Python library. First, we build a data set of unique *nodes* as start or end points of the original edges, with their coordinates, and give each node a unique id as an artificial *osmid* attribute. Second, we build a data set of *edges*, looking up the osmid of their start and end points of every road segment, that we store as u and v attributes, respectively, and a unique *key* in case multiple edges connect the same pair of nodes. To resemble a *simplified* osmnx network, we drop the geometry of straight edges and maintain only the geometry of curved edges (i.e. having additional coordinates not captured by its nodes). We also add the reverse of every edge, since the original road network was non-directed.

After map-matching, we count the number of unique trajectories per road segment, and normalize these values to obtain the estimated relative vehicle flows within each city.

C. Evaluating suitability of trajectories

We compare our estimated flows to three ground truth data sets to evaluate the suitability of the generated trajectories. We join the midpoint of every ground truth road segment to the nearest (within 10 m) map-matched road segment. We then calculate the Spearman's ρ correlation between the observed traffic flow and the flow estimated from the mobile phone geolocation data, using a significance threshold of 0.01. We use varying sets of trajectories as input for the flows to evaluate how correlations are affected by the sample size and characteristics (i.e., number of points, average speed and resolution) of trajectories, to get insight into the most suitable trajectories for estimating flows. A smaller data sample has advantages in terms of map-matching and analysis efficiency, and reduced mobile data costs.

Robustness and sample size. To evaluate the robustness of the data, we take 10 random samples of 100,000 trajectories as input, test how successfully they are map-matched, and how resulting flows correlate with the ground truth data. To test the effect of sample size, we gradually increase the sample from 100,000 to 1,000,000 trajectories, concatenating the random samples and dropping duplicates, to see how correlations change.

Trajectory characteristics. To evaluate effects of trajectory characteristics, we use samples of 100,000 map-matched trajectories that fulfil various quality criteria: have a minimum number of points, gradually increasing from 2 to 10; a certain average speed, from 0 to 20 km/h, i.e., as a potential means to filter out non-motorized traffic; a maximum average distance between points, from 1000 to 100 m, i.e., as a measure of resolution. We then test how different thresholds correlate with the ground truth.

4. RESULTS

From our raw data, we generated 85 million trajectories in Stockholm and 46 million in Gothenburg. On average, trajectories span 0.9 km and last 1.9 min, with a resolution of 246 m between points. From these trajectories, we have estimated flows for 204,032 road segments in Stockholm and 127,290 in Gothenburg.

Robustness and sample size. All 10 random samples of 100,000 trajectories resulted in significant and comparably strong correlations (*rho* between 0.633 and 0.660, median 0.647, p-value < 0.01), and regardless of the sample, 73% of trajectories could successfully be map-matched to the road network, indicating stable outcomes when using different random samples of this size. Unmatched trajectories include trajectories at sea (e.g. ferries), at the borders of the study area, and to some extent at highways (e.g. due to the large spacing between nodes at these roads).

When increasing sample size from 100,000 to 1,000,000 input trajectories, we observed a negligible increase in correlation strength from 0.644 to 0.656 (p-value < 0.01). Simultaneously, the share of road segments that received an estimate increased from 54% to 81%, but this increase flattens with increasing sample size. Understanding these effects and how they hold for smaller and larger sample sizes warrants further research.

Trajectory characteristics. When selecting trajectories with a minimum amount of trajectory-points as input for estimating flows, correlation strengths varied between 0.478 and 0.636 (p-value < 0.01), but no pattern could be identified.

When setting a minimum threshold for trajectories' average speed (figure 1), comparison with the Stockholm data set showed decreasing correlations with increasing speed, from 0.600 when including trajectories of any speed to 0.390 when using a 20 km/h threshold (p-value < 0.01). For Gothenburg's local and major roads, we observed a slight increase from 0.574, to 0.636 at 7.5 km/h (p-value < 0.01), before decreasing again for higher speeds. For Gothenburg's highways, we observed a rise in correlation strengths from moderate, to strong when selecting trajectories over 10 km/h (0.629, p-value < 0.01), to very strong when selecting trajectories over 20 km/h (0.809, p-value < 0.01). This suggests that average trajectory speed is a promising indicator to improve estimates at highways, while it may not be a useful sole indicator to distinguish vehicle trajectories from other modes within the city, given speed variance between road types, and due to multi-modal trips, crossings, traffic lights, congestion, etc.

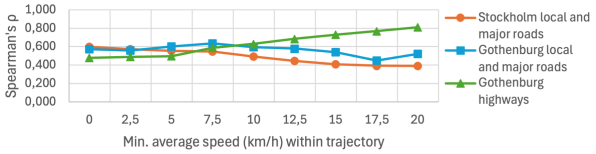

Fig. 1. Correlation strength differing with average speed.

When decreasing the allowed average distance between trajectory-points (i.e. increasing resolution), we found a decline in correlation strength with all three ground truth data sets for the most high-resolution trajectories. At high resolutions of 200 and 100 m on average, correlations are moderate for Stockholm, weak or insignificant for Gothenburg's local and major roads, and insignificant or even negative for Gothenburg's highways. This might indicate that high-resolution trajectories are not represented in every road type, requiring further investigation.

5. CONCLUSION

This paper introduces a methodology to estimate traffic flows from empirical vehicle trajectories based on sparse mobile phone geolocation data. We intend to open up new possibilities for data-driven transport, environmental and planning research and practice by providing extensive and frequent traffic flow estimates in entire metropolitan areas. Our approach yields strong correlations with ground truth data. For highway locations in Gothenburg, selecting trajectories of at least 20 km/h on average results in very strong correlations, while no subset of trajectories could be identified that outperforms a random sample of the data for the other ground truth data sets. Future work remains to understand where, when, and why deviations remain and how estimates within cities can be improved further, for instance by exploring other indicators than speed to identify vehicular trajectories, by ensuring a larger share of trajectories are map-matched, or by investigating how sparse mobile phone geolocation data can serve to estimate flows for different times, days and seasons.

Funding. This research is funded by Chalmers University of Technology Transport Area of Advance under grant agreement no. 2024-0299.

Disclosures. The authors declare no conflicts of interest.

Data availability. Estimated flows and ground truth data underlying the results presented in this paper are available on Zenodo. Geolocation points and trajectories cannot be made public due to privacy restrictions.

REFERENCES

1. J. C. Herrera, D. B. Work, R. Herring, *et al.*, "Evaluation of traffic data obtained via GPS-enabled mobile phones: The Mobile Century field experiment," Transp. Res. Part C: Emerg. Technol. **18**, 568 (2010).
2. C. Bongiorno, Y. Zhou, M. Kryven, *et al.*, "Vector-based pedestrian navigation in cities," Nat. Comput. Sci. **1**, 678 (2021).
3. L. Zhu, J. R. Holden, and J. D. Gonder, "Trajectory Segmentation Map-Matching Approach for Large-Scale, High-Resolution GPS Data," Transp. Res. Rec. J. Transp. Res. Board **2645**, 67 (2017).
4. G. Boeing, "Modeling and Analyzing Urban Networks and Amenities With OSMnx," Geogr. Anal. p. gean.70009 (2025).

The promise and pitfalls of mobile phone data for recreational fishing: Evaluating potential, biases, and representativeness

KATARZYNA SILA-NOWICKA,[1,*] RICHARD BEAUCHAMP,[1] MURRAY FORD,[1] IAN TUCK[2] AND BRUCE HARTILL [2]

[1]*School of Environment, the University of Auckland, New Zealand*
[2]*Ministry of Primary Industries, Fisheries New Zealand, New Zealand*
*Katarzyna.sila-nowicka@auckland.ac.nz

Received XX Month XXXX; revised XX Month, XXXX; accepted XX Month XXXX; posted XX Month XXXX (Doc. ID XXXXX); published XX Month XXXX

1. INTRODUCTION

Marine recreational fishing is a widespread activity that contributes to national economies, supports wellbeing, and strengthens people's connection to the marine environment. However, it can also place significant pressure on fish stocks, making effective monitoring and management essential [1]. Understanding the extent and spatial distribution of recreational fishing is critical to evaluating its ecological impact [2]. In New Zealand, current monitoring approaches—such as boat ramp surveys, telephone interviews, and aerial surveys—are often costly, sporadic, and reliant on self-reported data, limiting their coverage and reliability [3].

As a result, there is growing interest in alternative data sources that offer broader, timelier, and more cost-effective insights. Mobile phone location data (MPD) has emerged globally as a promising tool for tracking human movement [4], but its application to recreational fishing remains largely unexplored [5]. While MPD offers potential for mapping fishing activity at fine spatial and temporal scales, it is not without limitations. Biases related to phone ownership, app usage, and location service settings can result in underrepresentation of certain demographic or geographic groups. Technical issues, such as signal loss at sea and inconsistent data collection, further affect data completeness. Addressing these challenges—particularly those related to representativeness, bias, and privacy—is essential to ensure MPD can be used responsibly and effectively in fisheries research and management. Our paper directly addresses this first set of challenges to test the feasibility of using mobile phone data to study recreational fishing. It further explores how such biases emerge, how representative the data really are, and how methodological and design choices can help accommodate and mitigate these limitations.

2. DATA AND CASE STUDY

This study used de-identified mobile phone data from a third-party company to examine maritime activity around the New Zealand coast between 1 January 2020 and 31 December 2023. The dataset consists of over 400 million location records from mobile apps, focused on maritime zones and adjacent coastal areas (Fig. 1).

3. METHODS

We applied a two-stage spatial modelling approach to mobile phone data to detect recreational fishing activity and identify likely fishing zones. First, raw MPD records were de-duplicated, simplified, and classified by journey origin and vessel type using coastal boundary and origin polygons. Then, vessel movement characteristics—such as speed and turn angles—were analysed to detect stationary behaviour and spatial clusters, revealing probable fishing activity. These patterns were validated against known fishing areas and survey data, while temporal analysis highlighted seasonal and weather-related trends.

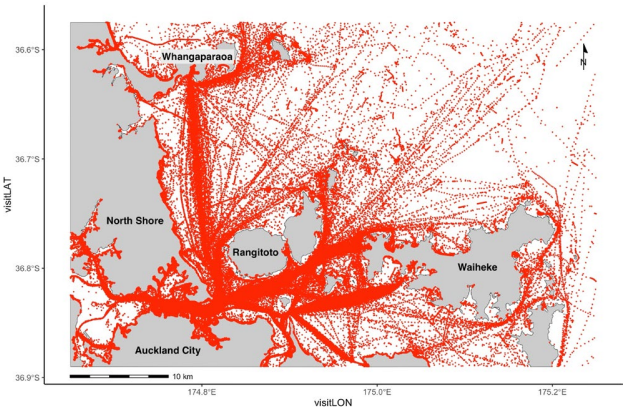

Fig. 1. MPD captured within Auckland area of interest from 1 March 2020 to 31 March 2020.

To ensure the reliability and representativeness of mobile phone data for studying marine activity in New Zealand, we designed a range of calibration and validation techniques. MPD can be cross-checked against CCTV footage at boat ramps, satellite imagery, AIS vessel tracking data, and ramp surveys to assess its accuracy in detecting vessels and people (Fig. 2). Using geospatial and AI-based analysis, we can evaluate spatial and temporal coverage, identify biases, and correct for them using statistical and machine learning models. We aim to account for variations in phone use, app coverage, reception, and demographic patterns, producing a spatially and temporally varying representativeness index. Together, these methods will enhance the robustness of MPD as a tool for understanding marine mobility. instructions.

Using AI-driven trajectory analysis and debiased mobile phone data, we develop predictive models to estimate

recreational fishing effort and the number of people on vessels across New Zealand's coastal waters. By integrating these metrics with the National Institute of Water and Atmospheric Research's ramp survey data and habitat suitability models, we aim to identify likely fishing locations, infer target species, and validate catch estimates at multiple spatial scales. The resulting tools and baseline data can support fisheries management by multiple government and industry users.

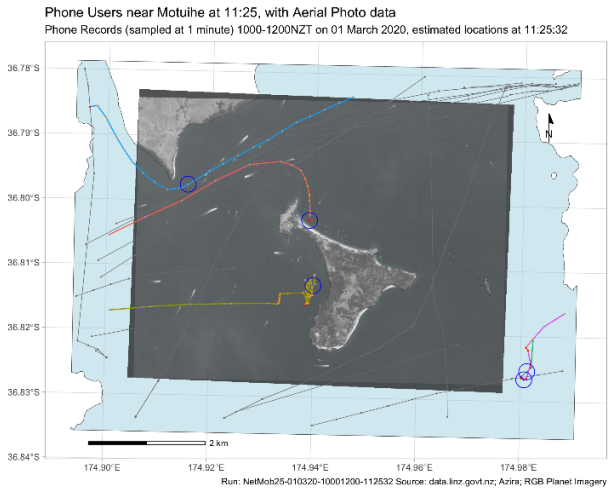

Fig. 2. Phone user journeys passing near Motuihe, Hauraki Gulf, on 1 March 2020, close to the time that the satellite image was taken.

4. PRELIMINARY RESULTS

Our analysis of mobile phone data successfully identified recreational fishing patterns with high spatial and temporal resolution, revealing activity at the level of specific hotspots as well as broader seasonal and daily trends. By classifying journeys based on their origin—such as boat ramps, marinas, or ferry terminals—we could infer likely vessel types, distinguishing between recreational boats, ferries, and other craft. Movement characteristics such as speed, turning angle, and stop duration enabled us to detect distinct journey types, including direct transits, fishing-focused outings, and mixed-use trips. Stop detection techniques, combined with spatial clustering, revealed where vessels were likely stationary for extended periods, indicating probable fishing activity. These fishing stops were concentrated in well-known inshore fishing areas, but we also identified less-documented locations. Differences in fishing locations were observed based on journey type and vessel origin, suggesting varying patterns of use across recreational fishing communities.

Using advanced movement analytics models, we were able to infer vessel behaviour and likely fishing zones with considerable detail. However, while these methods offer powerful new insights, the issue of data bias remains unresolved. We are only beginning to address challenges related to representativeness, particularly around who is captured in the data and under what conditions, highlighting the need for ongoing work to ensure more equitable and accurate use of MPD in fisheries research. As a result, fishing activity inferred from MPD may not fully capture the diversity or extent of recreational fishing, underscoring the importance of validation and debiasing techniques.

5. CONCLUSIONS

Mobile phone data offers valuable insights into maritime activity by linking phone locations over time to form vessel journeys. Journey simplification techniques reduced data volume while retaining essential movement characteristics, enabling accurate identification of vessel routes and stopping points, as verified through satellite imagery. Classifying journeys by their origin (e.g. boat ramps or ferry terminals) allowed inference of vessel types, which exhibit distinct behaviours. Movement analysis—using speed, turn angles, and spatial clustering—helped reveal likely fishing activity, aligning well with known fishing areas and existing survey data.

This research demonstrates the untapped potential of MPD for monitoring recreational fisheries, a sector largely invisible to conventional tracking systems like AIS. By adapting movement analysis methods, we showed that MPD can deliver detailed, high-resolution insights into fishing effort and behaviour. The approach provides a scalable, cost-effective complement to traditional monitoring tools, with applications in fisheries management, marine spatial planning, and conservation. In New Zealand, MPD could improve the accuracy and timeliness of recreational harvest estimates under the Quota Management System and inform responses to emerging challenges such as kina barrens or habitat restoration.

While differences in phone use, app coverage, and signal reception introduce bias into MPD, this study developed new methods to help address these limitations. Techniques such as journey classification, clustering, and the creation of a representativeness index show promise in correcting for uneven data coverage. However, these methods must be rigorously tested and validated before MPD can be reliably used to inform policy and management decisions.

Funding. This research was made possible through generous funding from the Ministry for Primary Industries through Fisheries New Zealand (FNZ).

References

1. M. Rufino, T. Mendo, J. Samarao, *et al.* "Estimating fishing effort in small-scale fisheries using high-resolution spatio-temporal tracking data (an implementation framework illustrated with case studies from Portugal", Ecol. Indic. **154**, 110628 (2023).
2. S. D. Bower, Ø. Aas, R. Arlinghaus *et al.*, "Knowledge gaps and management priorities for recreational fisheries in the developing world," (2020).
3. B. W. Hartill and T. T. Edwards "Comparison of recreational harvest estimates provided by onsite and offsite surveys: detecting bias and corroborating estimates" Can. J. Fish. Aquat. Sci. **72(9),** 1379–1389 (2015)
4. K. Siła-Nowicka, J. Vandrol, T. Oshan *et al.,* "Analysis of human mobility patterns from GPS trajectories and contextual information" Int. J. Geogr. Inf. Sci. **30(5),** 881–906 (2016).
5. R. Beauchamp, "Feasibility of estimating the extent of recreational fishing using anonymous smartphone location data," M.Sc. thesis, Univ. of Auckland (2025).

Modeling Mobility with Graph Neural Networks to Predict Neighbourhood Culture

THIAGO H SILVA[1,2*] AND DANIEL SILVER[2*]

[1] University of Toronto, Toronto, Canada
[2] Universidade Tecnológica Federal do Paraná, Curitiba, Brazil
[*] th.silva@utoronto.ca, dan.silver@utoronto.ca

1. INTRODUCTION

Urban neighbourhoods exhibit complex dynamics shaped by internal characteristics and external connections [1]. While early scholars recognized these relational aspects [2], quantitative modelling has been limited by methodological constraints. Our work bridges this gap by developing a graph neural network (GNN) framework that integrates three critical dimensions of neighbourhood change: mobility patterns using Yelp data [3], demographic attributes, and cultural dimensions derived from the Scenes Theory [4].

Unlike traditional mobility data derived from GPS traces or transportation logs, our model leverages digital co-visitation inferred from Yelp reviews, capturing a form of behavioural or social mobility that reflects routine consumer activity and lifestyle affinities [5].

This study offers two primary contributions. First, we demonstrate that GNNs effectively capture relational neighbourhood dynamics by representing cities as mobility (co-visitation) graphs. Second, we show that user-generated Yelp data can approach the predictive power of traditional census data in predicting cultural change, particularly in rapidly transforming urban environments.

2. METHODS

As shown in Figure 1, our approach follows a four-stage analytical pipeline that transforms raw Yelp data into predictive insights about neighbourhood cultural evolution.

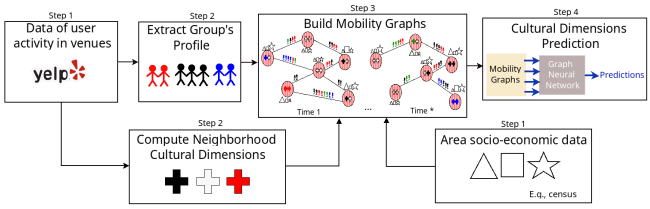

Fig. 1. Research workflow showing the four-stage process from data collection to cultural dimension prediction. Stage 1 extracts Yelp data (venue and user) and census data. Stage 2 identifies group profiles through topic modelling and clustering. Stage 3 constructs mobility graphs. Stage 4 implements the GNN prediction model. [Figure from [6]]

A. Data Preparation

We analyzed Yelp data (2011–2018) from Calgary, Montreal, and Toronto, focusing on commercial Forward Sortation Areas (FSAs) with at least 30 venues. Table 1 summarizes our dataset.

B. Model Components

Three key components formed our analytical framework:

Table 1. Overall statistics of the studied dataset.

	# Venues	# Reviews	# Users	# Categories
Calgary	7,736	97,650	34,645	774
Montreal	6,602	155,966	58,660	639
Toronto	18,906	525,435	169,516	903

Table 2. Graph Configuration Scenarios. Full model refers to Area info + mobility + group profile.

Scenario	Features		
	Vertex (D)	Edge (M)	Edge (C)
Full model	X	X	X
Area info + mobility	X	X	
Area info + group profile	X		X
Area info	X		
Mobility + group profile		X	X
Mobility		X	
Group profile only			X
None (baseline)			

Group Profiles: Using Latent Dirichlet Allocation (LDA) [7], we identified latent topics from venue categories visited by users. K-means clustering then grouped users with similar visitation patterns [8].

Mobility Graphs: Figure 2 illustrates our graph construction approach, where nodes represent FSAs and edges reflect co-visitation patterns. Edge weights quantify the number of users reviewing venues in both connected areas. In our mobility graph models, we used three types of features: *Vertex features (D):* Demographic and socioeconomic characteristics from census data and urban cultural dimensions; *Edge mobility features (M):* Weighted connections based on user mobility between areas; and *Edge group features (C):* Composition of user groups moving between neighbourhoods.

Cultural Dimensions: Following [4], we measured 15 scene characteristics (e.g., tradition, transgression) through systematic coding of venue types in each FSA.

C. GNN Architecture

We implemented a DeeperGCN model [9] with five GENConv layers, batch normalization, and residual links. The architecture processed different graph configurations. Table 2 summarizes the eight graph configurations we evaluated, systematically combining feature types (D, M, C) to assess their relative contributions. The full model incorporated all three feature types (D+M+C), while other configurations selectively omitted certain features to isolate their effects.

Comparative baselines included XGBoost [10] and Random

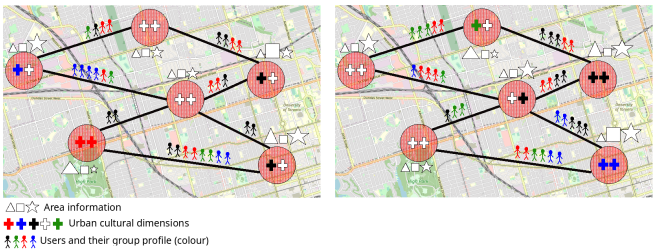

Fig. 2. Mobility graph structure illustrating node attributes (cultural dimensions and demographics) and edge attributes (mobility weights and group profiles) between neighbourhoods. [Figure from [6]]

Forest [11] models using only vertex features. All models were trained with historical data, including past cultural characteristics, and aimed to predict future urban cultural dimensions. More details on the methodology can be found in [6].

3. RESULTS AND DISCUSSION

Figure 3 shows the performance (RMSE) of different GNN-based graph scenarios across Calgary, Montreal, and Toronto. As expected, the "None" scenario, using only the base graph structure, produces the worst results across all cities. Similarly, adding only mobility data offers no significant improvement, indicating that neighbourhood connectivity alone is insufficient to predict cultural patterns.

The best predictions come from incorporating socioeconomic area information. Group profiles derived from Yelp activity perform nearly as well, especially in Calgary and Montreal, with Toronto showing only a marginal RMSE difference (∼0.001). Combining multiple data types (e.g., area info + group profile) does not yield better results than using area info alone.

Interestingly, while mobility alone contributes little predictive power, its inclusion with group profiles marginally improves prediction in certain areas, suggesting context-dependent value.

Compared to non-GNN baselines, the Naive model performs significantly worse, while tree-based models (Random Forest, XGBoost) using area data match GNN performance. Thus, simpler models may suffice when socioeconomic data is available; otherwise, group profiles offer a robust alternative.

Areas with higher prediction errors (e.g., Toronto's Parkdale and Leslieville) represent neighbourhoods undergoing rapid cultural transformation [12]. This suggests our model can identify areas where current trajectories diverge from historical patterns.

4. CONCLUSION

Our GNN framework advances urban analytics by simultaneously modelling mobility patterns, neighbourhood attributes, and cultural evolution. The approach demonstrates that while traditional census data remains valuable, digital traces from platforms like Yelp can effectively supplement official statistics. While our approach captures latent mobility through digital traces, it does not explicitly incorporate temporal dynamics or real-time movement. Future work could explore integrating temporally resolved location data or commuting flows to enrich the mobility component. Future extensions could also incorporate review text analysis and multi-city comparisons to enhance predictive capabilities.

ACKNOWLEDGMENTS

FAPESP (grant 2023/00148-0), and CNPq (grants 314603/2023-9, 441444/2023-7, 409669/2024-5, and 444724/2024-9).

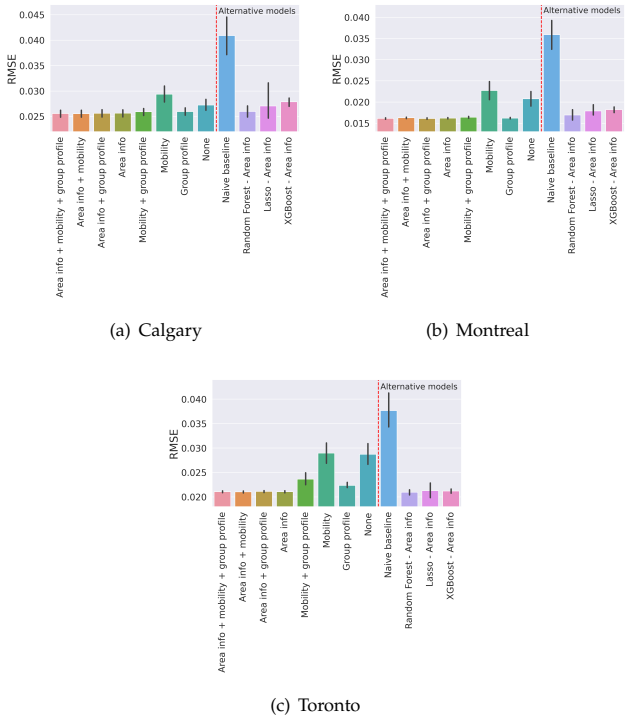

(a) Calgary (b) Montreal

(c) Toronto

Fig. 3. Comparison of prediction errors across model configurations in the cities studied, demonstrating the relative importance of different feature types. Error bars show 95% confidence intervals. Scenario None, a graph without census information, group profile, and mobility information, presents the worst results for all cities. Note on group profiles: for Calgary and Montreal, we found five group profiles considering five topics; for Toronto, we found four group profiles considering seven topics. Recall that the GNN scenarios' names are represented in Table 2. [Figure from [6]]

REFERENCES

1. A. Abbott, *Department and discipline: Chicago sociology at one hundred* (University of Chicago Press, 2017).
2. H. W. Zorbaugh, *The gold coast and the slum: A sociological study of Chicago's near north side* (University of Chicago Press, 1983).
3. Yelp, "Yelp dataset," (2020). https://www.yelp.com/dataset.
4. D. A. Silver and T. N. Clark, *Scenescapes: How qualities of place shape social life* (University of Chicago Press, 2016).
5. T. H. Silva, A. C. Viana, F. Benevenuto, *et al.*, "Urban computing leveraging location-based social network data: A survey," ACM Comput. Surv. **52**, 17:1 (2019).
6. T. H. Silva and D. Silver, "Using graph neural networks to predict local culture," Environ. Plan. B: Urban Anal. City Sci. **52**, 355 (2025).
7. D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," J. Mach. Learn. Res. **3**, 993 (2003).
8. D. A. Silver and T. H. Silva, "Complex causal structures of neighbourhood change: Evidence from a functionalist model and yelp data," Cities **133**, 104130 (2023).
9. G. Li, C. Xiong, A. Thabet, and B. Ghanem, "Deepergcn: All you need to train deeper gcns," arXiv preprint arXiv:2006.07739 (2020).
10. T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proc. of KDD*, (2016), pp. 785–794.
11. P.-N. Tan, M. Steinbach, and V. Kumar, *Introduction to data mining* (Pearson Education India, 2016).
12. D. Silver, P. Adler, and M. S. Fox, "Towards a model of urban evolution—part i: Context," Urban Sci. **6**, 87 (2022).

Mapping Urban Mobility with Interest Networks: Cross-Platform and User Behavior Insights

GUSTAVO H. SANTOS[1,*], MYRIAM DELGADO[1], AND THIAGO H. SILVA[1,2]

[1] *Universidade Tecnológica Federal do Paraná, Curitiba, Brazil*
[2] *University of Toronto, Toronto, Canada*
[*] *gustavohenriquesantos@alunos.utfpr.edu.br, {myriamdelg,thiagoh}@utfpr.edu.br*

1. INTRODUCTION

Location-Based Social Networks (LBSNs) provide valuable data for modeling urban behavior through Interest Networks (iNETs), which are undirected graphs with nodes representing urban areas and edges linking co-visited regions. We construct iNETs by connecting regions based on shared user interactions, enriching the studied regions with contextual attributes such as socioeconomic and cultural data to understand user interest. The work presents the key contributions:

Validation across platforms: We compare iNETs from Google Places and Foursquare, two different LBSN platforms, taken from cities in different countries, showing that coarser granularities yield consistent patterns (high correlation between iNETs), while finer levels expose platform-specific behaviors. To reconcile differences between platforms in finer granularities modeling, we introduce *Urban Preference Zones (UPZones)*, behaviorally coherent clusters within each LBSN.

Understanding user interests: Geographic proximity and venue similarity emerge as the primary drivers of user mobility, while socioeconomic factors, especially education, play a modest yet notable role. These findings are based on a case study conducted in Curitiba.

Proposed tool: `h3-cities`[1] facilitates multi-scale spatial analysis with an interface that allows users to select the resolution and city to obtain H3 hexagons across the world.

Prior work has explored LBSN data complementarity across platforms like Instagram and Foursquare [1], with applications in urban dynamics [2] and neighborhood studies [3]. Some studies validate LBSN data against official sources [4]. We focus on cross-platform behavioral modeling discrepancies, extending multi-scale urban analysis frameworks [5] and clustering approaches like Livehoods [6] with our grid-based UPZones method.

2. DATA

LBSN Datasets: We analyze two primary datasets: *Foursquare (FS)*: 7.5M check-ins from 1.4M users at 2.7M global locations; *Google Places (GP)*: 11M reviews from 5M users for 3M venues worldwide; **Contextual Data:** Urban regions are enriched with: i) Socioeconomic indicators (income, education, employment, population size, racial composition); ii) Political orientation (presidential election data); iii) Venues and Cultural characteristics (Scenes Theory vectors [7]). **Spatial Granularities** We employ: i) Administrative divisions (neighborhoods, ZIP codes, census tracts); ii) Multi-resolution hexagonal grids (h6-h9, 0.11-36.12 km^2) via `h3-cities`. We analyze the 20 most popular U.S. counties (by review count) and the 20 most popular U.S. cities, and conduct case studies in London, UK, and Curitiba, Brazil.

3. METHODOLOGY

Our methodology develops a three-phase (see Figure 1) framework for urban interest modeling. First (Figure 1a), we gather user reviews by region, e.g., neighborhoods or H3 cells. Then (Figure 1b), we identify users who have visited the same pair of regions. Finally (Figure 1c), we aggregate this data to build Interest Networks (iNETs) as weighted graphs where nodes are urban regions and edges capture co-visitation patterns, with edge weights reflecting the number of users interacting with both regions (self-loops map repeated visits in the same region).

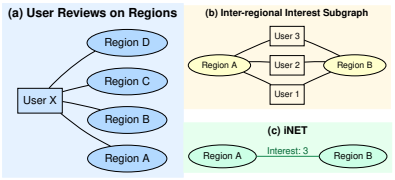

Fig. 1. iNETs modelling explored in this study

Fig. 2. Formation of Urban Preference Zones (UPZones)

For cross-platform validation, we systematically compare iNETs derived from Google Places and Foursquare data across multiple spatial resolutions, including neighborhoods, boroughs, census tracts, and zip codes; and at varying granularities, from coarse (h6 hexagons, 36.12 km^2) to fine (h9 hexagons, 0.11 km^2). We employ correlation metrics (Pearson and Spearman coefficients) to assess edge weight similarity and rank the nodes by their eigenvector centrality, using Kendall's Tau to evaluate node ranking consistency. We use Pearson to capture linear relationships in edge weights, Spearman to assess monotonic trends, and Kendall's Tau to evaluate the consistency of region importance rankings based on network centrality.

To reconcile platform-specific behaviors, we introduce Urban Preference Zones (UPZones), which are defined separately for each LBSN. Observing that the strongest connections at fine spatial scales (e.g. h9) occur between neighboring cells, we filter the iNET (Figure 2a) to include only adjacent h9 cells (Figure 2b). Next, we apply the Leiden algorithm (Figure 2c), with $\gamma = 1$, to detect behaviorally coherent and spatially contiguous regions, i.e., densely connected clusters of adjacent cells, referred to as UPZones (Figure 2d). We characterize each UPZone's functional identity (e.g., entertainment, retail) using TF-IDF on its venue category frequencies. To compare UPZones across the two LBSNs, we first align them using the Hungarian algorithm and then measure their structural similarity with Normalized Mutual Information (NMI) and the Rand Index.

To understand what drives LBSN user interest, we correlate the edge weights of iNETs with differences and similarities between areas, both calculated based on several socio-economic information. Geographic distance reflects mobility patterns, such

[1] Proposed tool: https://h3-cities.vercel.app

as a tendency to visit nearby destinations. Population size captures contrasts between residential and business-dense regions. Economic alignment is measured through mean household income. The absolute difference in literacy rates measures educational alignment. Demographic affinity is assessed through racial composition, calculated as half the sum of absolute differences across racial groups [8]. Political orientation is represented by differences in presidential voting, highlighting ideological alignment. Venue category similarity, measured with cosine similarity of category frequency vectors, reflects overlap in available experiences. Finally, cultural alignment is informed by Scenes Theory [7], operationalized as the Euclidean distance across 15 cultural dimensions (see [9] for more details).

4. RESULTS

Our analysis reveals several key findings about urban interest patterns and platform comparability. The cross-platform validation demonstrates for U.S. regions an inverse relationship between spatial granularity and consistency (Fig. 3), with coarser resolutions (h6) showing strong correlations (Pearson's $r \gtrsim 0.75$) between Google Places and Foursquare iNETs, while finer resolutions (h9) exhibit greater divergence ($r \lesssim 0.6$)); administrative divisions like neighborhoods and census tracts typically aligning with h7-h8 granularities and zipcodes and boroughs matching h6 characteristics. At finer resolutions, the iNETs diverge more significantly because a shared user interest in a large area gets diluted across many smaller regions, revealing distinct, platform-specific user behaviors and purposes at a micro-level. We notice that it is not an artifact of U.S. regions, as this pattern holds across other studied cities (Curitiba and London).

The UPZone clustering successfully identifies semantically meaningful urban areas (Fig. 4a), as evidenced by the London case study at h9 granularity, where detected zones correspond to known districts like Soho (entertainment hub) and Mayfair (luxury retail). The quantitative evaluation shows moderate agreement between platforms (NMI = 0.6364, Rand Index = 0.7378), suggesting that while platform-specific behaviors exist, underlying urban interest patterns remain detectable even for finer granularities. We also evaluate UPZones constructed using h8 cells, but the larger created zones loose some of the detail obtained through the h9 modeling. Finally, we show how UPZones can help urban planners familiar with the city understand the complex links among these densely connected areas. This is achieved by modeling an iNET where each UPZone functions as a node (UPZ-iNET in Fig. 4b).

Investigating user interest, we consider a case study on iNETs of both GP and FS in Curitiba. Table 1 highlights Geographic distance as the strongest predictor of user interest, followed by education levels and venue similarity.

Table 1. Spearman's ρ for predictors of user interest (Curitiba)

Predictor	GP	FS
Geographic distance	-0.42	-0.40
Education level	-0.33	-0.23
Venue similarity	0.33	0.12
Racial composition	-0.26	-0.14
Political polarization	-0.26	-0.14
Income difference	-0.17	-0.26
Population size	-0.14	-0.03
Cultural affinity	-0.28	0.00

5. CONCLUSION

In this work, we have explored Interest Networks (iNETs) to model urban mobility using LBSN data, showing that spatial granularity strongly impacts cross-platform consistency: coarser resolutions yield stable patterns, while finer ones reveal platform-specific behaviors. To reconcile these differences, we introduce Urban Preference Zones (UPZones), coherent spatial clusters that go beyond administrative boundaries. Our analysis reveals that geographic proximity, education levels and venue similarity are the main drivers of user mobility (in the case study), with socioeconomic and political factors playing secondary roles.

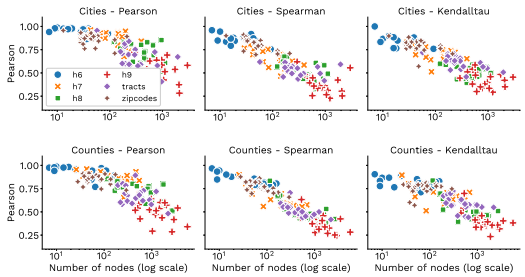

Fig. 3. GP and FS iNETs similarity across granularities

(a) UPZones modeled in London

(b) UPZ-iNET: Interest network of urban preference zones

Fig. 4. UPZones and UPZ-iNET in London

6. ACKNOWLEDGEMENTS

FAPESP (grant 2023/00148-0) and CNPq (grants 314603/2023-9, 441444/2023-7, 409669/2024-5, and 444724/2024-9).

REFERENCES

1. T. H. Silva, P. O. S. Vaz de Melo, J. M. Almeida, *et al.*, "A comparison of foursquare and instagram to the study of city dynamics and urban social behavior," in *Proc. of UrbComp*, (Chicago, USA, 2013).
2. P. Martí, L. Serrano-Estrada, and A. Nolasco-Cirugeda, "Social media data: Challenges, opportunities and limitations in urban studies," Comput. Environ. Urban Syst. **74**, 161 (2019).
3. A. Nolasco-Cirugeda and C. García-Mayor, "Social dynamics in cities: Analysis through lbsn data," Procedia Comput. Sci. **207**, 877 (2022).
4. T. H. Silva, P. O. V. de Melo, J. M. Almeida, *et al.*, "A large-scale study of cultural differences using urban data about eating and drinking preferences," Inf. Syst. **72**, 95 (2017).
5. D. Q. Wu, J. Tan, F. Guo, *et al.*, "Multi-Scale Identification of Urban Landscape Structure Based on Two-Dimensional Wavelet Analysis: The Case of Metropolitan Beijing, China," Ecol. Complex. **43**, 100832 (2020).
6. J. Cranshaw, R. Schwartz, J. Hong, and N. Sadeh, "The Livehoods Project: Utilizing Social Media to Understand the Dynamics of a City," in *Proc. of ICWSM*, (2012).
7. D. A. Silver and T. N. Clark, *Scenescapes: How qualities of place shape social life* (The University of Chicago, 2016).
8. P. Huang and C. T. Butts, "Rooted america: Immobility and segregation of the intercounty migration network," Am. Sociol. Rev. **88**, 1031–1065 (2023).
9. G. H. Santos, F. R. Gubert, M. Delgado, and T. H. Silva, "Modeling interest networks in urban areas: A comparative study of google places and foursquare across countries," J. Internet Serv. Appl. **16**, 25–42 (2025).

Evaluating Deep Learning for Transport Mode Detection on Crowdsourced Smartphone GPS Data

PAUL FERMÉ[1,*], ABDELLAH EL HASSANI[2], AURÉLIE PIRAYRE[3], AND GUILLAUME SABIRON[4]

[1,2,4] IFP Energies nouvelles, Rond-point de l'échangeur de Solaize, BP 3, 69360 Solaize, France
[3] IFP Energies nouvelles, 1 et 4 avenue de Bois-Préau, 92852 Rueil-Malmaison, France
[*] paul.ferme@ifpen.fr

1. INTRODUCTION

With the widespread adoption of connected devices, it is possible to accurately track individual mobility patterns. In particular, the automatic detection of transportation modes [1] enables a range of applications, including urban transportation planning [2], mobility behavior analysis [3], and carbon footprint estimation [4].

Historically, the task of transport mode classification was addressed using heuristic methods based on statistical rules [5–8]. This was followed by machine learning approaches [9–12], which relied on carefully engineered features. More recently, deep learning methods [13–15] have been employed to automatically extract relevant features from mobility time series data using deep neural networks.

Although the latter approaches have shown higher efficiency on publicly available mobility datasets such as GeoLife, it remains unclear whether these deep learning models are resilient to automatically collected, uncontrolled data. In this study, we use data collected through the Geco Air application to investigate this issue. We demonstrate that a machine learning approach based on carefully engineered features outperforms significantly the deep learning method proposed in [15] on this dataset. However, hybrid models that combine both approaches yield a small but measurable performance improvement.

2. GECO AIR APPLICATION

Geco Air[1] is a smartphone application that promotes eco-driving and sustainable transportation practices. It estimates users' carbon footprints and pollutant emissions based on travel data, using automatic transport mode detection. Users can manually confirm or correct the predicted mode, allowing the collection of a large set of labeled trips. In this study, we focus on trips collected and labeled in 2023; see Table 1 of [12] for the detailed distribution. Overall, our dataset includes ten transportation modes—one of the most comprehensive in the literature [1].

3. GEOLIFE AND GECO AIR BENCHMARKING

We focus on the deep learning architecture described in [15]. Although it is not as accurate as the model in [14], it is significantly more efficient in terms of computation time—a critical factor for deployment within the Geco Air application. We apply this architecture to the Geco Air dataset instead of the commonly used GeoLife dataset [16], which consists of GPS trajectories collected via user-controlled start-stop recordings on GPS loggers and GPS-enabled phones. To address the pronounced class imbalance in the Geco Air data, we employ undersampling, oversampling, and the Synthetic Minority Over-sampling Technique (SMOTE) [17] over fixed-size triplegs. This dataset was divided into training, validation, and test sets using a stratified approach with 64%, 16% and 20% of the data allocated to each set respectively, on which we have trained the structure described in [15].

[1] https://www.gecoair.fr/en/

We provide a summary of results per mode based on precision, recall and the F_1-score in Table 1.

Table 1. Evaluation metrics of Deep Learning Structure from [15] applied to Geco Air data

Class	Precision	Recall	F_1	GeoLife F_1
Walking	0.72	0.85	**0.78**	0.89
Bicycle	0.94	0.80	**0.87**	0.80
Car	0.87	0.74	**0.80**	0.74
Bus	0.51	0.42	**0.46**	0.74
Subway	0.39	0.41	**0.40**	0.82
Train	0.50	0.76	**0.60**	0.92
Average	0.65	0.66	**0.65**	0.82

We observe significantly lower performance on bus, subway, and train in comparison to results obtained using the GeoLife dataset, where we reproduced the findings of [15]. For the other modes, performance remains more or less consistent across both datasets.

4. CLASSIFICATION ALGORITHM

The core component of our machine learning approach is the careful selection of input features, which include: total distance and duration of the trip; mean and standard deviation of speed, acceleration (both positive and negative), and elevation; Geographic Information System-based heuristics such as proximity to railroads, subway lines, bus lines, and airports; number of stops, and whether these stops occur near train stations or bus stops (novel feature compared to [12]); signal quality indicators; and smartphone-based motion activity from the device's core motion sensors.

To integrate the deep learning architecture into our algorithm, we adopt a hierarchical structure that handles, in a second stage, transport modes that are particularly difficult to distinguish. For this purpose, we train a dedicated deep learning model specifically on these ambiguous modes, using trips collected and labeled in 2024. The output of this model is then incorporated as an additional feature in our standard machine learning pipeline, forming our hybrid approach (see Figure 1).

For these hierarchical classification tasks, we use compact neural networks consisting of two hidden layers with 64 and 32 neurons, respectively. The ReLU activation function is applied to the hidden layers, while a Softmax activation is used in the output layer for multi-class classification.

5. RESULTS

Our dataset was divided into training, validation, and test sets using a stratified approach with 64%, 16% and 20% of the data

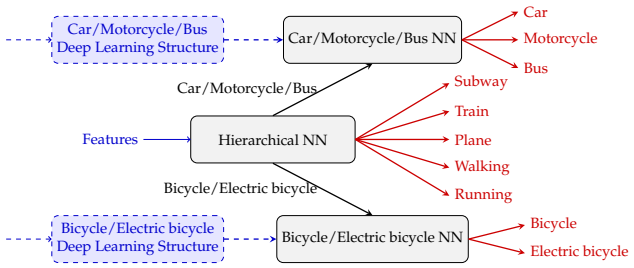

Fig. 1. Hierarchical structure of the proposed classification algorithm.

allocated to each set respectively. A summary of the results for each hierarchical level, with and without the inclusion of the deep learning feature, is presented in Tables 2,3 and 4.

Table 2. Evaluation metrics of Hierarchical Neural Network

Class	Precision	Recall	F_1
Car/Motorcycle/Bus	0.97	0.99	**0.98**
Subway	0.78	0.61	**0.69**
Train	0.89	0.88	**0.88**
Plane	0.91	0.86	**0.88**
Walking	0.92	0.71	**0.80**
Running	0.78	0.65	**0.71**
Bicycle/Electric bicycle	0.94	0.90	**0.92**
Average	0.88	0.80	**0.84**

Table 3. Evaluation metrics for Car/Motorcycle/Bus classification

Class	Prec.	Rec.	F_1	Deep Learning F_1
Car	0.98	0.98	**0.98**	**0.98**
Motorcycle	0.76	0.71	**0.73**	**0.72**
Bus	0.53	0.57	**0.55**	**0.57**
Average	0.76	0.75	**0.75**	**0.75**

Table 4. Evaluation metrics for Bicycle/Electric Bicycle classification

Class	Prec.	Rec.	F_1	Deep Learning F_1
Bicycle	0.83	0.78	**0.80**	**0.82**
Electric Bicycle	0.58	0.65	**0.61**	**0.63**
Average	0.70	0.72	**0.71**	**0.72**

The first step of the hierarchical classification, as shown in Table 2, performs well—particularly for modes such as subway and train where the deep learning approach previously failed. It achieves a weighted F_1-score of 0.96, taking into account the class imbalance present in the data. The classification performance of car, bus, and motorcycle described in Table 3 does not improve with the addition of the deep learning feature. In contrast, for bicycle and electric bicycle, adding the deep learning feature yields a small performance improvement; see Table 4.

6. CONCLUSION AND FUTURE WORK

While deep learning approaches have shown strong results on datasets like GeoLife, their robustness to automatically collected, uncontrolled data remains uncertain. In particular, on Geco Air data, we demonstrate that a machine learning method with carefully engineered features significantly outperforms the deep learning model from [15]. Hybrid models integrating both approaches demonstrate a small but measurable performance improvement, marking an encouraging early trend.

Future work will explore alternative deep learning architectures, particularly [14], to better assess their practical utility beyond computational efficiency for real-world deployment.

REFERENCES

1. P. Sadeghian, J. Håkansson, and X. Zhao, "Review and evaluation of methods in transport mode detection based on gps tracking data," J. Traffic Transp. Eng. (English Ed. **8**, 467 (2021).
2. M. D. Meyer and E. J. Miller, *Urban transportation planning: A decision-oriented approach* (2001).
3. T. Gärling, D. Ettema, and M. Friman, *Handbook of sustainable travel* (Springer, 2014).
4. H. S. Matthews, C. T. Hendrickson, and C. L. Weber, "The importance of carbon footprint estimation boundaries," (2008).
5. S. Y. A. Tsui and A. S. Shalaby, "Enhanced system for link and mode identification for personal travel surveys based on global positioning systems," Transp. Res. Rec. **1972**, 38 (2006).
6. P. Stopher, C. FitzGerald, and J. Zhang, "Search for a global positioning system device to measure person travel," Transp. Res. Part C: Emerg. Technol. **16**, 350 (2008).
7. C. Chen, H. Gong, C. Lawson, and E. Bialostozky, "Evaluating the feasibility of a passive travel survey collection in a complex urban environment: Lessons learned from the new york city case study," Transp. Res. Part A: Policy Pract. **44**, 830 (2010).
8. F. G. de Cossío and G. Sabiron, "A robust approach for transportation mode detection using smartphone-based gps sensors and road network information," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, (IEEE, 2023), pp. 959–965.
9. S. Dodge, R. Weibel, and E. Forootan, "Revealing the physics of movement: Comparing the similarity of movement characteristics of different types of moving objects," Comput. Environ. Urban Syst. **33**, 419 (2009).
10. Y. Zheng, Y. Chen, Q. Li, *et al.*, "Understanding transportation modes based on gps data for web applications," ACM Trans. on Web (TWEB) **4**, 1 (2010).
11. H. Mäenpää, A. Lobov, and J. L. M. Lastra, "Travel mode estimation for multi-modal journey planner," Transp. Res. Part C: Emerg. Technol. **82**, 273 (2017).
12. P. Fermé and G. Sabiron, "Transport mode detection using crowd-sensed data: A case study with the Geco Air application," in *NetMob 2024 – Book of Abstracts*, (NetMob, Washington, DC, USA, 2024), pp. 118–119.
13. Y. Endo, H. Toda, K. Nishida, and J. Ikedo, "Classifying spatial trajectories using representation learning," Int. J. Data Sci. Anal. **2**, 107 (2016).
14. J. James, "Travel mode identification with gps trajectories using wavelet transform and deep learning," IEEE Trans. on Intell. Transp. Syst. **22**, 1093 (2020).
15. H. Moreau, A. Vassilev, and L. Chen, "The devil is in the details: An efficient convolutional neural network for transport mode detection," IEEE Trans. on Intell. Transp. Syst. **23**, 12202 (2021).
16. Y. Zheng, X. Xie, W.-Y. Ma *et al.*, "Geolife: A collaborative social networking service among user, location and trajectory." IEEE Data Eng. Bull. **33**, 32 (2010).
17. N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "Smote: synthetic minority over-sampling technique," J. artificial intelligence research **16**, 321 (2002).

Modeling Urban Air Quality Using Taxis as Sensors

ANASTASIOS NOULAS[1], YASIN ACIKMESE[2], CHARLES QC LI[6], MILAN Y. PATEL[5], SHAZIA'AYN BABUL[3], RONALD C. COHEN[5], RENAUD LAMBIOTTE[3], AND MARTA C. GONZÁLEZ[4]

[1] *The Alan Turing Institute, UK , anoulas@turing.ac.uk*
[2] *Firefly, USA, yasin@fireflyon.com*
[3] *Mathematical Institute, University of Oxford, UK, renaud.lambiotte@maths.ox.ac.uk*
[4] *Departments of City and Regional Planning and Civil and Environmental Engineering, University of California, Berkeley, USA, martag@berkeley.edu*
[5] *Department of Chemistry, University of California, Berkeley, USA, milan_patel@berkeley.edu , rccohen@berkeley.edu*
[6] *College of Environmental Design, University of California, Berkeley, USA, liqianchuan9321@berkeley.edu*

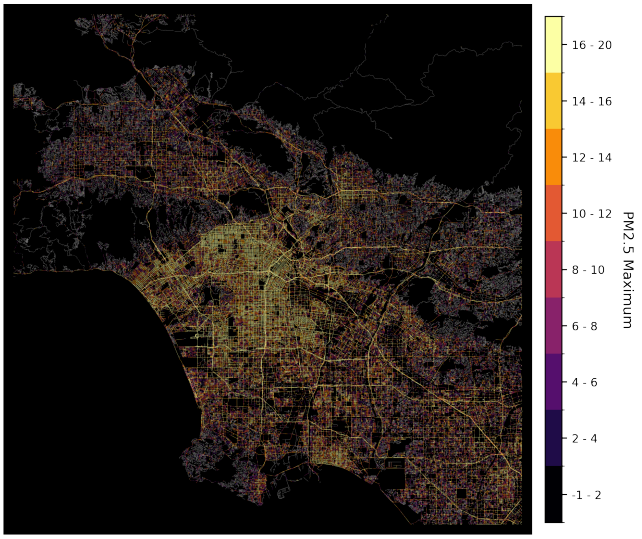

Fig. 1. Urban Street Network visualization of the Los Angeles area. Color intensity is set according to the max PM2.5 record in each street segment. Areas of the network without available records have been colored in gray.

Poor environmental and atmospheric conditions in urban areas are among the leading causes of premature death and chronic respiratory conditions such as asthma and lung cancer [1]. This affects billions of people residing in urban and peri-urban environments worldwide [2]. A major concern for citizens, urban authorities, and policymakers is the ability to accurately and efficiently monitor urban air quality [3]. One of the most widely adopted approaches to large-scale air quality monitoring involves citizen-led projects [4], which have emerged to address limitations in scalability and maintenance costs associated with government-operated sensing infrastructure. These citizen initiatives have been inspired by the success of crowdsourcing efforts over the past two decades, which have enabled data collection across domains including mapping, urban social activity, and human mobility [5].

Despite their critical value, these projects face significant challenges, including the need for sustained citizen participation, along with deployment and operational hurdles [6]. In recent years, the concept of leveraging taxi fleets as urban sensing platforms has gained attention due to their advantages in scale, spatial resolution, and cost-effectiveness [7, 8]. A growing body of work [9] has shown that taxis and other vehicles can serve as effective platforms for air quality monitoring. Operating continuously across time and space within cities, taxis and ride-sharing vehicles offer a high degree of spatiotemporal coverage, enabling the collection of environmental data with fine granularity [10]. Unlike dedicated monitoring vehicles deployed by governments, research institutions, or private organizations [11], taxis do not require explicit commissioning to traverse city streets—a process that adds substantial logistical and economic burden.

Here, we investigate the use of taxi fleets for urban air quality sensing through data collected by Firefly, a mobile advertising and technology platform. While there have been previous deployment scenarios that have utilized taxis as a sensing platform [12, 13], in this work we take a step further introducing a dataset of a much larger scale. We analyze over 100 million PM2.5 sensor readings collected from more than 3,000 vehicles operating in six major U.S. cities over a 12-month period (2022–2023). We demonstrate how existing vehicle fleets can be transformed into a large-scale air quality sensing network when supported by a sustainable business model [14]. Using vehicle-based sensing, we capture the temporal variation in PM2.5 levels, including seasonal patterns and major environmental events [15]. Across the spatial dimension, we show how taxis can provide highly granular, street-level air quality measurements while maintaining broad citywide coverage [16]. We further investigate new possibilities in urban air quality modeling made feasible by large-scale mobile sensing. In particular, we assess the predictability of pollution levels based on urban indicators such as traffic congestion, built environment characteristics, and human mobility. Our analysis explores the trade-offs between localized prediction accuracy and the availability of supporting PM2.5 data [17]. Visualizing PM2.5 concentrations at the city level, we identify distinct seasonal and geographic trends. We show that models using traffic congestion and wind speed data alone can effectively capture a city's characteristic pollution dynamics [18].

Our findings suggest that taxi fleets can serve as operationally efficient, near-real-time air quality monitoring systems with the potential for global scalability. We expect these data to further advance research into the interplay between vehicular traffic, urban activity, and pollution levels [19–25], and to support the development of mobility models and simulators [26**?** –28] that integrate pollution dispersion with urban dynamics. Moreover, while the link between air quality and socioeconomic disparities [29] are beyond the scope of this study, we hope these data can support future research addressing this urgent public health issue.

REFERENCES

1. R. Li, R. Zhou, and J. Zhang, "Function of PM2.5 in the pathogenesis of lung cancer and chronic airway inflammatory diseases," Oncol. letters (2018).
2. UNESCO, "Urban air pollution: The silent killer," (2022). Available at https://www.unesco.org.

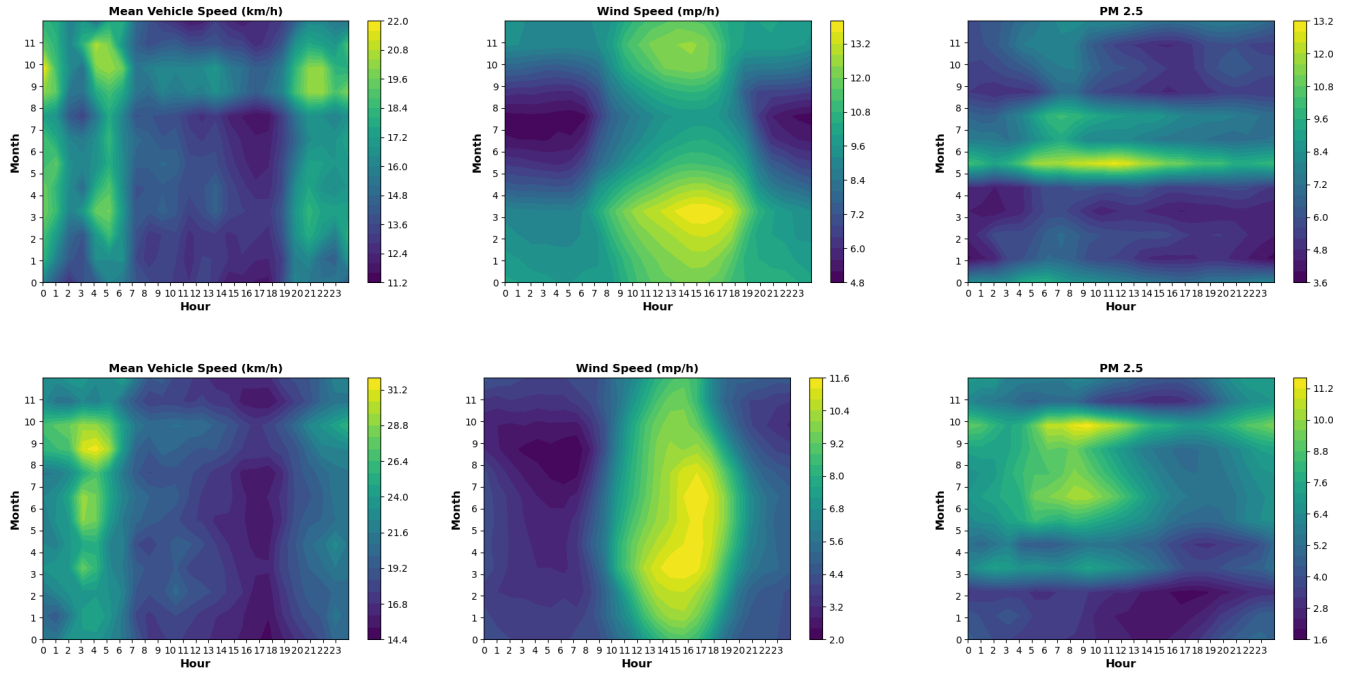

Fig. 2. Contour plots representing intensity levels of mean vehicle speed (km/h), mean wind speed (km/h) and mean PM2.5 levels in the Chicago (top) and Los Angeles areas (bottom). Signal variations are shown for different months (1-12) and daily hours (0-23).

3. W. H. Organization, "Who global air quality guidelines," (2021). Available at https://www.who.int/publications/i/item/9789240034228.
4. N. Castell, F. Dauge, P. Schneider, *et al.*, "Can commercial low-cost sensor platforms contribute to air quality monitoring and exposure estimates?" Environ. international (2017).
5. M. Haklay, "Citizen science and volunteered geographic information: Overview and typology of participation," in *Crowdsourcing geographic knowledge,* (Springer, 2013), pp. 105–122.
6. B. Maag, Z. Zhou, W. Apolinarski *et al.*, "sensebox: A community-driven approach for environmental monitoring and education," Sensors **18**, 3714 (2018).
7. K. O'Keeffe, A. Anjomshoaa, S. Strogatz, *et al.*, "Quantifying the sensing power of vehicle fleets," PNAS (2019).
8. D. Hasenfratz, O. Saukh, S. Sturzenegger *et al.*, "Participatory air pollution monitoring using smartphones," in *Mobile Sensing: From Smartphones and Wearables to Big Data,* (2012).
9. Y. M. Idir, O. Orfila, V. Judalet, *et al.*, "Mapping urban air quality from mobile sensors using spatio-temporal geostatistics," Sensors **21**, 4717 (2021).
10. L. Aplin *et al.*, "Moving urban monitoring forward: From static sensors to mobile platforms," Environ. Res. Lett. **16**, 110201 (2021).
11. I. Testi, A. Wang, S. Paul, *et al.*, "Big mobility data reveals hyperlocal air pollution exposure disparities in the bronx, new york," Nat. Cities **1**, 512 (2024).
12. Y. Sun, P. Brimblecombe, P. Wei, *et al.*, "High resolution on-road air pollution using a large taxi-based mobile sensor network," Sensors **22**, 6005 (2022).
13. Y. Wu, Y. Wang, L. Wang, *et al.*, "Application of a taxi-based mobile atmospheric monitoring system in cangzhou, china," Transp. research part D: transport environment **86**, 102449 (2020).
14. Y. Cheng, X. Wu, and Y. Lin, "A scalable and low-cost sensing framework for real-time urban air quality monitoring," Sensors **21**, 4032 (2021).
15. J. Zhou, H. Li, and Y. Zhang, "Urban mobile sensing: Fine-grained air quality monitoring with taxi fleets," Environ. Pollut. **283**, 117065 (2021).
16. M. Kamionka and K. Kasztelnik, "Mobile air quality monitoring: Assessing the spatial representativeness of urban air quality data," Atmosphere **11**, 946 (2020).
17. X. Song *et al.*, "Airpred: A spatio-temporal deep learning framework for air pollution forecasting," J. Environ. Informatics **40**, 91 (2022).
18. Y. Zheng, F. Liu, and H.-P. Hsieh, "U-air: When urban air quality inference meets big data," in *Proceedings of the 19th ACM SIGKDD,* (2013), pp. 1436–1444.
19. Y. Miao, W. C. Porter, K. Schwabe, and J. LeComte-Hinely, "Evaluating health outcome metrics and their connections to air pollution and vulnerability in southern california's coachella valley," Sci. The Total. Environ. **821**, 153255 (2022).
20. Y. Shu, J. Hu, S. Zhang, *et al.*, "Analysis of the air pollution reduction and climate change mitigation effects of the three-year action plan for blue skies on the "2+26" cities in china," J. Environ. Manag. **317**, 115455 (2022).
21. S. Sun, L. Sun, G. Liu, *et al.*, "Developing a vehicle emission inventory with high temporal-spatial resolution in tianjin, china," Sci. The Total. Environ. **776**, 145873 (2021).
22. H. Wang, X. He, X. Liang, *et al.*, "Health benefits of on-road transportation pollution control programs in china," Proc. National Acad. Sci. **117**, 25370 (2020).
23. P. Wanke, Z. Chen, X. Zheng, and J. Antunes, "Sustainability efficiency and carbon inequality of the chinese transportation system: A robust bayesian stochastic frontier analysis," J. Environ. Manag. **260**, 110163 (2020).
24. E. J. Westmoreland, N. Carslaw, D. C. Carslaw *et al.*, "Analysis of air quality within a street canyon using statistical and dispersion modelling techniques," J. Environ. Monit. **9**, 657 (2007).
25. Y. Zhou and J. I. Levy, "Impact of traffic-related air pollution on urban health," Curr. Environ. Health Reports **5**, 323 (2018).
26. M. Cai, H. Yin, M. Xie, and V. J. Gan, "Investigating the impact of urban morphology on air quality using data-driven models," Sci. Total. Environ. **776**, 145964 (2021).
27. J. Kim, J. J. Baik, and R. D. Sharman, "Evaluation of urban air quality modeling systems for predicting traffic-related air pollution," Atmospheric Environ. **224**, 117222 (2020).
28. G. Santos and H. Behrendt, "Using transportation models to estimate emissions and air quality impacts," Transp. Policy **102**, 1 (2021).
29. D. Fecht, P. Fischer, L. Fortunato, *et al.*, "Associations between air pollution and socioeconomic characteristics, ethnicity and age profile of neighbourhoods in england and the netherlands," Environ. Pollut. **198**, 201 (2015).

Measuring and mapping patterns of space utilization for characterizing human mobility in epidemiological contexts using a novel digital data set

Daniel T Citron,[1,†] Shankar Iyer,[2] Daniel J Weiss[3], Robert C Reiner[4], David L Smith[4]

[1]*Department of Population Health, NYU Grossman School of Medicine, NY, NY, USA*
[2]*Meta, Inc., Menlo Park, CA, USA*
[3]*Curtin University, School of Population Health, Bentley, WA, Australia*
[4]*Telethon Kids Institute, Nedlands, WA, Australia*
[5]*U. of Washington, Institute of Health Metrics and Evaluation, Seattle, WA, USA*
[6]*U. of Washington, Department of Health Metrics Sciences, Seattle, WA, USA*
**daniel.t.citron@gmail.com*

BACKGROUND

Data sets which describe human movement patterns are crucial for understanding spatial patterns of disease transmission. Digitally-collected user location history data offers an unprecedented opportunity to quantify patterns of movement behavior at a broad scale. Many early digital datasets describing movement patterns have been formatted as origin-destination matrices, counting the number of users who travel from one location to another, which may be valuable in some contexts but are less useful for modeling exposure risk where transmission depends on the time spent in contact with environmental hazards. In collaboration with the Data for Good initiative at Facebook, we developed a new dataset, Activity Space Maps, which aggregates human mobility data in order to quantify how much time an individual from one location spends in another location.

METHODS

Activity Space Maps are derived from de-identified Facebook mobile app user location history data, sampling users who have opted into Meta's location history data collection services. Each sampled user yields a trajectory of locations, logged as level-12 bing tiles (approximately 5km x 5km patches near the equator) over the past 28 days. Each trajectory is first aggregated in space and time to determine the modal daytime (8AM-6PM) and nighttime (8PM-6AM) tile over each 24-hour period. We interpret the modal nighttime tile as a "home tile" where the individual is likely to have stayed overnight. We exclude users who did not log a nighttime modal location at least 7 out of the past 28 day calculation period. We also exclude all travel that crossed international borders. Next, records are aggregated across people and time: for each home tile, we sum the total number of person-days and person-nights spent at each other visit tile. All home tile locations with fewer than 10 residents are excluded from the dataset. After normalizing, the result are two joint probability distributions of how the residents of each home tile distributed their daytime and nighttime periods at each other visit tile. This aggregation procedure is designed to obfuscate any individual's trajectory and preserve user privacy and anonymity. Activity space maps data are available for non-commercial use by Facebook Data for Good partners.

We generated a set of example maps illustrating the distribution and coverage of Activity Space Maps in Italy, Thailand, and Colombia. Next, we mapped the estimated fraction of time spent at home, as a proxy for how often individuals travel away from their home locations. Lastly, we demonstrated the use of Activity Space Maps in conjunction with a spatial model of malaria transmission in Nigeria[1,2]. Malaria is a vector-borne infectious disease that is only spread through contact with vectors in the environment, meaning that the number of nighttime periods spent in an environment with viable malaria vectors directly informs an individual's exposure risk.

RESULTS

Overview

Figure 1 provides an example set of Activity Space Maps showing the geographical distribution of nighttime periods spent for residents of Rome, Italy, Bangkok, Thailand, and Bogota, Colombia in the month of September, 2021. Most nights are spent at home, but the broader spatial spread of the distribution does show overnight travel throughout each country.

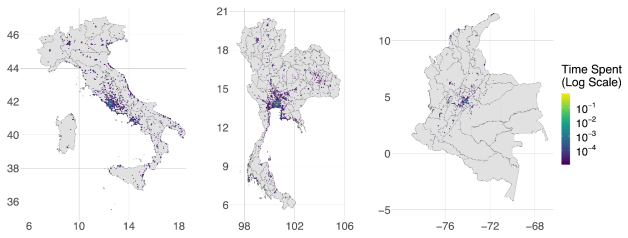

Fig. 1: Activity Space Maps from September 2021. Distribution of

nighttime spent for residents of (left to right) Rome, Italy, Bangkok, Thailand, and Bogota, Colombia.

Figure 2 illustrates the mean fraction of time spent at home for each home tile mapped for each country in Figure 1. Mean fraction of nights at home are 90% in Italy, 92% in Thailand, and 87% in Colombia. Mean fractions of days at home are 68% in Italy, 76% in Thailand, and 74% in Colombia.

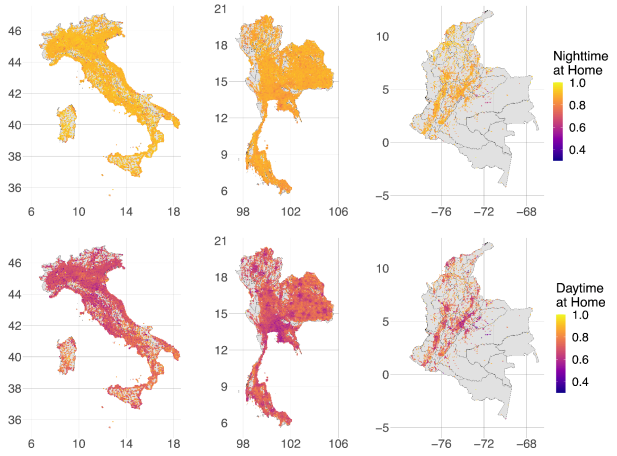

Fig. 2: Fraction of nighttime and daytime periods spent at home for each plotted tile.

Spatial Epidemiology

Malaria is a vector-borne parasitic disease. Exposure risk of malaria is highly dependent on local ecological conditions, and whether or not one travels to and spends time in a location with infectious vector mosquitoes is important for assessing exposure risk. This arises in travel between urban and rural areas, where people living in urban areas are less likely to be at risk of malaria exposure unless they travel to rural areas.

We used Activity Space Maps to perform a source-sink analysis of spatial malaria transmission in the geographical area around Lagos, Nigeria. Using a mathematical model, we combined spatial malaria prevalence estimates from the Malaria Atlas Project[3] (Figure 3, upper left) and Activity Space Maps (Figure 3, upper right) to estimate the relative disease burdens of malaria transmission risk which occurs while traveling away from home (Travel Fraction, Figure 3, lower left) and that which occurs locally at home (Local Residual Fraction, Figure 3, lower right). We note that the areas around central Lagos, circled in red, have a lower malaria prevalence than less developed areas further to the north. At the same time, our analysis suggests an elevated rate of malaria attributable to travel and a lower rate of malaria attributable to local transmission.

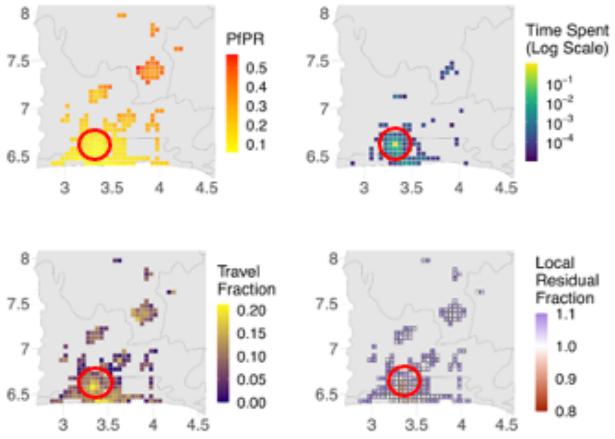

Fig. 3: Example source-sink analysis of malaria transmission near Lagos, Nigeria.

DISCUSSION

We demonstrated how Activity Space Maps augment existing disease models by quantifying the relative risks of exposure at home and away from home. This use case may also be extended to analyzing other environmental exposures such as extreme heat, contaminated water, or even airborne diseases in areas with high population density. This sort of risk evaluation is made possible through Activity Space Maps and would be more difficult to estimate with origin-destination matrix data.

Limitations of this dataset include its likely biased sampling among Facebook mobile app users. Facebook adoption varies significantly across and between different countries, and is more likely accessible to individuals affluent enough to afford a smart phone and access mobile app services. Validating the representativeness of this dataset externally is an important next step in its development.

Future development of this dataset will also include leveraging geostatistical tools for imputing missing Activity Space Maps data in undersampled areas. Using spatial covariates related to demographic, socioeconomic, and environmental data, it will be possible to impute a more globally representative version of these data for applications in resource-limited settings.

REFERENCES

1. Ruktanonchai NW, et al. Identifying malaria transmission foci for elimination using human mobility data. PLoS Comput Biol. 2016;12:e1004846.
2. Citron, D.T., Guerra, C.A., Dolgert, A.J., Wu, S.L., Henry, J.M. and Smith, D.L., 2021. Comparing metapopulation dynamics of infectious diseases under different models of human movement. PNAS, 118(18).
3. Gething PW, Patil AP, Smith DL, Guerra CA, Elyazar IR, Johnston GL, et al. A new world malaria map: Plasmodium falciparum endemicity in 2010. Malar J. 2011;10:378.

Decomposing International Migration Flows with Mobile Trace Data: Trends, Seasonality, and Shocks

WENLAN ZHANG[1,*] AND DOUGLAS LEASURE[1]

[1] Leverhulme Centre for Demographic Science, University of Oxford, 242-43 Park End Street, Oxford, OX1 1JD
[*] wenlan.zhang@demography.ox.ac.uk

1. INTRODUCTION

Understanding international migration is essential to studies of population dynamics, socio-economic change, and global connectivity. However, traditional migration statistics are often annual, delayed, and inconsistent across countries, limiting their ability to capture short-term or cyclical patterns such as student flows or pandemic-driven shifts [1, 2].

Digital trace data, particularly from mobile phones, offer near–real-time insights into human mobility, but have seen limited use in international migration research due to data fragmentation, roaming issues, and lack of standardisation [3].

The Meta International Migration Dataset [4] addresses these challenges by providing harmonised, monthly estimates of cross-border flows across 181 countries (2019–2022), derived from aggregated, privacy-protected Facebook mobility data and corrected for demographic biases.

In this work, we apply time-series decomposition techniques to this dataset to uncover seasonal migration patterns. Our results illustrate how mobile-derived data can capture temporal dynamics often missed in conventional administrative sources.

2. DATA

A. Meta Global Migration Dataset

Our analysis uses the *International Migration Flows* dataset released by Meta's Data for Good initiative [4], which provides monthly estimates of international migration inflows across over 32,000 country pairs in 181 countries from January 2019 to December 2022. Constructed from aggregated, privacy-preserving data on Facebook users, the dataset estimates country of residence based on IP geolocation and self-reported location, applying a segment-based algorithm to identify long-term residence spells while filtering out short-term or noisy movements. A user is defined as a migrant if they reside in one country for the majority of a 12-month period and in another country for the majority of the subsequent 12 months, consistent with the UN definition of long-term international migration [5, 6]. These transitions are aggregated by origin–destination and month, with calibration weights applied to correct for Facebook penetration and demographic biases. The result is a globally harmonised, high-frequency digital trace dataset of estimated total migration flows, made publicly available via the Humanitarian Data Exchange [7].

B. Eurostat Migration Statistics

To contextualise patterns observed in the Meta dataset, we incorporate official statistics from Eurostat, the statistical office of the European Union. Eurostat provides annual data on international migration and mobile students, disaggregated by factors such as age, citizenship, and reason for migration. Although these data lack the temporal granularity of the Meta dataset, they offer valuable structural insights and serve as a benchmark for assessing cross-national comparability. In particular, information on student mobility helps to interpret seasonal regularities in migration patterns observed in digitally derived data.

3. METHODOLOGY

To analyse cyclical dynamics in international migration flows, we applied *Seasonal-Trend decomposition using LOESS (STL)* to monthly time series of migration inflows. STL is a non-parametric technique that decomposes a time series Y_t into three additive components:

$$Y_t = T_t + S_t + R_t \tag{1}$$

where T_t represents long-term trends, S_t captures recurring seasonal patterns, and R_t accounts for residual fluctuations including irregular shocks or noise [8].

To quantify the contribution of each component, we compute normalised variance-based metrics:

$$\text{Trend strength} = \frac{\text{Var}(T_t)}{\text{Var}(T_t) + \text{Var}(S_t) + \text{Var}(R_t)} \tag{2}$$

$$\text{Seasonal strength} = \frac{\text{Var}(S_t)}{\text{Var}(T_t) + \text{Var}(S_t) + \text{Var}(R_t)} \tag{3}$$

$$\text{Residual strength} = \frac{\text{Var}(R_t)}{\text{Var}(T_t) + \text{Var}(S_t) + \text{Var}(R_t)} \tag{4}$$

These strengths sum to one, allowing for direct comparison across countries. A high seasonal strength reflects strong calendar-based regularities, while high trend strength indicates structural or pandemic-related shifts. Residual strength captures short-term noise or event-driven anomalies. This decomposition enables systematic comparison of the temporal structure of migration flows across different national contexts.

4. PRELIMINARY RESULT

We applied STL decomposition to monthly international migration inflow time series across multiple countries from 2019 to 2023, yielding three components: long-term trend (T_t), seasonal variation (S_t), and residuals or anomalies (R_t). Figure 1 illustrates these components for a selection of eight countries.

To quantify the relative importance of each component, we computed variance-based strength metrics for trend, seasonal, and residual signals (see Equations 2–4 in Methods). Figure 2 presents the distribution of component strengths across all countries. These metrics reveal substantial cross-country heterogeneity: in some cases, trend dynamics dominate, suggesting structural or policy-driven migration changes, while in others, seasonal patterns are more prominent, particularly where academic or labour-related cycles influence mobility. Residuals are generally limited but exhibit spikes during episodes of geopolitical or economic disruption.

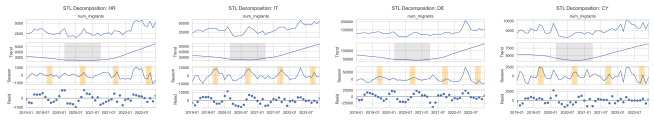

(a) Country HR **(b)** Country IT **(c)** Country DE **(d)** Country CY

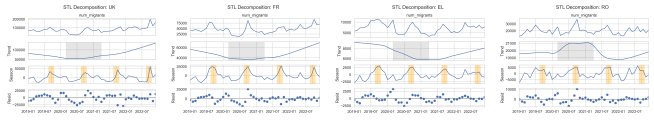

(e) Country UK **(f)** Country FR **(g)** Country EL **(h)** Country RO

Fig. 1. STL decomposition of migration inflows for selected countries. The COVID-19 period is highlighted in grey (trend), and September–October in orange (seasonality).

A. Trend Dynamics: Impact of the COVID-19 Pandemic

Across most countries, the trend component exhibits a sharp decline in migration inflows beginning in early 2020, aligning with the onset of the COVID-19 pandemic and the implementation of global travel restrictions (shaded in grey from 1 March 2020 to 30 March 2021). Romania (RO), however, deviates from this pattern, showing a sustained increase in migration between June 2020 and June 2021. This divergence may be attributed to Romania's relatively lighter travel restrictions during the summer of 2020 compared to Western Europe. Given Romania hosts one of the largest diasporas in the EU, the rise in inflows may reflect temporary or permanent returns motivated by family reunification, holidays, or employment disruptions abroad—particularly in sectors like hospitality, agriculture, or construction. This atypical behaviour also resulted in a higher share of unexplained variation: 43% of the time series variance was captured by the residual component, complicating the interpretability of the trend signal.

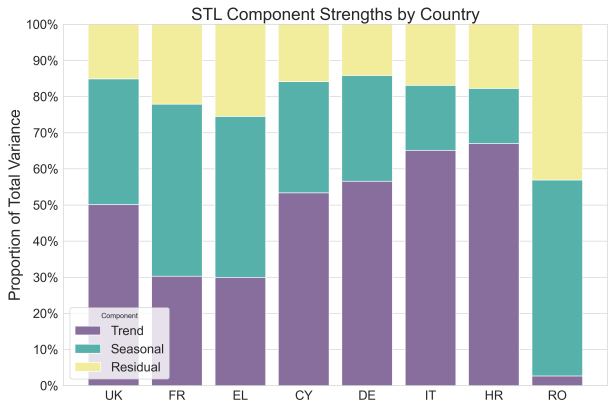

Fig. 2. STL trend components for selected countries, showing the relative contributions of trend, seasonality, and residual components. The bar chart illustrates how these components vary by country.

B. Seasonality and Academic Mobility

To explore cyclical migration behaviour, we computed the average monthly share of migration inflows for each country, aggregated over the period 2019–2023. Figure 3 reveals distinct seasonal peaks in many countries, typically aligned with academic calendars and labour demand cycles. Notably, countries such as France (FR), the United Kingdom (UK), and Cyprus (CY) show clear peaks in September and October, consistent with the start of university semesters and international student arrivals.

These patterns align with official education statistics. The UK accounts for 20% of global international student flows (OECD, 2022), while Cyprus (CY) has the third-highest share of international students in the EU (20%) and the highest enrolled in bachelor's programmes (28%) (Statistics Explained, 2024). France (FR) hosts 16% of EU-based international tertiary students, many from African countries. These inflows likely contribute to the strong seasonal peaks observed in these countries. Some cases deviate from this pattern: Greece (EL) shows a high seasonal strength (45%) despite a low international student rate (3%), possibly due to overlapping academic and tourism-driven migration from July to November. Conversely, Germany (DE), which hosts 24% of international students in the EU, shows a lower seasonal strength—likely because the large refugee-related spike in May–June 2022 increased the trend component, dampening the seasonal signal.

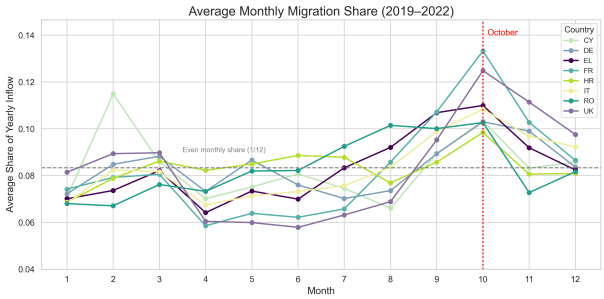

Fig. 3. Monthly migration inflow shares averaged across four years (2019–2023).

C. Residuals and Geopolitical Disruption

The residual component (R_t) of the STL decomposition captures irregular, non-cyclical fluctuations not explained by trend or seasonal structure. In most countries, residual variation is relatively low, suggesting stable and predictable migration patterns. The UK and Cyprus (CY) show particularly low residual strength (15%), indicating inflows closely tied to recurring seasonal drivers like academic calendars. Germany (DE) also exhibits a low residual share (14%), but this masks a major irregular event: the sharp increase in migration in May 2022 due to Ukrainian refugee arrivals. STL appears to have absorbed this sustained inflow into the trend and seasonal components, illustrating a limitation of the method that prolonged shocks may be smoothed into structural signals, reducing their visibility as anomalies. In contrast, Romania (RO) stands out with the highest residual strength (43%), suggesting greater volatility potentially linked to idiosyncratic or short-term migration drivers.

REFERENCES

1. F. Rampazzo, J. Bijak, A. Vitali, *et al.*, Int. Migr. Rev. **59**, 119 (2024).
2. OECD, "International migration outlook," (2023).
3. P. Blanchard and S. Rubrichi, arXiv (2024). Https://arxiv.org/abs/2406.15216.
4. G. Chi, G. J. Abel, D. Johnston, *et al.*, Proc. National Acad. Sci. **122**, e2409418122 (2025).
5. United Nations Statistical Division, *Recommendations on Statistics of International Migration, Revision 1* (United Nations Publications, 1998).
6. International Organization for Migration (IOM), "Fundamentals of migration," (2021).
7. Meta Data for Good, "International migration flows - meta," (2024).
8. R. B. Cleveland, W. S. Cleveland, J. E. McRae, and I. J. Terpenning, J. Off. Stat. **6**, 3 (1990).

How People Navigate in Crowds Indoors: Evidence from Trajectories in UCL PEARL

GUANGSHENG DONG[1,12], YOUNAN HUANG[1], HANZHEN FANG[1], XIANGHUI ZHANG[1], BEKY STODDART[2], JESSICA IN[3], GEORGE PROFENZA[4], STEVE MAYO[2], RALPH STOKELD[2], BARBARA PIZZILEO[2],SHANE GILL[2], MATTHEW VILE[2], BEN PAVELEY[2], ISABELLE SJOVALL[5], UEMEE JUNG[6], JULIETTE DUPERTUYS[6], EMRE YAVUZ[6], SAM MCELHINNEY[7], PAUL WEBSTER[8], STEPHEN LAW[9], ANTONIA HAMILTON[10], NICK TYLER[2,11], STEPHEN GAGE[3], SEAN HANNA[3], HUGO SPIERS[6], FIONA ZISCH[3], TAO CHENG[1,*]

[1]*SpaceTimeLab for Big Data Analytics, Department of Civil, Environmental and Geomatic Engineering, University College London, London , UK*
[2]*Person Environment Activity Research Laboratory, University College London, London, UK*
[3]*The Bartlett School of Architecture, University College London, London, UK*
[4]*Opificio, London, UK*
[5]*Research institutes of Sweden, Sweden*
[6]*Institute of Behavioural Neuroscience, Department of Experimental Psychology, Division of Psychology and Language Sciences, UCL, London, UK*
[7]*University for the Creative Arts, Canterbury, UK*
[8]*Ubisense Limited, Cambridge, UK*
[9]*Department of Geography, University College London, London, UK*
[10]*Institute of Cognitive Neuroscience, University College London, London, UK*
[11]*Pedestrian Accessibility Movement Environment Laboratory, Department of Civil, Environmental and Geomatic Engineering, University College London, London, UK*
[12]*State Key Laboratory of Information Engineering in Surveying, Mapping and Remote Sensing, Wuhan University, Wuhan, China*

1. INTRODUCTION

Outdoor navigation has been revolutionized by GPS-based navigation, yet most indoor environments — such as train stations, shopping malls, and hospitals — remain poorly served by comparable localization infrastructure. Regular occupants may develop mental maps of familiar interiors, but first-time visitors to large or complex buildings often experience disorientation, leading to delays, frustration, and, in critical settings, elevated safety risks. Effective indoor navigation not only enhances everyday convenience but also underpins efficient healthcare delivery, smooth intermodal transportation, focused learning and work, and rapid, safe evacuation during emergencies.

A growing body of work shows that outdoor GPS-derived mobility patterns reliably reflect users' psychological states and personality traits. Zhang et al. (2022) demonstrated a bidirectional association between movement features (e.g., daily travel radius) and depression severity. Meyerhoff et al. (2021) found that changes in GPS trajectories, call logs, and app-use patterns corresponded with fluctuations in depression, generalized anxiety, and social anxiety. Balliu et al. (2024) used home-stay duration and location-entropy variability to predict depression severity up to three weeks in advance. In the personality domain, Liu et al. (2024) introduced a "Trajectory Trait Scale" by analyzing over one million car-GPS trips to quantify extraversion, openness, neuroticism, and conscientiousness.

Despite these advances in outdoor settings, analogous studies of indoor spatial behavior are scarce, primarily due to challenges in data collection, and privacy concerns. To address this gap, we have built a full-scale, configurable indoor environment within Person–Environment–Activity Research Laboratory (PEARL) and recruited over 100 participants. Using a high-precision indoor-positioning system, we recorded each participant's trajectory as they completed a navigation task, alongside demographic surveys and standardized psychological assessments. Through systematic analysis of how indoor movement patterns vary by age, gender, neurodiversity, and other individual attributes, this study seeks to elucidate the relationships between navigation behaviors and psychological characteristics, thereby generating insights for the assessment and design of user-friendly architectural and urban environments.

2. Data and Methods

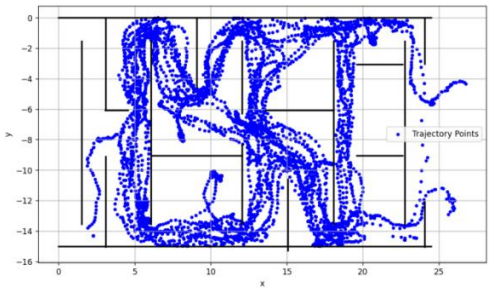

Fig. 1. Example of indoor trajectories.

We conducted our experiments in PEARL, a 40,000 m^3 adaptable space replicating various real-world environments (e.g., high street, park, station, ward, supermarket) in UCL. We recruited 110 participants, each of whom wore an indoor-positioning device and completed a predefined sequence of navigation tasks within the experimental layout. From these trials, we captured indoor trajectories at 5 Hz. Figure 1 shows the experimental layout and sample paths. After navigation, participants completed a survey collecting demographic information — age, gender, navigation ability, GPS reliance, as well as their psychological responses before and after the task (e.g., curiosity and enjoyment).

To prepare the raw trajectory data for analysis, we developed a bespoke preprocessing pipeline that accounts for the high sampling rate and spatial confinement of indoor tracking. Steps included: digital mapping of the facility; spatio-temporal registration; noise filtering; identification and

* tao.cheng@ucl.ac.uk

removal of redundant and spurious points; and trajectory segmentation. From the cleaned data, we extracted movement metrics — e.g., total distance, range, stayarea count, revisit index, location variance, radius of gyration, location entropy, and average deviation from line — to quantify each participant's movement intensity, spatial distribution, behavioral diversity, and variability. We then applied ST-DBSCAN to each individual's trajectory to detect stay points, and clustered all participants' stay points to identify common hotspots. This reduced each trajectory to a sequence of transitions between hotspots, enabling computation of revisit counts and other movement metrics. Finally, we assessed associations between movement metrics and demographic measures using Spearman's rank correlation.

3. Results

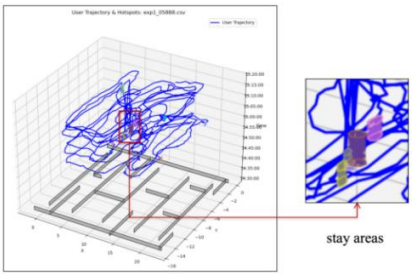

Fig. 2. Stay points identified in an individual's trajectory.

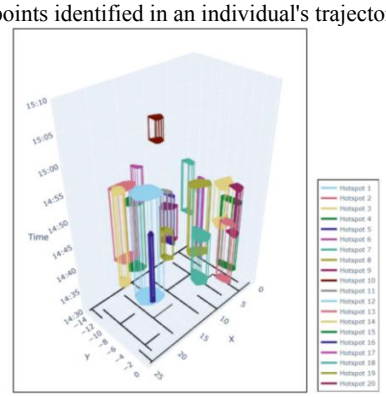

Fig. 3. Hotspots identified from group trajectories.

Spatio - temporal clustering can identify users' behavior patterns. Fig. 2 illustrates a representative trajectory with overlaid stay areas detected by ST-DBSCAN: the participant exhibited high - density looping and overlapping paths, particularly forming prolonged dwell clusters in the central and right - hand zones of the layout. The close - up visualization of these stay areas confirms that our spatio - temporal clustering robustly captures sustained pauses. Aggregating every participant's stay points and reclustering yielded twenty group hotspots (Fig. 3), predominantly located along the central corridor and in corner niches — many co - located with interactive installations — indicating collectively attractive regions.

Temporal patterns within several hotspots further suggest that participants concentrated their visits during specific task intervals, likely reflecting the prescribed navigation sequence. By encoding each trajectory as a sequence of transitions between these hotspots, we computed revisit frequencies: most participants revisited at least one hotspot more than twice, with marked inter-individual variation. For instance, some oscillated frequently between hotspot 3 and hotspot 7, exhibiting rapid alternation, whereas others demonstrated spatial attachment by dwelling at a single hotspot for extended periods.

We then examined how these movement patterns related to demographic measures using Spearman's rank correlations (Fig. 4). Notably, GPS reliance and experiment enjoyment(Enjoyment_exp) displayed inverse associations with movement metrics, whereas self-reported navigation ability consistently aligned with more efficient trajectories—fewer stay areas and broader range. Specific correlations of interest included a positive relationship between experiment enjoyment (Enjoyment_exp) and stay areas count (r=0.204), indicating that participants who reported greater enjoyment visited more distinct zones; a negative correlation between interest and average duration (r=-0.277), suggesting that more interested individuals paused for shorter intervals; a negative association between crowded and total distance (r=-0.306), implying that crowding curtailed exploration; and a negative link between curiosity and speed variability(speed SD) (r=-0.231), reflecting steadier speeds among more curious participants. Measures of anxiety (Anxiety_party/ Stress_nav) exhibited an inverse relationship with movement range, aligning with findings from outdoor trajectory research.These findings underscore how individual traits and subjective experiences shape indoor navigation behaviors.

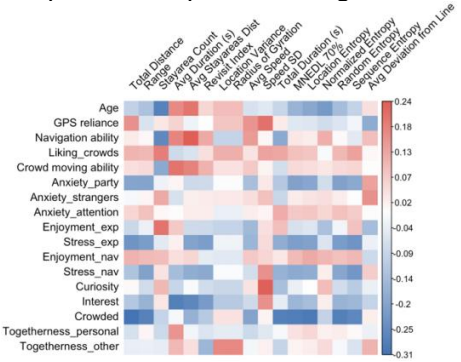

Fig. 4. Spearman correlation matrix between movement metrics and psychological, behavioural, and demographic measures.

References

1. Zhang Y. et al. (2022). Longitudinal relationships between depressive symptom severity and phone-measured mobility. JMIR Mental Health, 9(3): e34898.
2. Meyerhoff, J. et al. (2021). Evaluation of changes in depression, anxiety, and social anxiety using smartphone sensor features: longitudinal cohort study. Journal of medical Internet research, 23(9), e22844.
3. Balliu B. et al. (2024). Personalized mood prediction from smartphone-collected behavior patterns. NPJ Digital Medicine, 7(1): 49.
4. Liu Y. et al. (2024). Profiling mobility patterns and driving behaviors via trajectory trait. The Innovation Geoscience, Article 100114-1.

Enhancing mobile network data with GPS and satellite imagery for accurate population detection

MARTA ALONSO[1], MIGUEL BAENA[1*], AN VO QUANG[2], OLIVA G. CANTÚ-ROS[1], ANA BURGIN[2], ANTOINE MASSE[2]

[1]*Nommon Solutions and Technologies, Plaza Carlos Trías Bertrán 4, 2ª planta, 28020 Madrid, Spain*
[2]*IGN FI, rue Biscornet 7, 75012 Paris, France*
*miguel.baena@nommon.es

Received 1 July 2025; revised XX Month, XXXX; accepted XX Month XXXX; posted XX Month XXXX (Doc. ID XXXXX); published XX Month XXXX

1. INTRODUCTION

In recent years, the need for precise dynamic population mapping solutions became more and more important to convey with the changing patterns of spatial dynamics. Traditional population mapping often relies on census data, which lacks temporal and behavioural information [1]. Distinguishing between indoor and outdoor activities allows for real-time and situation-specific population density estimates, improving accuracy. This information could be key for experts in urban planning, emergency management, public health crises like pandemics [2], and other aspects of smart cities, which rely on real-time data to manage traffic, public transport, and services. In a context of increasing Innovative Air Mobility (IAM) operations, dynamic data of pedestrians provide key insights to plan flights in order to reduce ground risk and other impacts of this new mobility, such as noise, visual pollution and privacy concerns [3]. Mobile Network Data (MND), passively collected from mobile antennas have proven to be highly effective in accurately capturing population presence and mobility patterns. By enhancing this rich data source with GPS data for spatial accuracy and validating the results with satellite imagery pedestrians detected, we provide a procedure for indoor and outdoor population detection.

2. METHODOLOGY

A. Population presence calculated with MND and GPS

Mobile network data has the disadvantage of low spatial resolution; location is subject to the coverage areas of the mobile network antennas, which vary from some hundreds of metres to kilometres. GPS data is spatially highly precise; however, the temporal resolution and sample size is at the moment lower. Then, several months of GPS data in the area are used to calculate the probability location distribution of the population, for different times of the day, within the MND coverage areas [4].

The differentiation between indoor and outdoor activities is performed by considering the buildings' vectorised shape and comparing it with the people located in each zone of the high-resolution zoning system. For this case, we considered a square grid with a 15 m side.

While the integration of GPS data allows for the creation of high-resolution probability maps of population presence, our analysis revealed systematic biases. In particular, we observed abnormally high presence probabilities at night, likely stemming from biases inherent in GPS data collected through mobile applications. To address this issue and generate a more accurate spatial distribution of population presence, we implemented both temporal and spatial calibration techniques to adjust the observed probabilities accordingly.

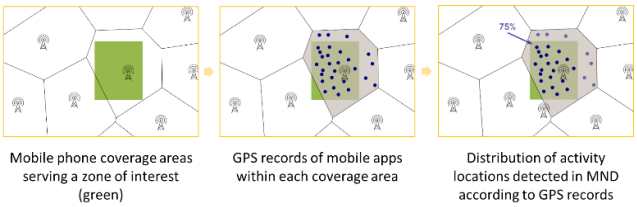

Fig. 1. Description of the process to disaggregate MND population detected using GPS

The adjustment of the dynamic probability maps was carried out through comparison with a ground truth proxy derived from the mobile network indicators themselves. During nighttime hours (00:00–01:00), we used presence indicators associated with the activity type "HOME" as reference, restricting the analysis to purely residential areas where other types of activity can reasonably be considered negligible. Similarly, between 09:00 and 12:00, the same approach was applied in exclusively office areas, with no residential or leisure spaces present. In this case, the ground truth was based on indicators associated with the "WORK" activity. This calibration strategy based on time slots and urban typologies allows us to scale the GPS-derived probability estimates and correct for systematic biases resulting from the dynamic and uneven nature of GPS data coverage.

B. Pedestrian detection with satellite imagery

We developed a remote sensing pipeline capable of detecting individuals from very high-resolution satellite imagery. Our approach leverages Pléiades Neo images (15 cm resolution) acquired over central Madrid in February and July 2023, allowing us to evaluate seasonal effects on detectability, particularly due to variations in shadow length and occlusions caused by urban structures. We conducted manual annotation of 4,500+ pedestrian instances across 21 orthoimages from the Madrid council geoportal, based on visual shadow identification.

For model training, we used a High-Resolution Network (HRNet) architecture, chosen for its ability to preserve fine spatial features through parallel multi-scale representations. The model was trained on the annotated orthophotos and

applied to the two Pléiades Neo scenes. To improve precision, we integrated Urban Atlas land cover data to spatially filter detections in likely pedestrian areas and applied a post-processing pipeline to clean artifacts and extract geolocated point features.

The detection model achieved an estimated 70.41% detection rate for the February image and 80.86% for the July image, after post-processing. This performance difference reflects the seasonal trade-off: winter scenes provide longer and more visible shadows, which are helpful for detection, but are also more affected by occlusions caused by tall buildings and narrow streets. In contrast, the July image showed less occlusion, leading to better visibility in open areas, although shorter and softer shadows slightly reduced contrast.

Fig. 2. Detected pedestrians (yellow points) on Pléiades Neo imagery at 1/122 scale in February 2023

False positives were primarily associated with urban clutter such as rooftops, vehicles, and street furniture. These were significantly reduced by integrating Urban Atlas land cover data and applying a morphological cleaning step during post-processing.

3. VALIDATION

The validation process, based on the direct comparison between presence indicators and satellite-derived person counts, were specifically conducted in selected urban squares of Madrid and shows a reasonable level of correlation with the MND+GPS-based presence estimates. Urban squares can be considered as areas where the intrinsic limitations of both data sources are less pronounced.

Table 1 summarises the results obtained employing both methodologies. We used the February data for validation, despite having less detection rate, the capability for detection is better due to longer pedestrian shadows. By comparing in urban squares, the area not visible due to building shadowing is minimised.

Table 1. Comparison between MND+GPS and satellite imagery for February 2023

Urban square	MND+GPS	Satellite imagery
Sol	392	201
Callao	72	70
Plaza Mayor	266	221

The interpretation of results must, however, consider key sources of bias. On the one hand, the typology of mobile applications generating GPS traces is unknown, which may lead to an overrepresentation of certain behaviours, such as outdoor movements captured by navigation or fitness apps, potentially skewing the estimated probability of presence. On the other hand, Earth Observation (EO) data is subject to shadowing effects that systematically underestimate presence levels, with a degree of distortion that varies depending on the urban morphology and season. Furthermore, EO imagery is typically available for a single moment during the day, limiting the temporal granularity of the validation and complicating the assessment of dynamic population patterns throughout the day.

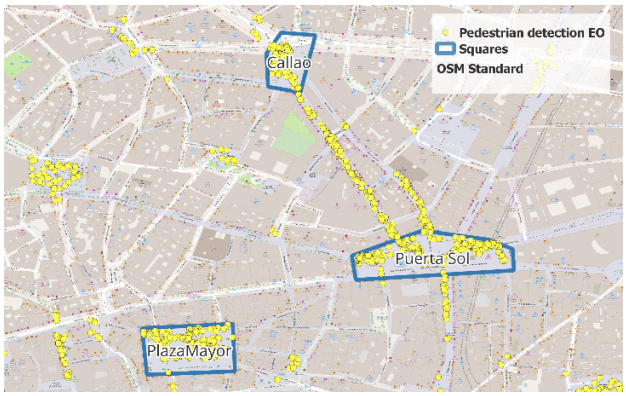

Fig. 3. Pedestrian detected at 11 am in February 2023

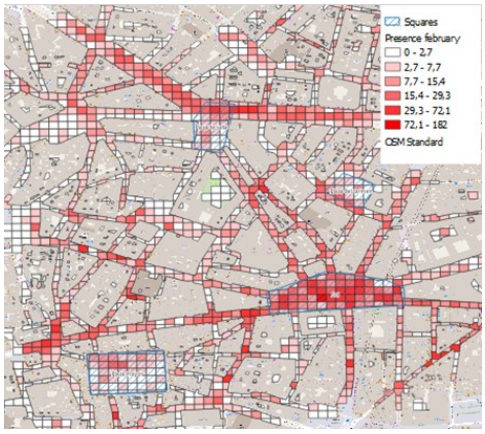

Fig. 4. Presence indicators obtained with the methodology at 11 am in February 2023

Acknowledgment. This project has received funding from the SESAR 3 Joint Undertaking (SESAR 3 JU) under grant agreement No 101114858. The JU receives support from the European Union's Horizon Europe research and innovation programme and the SESAR 3 JU members other than the Union.

References

1. Picornell, M., Ruiz, T., Borge, R., García-Albertos, P., de la Paz, D., & Lumbreras, J. (2019). Population dynamics based on mobile phone data to improve air pollution exposure assessments. Journal of Exposure Science & Environmental Epidemiology, 29(2), 278-291.
2. Romanillos, G., García-Palomares, J. C., Moya-Gómez, B., Gutiérrez, J., Torres, J., López, M., Cantú-Ros, O.G., & Herranz, R. (2021). The city turned off: Urban dynamics during the COVID-19 pandemic based on mobile phone data. Applied Geography, 134, 102524.
3. Krstić Simić, T., Ganić, E., Mirković, B., Baena, M., LeGriffon, I., & Barrado, C. (2024). U-Space Social and Environmental Performance Indicators. Drones, 8(10), 580.
4. Gonzalez AB, Burrieza-Gal'an J, Diaz JJ, de Castro IP, Wilby MR, Cantu-Ros OG. Using App Usage Data from Mobile Devices to Improve Activity-based Travel Demand Models. IEEE Transactions on Big Data. 2024 Feb 1(01):1-1.

The (un)predictability of the earthquake-induced mobility

Stijn van Hannen[1], Bilgeçağ Aydoğdu[1,*], Subhi Güneş[2], Semiha Nur Yağcıklı[2], and Albert Ali Salah[1]

[1] Utrecht University, The Netherlands
[2] Turkcell Technology, Türkiye
[*] b.aydogdu@uu.nl

1. INTRODUCTION

Earthquakes are sudden-onset and difficult-to-predict disasters that continue to cause mass displacement, damage, and loss of life. Despite the difficulty in predicting the disaster itself, models that can extrapolate crisis-induced mobility in real time from historical mobility data can provide a crucial window of preparation for governments and humanitarian organizations. In this study, we investigate the (un)predictability of earthquake-induced mobility following the devastating earthquake that hit southeastern Turkey and northern Syria in 2023, causing large numbers of casualties and widespread displacement from the region.

We propose a graph convolutional network model for mobility prediction in disaster contexts. Our analysis reveals that while the model effectively predicts mobility patterns before the earthquake, it struggles significantly with predictions during the earthquake period. We examine the spatial distribution of areas where the model exhibits the greatest prediction errors both before and during the earthquake to understand the spatial correlations of the earthquake's impact on prediction accuracy.

2. METHODOLOGY

In this research, the CDR datasets we use are collected between January 1 and March 15, 2023, and prepared in a fine-grained mobility format within a similar anonymization and aggregation framework [1] to previous data challenges such as Data for Development (D4D) [2] and Data for Refugees (D4R) challenges [3]. The CDR sample includes around 127,700 individuals in the affected region and 500,000 individuals across Turkey. Nearly half of the individuals are Syrian, and the rest are Turkish users. Therefore, the crisis-induced mobility represents both the patterns of secondary displacement of Syrians and first displacement of Turkish populations. The data is binned at an hourly level and aggregated at cell tower level to represent the inflows and outflows to each cell in each hour [4]. Additionally, we incorporate road network data to construct our novel Road Connectivity graphs based on centrality measures of transportation infrastructure.

We use Graph Convolutional Networks (GCNs) to predict inflows of people at an hourly and cell tower level. Our model predicts one hour ahead using the past five hours as input features. To calculate the inflows on CDR, we measured the trips for each person and aggregated the number of incoming people for each cell tower. The graph is constructed using the approximate cell tower locations as nodes. We use a novel way to represent the edges using the road network connectivity strength, which is the sum of closeness centrality scores of road junctions that connect between neighboring cell tower areas, as the weights between neighboring cell towers [4]. The GCNs have emerged as powerful tools in mobility research [5] [6], but these models typically struggle to capture mobility patterns induced by unexpected disruptive events such as earthquakes. To address this limitation, some studies have attempted to incorporate disaster-specific modules into their architectures, such as Tang et al. [7], who used GCNs to predict flood-induced mobility in the city ZhengZhou. In their work, they [7] introduced a novel heterogeneity component, which captures global meta-information about nodes in the graph based on centrality measures like degree, closeness, and betweenness to predict the irregular mobility patterns during floods. In our research, we refer to this type of additional component as the "disaster component", which encompasses contextual information about the disrupting event and regional characteristics that influence how populations respond to the crisis.

Earthquakes present particularly challenging prediction scenarios because the spatial patterns of displacement can vary significantly even among neighboring regions. The level of damage caused by the earthquake in an area depends on diverse local factors including terrain type, soil composition, and infrastructure quality, creating complex spatial heterogeneity that differs from other disaster types. As a result, the mobility ensuing the earthquake depends not only on the damage in the residential areas, but also on the damage in the existing roads, and in the other potential important places, where the individual can seek shelter. Our proposed model addresses these challenges by incorporating both a disaster component adapted from the work of Tang and others [7], and the novel Road Connectivity graph For capturing temporal dependencies in the displacement patterns, we employ a Gated Recurrent Unit (GRU) architecture that can effectively model the sequential nature of mobility changes following the earthquake event. We use 80% of pre-earthquake data for training, with the remaining 20% of pre-earthquake data plus during-earthquake data for validation, and post-earthquake data for testing.[1]

3. RESULTS AND DISCUSSION

Our analysis reveals significant variations in model performance when predicting earthquake-induced mobility patterns. Under normal conditions (pre-earthquake period), the proposed GCN model with Road Connectivity graphs and disaster component achieved strong predictive performance with R^2 score of 0.77 for hourly predictions, substantially outperforming baseline models including ARIMA and Historical Average [4]. During the day of the earthquake, the performance of predictions drop significantly both for proposed model, as well as for baseline models, and various other models we tested with different representations of inter-regional relations.[2] For our proposed model at hourly level we observed R^2 score of 0.47 during the earthquake day.

[1] More details on the methodology, the model, and the training can be found in the thesis of Van Hannen [4]

[2] In addition to road-connectivity, we also tried spatial adjacency, distance and travel time-based graphs to represent inter-regional relations.

We investigated the spatial distribution of performance losses measured by R^2 during the earthquake day (06/02/2023) compared to a baseline day a week before (30/01/2023) to understand the spatial correlates of crisis-specific uncertainty in mobility. Previously, we found out that the existing refugee camps and border areas appeared as significant destinations to the displaced people after the earthquake [8]. In this study, our goal is to test whether the fall in model performance could be understood by such mobility dynamics.

Figure 1 illustrates the spatial distribution of the difference between the R^2 calculated during the day of the earthquake (06/02/2023) and the baseline day. Cells with zero incoming people on either the baseline day or earthquake day were excluded from the spatial analysis and marked as "all zeros". Our analysis show that the GCN-based model struggled to capture the reduced inflows to the city centers after the earthquake, and increased inflows to the rural areas surrounding major cities such as Gaziantep, Şanlıurfa and Malatya. In addition, the model underestimated the inflows of people during the earthquake day to certain areas across the Türkiye-Syria border, such as the border gates. In Figure 2, we show the model predictions for the inflows of people to the areas around the Ceylanpınar border gate (see Figure 1 for the location) in comparison with the actual values. On the earthquake day, the model struggled to accurately predict the earthquake-induced inflows of people into this region, compared to its performance on the baseline day. These findings validate our findings in a previous study [8], where we found out that a portion of Syrian population moved towards the border after the disaster, and the Turkish population moved from urban centers to rural areas as a coping strategy [8].

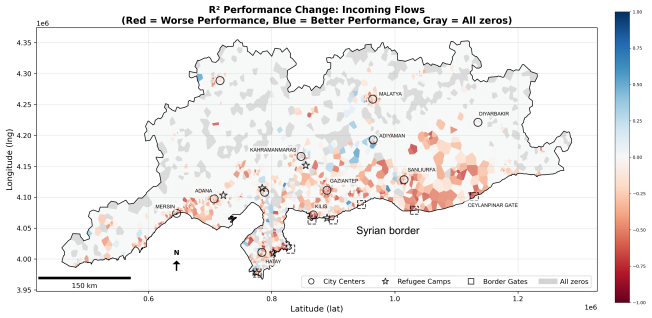

Fig. 1. The figure shows the spatial distribution of the difference of R^2 on the earthquake day compared to the baseline day. We labeled border gates, refugee camps, and city centers for demonstrating the potential contributors to the low performance of the model.

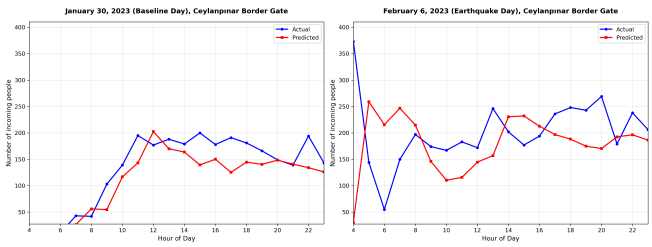

Fig. 2. This figure shows the actual and predicted inflows to the cell towers providing service around Ceylanpınar border gate on the earthquake and baseline days, respectively.

Our study demonstrates that while GCN with disaster component and road connectivity can effectively model routine mobility patterns, but earthquake-induced displacement presents fundamental challenges. The dramatic performance losses during the earthquake period reveals that crisis-induced mobility operates under significantly different dynamics than normal conditions. Our findings highlight that more successful predictive models should incorporate distinct mobility patterns and behaviors of specific sub-populations, such as Syrian refugees, rather than treating all displaced populations as homogeneous. While our models cannot predict the immediate chaos following earthquake onset, understanding these failure modes is important for developing more robust predictive models, and improving humanitarian response planning in disaster-prone regions.

4. CONCLUSION

This study examined the predictability of earthquake-induced mobility using graph convolutional networks with road connectivity and disaster components. While the model performed well under normal conditions, it struggled significantly during the earthquake period, particularly in border regions and rural areas receiving unexpected displacement flows. These findings highlight the fundamental challenges in predicting crisis-induced mobility and suggest that future models must incorporate adaptive mechanisms to account for the shift from routine to crisis-driven displacement patterns.

Funding. This work was funded by the European Union Horizon 2020 Research and Innovation funding Program, project number 870661.

Acknowledgment. The authors thank the European Union for funding this research through the Horizon 2020 program. We also acknowledge Turkcell Technology for providing the mobile phone data that made this analysis possible.

Disclosures. The authors declare no conflicts of interest.

Data Availability Statement. Data underlying the results presented in this paper are not publicly available due to privacy and commercial restrictions but may be obtained from the authors upon reasonable request and subject to appropriate data sharing agreements.

REFERENCES

1. B. Aydogdu, A. A. Salah, O. Ones, and B. Gurbuz, HumMingBird Proj. Deliv. **6** (2021).
2. V. D. Blondel *et al.*, arXiv preprint arXiv:1210.0137 (2012).
3. A. Salah, A. Pentland, B. Lepri, and E. Letouzé, *Guide to Mobile Data Analytics in Refugee Scenarios* (Springer, Cham, Switzerland, 2019).
4. S. van Hannen, "Introducing spatio-temporal graph networks to displacement predictions in the context of the 2023 Turkey-Syria earthquake," Msc thesis, Utrecht University, Utrecht, Netherlands (2025). MSc Artificial Intelligence, Faculty of Science, Department of Information and Computing Sciences.
5. F. Terroso-Sáenz and A. Muñoz, Appl. Intell. **52**, 4144 (2022).
6. W. Wang, J. Chen, Y. Zhang, *et al.*, ACM Trans. on Internet Technol. **21**, 1 (2021).
7. J. Tang, J. Wang, J. Li, *et al.*, Comput. Environ. Urban Syst. **113**, 102160 (2024).
8. B. Aydoğdu, D. Danış, Bilgili, *et al.*, "A novel activity space approach to discover displacement patterns via mobile phone data: An analysis of the 2023 türkiye-syria earthquakes," (2025). Manuscript under review.

Mixture models uncover polycentric urban structure from GPS mobility data

ANDREW RENNINGER[1*], EGOR KOTOV[2], AND CARMEN CABRERA[3]

[1] Centre for Advanced Spatial Analysis, University College London
[2] Geographic Data Science Lab, Liverpool University
[3] Laboratory of Migration and Mobility, Max Planck Institute for Demographic Research
[*] andrew.renninger.12@ucl.ac.uk

1. INTRODUCTION

A central debate in urban theory concerns the benefits of a monocentric agglomeration against the advantages of a polycentric conurbation [1]. In the monocentric model, a central business district concentrates jobs, a thick and liquid labour market where workers and firms and match and rematch without burdensome changes to residential or commercial location. This unified market also enables returns to scale by reducing the need for infrastructure while improving the ease of communication. However, a centralized city also incurs high congestion costs: as cities grow, radial commutes burden transport networks and land rents grow near the core; decentralization becomes a practical necessity beyond a certain metropolitan size, leading to the emergence of hubs and clusters throughout the city. The modern city is thus often characterized by a tension between centralizing forces and decentralizing forces.

Human mobility patterns offer a powerful lens to understand a city's dynamic structure, often revealing complexities hidden in static maps of population or land use. Tracking where and when people travel – their daily moves across the city – yields a map of functional urban connectivity, effectively the "revealed preference" of residents using and experiencing the city. Here we exploit GPS mobility data for urban residents in cities across Germany, France and the United Kingdom to develop a method for understanding how cities balance the forces of centralization and decentralization. In doing so, we find regularities in how people distribute activities in space with implications social mixing, productivity, and sustainability.

2. DATA AND METHODS

Data. We analysed two months of anonymized GPS traces for 15 million residents across 226 functional urban areas in France, Germany, and the United Kingdom. For every device the modal cell between midnight and 6:00AM is its *home*. We use the infostop algorithm to detect discrete visits to locations [], and every distinct stop defines a *trip* of length d_{ij}, the distance between home and stop. We then pool these trips according to neighborhoods, `H3` resolution 7 hexagons ($\approx$ 1 km), so that each neighborhood is assigned a vector of trip distances. Because these length distributions are comprised of trips generated from different processes, which include leisure activities and commutes, we employ finite *mixture models* to separate shorter-range routine mobility from longer-range travel within a single, interpretable framework.

Methods. For each home hexagon we model the empirical distribution of trip lengths with a three-component Gamma mixture

$$f(d) = \sum_{c\in\{p,m,d\}} \pi_c \, \Gamma(d \mid \alpha_c, \beta_c),$$

representing *proximal*, *medial*, and *distal* mobility scales. Mixture parameters $(\pi_c, \alpha_c, \beta_c)$ are obtained by maximum-likelihood expectation–maximization and the components are ordered by increasing mean $\hat{\mu}_c = \alpha_c/\beta_c$. The resulting means $(\hat{\mu}^p, \hat{\mu}^m, \hat{\mu}^d)$ represent the typical trip distance within each component, or scale, while the weights (π^p, π^m, π^d) quantify the share of all trips undertaken in that scale for the neighborhood in question. To understand how these values vary with respect to urban spatial structure, we estimate an urban center for each city by taking the centroid of the busiest location by GPS pings and computing each neighborhood's distance to that centroid. (In alternative specifications we use two-component and four-component models, but three-components offer the best fit by BIC.)

3. RESULTS

We demonstrate our approach on London in Fig. **1a** and **b**, where we can see that as a neighborhood gets farther away from the center—which in London is near Piccadilly Circus—the distal mean $\hat{\mu}^d$ grows, suggesting that these communities need to travel farther to access important services or jobs. The proximal $\hat{\mu}^p$ distances do not systematically vary by distance to the center, suggesting that local activity operates at a similar scale no matter than neighborhood, and medial $\hat{\mu}^m$ distances scale but less so than $\hat{\mu}^d$. We document an interesting fact in London: because our modeling approach is agnostic to what kinds of trips it decomposes into constituent distributions, it detects a distal component in central locations, which represents trips away from the center. These trips constitute smaller proportions for their neighborhoods relative to suburban ones, but it suggests urban residents might gain flexibility and leisure time.

Fitting mixture models to $\sim$ 20,000 neighborhoods across $\sim$ 200 cities in Europe, we see a number of important regularities, shown in Fig. **1c**. Across all cities, as in London, $\hat{\mu}^p$ is typically location invariant while $\hat{\mu}^m$ and $\hat{\mu}^d$ vary more. The insets in Fig. **1c** show the slope of the best fit line for each city, with brigher colors corresponding to larger cities; we see a clear stratification where larger cities have both lower slopes and lower intercepts. This indicates that smaller cities are more dependent on their urban centers while larger cities fragment into subcenters.

An advantage of this method is that it lets us assign individual trips to components and examine the composition of them. Assessing visit locations using FourSquare points of interest confirms intuitions: the proximal component is dominated by residential areas, medial by leisure and consumption amenities, and distal by employment centers. Assigning users socioeconomic strata using neighborhood income data, we compute residents' exposure to different strata. We show 50% of intergroup interaction across cities comes from distal components, with implications for social contact given remote work trends; if centralized work declines, intergroup exposure may decrease.

REFERENCES

1. A. Bertaud, Town Reg. Plan. **79**, 2 (2021).

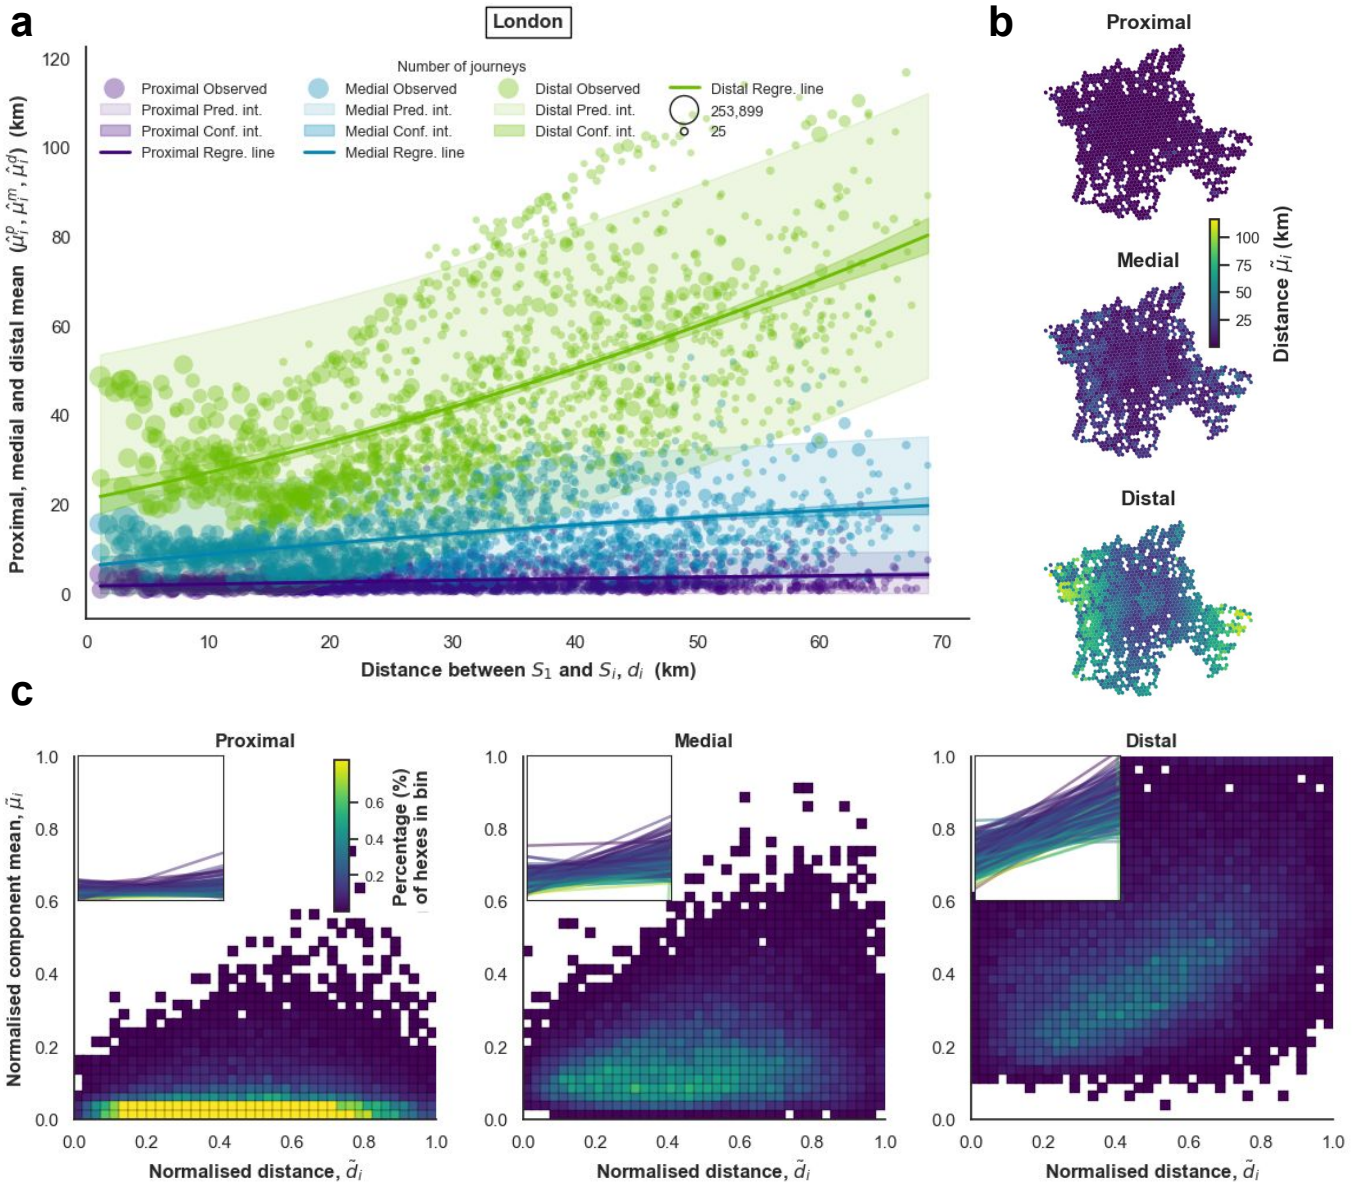

Fig. 1. **a** Scatter for London showing proximal (purple), medial (blue) and distal (green) component means against distance to the city center; curves are quadratic fits with 95% confidence/prediction ribbons; the proximal component is flat, indicating that people typically stay close to home, and the medial component grows, indicating that in sprawling suburbs, even local trips becoming extended. **b** Maps of the same three component means across London on a common color scale, showing that the most central locations have a distal component representing trips away from the CBD. **c** Normalized relationships pooling 224 European cities; density gradients steepen from proximal to distal and rise with city size, shown in the insets with linear fits per city, colored by rank size. In these insets, slope declines with city size, indicating that smaller cities have a steeper gradient—and suggesting that smaller cities are more dependent on the CBD.

Activity Space Segregation Through the Lens of Income Outliers in London

OLENA HOLUBOWSKA [1], ANIRUDH GOVIND [1], AND ANDREW RENNINGER [2]

[1] Division of Geography and Tourism, Department of Earth and Environmental Sciences, KU Leuven, Belgium
[2] Centre for Advance Spatial Analysis, University College London, UK
[*] Corresponding author: olena.holubowska@kuleuven.be

1. INTRODUCTION

Assessing segregation in residential areas and activity spaces, as well as the relationship between them, is an important step toward addressing sociospatial inequalities in cities. Segregation research typically relies on models of hypothetical 'average' individuals that are assumed to adequately represent larger groups, such as those residing in the same neighborhoods. Consequently, these studies use neighborhood-scale demographic data to model these hypothetical 'average' individuals [1, 2]. However, such approaches overlook the heterogeneity of neighborhoods. For example, some people may have income levels that differ substantially from their neighbors. These individuals, referred to here as *income outliers*, may have different mobility patterns and thus experience segregation differently from their neighbors, both in their residential areas and in their activity spaces.

This paper investigates whether such income outliers engage with city spaces differently from the norms of their residential neighborhoods through their daily movements. Using granular GPS trace data, we model individual mobility patterns of approximately $\sim$ 10% of London's residential population. We examine how their mobility translates into residential and activity space segregation, or conversely, the absence of segregation, suggesting integration.

2. DATA AND METHODS

A. Data

We used anonymized GPS mobility traces for people living in Greater London over a 1-month period. These traces allow us to infer each individual's home location and identify the places they visit during daily routines. To estimate individual income levels and median neighborhood income, we rely on two spatial proxies: (1) housing transaction prices for properties traded on the open market and (2) social housing in the form of council estates.

Housing Transactions. We obtained full transaction-level property sale data from the UK's HM Land Registry[1]. This dataset includes all residential property transactions in the UK by address and geocoded postcode locations.

Council Estates. In the absence of a centralized dataset of council estates, we created a heuristic approximation. We identified properties owned by local authorities, filtered for names commonly associated with residential buildings (e.g., "estate", "house", "court"), and removed civic or commercial properties. To improve specificity, we retained only ownership classifications consistent with council estate typologies. This yielded a set of point-level indicators representing locations across London likely to be social housing.

[1] https://www.gov.uk/government/statistical-data-sets/price-paid-data-downloads

B. Estimating Individual and Neighborhood Income

These two datasets allowed us to partition the city into small sub-neighborhood units that capture variation in property transaction prices and ownership (see Figure 1. Detected home locations were matched to these units to approximate individual income levels as a function of property tranaction prices and ownership. Lower Layer Super Output Areas (LSOAs) were used as operation definitions of neighborhoods and neighborhood incomes were calculated as the median transaction price within each LSOA.

Using this framework, we classify individuals into three income groups:

- **Low-income outliers**: Individuals whose inferred home is located in or near a council estate or whose home value is in the lowest 20% of the distribution for its LSOA.
- **High-income outliers**: Individuals whose estimated home value falls in the highest 20% of the distribution for its LSOA.
- **Neighborhood average**: All other individuals whose home income aligns with their neighborhood median in relative terms.

This three-part classification allows us to compare the activity spaces of income outliers with those of individuals with typical income levels for their area.

C. Assessing Spatial Integration

Spatial integration involves both the physical spaces people move through and the social groups they encounter within those spaces. To assess this dual dimension of integration—places and people—we ask three main questions:

- To what extent are individuals from different income groups exposed to one another in the activity spaces they frequent, and how does this compare to their exposure at home? (Question A)
- How much do the activity spaces of income outliers geographically overlap with each other (i.e., low vs. high) and with other income classes? (Question B)
- What is the variation in income levels of each individual's residential and visited locations? (Question C)

To address these questions, we derive individuals' activity spaces from the aforementioned dataset of GPS traces. We link these activity locations to the derived income classes. We then quantify the intergroup exposure of individuals in different income classes as their co-presence in the same locations within the same hours of the day, across all time points and locations (Question A).

We then construct convex hulls around each individual's activity locations to delineate activity space extents. We use these to calculate the extent of overlap in the activity space, space,

normalized by the total activity space area of each income class, between individuals of different income classes (Question B).

Finally, we differentiate between residential and visited locations, per user, to compare their variation in income levels. We compute mean property values across all visited locations with the property value of their residential location (Question C).

3. RESULTS

A. Intergroup Exposure in Activity Spaces

We begin by examining the intergroup exposure between individuals from different income classes based on the locations they visit. Figure 1 B shows the average proportion of people in locations with income classes different from their own, for each hour of the day. We focus on the outliers, i.e., lower-income and higher-income, in turn.

Lower-income individuals encounter higher-income individuals at higher rates throughout the day, particularly during the midday and afternoon hours. However, the converse does not hold true: higher-income individuals experience relatively lower exposure to lower-income individuals, with only modest increases during peak hours. These results reveal a clear pattern of asymmetric exposure between people from income classes.

Those in the average-income class tend to have a more balanced exposure, but with a skew towards contact with the higher-income classes. This suggests that average-income individuals move through more socially-mixed environments, while higher-income users remain comparatively insulated.

B. Spatial Overlap Between Income Groups

Using convex hulls of individual activity spaces, we measure the geographic overlap between groups. As shown in Figure 1 C, the geographic overlap is substantial. For example, 94.5% of the average-income classes' space overlaps with lower-income individuals, and 98.2% with higher-income individuals. Again, the converse is not true. Higher-income users share only 82.9% of their activity space with lower-income users. This asymmetry suggests a form of spatial privilege, where some groups move through areas shared with others, but not vice-versa.

C. Income Profile of Visited Locations

Finally, we compare the average income of residential and visited locations across income classes to assess whether users tend to inhabit spaces with income profiles similar to their own. Figure 1 D shows that individuals from lower income areas consistently visit neighborhoods with higher incomes than their own. This shows that their mobility patterns exposes them to social heterogeneity. In contract, higher-income individuals tend to stay within similarly affluent areas. This reflects a form of spatial insulation wherein people remain in socially homogeneous areas despite their mobility spanning wide geographic distances. The average income group falls in between these two groups. They visit areas that are, on average, modestly more affluent than their home neighborhoods. However, the variation within this group is relatively small.

The results show that individuals from low-income areas consistently visit higher-income neighborhoods than their own. The gap between home and visited area income is especially large for this group, suggesting that their daily mobility exposes them to more affluent environments than their residential context would imply.

By contrast, high-income individuals tend to stay within similarly affluent areas, with minimal difference between the income level of their home and that of their visited locations. This reflects a form of spatial insulation: while their mobility spans wide geographic distances, it remains socially homogeneous. The average-income group falls in between—visiting areas that are, on average, modestly more affluent than their home neighborhoods. However, the variation within this group is relatively small.

Together, these patterns illustrate a broader asymmetry in urban experience between income outliers and the rest of the population. Lower-income individuals move through more economically diverse environments, while higher-income individuals remain spatially and socially segregated despite traversing the city to similar distances. In this way, daily mobility may reinforce—not reduce—existing inequalities in exposure to different economic conditions.

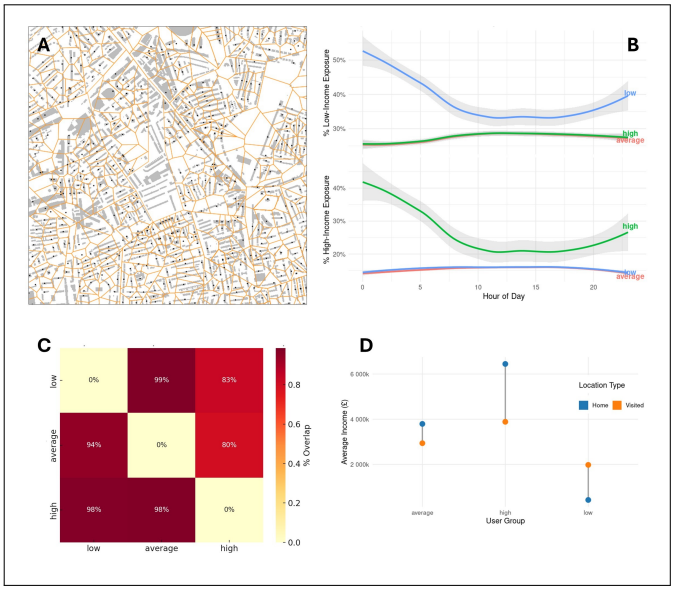

Fig. 1. Results: A) Voronoi polygons drawn according to the nearest home sale or council estate, B) Temporal patterns of intergroup exposure in activity spaces across income groups This figure shows the percentage of low-income (top) and high-income (bottom) individuals encountered in visited locations throughout the day, by users from each income group. Each line represents one group's average exposure to individuals from another income group, with shaded areas indicating confidence intervals. C) Asymmetric spatial overlap of activity spaces across income groups This heatmap shows the percentage of each group's activity space that overlaps with the activity space of another income group, based on convex hull intersections. Rows indicate the focal group, and columns the comparison group. D) Income composition of home vs. visited locations across user groups This figure compares the average income of individuals' home neighborhoods (blue) with the average income of the areas they visit (orange), broken down by income group. Error bars represent standard deviations across individuals.

The data that support the findings of this study are available from the corresponding author upon request.

REFERENCES

1. D. Wang, F. Li, and Y. Chai, Urban Geogr. **33**, 256 (2012).
2. Z. Zheng, F. Chen, J. Lu, and S. Zhou, Trans. Urban Data, Sci. Technol. **2**, 19 (2023). Original work published 2023.

Spatial Modeling of NOx Emissions from Vessel Trajectories Using AIS Data

JIZHOU TANG,[1,2] XIAOWEI GAO,[1,3] XIANGHUI ZHANG,[1,4] TAO CHENG[1]

[1]*SpaceTimeLab, Department of Civil, Environmental and Geomatic Engineering, University College London, London, UK*
[2]*Civil Engineering, Imperial College London, London, UK*
[3]*Department of Earth Science and Engineering, Imperial College London, London, UK*
[4]*Department of Geography and Planning, University of Liverpool, Liverpool, UK*

* Corresponding author: tao.cheng@ucl.ac.uk

Received XX Month XXXX; revised XX Month, XXXX; accepted XX Month XXXX; posted XX Month XXXX (Doc. ID XXXXX); published XX Month XXXX

1. INTRODUCTION

Shipping significantly contributes to nitrogen oxides (NOx) pollution in coastal regions, making detailed spatial emission assessments essential. Automatic Identification System (AIS) data offer comprehensive, high-resolution vessel trajectories, providing a unique opportunity for fine-grained emission modeling[1]. In this study, we develop a two-step spatial approach to map NOx emissions using AIS vessel trajectories, leveraging methods familiar in broader mobility analytics. First, vessel tracks are projected onto a spatial grid to quantify maritime traffic density by location. Second, distances traveled within each grid cell are converted into NOx emissions using a vessel-type-specific emission factor model. We demonstrate this methodology using 2015 AIS data for UK waters, producing a detailed NOx emission map at 5 km resolution. Our key contributions include: (i) a trajectory-to-grid mapping approach for AIS data, and (ii) an emission estimation model calibrated by vessel type. The resulting spatial inventory identifies critical shipping corridors as NOx hotspots, providing actionable insights for environmental policy and maritime management. By framing AIS data within the broader mobility analytics context, our study bridges maritime transportation research and mobility data science. The resulting emission distribution highlights major shipping corridors as NOx hotspots, providing insights for environmental policy and maritime traffic management

2. Data and Methodology

A. AIS Trajectory Data:

We utilize the *Anonymised AIS Derived Track Lines 2015* dataset from the UK Marine Management Organisation (MMO)[2]. This dataset contains full-year AIS vessel transit lines around the UK, derived from sequential AIS-A/B position reports (over 1 billion positions covering >20 million nautical miles)[3]. Each transit line record includes attributes such as start/end time, geometry, and a Ship Type Group (STG) code classifying the vessel into one of 11 categories (0: Unknown, 1: Non-port service, 2: Port service, ..., 6: Passenger, 7: Cargo, 8: Tanker, etc.)[2]. The AIS trajectories were processed into spatially referenced vessel track datasets suitable for geospatial analysis. To ensure consistency and comparability in spatial computations, all trajectory data were transformed into the EPSG:4258 geographic coordinate reference system.

B. Trajectory-to-Grid Mapping:

To quantify spatial activity, we projected all vessel trajectories onto a uniform grid of 5 × 5 km cells covering the study area. Each ship's polyline trajectory was intersected with the grid using a spatial join, yielding the portions of each track falling within each grid cell. We then calculated the total trajectory length (in km) traversed in each cell by summing the lengths of all track segments within that cell. This effectively "bins" the vessel traffic by location, producing a map of total kilometers traveled per grid cell (akin to a shipping traffic density map). Next, we associated each trajectory segment with its vessel type and assigned a conversion factor to translate distance traveled into NOx emitted. We denote this conversion factor as *EF* (in grams NOx per km), which varies by ship type.

C. Emission Conversion Model:

Our emission model transforms the length of ship track in each cell into NOx mass using a multiplicative factor specific to each vessel category. The conversion factor for a given ship type is derived from engine emission parameters obtained in prior studies. In particular, we compute *EF* for each ship type by combining three components: (1) the average NOx emission rate per unit engine work (g NOx per kW·h), (2) the typical engine power (kW) for that vessel class (averaging main and auxiliary engines), and (3) the average operating speed (km/h). By multiplying the first two factors and

dividing by the speed, we obtain the desired rate of NOx emitted per distance traveled. Formally, for each ship type *s*:

$$EF_s = \frac{E_{kWh,s} \times P_s}{v_s}$$

Where $E_{kWh,s}$ is the NOx emission factor per engine work (g NOx/kW·h), P_s is the average engine power, and v_s is the average speed for ship type *s*. These values were drawn from authoritative sources: Entec UK's shipping emission report for emission rates, a Hong Kong marine emission inventory study for engine power data[4], and an ICF International report for typical vessel speeds[5]. For example, large cargo ships have higher P_s and thus a larger *EF* than small craft. A lookup table of *EF* by ship type was prepared (from an internal CSV of literature values) and each trajectory segment was assigned the appropriate *EF* based on its type.

Finally, for each grid cell, we calculate total NOx emissions by summing $L_i \times EF_{s(i)}$ over all trajectory segments i in that cell (where L_i is the length of segment *i* and *s(i)* its ship type). This yields the annual NOx emissions per cell (in grams or tonnes). Notably, our model is a simplified approximation; it assumes average operating conditions (we used values averaged across "at sea" and "in port" modes) and does not explicitly account for factors like engine load or specific speed variations. However, it provides a reasonable first-order estimate of spatial emission distribution given the available data and time constraints.

3. Results and Discussion

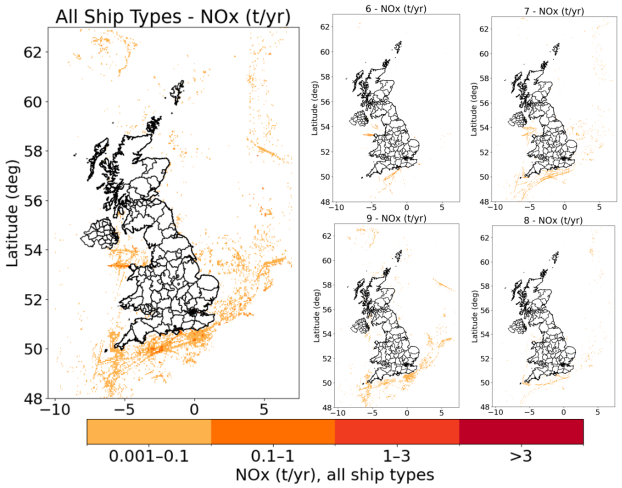

Fig. 1. Estimated spatial distribution of ship-based NOx emissions (annual total for 2015) over the UK region.

The spatial distribution of estimated annual NOx emissions for 2015, shown in Figure 1, highlights clear hotspots along major shipping lanes. High-emission areas (dark orange and red) prominently appear in heavily trafficked regions, especially in the English Channel and routes across the southern and eastern UK coasts. Lower emission intensities (light orange) are evident in less frequented areas, including waters around northern Scotland and parts of the Irish Sea. The results distinctly illustrate that passenger vessels (Type 6), cargo ships (Type 7), tankers (Type 8), and fishing vessels (Type 9) have notably different spatial footprints. Particularly, cargo and tanker vessels dominate the emissions along the busiest shipping corridors, while fishing vessels contribute emissions near coastal and regional fishing areas. The spatial patterns obtained closely match known vessel traffic density maps[3]. Despite simplified assumptions about emission factors, the AIS-based model effectively captures the spatial variability of maritime NOx emissions, enabling the identification of regional pollution hotspots. Such detailed spatial mapping can significantly aid targeted environmental monitoring and regulatory planning.

4. Conclusion

We presented a spatial methodology to map NOx emissions from maritime traffic by integrating AIS trajectory data with an emission estimation model. By projecting vessel tracks onto a 5 km spatial grid and applying vessel-type-specific emission factors, we generated a detailed NOx emission inventory for UK waters. The results clearly identified emission hotspots along major shipping corridors, notably the English Channel, highlighting maritime traffic's significant role in regional air pollution.

Beyond the marine domain, this fundamental methodology—mapping movement trajectories onto spatial grids and integrating them with domain-specific models—can be extended to other mobility analysis scenarios. For example, similar techniques can be applied to urban traffic data (via vehicle GPS or mobile phone trajectories) to map road traffic emissions or congestion patterns, to crowd mobility data to understand crowd dynamics and optimise event management, or to human commuting data combined with air quality measurements to assess crowd exposure to pollution. Our work aligns with the NetMob community's vision of leveraging personal mobility data to address pressing societal challenges. By bridging marine and terrestrial mobility analysis, this study provides a transferable framework for trajectory-based spatial modelling, offering a foundation for policy-making and planning across diverse domains of mobility and transportation.

References

[1]. Ricardo Energy & Environment, *The contribution of shipping emissions to pollutant concentrations and nitrogen deposition across the UK* (2020)

[2]. Marine Management Organisation (MMO). Anonymised AIS Derived Track Lines 2015 (UK Vessel Transit Data). Published Open Data, June 2017 (Available via data.gov.uk, accessed 2025).

[3]. ABPmer, "View the new UK 2015 national dataset of marine vessel traffic," *ABPmer Blog*, 2017. Available: https://www.abpmer.co.uk/blog/view-the-new-uk-2015-national-dataset-of-marine-vessel-traffic/ [Accessed: 6 July 2025].

[4]. Institute for the Environment, HKUST. *Study on Marine Vessels Emission Inventory for Hong Kong – Final Report.* Submitted to Hong Kong Environmental Protection Department, 2012

[5]. ICF International. *Current Methodologies in Preparing Mobile Source Port-Related Emission Inventories – Final Report.* Prepared for U.S. Environmental Protection Agency, April 2009

Robustness of Stop Detection in Sparse Location Datasets

FRANCISCO BARRERAS[1,*] AND THOMAS H. LI[2]

[1] Department of Computer and Information Science, University of Pennsylvania, 3333 Chestnut Street, Philadelphia, PA 19104
[2] Graduate School of Business, Stanford University, 655 Knight Way, Stanford, CA 94305
[*] fbarrer@sas.upenn.edu

1. INTRODUCTION

Commercial GPS datasets have become central to human mobility research, powering applications in epidemiology, disaster response, transportation, urban planning, and behavioral science. These datasets provide granular, high-frequency location signals from smartphones, but they are also sparse, bursty, and noisy—particularly when pings are collected opportunistically via apps or under privacy-preserving constraints. To extract useful signals from these raw data, researchers apply pre-processing algorithms to identify stops, trips, and frequently visited locations. However, many of these algorithms, including variants of DBSCAN [1, 2], rule-based methods like Project Lachesis [3], and grid-based heuristics, were not designed with sparse or bursty signals in mind. Their performance in such settings is rarely tested, and robustness to signal degradation remains poorly understood.

Evaluating robustness in these settings is nontrivial. Real datasets typically lack ground truth: it is unknown whether a detected stop reflects a real event, or if a missed stop is due to noise, sparsity, or poor parameterization. Even synthetic benchmarks often assume uniformly sampled trajectories, which overlook the temporal irregularities that characterize real data. As a result, existing evaluations often prioritize computational efficiency over behavioral accuracy [4, 5], leaving open questions about how well these algorithms perform under the kinds of signal degradation found in commercial datasets.

The implications of this research touch on critical societal issues that depend on reliable human mobility data, as large-scale location datasets are increasingly utilized for urban planning, public health interventions, and transportation policy. Flaws in pre-processing algorithms—particularly those caused by sparsity and incompleteness—can lead to biased metrics and misguided decisions that directly impact public resources, from skewing disease spread understanding to misallocating infrastructure funding. By rigorously evaluating and improving the robustness of these algorithms, this work aims to mitigate such risks and foster more effective decision-making processes in areas critical to societal well-being.

2. METHODS

We address this challenge by introducing a synthetic benchmarking framework designed to test the robustness of stop-detection algorithms under realistic conditions. Our framework builds on the Garden City mobility model [6], which uses an Exploration and Preferential Return (EPR) model to simulate human mobility patterns at both macro and micro scales. Agents generate full, minute-level "ground-truth" trajectories as they move through a synthetic city composed of homes, workplaces, and public locations. These trajectories are then sparsified using a self-exciting point process that captures the bursty nature of GPS pings, and measurement noise is added to mimic GPS error (see Figure 1). These sparse trajectories, alongside their ground-truth counterparts, allow us to quantify the accuracy and robustness of stop-detection algorithms used in the human mobility literature.

Using this framework, we evaluate the robustness of two stop-detection algorithms—ST-DBSCAN [1] and Lachesis [3]. By comparing the detected stops to the ground-truth data, we can compute error metrics to quantify algorithm performance. Crucially, this synthetic approach allows us to generate thousands of trajectories with control over key parameters, enabling us to identify consistent patterns in algorithm performance under different sparsity conditions.

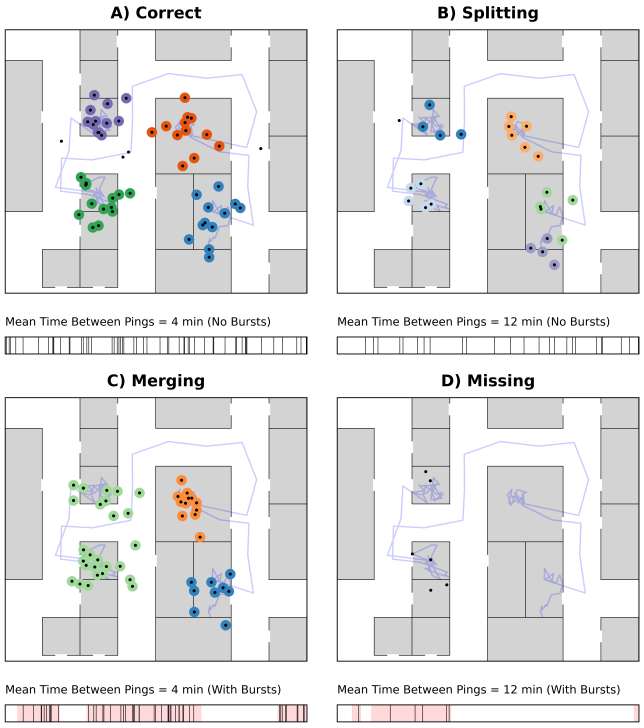

Fig. 1. Each subplot shows sparse trajectory data with detected stops (colored circles) overlaid on a spatial map (top) and corresponding temporal ping distribution (bottom "barcode" plot). In the spatial map, small black dots represent individual GPS pings, while the light blue line indicates the ground-truth trajectory. In the barcode plots, vertical ticks denote ping timestamps and red shading indicates periods of bursts (i.e., ping collection is "turned on"). **(A) Correct**: Optimal stop detection. **(B) Splitting**: False fragmentation of the bottom-right stop into multiple clusters. **(C) Merging**: False consolidation of the left two stops. **(D) Missing**: Failure to detect stops.

Our selected results first focus on identifying three key types of detection errors outlined in Figure 1: merging neighboring stops, splitting larger or longer stops, and completely missing stops. We characterize the sparsity patterns and algorithm pa-

rameter regimes under which these failure modes occur using controlled synthetic trajectories. We also provide results that show that these failure modes persist under more realistic mobility traces—analyzing their frequency and correlation with features of the signal such as ping density and burst structure. In future results, we will also show that careful parameter tuning can mitigate some—but not all—of these issues, and provide practical guidance on how to select parameters in the presence of sparse and noisy data.

3. SELECTED RESULTS

To measure the propensity to merge different stops, we use a trajectory of four stops close to one another (Fig. 1). As shown in Figure 2, we note that the TA-DBSCAN algorithm is relatively robust to merging: the percentage of stops remains roughly constant with the ping frequency. Second, we observe that LACHESIS exhibits less merging as the trajectory dataset becomes more sparse. This could be explained by the sequential nature of the algorithm in combination with the strict size diameter cutoff (the "`max_roam`" parameter). Since human movement is continuous, a higher frequency results in each ping being on average closer to its predecessor (all other factors equal).

Fig. 2. Determinants of Stop Merging. The decrease in ping frequency (higher mean time between pings) results in Lachesis merging less stops, while TA-DBSCAN appears to be robust.

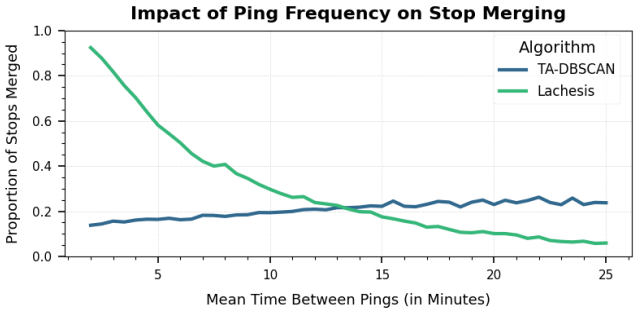

To examine the propensity of algorithms to split stops, we specify a movement pattern of five stops, chosen so they are far from one another (disallowing any merging). In this case, we investigate the effect that gaps between bursts of pings can have on long-dwell stops being split into many (but correct) stops. Thus, we modulate the mean number of minutes in between bursts—where smaller values are more complete and larger values are more incomplete and "bursty"—while keeping the expected completeness constant by commensurately modulating the mean length of bursts. In short, the long gaps add up to about the same amount of time, but they are more or less. Figure 3 reveals that, for the same trajectories, DBSCAN has a much lower propensity for splitting than LACHESIS. The Lachesis algorithm's strict "`max_roam`" parameter prevents large stops from being classified together if the parameter is too small relative to the stop size. Both stops exhibit a large fraction of stops missing (not recovered). A clear trend exists in both algorithms where the splitting is reduced as the mean time between bursts increases. We note that there is a merging-splitting tradeoff between the two algorithms: DBSCAN tends to merge but does not split, whereas Lachesis splits but not does merge.

These previous results were produced using artificial trajectories in order to isolate the determinants of the detection errors. We now turn to more realistic trajectories that are constructed us-

Fig. 3. Determinants of Stop Splitting.

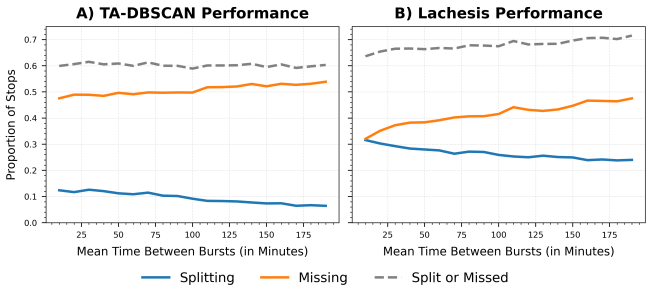

Table 1. Mean Stop Detection Metrics by Algorithm

Algorithm	Missed	Merged	Split	Precision	Recall	F1
Oracle	0.498	0.038	0.001	0.997	0.449	0.611
DBSCAN	0.679	0.141	0.003	0.992	0.419	0.575
Lachesis	0.620	0.052	0.149	0.997	0.362	0.524

ing the EPR model and heterogeneous sparsity patterns. For 100 agents, we simulate three weeks worth of ground-truth trajectory data and sample each agent's trajectory many times to obtain different sparsity patterns. DBSCAN and Lachesis are deployed on these synthetic datasets and their outputs are analyzed, alongside a grid-based "oracle" stop detection that provides a benchmark for the best possible results given the dataset[1]. Table 1 provides the mean performance metrics of all algorithms. The results show DBSCAN having superior performance over Lachesis, with F1 and Recall comparable to the Oracle benchmark. Crucially, in this more general sample of trajectories, Lachesis seems to miss substantially less stops than DBSCAN, which could make it preferred when it is more important to know if a visit occurred than to know exactly for how long.

Funding. This work was possible with the support of the National Science Foundation, NSF award 2341932.

Data availability. Data for the experiments in this paper was generated *ad-hoc* using the garden-city sandbox [6]. The code generating the data and conducting the experiments is available at https://github.com/Watts-Lab/nomad/blob/main/examples/research/robustness-of-algos-paper.ipynb.

REFERENCES

1. D. Birant and A. Kut, Data & knowledge engineering **60**, 208 (2007).
2. M. Ester, H.-P. Kriegel, J. Sander, *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise." in *kdd*, , vol. 96 (1996), pp. 226–231.
3. R. Hariharan and K. Toyama, "Project lachesis: parsing and modeling location histories," in *International Conference on Geographic Information Science*, (Springer, 2004), pp. 106–124.
4. U. Aslak and L. Alessandretti, arXiv preprint arXiv:2003.14370 (2020).
5. W. Chen, M. Ji, and J. Wang, Int. J. Online Eng. **10** (2014).
6. T. H. Li and F. Barreras, arXiv preprint arXiv:2412.00913 (2024).

[1]For example, some stops might be missed simply because they contain no pings.

Data Challenge

posters

Persona-Aware Mobility Management: Detecting Sustained Demand, Speed, and Coverage Shifts

KAVYA HARINARAYANA BHAT[1*], VEENA MENDIRATTA[2*]
Northwestern University, Evanston, IL, USA

[1*]kavyabhat2025@u.northwestern.edu [2*]veena.mendiratta@northwestern.edu

1. INTRODUCTION

This paper presents exploratory data analysis for urban mobility and develops algorithms to manage city transport operations driven by near-real-time data intelligence. Using weighted trip records and GPS data, we build two independent modules. **Traveler personas** characterize who moves how and supports targeted policy. **Real-time anomaly detection** at the IRIS level at 15-minute resolution outputs actionable alerts mapped to playbooks for operations. Together, these modules enable proactive management of multimodal systems with strong spatial–temporal variability.

2. APPROACH

To keep the scope manageable, we analyze weekday mobility at IRIS, 15-minute intervals by mode using three data sources: individuals' dataset, trips dataset, and GPS dataset.

Module A – Traveler personas (segmentation for policy)

- Features (weekday aggregates): trips/day and variability; home vs. top-place dwell; daily distance; radius of gyration (mean/standard deviation) as activity-space size/variability.
- Two-stage clustering (concise). Step-1 (temporal split): standardize and run K-Means on temporal features to separate low/regular/high activity patterns. Step-2 (spatial refinement): within each Step-1 group, standardize and run K-Means on spatial features to distinguish local-frequent vs long-range commuters.
- Silhouette and profile interpretability guide k. We retain four personas and overlay Navigo and bike access, car availability, household size to connect segments to policy (peak reliability, school-hour boosts, last-mile programs).

Module B - Real-time anomaly detection (15-minute operations)

- Signals per slot: weighted departures and arrivals per IRIS by mode; mode shares; GPS speed reliability (minute-median SPEED pooled); coverage (Coverage per slot is the number of unique IRIS with ≥1 weighted start or end).
- Baseline and residuals: robust weekday median by IRIS, mode, slot with hierarchical fallback of IRIS, ZoneClass, citywide-by-mode; variability via robust σ (IQR/1.349).
- Detectors used (sustained changes): CUSUM (CUmulative SUM) and Page–Hinkley (PH) on residuals with level shifts (up/down); Weighted Kolmogorov–Smirnov (KS) on SPEED comparing the last 60 min "before" vs the next 15 min "after" for distributional slowdowns; reporting change in median SPEED; coverage CUSUM for pulses/lulls in unique IRIS. Prioritization and outputs: alerts scored by effect size, persistence and spatial spread (via a department-prefix neighborhood proxy) and mapped to concise playbooks (dynamic dispatch/headway control, bus/traffic signal priority, egress management, etc.). We publish a ranked demand-alerts table and a cross-signal ticker.

3. TRAVELER PERSONAS

We segmented weekday travelers into four interpretable personas using policy-relevant trip/trajectory aggregates (trips per day, dwell at home vs. top place, daily distance, radius of gyration). The goal is to summarize who moves how so that peak reliability, school-hour support, and last-mile programs can be targeted. The clustering in Figures 1 and 2 shows that almost 55% (Persona 0_0) are moderate-trip long-distance commuters, while Persona 1_1 (27%) are frequent long-range travelers, both requiring reliable transit. Persona 1_0 (6%) are local movers with many short trips, pointing to a need for micro-mobility. Persona 0_1 (12%) are low-frequency, place-anchored travelers, suggesting value in strengthening first/last-mile links. Figures 3, 4, 5 and 6 show persona characteristics. In Figure 4 the x-axis bars represent

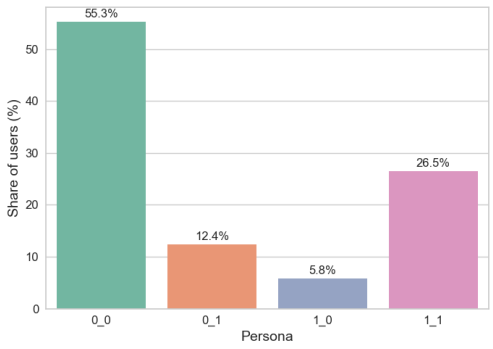

Fig. 1. Persona distribution in the population

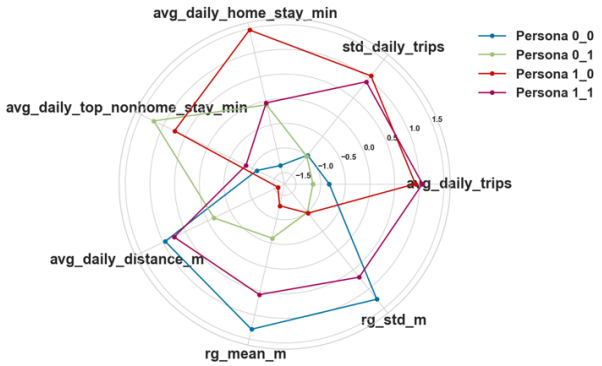

Fig. 2. Feature profile (z-score radar)

household types: Alone, Couple no children, Couple with children, Flat-sharing, Non-Family, Single-parent, With family, With parents. Personas, 1_0 and 1_1 over-index in prime working ages and family households, while 0_0 and 0_1 skew to smaller households with more students/older adults; this mirrors their public-transport use, which is highest for 1_0 and 0_1, moderate for 0_0, and more mixed for 1_1. **Actions:** concentrate peak-hour capacity and school-aligned timetables on 1_0/1_1 corridors; pair 1_1 zones with family bundles and park and ride; and in 0_0/0_1 areas emphasize discounted youth/low-income passes plus dense first/last-mile micromobility where car ownership is lower. **Net effect:** resources shift toward peak, family-heavy flows while off-peak and access products are targeted to personas 0_0 and 0_10_0 and 0_1, improving load-factor efficiency and equity.

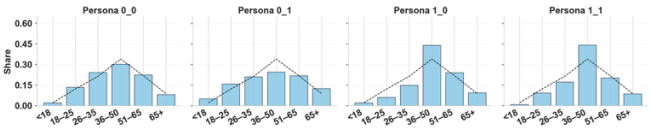

Fig. 3. Age distribution by persona

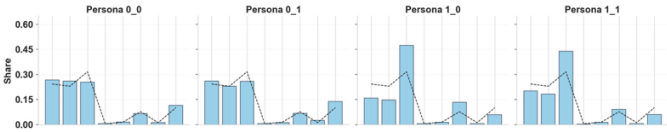

Fig. 4. Household type by persona

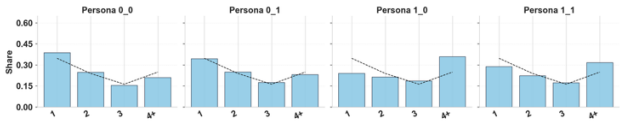

Fig. 5. Household size by persona

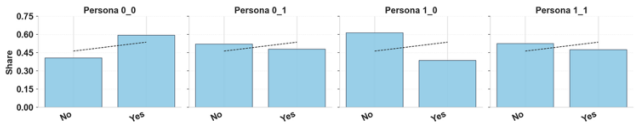

Fig. 6. Public transport (NOVIGO) subscription by persona

4. ANOMALY DETECTION

The goal is to monitor weekday mobility at IRIS by 15-minute by mode to flag sustained deviations in demand, mode mix, speed, and coverage. Events are scored by effect size, persistence and spatial spread, and translated to a playbook (dispatch and headway control, egress management, etc.).

Figure 7 ranks sustained events by effect, persistence and spread, highlighting central-IRIS walking surges and outer-ring car overshoots. Action: egress staffing in core stations and short-term signal adjustments on feeder corridors. Figure 8 (75120 walking) shows observed demand outside the ±1σ band for ~1 h around 10:30, confirming a true change. **Action:** deploy egress flow control (gate opening, escalator reversal). Figure 9 reveals a distributional speed drop (KS = 0.289, median change ≈ −6.8m/s) indicating system-wide slowdown. Action: incident checks, bus/traffic priority, and rider guidance until recovery.

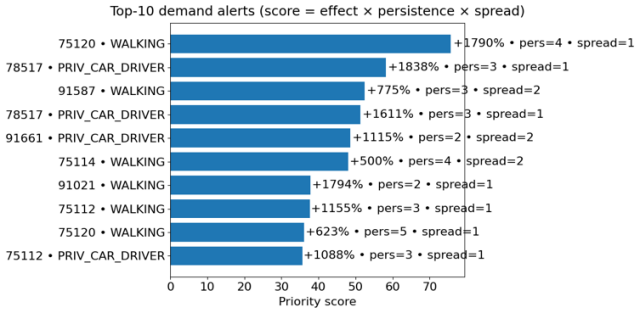

Fig. 7. Ranked demand anomalies by IRIS and mode

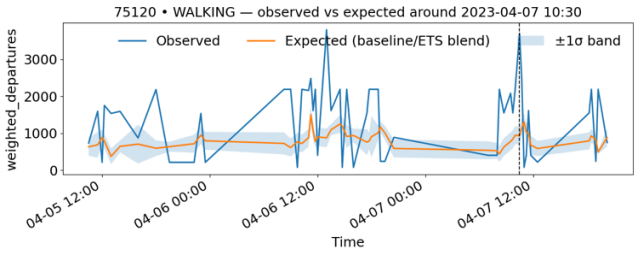

Fig. 8. 75120 • WALKING — Observed vs baseline (±1σ)

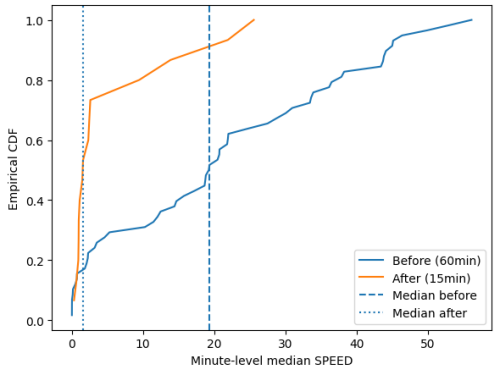

Fig 9. Reliability degradation: speed distribution before/after

5. CONCLUSIONS

Our analysis yields actionable insights for managing city transport networks with data-driven strategies. We convert IRIS 15-minute data into ranked, evidence-backed playbooks. Personas explain who drives demand in each area, while sustained-change detectors show provide actionable insights. We recommend allocating vehicles, staff, and priority windows by mode and corridor, and adjusting intra-day where anomalies persist. For implementation at scale, an AI/ML pipeline should continuously ingest data, maintain day-specific baselines and **forecast** near-term demand to pre-position resources. As future work, we will run similar detection on arrivals, integrate **context** for expectation down-weighting, and deploy a near-real-time loop (<5-min latency) so that capacity allocation and rider information are automated and adaptive**.**

Identifying First/Last Mile Deserts in the Paris Region Using Mobile GPS Data

MOHAMED NASSAR[1,*], REEM HUSSIN[1], MOHANNED ABDELFATAH[1], AND NOHA GAMAL[1]

[1] ITCS, Nile University, Cairo, Egypt
[*] M.Hossam2225@nu.edu.eg

1. INTRODUCTION

The First and Last Mile (FLM) segments of a journey are widely recognized as critical determinants of a public transport system's overall accessibility and equity [1, 2]. These hyperlocal segments—from a starting point like home to a transit hub, and from a hub to a final destination—are often overlooked in traditional network-level planning, yet they can represent the most significant barrier to mobility. Inefficiencies and inequities in FLM access can create profound social and economic disparities for urban residents.

This study addresses this challenge by developing a scalable methodology to systematically identify and, crucially, *diagnose* different archetypes of "FLM Deserts" within the Paris region. Using a large-scale GPS dataset from the NetMob25 Data Challenge [3] and building on methods for identifying transit legs from GPS traces [4], our objective is to move beyond simple mapping. We aim to understand the underlying causes of these deserts—whether time poverty, car dependency, or other barriers—in order to propose targeted policy interventions, such as the deployment of micromobility solutions, that can foster a more equitable transportation system [5].

2. METHODOLOGY

To produce a reliable map of FLM deserts, our methodology follows a robust, multi-stage data processing and analysis pipeline, building on established techniques for GPS-based mobility analysis [4].

Starting with the raw trip data, we first identified potential FLM segments using a two-step validation process. A temporal filter first flagged all non-public transport trips occurring immediately before or after a public transport leg (e.g., Métro, RER). These candidates then underwent a crucial geospatial validation, where GPS traces were used to confirm that the transfer was geographically plausible, defined as the end-point of the FLM trip and the start-point of the transit leg being within a 200-meter buffer. To ensure extreme edge cases did not skew our analysis, we then removed statistical outliers by filtering out trips in the 99th percentile of duration.

Finally, we performed a spatial join to link each of the remaining validated FLM trips to its originating IRIS neighborhood. Within each of these 5,264 zones, we calculated our key metrics—including weighted average duration and private car reliance—by applying the provided 'Weight_Day' survey weights. This final step ensures our neighborhood-level findings are statistically representative of the region's overall travel patterns.

In the next phase of this project, we will extend our diagnostic methodology by simulating policy interventions within an agent-based modeling framework. Each identified neighborhood will be populated with representative agents exhibiting empirically-derived mobility behaviors and constraints. By testing the introduction of targeted solutions—such as micromobility deployments, feeder bus expansions, or cycle infrastructure improvements—we aim to quantify the potential impact of these interventions on FLM access, travel times, and modal shifts. This simulation-driven extension will provide decision-makers with a predictive tool for evaluating trade-offs and prioritizing investments at the neighborhood level.

3. PRELIMINARY RESULTS

Our analysis, performed on an outlier-removed dataset of validated FLM trips, reveals that "FLM Deserts" are not a monolithic problem but manifest in distinct archetypes across the Paris region. By focusing on the 197 neighborhoods with the highest volume of observed FLM activity (representing over 800 real-world trips each), we identified a number of key patterns.

A. Diagnosing Desert Archetypes

While the average FLM journey in these significant zones is approximately 15 minutes, we found extreme variation, with the worst-performing neighborhood exhibiting a staggering average duration of 80 minutes. The primary challenges can be diagnosed as two distinct types, often appearing in different locations.

First, as shown in Fig. 1, we identified severe **"Time-Poverty Deserts."** These zones, often located in outer suburbs like **Vaux-le-Pénil**, are characterized by extremely long travel times coupled with low private car use. This points to a population forced to endure long walks or waits for infrequent feeder transit.

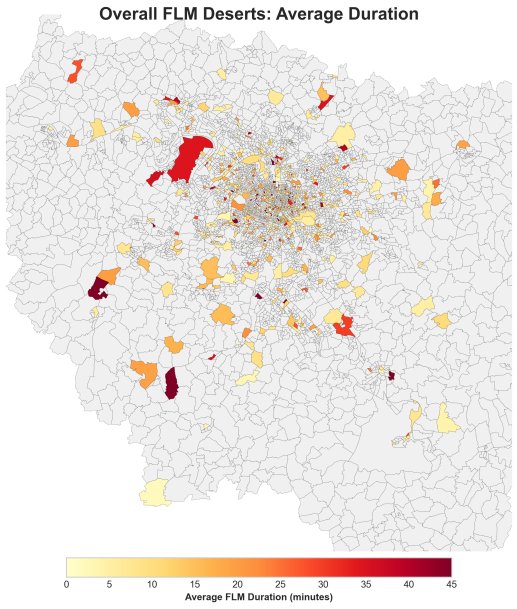

Fig. 1. Overall "Time-Poverty Deserts," showing the weighted average duration of all validated FLM trips. Darker red indicates a longer and more difficult journey to access public transport.

Second, Fig. 2 reveals distinct **"Car-Dependent Deserts."** These neighborhoods, frequently appearing around major economic hubs like **Puteaux** (near La Défense) and **Saint-Denis**, are defined by a high percentage of FLM trips requiring a private car. This indicates a critical lack of viable, non-car alternatives for accessing the main transit network in key employment zones.

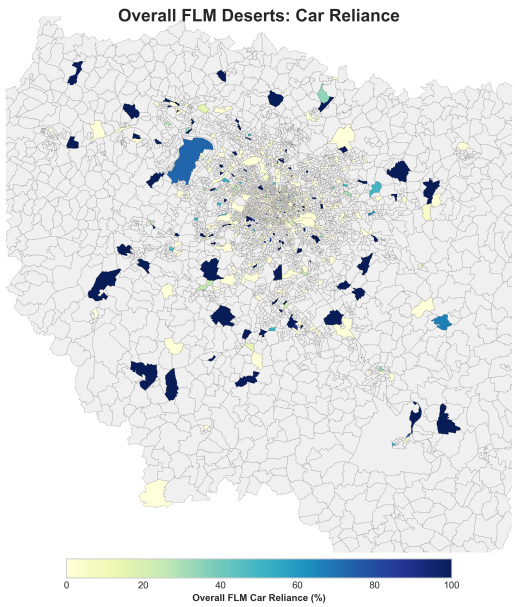

Fig. 2. Overall "Car-Dependent Deserts," showing the percentage of all FLM trips that rely on a private car. Darker blue indicates a higher dependency on private vehicles.

B. The First vs. Last Mile Asymmetry

A comparative analysis reveals a strong, directional asymmetry between the journey to and from public transport. The phenomenon of an inefficient First Mile was found to be far more pronounced in duration than that of an inefficient Last Mile. Our analysis identified zones where the average First Mile was over **22 minutes longer** than the Last Mile. We hypothesize this is due to several factors, including complex morning trip-chaining, lower off-peak frequency of evening transit services, and differing perceptions of safety after dark influencing mode choice for the journey home.

4. RECOMMENDATIONS & CONCLUSION

Our analysis, which identified and diagnosed distinct archetypes of FLM Deserts, allows for the formulation of targeted policy recommendations that address the specific nature of each mobility barrier.

1. **For "Time-Poverty Deserts,"** which are often located in outer suburbs like **Vaux-le-Pénil** and are characterized by long travel times but low car-reliance, the primary challenge is a lack of efficient transit options. We recommend that transport authorities:
    - Conduct feasibility studies to increase the frequency and expand the service area of feeder bus routes connecting these residential zones to the nearest RER/Transilien stations.
    - Co-deploy subsidized micromobility solutions, such as shared e-bikes and e-scooters, at key neighborhood locations to cover distances that are too long to walk but too short for new bus routes, a strategy supported by recent transport policy studies [5].
2. **For "Car-Dependent Deserts,"** often found around major economic hubs like **Puteaux** and **Saint-Denis**, the challenge is a lack of viable alternatives to the private car. We recommend:
    - Investing in high-capacity, secure bicycle parking at major transit hubs to eliminate a key barrier to bike-and-ride commuting.
    - Developing a network of physically protected cycleways that connect these hubs to surrounding business parks and residential areas, providing a safe and efficient alternative to driving.

In conclusion, this study demonstrates the necessity of moving beyond network-level analysis to diagnose hyperlocal mobility barriers. By distinguishing between different archetypes of FLM deserts—those defined by time poverty versus those defined by car dependency—we provide a scalable methodology that enables transport authorities to deploy targeted, context-aware, and cost-effective solutions. This approach is essential for building a truly equitable and accessible public transport system for all residents of the Paris region.

Data Availability Statement. Data underlying the results presented in this paper were provided as part of the NetMob25 Data Challenge. Derived data used for the analysis are available from the authors upon reasonable request.

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. E. Kåresdotter, O. Cats, and L. Reardon, J. Urban Technol. **29**, 3 (2022).
2. M. Alfaris, M. Bell, H. Raza, and R. A. Acheampong, Transp. Res. Rec. **2677**, 149 (2023).
3. NetMob25 Data Challenge, "Dataset provided for the netmob 2025 conference," (2025). Accessed as part of the official NetMob25 competition dataset.
4. M. Hosseini, A. Sadeghi-Niaraki, and K. Kim, Transp. Res. Procedia **70**, 229 (2024).
5. G. Oeschger, L. T. Truong, and G. Currie, Transp. Policy **136**, 13 (2023).

Segmenting Human Mobility Behavior Under Internet Quality Pressure in Paris

MOHAMED NASSAR[1,*], REEM HUSSIN[1], MOHANNED ABDELFATAH[1], AND NOHA GAMAL[1]

[1] ITCS, Nile University, Cairo, Egypt
[*] M.Hossam2225@nu.edu.eg

1. INTRODUCTION

While physical infrastructure has long been the focus of urban planning, the quality of digital infrastructure is now a critical determinant of economic opportunity and mobility behavior [1]. The behavior of city residents is increasingly influenced by the need to stay connected, yet the relationship between internet quality and physical movement remains underexplored. Understanding this dynamic is essential for designing cities that are not just "smart," but digitally equitable in how they serve residents.

This study addresses this gap by developing a scalable framework to model this interplay in the Paris region. Using large-scale GPS trajectories from the NetMob25 Data Challenge [2] fused with fine-grained RIPE Atlas internet quality measurements [1], we define two new behavioral metrics: the Internet Exposure Score (IES), quantifying the connectivity experienced during mobility, and the Bias Score, capturing preference for higher-connectivity zones. We then apply an unsupervised clustering pipeline (UMAP + HDBSCAN) to identify distinct digital-mobility personas and map their geographic concentrations across Paris[2, 3].

Beyond descriptive analysis, we embed these personas into a Predictive Agent-Based Modeling (ABM) framework that simulates how each group adapts to digital infrastructure changes. By testing scenarios such as continuous 5G coverage along the RER-D line, free Wi-Fi student corridors, and targeted subsidies in underserved zones, our results demonstrate both the heterogeneity of digital dependencies and the policy potential of connectivity-aware urban analytics.

2. METHODOLOGY

Our methodology transforms raw mobility data into actionable behavioral insights through a three-stage pipeline: data fusion and feature engineering, unsupervised clustering, and representative geographic profiling with predictive simulation.

First, acknowledging the challenges inherent in fusing heterogeneous data sources [4], we performed a rigorous spatio-temporal data fusion. We integrated individual-level GPS trip data from the NetMob25 dataset with high-resolution RIPE Atlas [1] internet quality measurements. From this unified dataset, we engineered novel behavioral features for each user, most notably an **Internet Exposure Score (IES)**, which captures the average internet quality experienced during travel, and a **Bias Score**, which measures the tendency to move toward areas with better connectivity.

To identify distinct behavioral archetypes, we employed an unsupervised machine learning pipeline similar to that used in other large-scale mobility studies [2]. After aggregating the engineered features for each individual, we applied **UMAP** to reduce the dimensionality of the feature space. Subsequently, we used **HDBSCAN** to discover natural, density-based clusters of users, allowing for the organic emergence of personas without predefining the number of groups.

Finally, we extended the analysis by embedding these personas into a **Predictive Agent-Based Modeling (ABM) framework**. Each persona was modeled as an autonomous agent with a utility function balancing efficiency, connectivity, and cost. Policy interventions—such as extending continuous 5G coverage along the RER-D line, deploying free Wi-Fi corridors for students, and subsidizing underserved areas—were simulated by modifying network connectivity attributes. The resulting behavioral adaptations were then quantified in terms of IES, Bias, and commuting efficiency, providing predictive capacity for scenario planning and strategic investment.

3. PRELIMINARY RESULTS

Our population-weighted analysis reveals that the population of the Paris region is not homogeneous in its digital behavior. Instead, it is composed of distinct personas with unique mobility patterns and sensitivities to internet quality. The clustering process successfully identified a spectrum of these behavioral archetypes, which we then mapped to reveal their geographic concentrations.

A. A Spectrum of Digital Personas

The unsupervised clustering algorithm identified over 20 distinct behavioral segments, visualized in the UMAP projection in Fig. 1. By analyzing the weighted characteristics of these groups, we defined several key personas that exist on a spectrum of digital dependency.

At one end are the **Digitally Demanding Professionals**, a group characterized by the highest *Bias Score*. This indicates they actively seek out and are influenced by the availability of high-quality internet during their often long-distance trips. At the other extreme are the **Top Tier Oasis Users**, who exhibit an exceptionally high *Internet Exposure Score*. These individuals, concentrated in continuous-coverage "digital oases," have low Bias not because they do not value connectivity, but because their environment already guarantees top-tier access. Between these extremes, we identified other crucial groups, including digitally-aware **Student Commuters**, whose mobility is centered around university zones, and more disconnected segments such as **Traditionalist Retirees**, who take shorter trips and exhibit minimal reliance on connectivity.

B. Geographic Concentrations of Behavior

The true power of this analysis comes from mapping the geographic distribution of these personas, as shown in Fig. 2. Rather than a simple model, the findings reveal a complex urban mosaic of digital behavior.

Our analysis pinpointed specific concentrations of these key personas. The **Digitally Demanding Professionals**, for instance, show a significant concentration in the suburban economic hub of **Évry-Courcouronnes**. In contrast, the **Top Tier Oasis Users** are most dominant in the affluent, dense central Paris districts of **Auteuil (16e)** and near the Seine in the **15e arrondissement**.

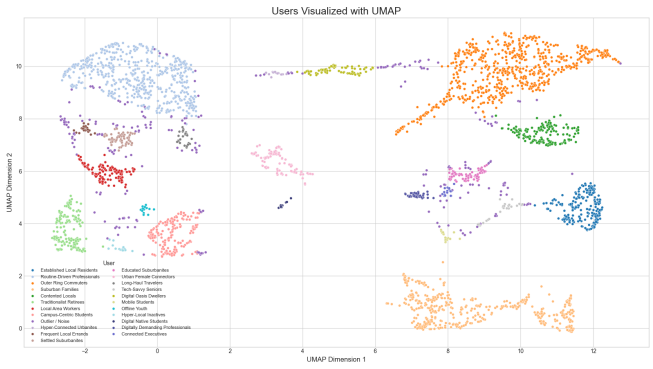

Fig. 1. UMAP visualization of user clusters. Each color represents a distinct behavioral segment identified by the HDBSCAN algorithm.

Similarly, **Student Commuters** are most prevalent in neighborhoods adjacent to university centers, such as **Croulebarbe (13e)**, while **Retirees** dominate peripheral zones with shorter trips and limited connectivity. This spatial analysis provides a critical link between abstract behavioral profiles and tangible locations.

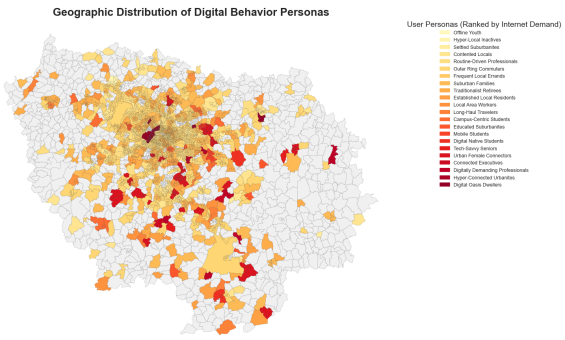

Fig. 2. Geographic distribution of dominant personas across Paris. Each IRIS zone is colored by its most prevalent persona, illustrating distinct spatial patterns of digital-mobility behavior.

C. Policy Simulation Results

To move beyond descriptive clustering, we embedded the identified personas into a **Predictive Agent-Based Modeling (ABM)** framework and simulated targeted interventions on the Paris transport network. Each persona was modeled as an adaptive agent balancing efficiency, cost, and connectivity. Three policy scenarios were tested:

1. **Continuous 5G coverage on the RER D line:** reduced detours by approximately **23%** among Digitally Demanding Professionals, increasing productive commuting time.

2. **Free Wi-Fi corridors on student routes:** raised the Internet Exposure Score of Student Commuters by around **17%**, narrowing digital inequality with central districts.

3. **Targeted subsidies in underserved areas:** improved baseline connectivity for Retirees, reducing their IES gap by roughly **12%**.

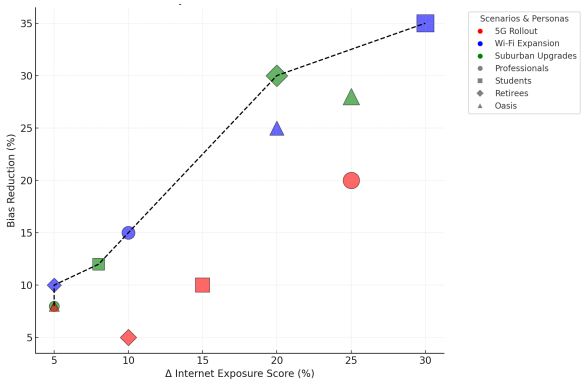

Fig. 3. Impact of simulated interventions on key personas. Shapes indicate the three policy scenarios relative to baseline: continuous 5G on RER D, free Wi-Fi student corridors, and targeted subsidies. Values show changes in Internet Exposure Score (IES) and detour reduction.

4. CONCLUSION

This study introduces a connectivity-aware paradigm for urban analytics. By defining two behavioral metrics—the Internet Exposure Score (IES) and Bias Score—and applying a UMAP+HDBSCAN pipeline, we identified more than twenty digital-mobility personas, with four dominant archetypes capturing the majority of the Paris population. Mapping their geographic concentrations revealed a clear mosaic of digital inclusion and exclusion across the metropolitan region.

Embedding these personas into an ABM framework allowed us to simulate concrete policy interventions, producing quantifiable reductions in detours, measurable increases in IES, and narrowed equity gaps. By combining clustering-based diagnostics with predictive simulation, our framework provides both researchers and policymakers with a scalable tool for connectivity-aware scenario planning. This approach enables investments that are not one-size-fits-all, but tailored to the distinct dependencies of urban populations.

Data Availability Statement. Data underlying the results presented in this paper were provided as part of the NetMob25 Data Challenge. Derived data used for the analysis are available from the authors upon reasonable request.

Disclosures. The authors declare no conflicts of interest.

REFERENCES

1. RIPE Atlas Project, ACM SIGCOMM Comput. Commun. Rev. **48**, 5 (2018). A detailed overview of the project and its capabilities.
2. M. Thuillier, D. D'Avolio, and D. Basso, Transp. Res. Part C: Emerg. Technol. **95**, 14 (2018).
3. P. Lemaire and T. Giraud, J. Urban Plan. Dev. **150**, 04023058 (2024).
4. NetMob25 Data Challenge, "Dataset provided for the netmob 2025 conference," (2025). Accessed as part of the official NetMob25 competition dataset.

Behavioral Segmentation of Older Adult Mobility in Metropolitan Areas Based on a High-Resolution GPS Data: A Case Study of Île-de-France

ABDULLAH ALLAMA[1], ISLAM AREF[2], AHMED MAKBOUL[3], ABDALLAH WAEL[3], RGHDA SALAH[4], AND NOHA GAMAL[5,*]

[1] Faculty of Computers and Data Science, Alexandria University, Egypt
[2] Faculty of Computer and Informatics, Zagazig University,Egypt
[3] Faculty of Computer Science, Helwan University, Egypt
[4] Faculty of Engineering, Zewail University, Giza, Egypt
[5] School of Information Technology and Computer Science, Nile University, Giza, Egypt.
[*] ngamal@nu.edu.eg

Abstract- **This study fundamentally challenges conventional assumptions about elderly mobility by demonstrating that residential location—not socioeconomic status—is the primary determinant of travel behavior among older adults in Île-de-France. Analyzing high-resolution GPS trajectory data from 292 individuals aged 65+ in the NetMob25 dataset, we applied advanced clustering techniques to identify three distinct mobility phenotypes: Extensive Car-Dependent Travelers (21.7%), Local Shopping-Focused Minimalists (22.9%), and Multimodal Active Agers (55.4%). Traditional demographic markers including professional category (p=0.212), education (p=0.090), and gender (p=0.823) show no significant association with mobility patterns, while geographic location demonstrates highly significant clustering ($\chi^2 = 56.64$, p < 0.0001). Central Paris shows 76.7% multimodal active elderly, dropping to 34.5% in outer suburbs. This geographic gradient, independent of social class, provides evidence that infrastructure accessibility shapes elderly travel patterns more than individual characteristics.**

Keywords: elderly mobility, clustering analysis, transportation behavior, spatial analysis

1. INTRODUCTION AND METHODOLOGY

With 21.6% of the EU population now aged 65+[1], metropolitan cities are particularly affected and efforts were made in the recent years to make cities more inclusive for people in all demographic groups [2]. Prior research emphasizes socioeconomic stratification as the primary driver of elderly mobility inequality. Yet, this study reveals that geographic location overwhelmingly determines mobility patterns, transcending class boundaries.

The NetMob25 dataset [3] provides unprecedented granularity through GPS coordinates captured every 2-3 seconds for 292 elderly individuals, combined with comprehensive demographic and infrastructure access data. We developed four-dimensional behavioral features capturing spatial (activity radius, unique locations), temporal (scheduling entropy), modal (transport diversity), and activity (purpose diversity) dimensions.

Multiple clustering algorithms (K-means, GMM, Hierarchical) converged on three distinct phenotypes, validated through bootstrap stability analysis (ARI = 0.589 ± 0.189) and comprehensive statistical testing.

2. RESULTS

A. Three Mobility Phenotypes

Extensive Car-Dependent Travelers (21.7%) demonstrate the largest activity radius (47.1 km), overwhelming car dependency (77% of trips), minimal public transit use (4%), and concentrate in outer suburbs (74.2% of group).

Local Shopping-Focused Minimalists (22.9%) show constrained mobility (12.3 km radius), shopping specialization (47% of trips), balanced modal mix, and mixed suburban distribution.

Multimodal Active Agers (55.4% - Majority) exhibit highest transport diversity, strong active transport use (49%), substantial public transit adoption (35%), minimal car dependency (17%), and concentrate in central Paris and inner suburbs (77.7% of group).

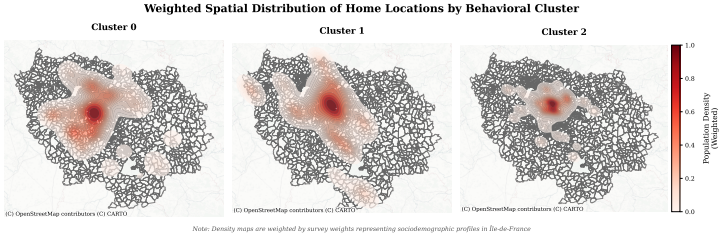

Fig. 1. Weighted spatial distribution of elderly mobility across Île-de-France.

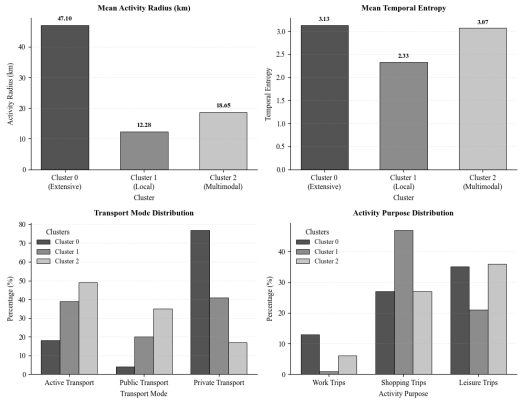

Fig. 2. Comparative behavioral profiles of the three elderly mobility phenotypes

The three-cluster solution reveals distinct behavioral profiles across spatial , temporal, modal, and activity dimensions. (as seen in Figure 2).Even though work trips is among the least frequent trips conducted by elderly people in Île-de-France as a

whole, the extensive car-dependent cluster were the cluster that commute to work the most.

B. The Irrelevance of Socioeconomic Status

Statistical analysis reveals significant associations between mobility phenotypes and residential location:

Table 1. Statistical Associations with Mobility Phenotypes

Factor	Test Statistic	p-value
Geographic Factors		
Residential Location	$\chi^2 = 56.64$	**< 0.0001*** **
Transit Pass Ownership	$\chi^2 = 48.92$	**< 0.001*** **
Socioeconomic Factors		
Professional Category (PRO_CAT)	$\chi^2 = 5.21$	0.212
Education Level	$\chi^2 = 7.89$	0.090
Gender	$\chi^2 = 0.38$	0.823
Mobility Impairment	$\chi^2 = 1.68$	0.497

C. Spatial Clustering Patterns

The geographic distribution reveals pronounced spatial segregation aligned with infrastructure availability:

Table 2. Mobility Phenotype Distribution by Residential Zone

Zone	Multimodal Active	Local Minimalist	Car-Dependent
Central Paris	76.7%	6.8%	16.4%
Inner Suburbs	75.5%	10.4%	14.2%
Outer Suburbs	34.5%	40.7%	24.8%

Transit infrastructure density provides the mechanism: central Paris offers 42.3 transit stops/km² versus 3.1 in outer suburbs, while transit pass ownership ranges from 14.5% (car-dependent) to 69.2% (multimodal active).

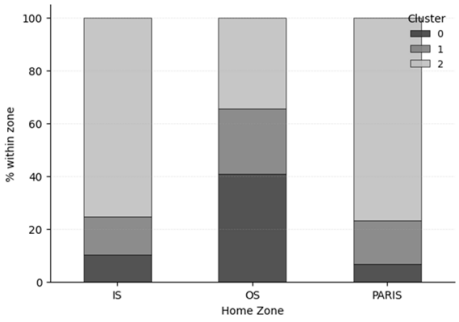

Fig. 3. Cluster distribution by residential zone

3. DISCUSSION

Our findings necessitate fundamental reconceptualization of elderly mobility determinants. The traditional framework emphasizing individual resources was limited in explain patterns where professional category—a proxy for socioeconomic status—shows weak association with mobility patterns (p=0.212). A retired executive in suburban Marne-la-Vallée exhibits similar car-dependent patterns to a working-class resident in the same area, while both would likely adopt multimodal behaviors in central Paris. Findings from Zhou et al. [4] in a study conducted on older people in Nanjing (a metropolitan city in China) affirm that the mode choice of older adults in this city is mainly driven by their socioeconomic status. The present analysis of Île-de-France shows that residential location is a strong determinant of mode choice regardless of socioeconomic status. A study conducted on two cases from the US and Japan [5] supports this conclusion by comparing the travel behavior of older adults from the United States—where urban spaces were designed to accommodate private motorized modes—and Japan. Lo & Mitra drew a similar conclusion to our study: that a denser environment and accessible public transit can reduce reliance on automobiles and encourage older adults to choose more active/public modes.

The dramatic center-periphery gradient (76.7% → 34.% multimodal) reveals that **mobility inequality is fundamentally spatial rather than social justice**. Infrastructure improvements would benefit all elderly residents regardless of socioeconomic status, suggesting geographic targeting proves more impactful than income-based interventions.

For central Paris and inner suburbs with high multimodal populations, policies should maintain existing infrastructure quality and focus on accessibility features. For outer suburbs with car-dependent patterns, interventions should prioritize new transit infrastructure, demand-responsive transport, and mobility hubs linking car parks with public transit. Transition zones require targeted last-mile connectivity improvements and integrated paratransit services.

4. CONCLUSION

This analysis demonstrates that 55.4% of elderly individuals maintain active, multimodal mobility when infrastructure permits, transcending all socioeconomic boundaries. The complete absence of associations between mobility phenotypes and professional category (p=0.212), education (p=0.090), or gender (p=0.823), coupled with overwhelming significance of residential location (p<0.0001), provides compelling evidence for spatial rather than social understanding of elderly mobility constraints for Île-de-France's older adults.

These findings demand reorientation from individual interventions to spatial solutions. Rather than assuming elderly immobility reflects personal deficits, planners have to recognize that appropriate infrastructure enables diverse mobility regardless of individual characteristics.

Acknowledgments: We thank the NetMob25 challenge organizers and survey participants.

REFERENCES

1. Eurostat, "Population structure indicators at national level," Report from European Commission (2021).
2. OECD, "Cities for all ages," Online Publication (2021).
3. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," arXiv preprint arXiv:2506.05903 (2025).
4. Y. Zhou, Q. Yuan, F. Ding, *et al.*, "Demand, mobility, and constraints: Exploring travel behaviors and mode choices of older adults using a facility-based framework," J. Transp. Geogr. **102**, 103368 (2022).
5. A. W. T. Lo and S. K. Mitra, "A tale of two aging nations: How do urban designs impact older adults' mobility in automobile versus transit societies?" Transp. Policy **133**, 80–92 (2023).

Beyond Averages: Inequalities in Mobility-Based $PM_{2.5}$ Exposure in Greater Paris

EGOR KOTOV[1,2,*], DANIELA PERROTTA[1], TOM THEILE[1], AND EMILIO ZAGHENI[1]

[1] Max Planck Institute for Demographic Research (MPIDR), Rostock, Germany
[2] Universitat Pompeu Fabra (UPF), Barcelona, Spain
[*] kotov@demogr.mpg.de

Accurately assessing personal exposure to air pollution requires accounting for daily mobility. Residence-based estimates, derived from home addresses and fixed monitors, miss heterogeneity generated by activity locations and travel. Evidence shows that even neighbours can experience markedly different $PM_{2.5}$ levels due to routines and modes of travel. Mobility-based approaches that integrate high-resolution movement with concentration fields address these issues and help mitigate the uncertain geographic context and neighbourhood-effect averaging problems [1, 2].

Averages conceal distributional differences and the share of people at high exposure. Adopting distribution-sensitive metrics from inequality research reveals *how exposure is shared across individuals*. This matters for policy because many concentration-response functions are convex: reducing the upper tail at a fixed mean can disproportionately lower health burden.

1. METHODS

Data and study window. We analyse the NetMob25 EMG 2023 survey for Île-de-France [3]: 3,337 residents with seven consecutive days of validated trips per participant and calibrated survey weights at the individual and day levels. Hourly $PM_{2.5}$ comes from the Copernicus Atmosphere Monitoring Service (CAMS) European reanalysis (surface, 0.1° resolution), covering October–December 2022 (validated) and January–May 2023 (interim) [4]. We link activity locations to hourly $PM_{2.5}$ to construct individual time-weighted mean concentrations and quantify inequality by sex, age, and socio-professional group using the survey-weighted mean, Gini, Theil T, and the Q90/Q10 ratio.

Linkage and exposure construction. Each trip destination is mapped to its CAMS raster cell. For a trip i of participant j, hours between arrival and the next trip start are assigned to that cell and hour; unobserved movements are ignored and in-transit exposure is not modelled. Let $PM_{2.5,ijh}$ denote the hourly concentration and w_{ij}^{day} the calibrated trip-day weight. Here, H_{ij} is the number of hourly intervals and Δt_{ijh} is the duration (hours) of interval h. The individual time-weighted mean concentration is

$$C_j = \frac{\sum_i \sum_{h=1}^{H_{ij}} \mathrm{PM}_{2.5,ijh}\, \Delta t_{ijh}\, w_{ij}^{\text{day}}}{\sum_i \sum_{h=1}^{H_{ij}} \Delta t_{ijh}\, w_{ij}^{\text{day}}}.$$

Population means use calibrated individual weights weight_indiv_j:

$$\mu = \frac{\sum_j \text{weight_indiv}_j\, C_j}{\sum_j \text{weight_indiv}_j}. \tag{1}$$

Inequality indices

We analyse the distribution of $\{C_j\}$ using four metrics: the survey-weighted mean (Eq. 1), the Gini coefficient, the Theil T index, and the Q90/Q10 ratio.

Gini. With $x_j = C_j$ and weights $w_j = \text{weight_indiv}_j$,

$$G = \frac{\sum_j \sum_k w_j w_k |x_j - x_k|}{2\,\mu\,(\sum_j w_j)(\sum_k w_k)} \in [0,1]. \tag{2}$$

Theil T.

$$T = \frac{1}{\sum_j w_j} \sum_j w_j \left(\frac{x_j}{\mu}\right) \ln\left(\frac{x_j}{\mu}\right). \tag{3}$$

For groups g with total weight W_g and mean μ_g,

$$T = \underbrace{\sum_g \frac{W_g}{\sum_j w_j} \frac{\mu_g}{\mu} \ln\left(\frac{\mu_g}{\mu}\right)}_{\text{between}} + \underbrace{\sum_g \frac{W_g}{\sum_j w_j} T_g}_{\text{within}}. \tag{4}$$

Tail spread (Q90/Q10). Let q_p be the survey-weighted p-quantile of $\{x_j\}$; the tail spread is $\text{Q90/Q10} = q_{0.90}/q_{0.10} \geq 1$ (computed from the survey design).

Diagnostics. We use complementary cumulative distribution functions (CCDFs, log scale) as a non-parametric visual summary of tail behaviour (plots shown for occupational groups).

Uncertainty

For all indices and diagnostics we use an individual-level bootstrap: (i) resample participants with replacement (keeping weight_indiv and their trip-day weights); (ii) recompute all metrics for each resample; (iii) form 95% percentile intervals.

2. RESULTS

Distributional comparisons and bootstrap inference The plot in Fig. 1 confirms that inequality varies across occupations. *Craftsmen/tradespeople/business owners* and *Unemployed/inactive* exhibit clearly higher Theil T and larger Q90/Q10 ratios than *Employees* and *Workers*. For example, Theil T for Craftsmen is around 0.16 with a narrow 95% interval and Q90/Q10 is around 4.6; Workers are around 0.09 and 2.7, respectively. Group means remain similar across occupations.

Figure 2 visualises whole distributions without any modelling for the occupational groups. The complementary cumulative distribution functions (CCDFs, log scale) display heavier right tails for *Craftsmen/tradespeople/business owners* and *Unemployed/inactive* than for *Employees* and *Workers*, with separation most evident in the upper decile. This pattern is consistent with the bootstrap intervals reported above.

Other dimensions (age and sex). By age, inequality indices are broadly similar across groups with overlapping intervals (Gini $\approx$ 0.28–0.30, Theil $T \approx$ 0.12–0.14, Q90/Q10 $\approx$ 3.4–3.9); mean concentration differs more clearly but only modestly: the 55–64 group is highest (about 12.5–13 $\mu g\,m^{-3}$) whereas 65+ is lowest (about 10–10.5 $\mu g\,m^{-3}$). By sex, the inequality indices are nearly identical; means differ only slightly, with women higher by roughly 0.5–0.7 $\mu g\,m^{-3}$, and intervals largely overlapping. Given the small and statistically non-discernible differences for age and sex, we do not include those plots here.

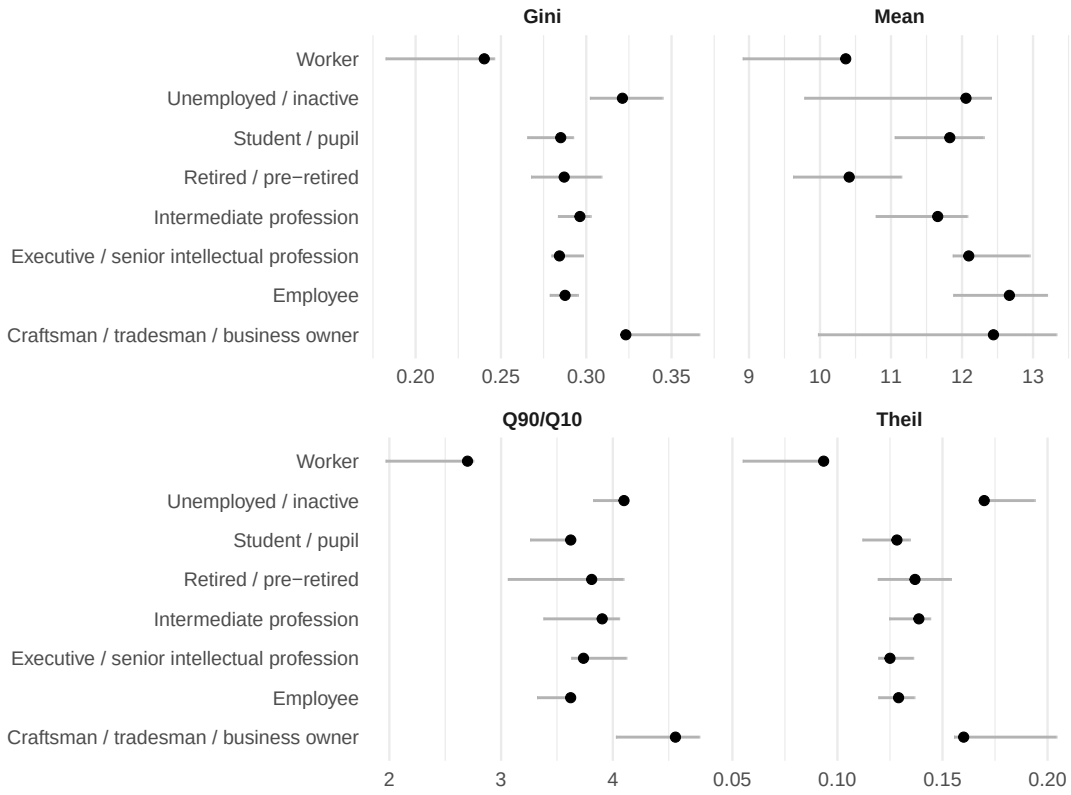

Fig. 1. Survey-weighted bootstrap 95% intervals for mean, Gini, Theil T, and Q90/Q10 by socio-professional category. Intervals for Craftsmen and Unemployed / inactive are shifted toward higher inequality while means remain similar across groups.

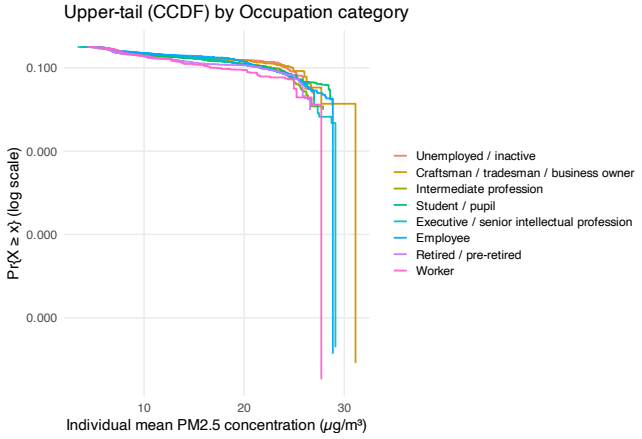

Fig. 2. CCDFs (log scale) of individual time-weighted mean $PM_{2.5}$ by socio-professional category. Heavier upper tails for Craftsmen/tradespeople and Unemployed/inactive indicate larger shares at high exposure despite similar means.

3. DISCUSSION

Our results show a clear pattern: different occupations have *very different spreads* of $PM_{2.5}$ exposure even though their *average* exposure is much the same. In Fig. 1, the occupation comparison highlights that *Craftsmen/tradespeople/business owners* and *Unemployed/inactive* have noticeably more unequal exposure profiles than *Employees* and *Workers*. Fig. 2 makes this concrete: a larger share of people in those two groups end up in the *high-exposure tail*, meaning they experience many more hours at higher pollution levels.

Looking only at averages can hide these pockets of high exposure. Alongside efforts to lower citywide averages, policies should also aim to shrink the high-exposure tail—e.g., by targeting specific places, times, and routines where certain occupations face worse air (workshops, delivery routes, peak-hour tasks), so that fewer people are pushed into the highest exposure levels.

REFERENCES

1. J. Ma, Y. Tao, M.-P. Kwan, and Y. Chai, "Assessing Mobility-Based Real-Time Air Pollution Exposure in Space and Time Using Smart Sensors and GPS Trajectories in Beijing," Ann. Am. Assoc. Geogr. **110**, 434–448 (2020).
2. M.-P. Kwan, "The Neighborhood Effect Averaging Problem (NEAP): An Elusive Confounder of the Neighborhood Effect," Int. J. Environ. Res. Public Health **15**, 1841 (2018).
3. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The NetMob25 Dataset: A High-resolution Multi-layered View of Individual Mobility in Greater Paris Region," (2025).
4. Copernicus Atmosphere Monitoring Service, "CAMS European air quality reanalyses," (2021).

Decoding Urban Interest: The Roles of Purpose, Transportation and Spatial Scale in City Mobility

GUSTAVO H. SANTOS[1,*], MYRIAM DELGADO[1], DANIEL SILVER[2], AND THIAGO H. SILVA[1,2]

[1] Universidade Tecnológica Federal do Paraná, Curitiba, Brazil
[2] University of Toronto, Toronto, Canada
[*] gustavohenriquesantos@alunos.utfpr.edu.br, myriamdelg@utfpr.edu.br, {dan.silver, th.silva}@utoronto.ca

1. INTRODUCTION

Understanding urban mobility is essential for better city design, yet traditional models often overlook why trips occur and how people travel. Using the high-resolution NetMob25 dataset, we show that decomposing trips by purpose and transport mode may reveal distinct, more explainable patterns of urban interest. Our models achieve strong explanatory power in some cases, e.g., explained deviance (D^2) ≈ 0.7 for purchase and walking, but struggle to explain trips in certain instances, such as purchase visits made with a Car. We also explore Interest Networks (iNETs) based on co-visitation patterns to model urban dynamics beyond individual locations. The iNETs reveal that shared interest between areas is predictable for some purposes, like leisure, but not for others, such as work, highlighting how particular trip purposes shape collective mobility. This effect weakens at finer spatial scales (IRIS), where even our best-performing variables lose contrast, underscoring the need for more granular (e.g, communes) and context-aware descriptors.

2. DATA

We study data from the NetMob25 dataset. It captures a week of detailed mobility from over 3,300 Île-de-France residents using high-frequency GPS devices and validated travel diaries with the observation weeks drawn between Oct 2022 and May 2023. Each of 80,000+ trips is annotated with purpose and main transport modes. Structured across individual profiles, trip records, and raw trajectories, the dataset includes calibrated weights to support population-level analysis by correcting for sociodemographic and temporal sampling biases. We apply these weights in the visit counts and co-visitation analyses.

At the commune scale, we examine every municipality in Île-de-France, including the Paris *arrondissements*. From the official French data portal[1], we get key socioeconomic indicators: population, number of families, median disposable income, share of bachelor holders, social disadvantage index, and employment rate. For the IRIS scale (areas around 2000 residents), we get population data through the INSEE [2] platform. For both the commune and IRIS scale, Overture Maps[3] venues enrich the context; applying Scenes Theory[1], we average culturally weighted venue categories to derive 15 cultural dimensions for each commune and IRIS (see [2] for methodological details). Finally, OpenStreetMap road data, extracted via pyrosm, yield driving road intersection density (nodes with degree ≥ 3), completing our multidimensional dataset.

3. METHODOLOGY

We adopt a two-phase approach to model interest in urban areas: Phase i) to study the relationship between area features and the interest, and Phase ii) to study the co-visitation pattern between areas. We define the interest score for a region r (commune or IRIS) as $I_r = \text{round}\left(\sum_{u \in U_r} w_u\right)$, where U_r is the set of visitors to r, and $w_u \in \Re$ is the population weight of visitor u. We exclude visit frequency, aiming to capture the breadth of interest and to avoid skew from high-frequency visitors. Additionally, visits are labeled by trip purpose and main transport mode (for example, work trips made by car, or all leisure trips, regardless of mode). This allows us to identify whether a region functions primarily as a leisure destination, a work hub, or is preferentially accessed via specific modes of transport.

In phase i), interest scores are modeled using a Generalized Linear Model (GLM) with a Poisson distribution to obtain interpretable and comparable coefficients. Other distributions, including Negative Binomial, are explored, but the Poisson specification provides higher explained deviance (D^2) and shows no significant overdispersion. Model performance is assessed through random 5-fold cross-validation, while accounting for potential spatial autocorrelation. All available features described in Sec. 2 are initially considered; however, due to multicollinearity (high VIF), only four predictors with VIF $\lesssim 5$ are retained: population (available at both commune and IRIS levels), total venues, driving-intersection density, and the composite "Theatricality" cultural score. Theatricality represents the magnitude of a vector composed of the cultural dimensions Glamour, Neighborliness, Transgression, Formality, and Exhibitionism, and reflects how a place presents itself—higher values typically occur in areas rich in landmarks, galleries, and entertainment venues.

While in the first phase, the proposed approach examines interest within areas, in Phase ii), it captures relationships between areas via shared visitors. This is helpful, for example, in guiding urban planners to promote more connections between regions of a city. To model this shared interest, we construct the iNETs: weighted undirected graphs where nodes represent regions (commune or IRIS), and edge weights are given by $w_{ij} = \text{round}\left(\sum_{u \in U_{ij}} w_u\right)$, with U_{ij} as the set of users who visited both i and j. Again, we use a Generalized Linear Model (GLM) with a Poisson distribution and a random 5-fold cross-validation. From the full set of candidate predictors, those with acceptable multicollinearity (VIF $\lesssim 5$) are retained: geographic distance between centroids, Euclidean distance of cultural vectors, cosine similarity of venue frequency distributions, and absolute differences in population and income (only population at IRIS-level, due to data availability). We also include average values of population, income, venues, and road density across the region pair as predictors. All predictors are z-scaled after aggregation or distance computation.

4. RESULTS AND DISCUSSION

To better contextualize the studied area (Ilê-de-France), we map dominant visit purposes (Fig. 1a). Paris (black borders) splits into leisure-focused arrondissements with tourist landmarks, work-oriented zones, and some residential areas. In terms of

[1] https://www.data.gouv.fr/
[2] https://www.insee.fr/
[3] https://overturemaps.org/

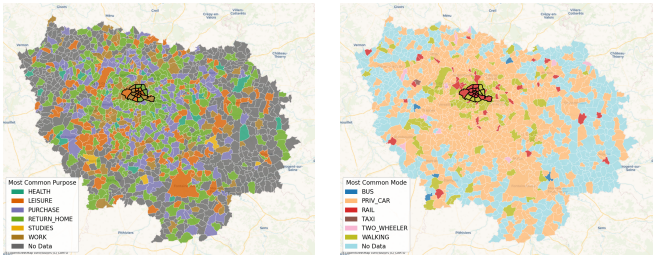

(a) Dominant visit purposes **(b)** Dominant visit transportation

Fig. 1. Analysis of visit patterns

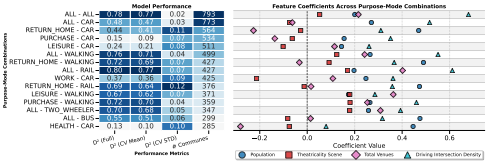

Fig. 2. Visitation GLM model performance and coefficients

transport (Fig.1b), Car dominates most of Île-de-France, while Paris relies mainly on rail and walking.

In Phase i), we apply the GLM to the commune-visitation data and obtain the results shown in Fig. 2. It reports the model's deviance-explained score (D^2 in the left) and the estimated coefficients (right) for each Purpose + Transportation combination, considering the 15 cases with the largest data volume, ordered by the number of communes exhibiting that type of visit.

Interest is well explained when considering all visits (ALL-ALL, i.e., purpose and transport mode are unspecified, with $D^2 = 0.77 \pm 0.02$), with road density as the main predictor. Among transport modes, Rail stands out ($D^2 = 0.77 \pm 0.07$ without purpose filter), where total venue count plays a stronger role compared to unfiltered visits. Return Home visits are well-explained across modes, especially walking and Rail ($D^2 = 0.69 \pm 0.04$ and $D^2 = 0.64 \pm 0.12$, respectively), driven by population and road density. Leisure and Purchase trips show similar explainability, in walking, with total venues and Theatricality emerging as stronger predictors. However, model explainability drops sharply for trips made by Car, the region's dominant mode (Fig. 1b), and varies by purpose: while Return Home by car achieves moderate explainability ($D^2 = 0.41 \pm 0.11$), Purchase trips are poorly explained ($D^2 = 0.09 \pm 0.11$). This highlights how trip purposes influence the way less constrained travel models, such as cars, interact with the urban form.

At the IRIS scale, the GLM explains almost none of the variation in interest ($D^2 \leq 0.03$), regardless of purpose–mode combination. Only a few cases, such as ALL–ALL, ALL–CAR, and ALL–RAIL, reach modest values ($D^2 \sim 0.12$). This sharp drop in explainability suggests that visits to micro-zones (with about four IRIS per commune) are less influenced by the general area characteristics we analyzed. Predictors that help differentiate broader areas lose contrast at this finer resolution—for example, intersection density can separate urban cores from suburbs, but is nearly uniform across adjacent IRIS within the same neighborhood. Under this uniformity, current predictors likely fail to capture the localized and purpose-specific factors that shape interest. Incorporating finer-grained features, segmenting by demographic group, or focusing on specific regions (e.g., central Paris) may offer better explanatory power, due to localized patterns (e.g., high tourism).

In Phase ii) of our study, these insights help understand the results of shared interests modeling across the city. Using iNETs,

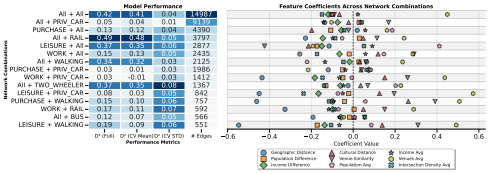

Fig. 3. iNETs GLM model performance and coefficients

we model networks of co-visitation where visits share the same purpose and transport mode. We then fit a separate GLM to each of these network combinations. The general iNET (ALL-ALL) has one of the strongest explainability ($D^2 = 0.41 \pm 0.04$, in Fig. 3, left), with shared interest between communes driven by proximity, venue density, and similarity in venue types (Fig. 3, right). A similar explainability holds for the Rail iNET, without considering purpose, and the Leisure iNET, without considering transport mode, where lower intersection density also plays a role. This is consistent with Fig. 1a, which highlights leisure visits to landmark-rich communes, such as those with castles or historical sites. In contrast, models for other purposes, such as Work, Purchase, and Health, explain little variation in shared interest. This divergence occurs because shared interest in leisure is likely driven by visitors seeking areas with comparable amenities and cultural scenes, whereas work co-visitation is tied to specific destinations, like two different office buildings, whose locations are uncorrelated by the analyzed area-level features.

In the IRIS iNETs, the GLM with area-level characteristics again fails to explain the interest ($D^2 \lesssim 0.03$). This occurs across all purpose–mode combinations, reinforcing earlier findings that the analyzed area-level features are insufficient at this finer scale.

5. CONCLUSION

Decomposing urban mobility by trip purpose and transport mode reveals behavioral nuances often missed in aggregated analysis. For communes, purchase/leisure - walking visits align well with area-level descriptors like venue density and cultural vibrancy; Return home and, separately, Rail trips align with population, while car-based and fine-scale patterns (e.g., IRIS level) are less predictable. Interest Networks (iNETs) further expose how trip purpose and transportation mode shape shared interest across regions, with leisure purpose and Rail transportation showing higher explainability than work or purchase-related visits with our current features. The understanding of how leisure areas are connected, for example, may guide urban planners in actions to promote more connections between different areas of the cities. Finally, these findings suggest that, because some cases show strong explanatory power while others do not, improving the spatial descriptors, by incorporating finer-grained data, spatial lag, demographic segmentation, land usage and studying temporal dynamics can deepen our understanding of urban interest across multiple scales and better capture the diversity of mobility behaviors.

6. ACKNOWLEDGEMENTS

FAPESP (grant 2023/00148-0) and CNPq (grants 314603/2023-9, 441444/2023-7, 409669/2024-5, and 444724/2024-9).

REFERENCES

1. D. A. Silver and T. N. Clark, *Scenescapes: How qualities of place shape social life* (The University of Chicago, 2016).
2. F. R. Gubert, G. H. Santos, M. Delgado, *et al.*, "Culture fingerprint: Identification of culturally similar urban areas using google places data," in *Proc. of ASONAM*, (Rende, Italy, 2025), pp. 286–297.

Urban Real Accessibility Evaluation Based on GPS Data: Integration of Space Syntax and Deep Learning

FU JIANYU,[1,*] JUE MA,[2] YOSHIHIDE SEKIMOTO[2]

[1]*Center for Spatial Information Science, The University of Tokyo, 4-6-1 Komaba, Meguro-ku, Tokyo 153-8505, Japan*
[2]*Center for Spatial Information Science, The University of Tokyo, 4-6-1 Komaba, Meguro-ku, Tokyo 153-8505, Japan*
** fu-jianyu@g.ecc.u-tokyo.ac.jp*

Received XX Month XXXX; revised XX Month, XXXX; accepted XX Month XXXX; posted XX Month XXXX (Doc. ID XXXXX); published XX Month XXXX

1. INTRODUCTION

Cities are evolving systems where rising travel demand makes accessibility a core determinant of urban equity and efficiency. Standard accessibility measures built on network geometry and facility distributions are useful but often diverge from how people actually move. A key gap is the lack of behavior-grounded, GPS-based accessibility that can be generalized beyond observed coverage while remaining interpretable. We address this by computing a trajectory-only Composite Accessibility Index (CAI) on pedestrian edges actually used in Grand Paris, combining reachability probability, observed travel-time cost, and detour index. We then derive space syntax metrics—integration, connectivity, choice—on the OSM walking network and train a lightweight MLP with syntax as inputs and CAI as target. The trained model is applied to the Paris commune, where trajectories are sparse, to predict citywide real accessibility from urban form[1-4].

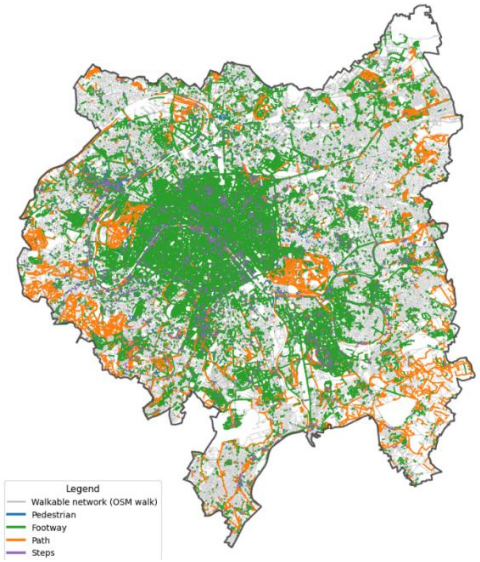

Figure 1. Pedestrian (walking) network of the Métropole du Grand Paris (source: OSM)

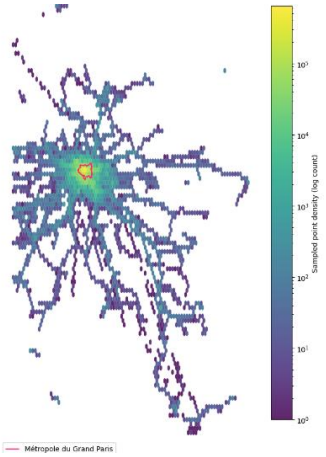

Figure 2. Trajectory Density(sampled)-Hexbin with Grand Paris Boundary

2. Data and Methods

A. Data Sources and Description

This study uses two datasets from the Paris region: GPS trajectories and the pedestrian transportation network. The GPS data were collected from 3,337 volunteers between October 2022 and May 2023, with each record containing latitude–longitude, timestamp, and (when available) speed. We retain only trajectory points that fall within the Métropole du Grand Paris administrative polygon (Figure 1) and reproject all data to EPSG:2154 before processing. The filtered trajectories are map-matched to the walking network to derive edge-level usage and travel-time signals used to compute the Composite Accessibility Index (CAI). The transportation network is extracted from OpenStreetMap for Grand Paris, restricted to walkable edges (pedestrian streets, footways, paths, sidewalks, steps, and streets permitting walking). We simplify the network, build an undirected graph, and compute space syntax metrics—integration, connectivity, and choice—for all edges. These two sources are sufficient for our pipeline: GPS yields behavior-grounded CAI on covered edges.

B. Methods and Models

1. Preprocessing and Map Matching

We use the OSM walking network for Métropole du Grand Paris and keep only GPS points inside that boundary. Trajectories and network are reprojected to EPSG:2154. After cleaning (time ordering, deduplication, gap-based trip segmentation), points are snapped to nearby edges via a spatial index; discontinuities are bridged with shortest paths on the pedestrian graph. For each trip we retain an ordered edge sequence, total path length, straight-line OD distance, and per-edge time shares (allocated by length proportion). Implausible speeds and one-point segments are removed. The result is a set of map-matched trips with edge-level usage and time, ready for CAI.

2 Composite Real Accessibility (CAI) from GPS

We compute CAI only on pedestrian edges covered by matched GPS trips. Three behavior-grounded components are derived per edge: (i) usage R, from the count of distinct trips; (ii) time efficiency T, comparing the mean observed per-pass time with a reference (free-flow or low quantile); and (iii) detour D, from the trip path-length to straight-line OD ratio, averaged back to edges. Each component is normalized to [0,1] and combined with fixed or learnable weights (default equal) into a single index:

$$CAI_e = \widehat{w_1} R_e + \widehat{w_2} T_e + \widehat{w_3} D_e$$

When coverage is sparse, optional graph Laplacian regularization can stabilize neighboring edges.

3 Space Syntax on the Walking Networkel

e compute three space-syntax features on the Grand Paris pedestrian network using its line graph (each street segment becomes a node; adjacency means walkable continuity). Connectivity counts incident segments at the two endpoints (local branching). Integration is line-graph closeness (shorter average shortest-path distance → more central). Choice is line-graph betweenness (how often a segment lies on shortest paths between other segment pairs → likely through-corridor). Paths use metric length as cost; all features are min–max normalized to [0,1] and used as predictors in the MLP.

4 Learning and Citywide Prediction

We train a lightweight MLP on edges where CAI is available, using space syntax features as inputs and CAI as the target. The feature vector for each edge is:

$$X_e = [, I_e, K_e, Ch_e]$$

The model outputs a prediction:

$$\hat{y}_e = f_\theta(X_e)$$

Parameters are estimated by minimizing mean-squared error, with an optional line-graph smoothness prior that encourages neighboring edges to have similar predictions:

$$\min_{\theta} \sum_{e \in E_{train}} \left(CAI_e - f_{\theta(X_e)} \right)^2 + \lambda \sum_{(e,e') \in E_L} \left(f_{\theta(X_e)} - f_{\theta\left(X_{\{e'\}}\right)} \right)^2$$

After spatial block cross-validation, we apply the trained model to edges in the Paris commune where GPS coverage is sparse to obtain citywide predictions:

$$\widehat{CAI}_e = f_{(\theta^*)(X_e)}$$

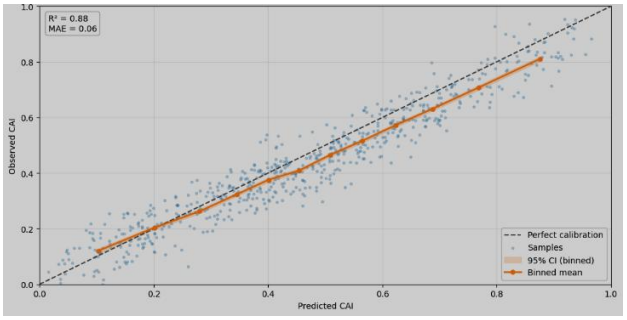

Figure 3. Calibration of predicted CAI.

Figure 4. Deep learning modeling and CAI prediction.

3. DISCUSSION

This study builds an edge-level Composite Accessibility Index (CAI) from GPS trajectories in Grand Paris and uses space syntax features with an MLP to predict CAI citywide. The GPS-derived CAI captures revealed accessibility on actually used pedestrian segments, emphasizing efficiency and directness. The syntax-based MLP explains a substantial share of CAI variance on covered edges (report R^2/MAE in the paper) and extrapolates to the Paris commune where trajectories are sparse. Largest errors appear along peripheral or low-sample corridors, indicating sensitivity to coverage and map-matching noise; residual clustering suggests factors beyond syntax. Diagnostics typically show integration and choice as main drivers, with connectivity shaping local detail. The approach links observed behavior with urban form and produces interpretable, city-scale predictions without POIs or time-slicing. Main limitations are dependence on GPS quality, lack of temporal stratification, and potential bias where walking is under-sampled.

References

1. M. Reneland, “Using GIS to calculate urban network accessibility,” (2006).
2. A. F. Stewart, “Advancing accessibility: public transport and urban space,” (2017).
3. I. Benenson, K. Martens, and Y. Rofè, Transp. Res. Rec. 2144, 28–35 (2010).
4. S. Sathisan and N. Srinivasan, Transp. Res. Rec. 1617, 78–83 (1998).

Optimizing ad-placement for socio-demographic segments using mobility data

SANTIAGO M. BARRÓN BUCOLO[1], TOMÁS DÍAZ[1], JEREMÍAS FIGUEIREDO PASCHMANN[1], JUAN KAPLAN[1], FRANCISCO NATTERO[1], AND MARIANO G. BEIRÓ[1,2*]

[1] *Universidad de San Andrés, Vito Dumas 284 (B1644BID), Victoria, Argentina*
[2] *Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), Buenos Aires, Argentina*
[*] *mbeiro@udesa.edu.ar*

1. INTRODUCTION

Optimizing ad placement is a central challenge at the intersection of data science, urban mobility, and personalized marketing. In this work we propose a predictive model that, using anonymized mobility data collected in the Île-de-France region (Paris and surroundings), estimates the probability that individuals belonging to specific socio-demographic segments will go through an urban area. Our aim is to generate probabilistic heatmaps that identify, for a socio-demographic segment, the areas most commonly visited, thus enabling strategic decision-making in geolocated ad-placement campaigns.

2. METHODOLOGY

A. Dataset Description

This study uses data from the 2023 EMG Survey, which includes one week of GPS tracking for 3,337 Île-de-France residents aged 16 to 80. The dataset comprises socio-demographic information, trip details, and raw GNSS traces. In this preliminary stage, we focus on the individual data and GNSS traces, limiting the analysis to GPS points within Paris.

B. Target Construction

The primary objective of this study is to estimate the probability that individuals will visit different areas of the city. As this information is not directly available in the dataset, we developed a pipeline to infer it from the GNSS data.

We focused on estimating visit probabilities for each individual in the dataset. For every user, we collected all GNSS traces within the city of Paris, and we fitted a Gaussian Mixture Model (GMM) to capture the spatial distribution of their movements. The optimal number of clusters was determined using the Bayesian Information Criterion (BIC), with a mean of 4.5 clusters across all users.

Subsequently, the city was divided into a 50x50 spatial grid and, for each user, we computed the probability of visiting each grid cell based on their personal GMM distribution. These probability values were then used as the target variable in the predictive modeling stage.

C. Feature Engineering

To identify meaningful subgroups and uncover latent behavioral profiles, we applied unsupervised techniques—GMM, Principal Component Analysis (PCA), and Latent Class Analysis (LCA)—to the full socio-demographic dataset. Accordingly, 41 engineered features were appended to each observation—10 cluster-membership probabilities from GMM, 20 principal components derived via Principal Component Analysis (PCA), and 11 cluster-membership probabilities from Latent Class Analysis (LCA).

In addition, minor preprocessing steps were applied to ensure data consistency. Missing values were handled by assigning default values per column, and categorical variables were label-encoded to ensure all features were represented as floats.

D. Dataset Split and Data Leakage Prevention

The dataset was split into 85% for training and 15% for testing. To prevent data leakage, all observations corresponding to a given user were assigned exclusively to either the training or the test set, ensuring no overlap between user data across splits.

E. Predictive Modeling and Training

We modeled the target with XGBoost, a gradient-boosted decision-tree method. After discretizing the spatial domain into grid cells, the mobility-probability task becomes a tabular prediction problem, for which XGBoost is a natural fit. Because the target is defined over latitude–longitude–derived cells, the model is well suited to distinguishing among these discrete spatial regions. Training used a squared-error objective, and performance in training was evaluated with root mean squared error (RMSE).

F. Prediction Validation

It is important to clarify that the goal of the project is not to replicate a user's exact mobility pattern. A model that perfectly reproduces an individual's trajectory would risk severe overfitting, merely memorizing their specific path during the measurement period. Instead, the objective is to generalize across users with similar socio–demographic characteristics. The aim is for the predicted heatmap to cover the regions a user *might* visit, without exactly reproducing their personal route. In doing so, the model captures broader behavioral tendencies of the target group rather than the unique patterns of any single individual, which is essential for scalable and privacy–aware applications such as targeted advertising.

This emphasis on *coverage* is particularly relevant in non–deterministic mobility: two people with *identical* observed features may select completely different routes on a given day. Under such uncertainty, a single "ground–truth" map is not uniquely defined, making many traditional validation metrics not a viable option. Therefore, to evaluate model performance, each user's *predicted heatmap* was compared against their GNSS points using **hit@K** as the main metric. Concretely, the 50×50 grid is ranked by the model's predicted probability; the top K cells form the model's high–confidence region. A user contributes a hit if a GNSS point falls inside this region, and *hit@K* is the fraction of that user's points that are covered. Finally, to evaluate the global performance of the model, the *hit@K* of every user is averaged, weighted by the amount of GNSS points logged by each user.

3. CASE STUDIES

To illustrate the practical value of our model, we present three case studies that simulate real-world advertising use cases. The

first case considers users with demographic profiles consistent with high-end automotive consumers, specifically, potential BMW buyers. The second profile targets individuals living in households with older adults, relevant for advertisement campaigns related to elder care services or products. The third profile targets educated parents living with underage children; the profile was designed to appeal to extra-curricular academical activities such as foreign language education.

To construct all three profiles, we set socio-demographic features to match typical attributes within each segment and then identified matching individuals in the test set. For the BMW-buyer profile, we drew on demographic insights reported by Hedges and Company [1]. We computed each individual's probability heat map, aggregated them within each case study, and applied a final min–max normalization.

As shown in Figure 1, the predicted mobility distributions differ substantially between the three socio-demographic profiles. The profile associated with potential BMW buyers exhibits a more spatially dispersed pattern, indicating broader movement across the city. This suggests that individuals in this segment may have greater access to private transportation and fewer mobility constraints. The highest-probability areas in the map are close to Boulevard Haussmann, Gare de Lyon and 6me arrondisement. In contrast, the mobility of individuals living in households with a person over 65 is more spatially concentrated around main train stations, likely reflecting limitations related to transportation mode or physical ability. In this case, the area with highest probability is Gare du Nord - Gare de l'Est. Finally the mobility of individuals with underage children is concentrated around the 3eme, 4eme, and 10eme arrondisements. These distinct patterns highlight the model's capacity to capture relevant behavioral differences across segments. From a marketing perspective, such insights can inform targeted ad-placement strategies by identifying high-affinity zones and eliminating low-yield areas where the target audience is unlikely to be present.

Upon evaluating the model over each case study, along with the entire test set, using the **hit@K** metric varying K (the K highest ranked cells in the model's prediction) the results are shown on Table 1. It is worth noting that, for each case study, the metric was obtained by evaluating the users with the highest survey weight, a variable indicating how many individuals in the total population each respondent represents. The test set was evaluated as discussed in Section 2.F.

Metric (%map)	Elderly	BMW	Children	Full
Hit@100 (4%)	0.589	0.633	0.514	0.432
Hit@200 (8%)	0.769	0.772	0.869	0.650
Hit@250 (10%)	0.990	0.772	0.928	0.734
Hit@300 (12%)	1.0	1.0	0.976	0.820
Hit@375 (15%)	1.0	1.0	1.0	0.9
Mean Prob.	0.777	0.810	0.780	0.668

Table 1. Hit@K results across different subsets of the test set. The mean probability is computed by averaging, for each user, the probabilities of the GNSS traces under that user's heatmap.

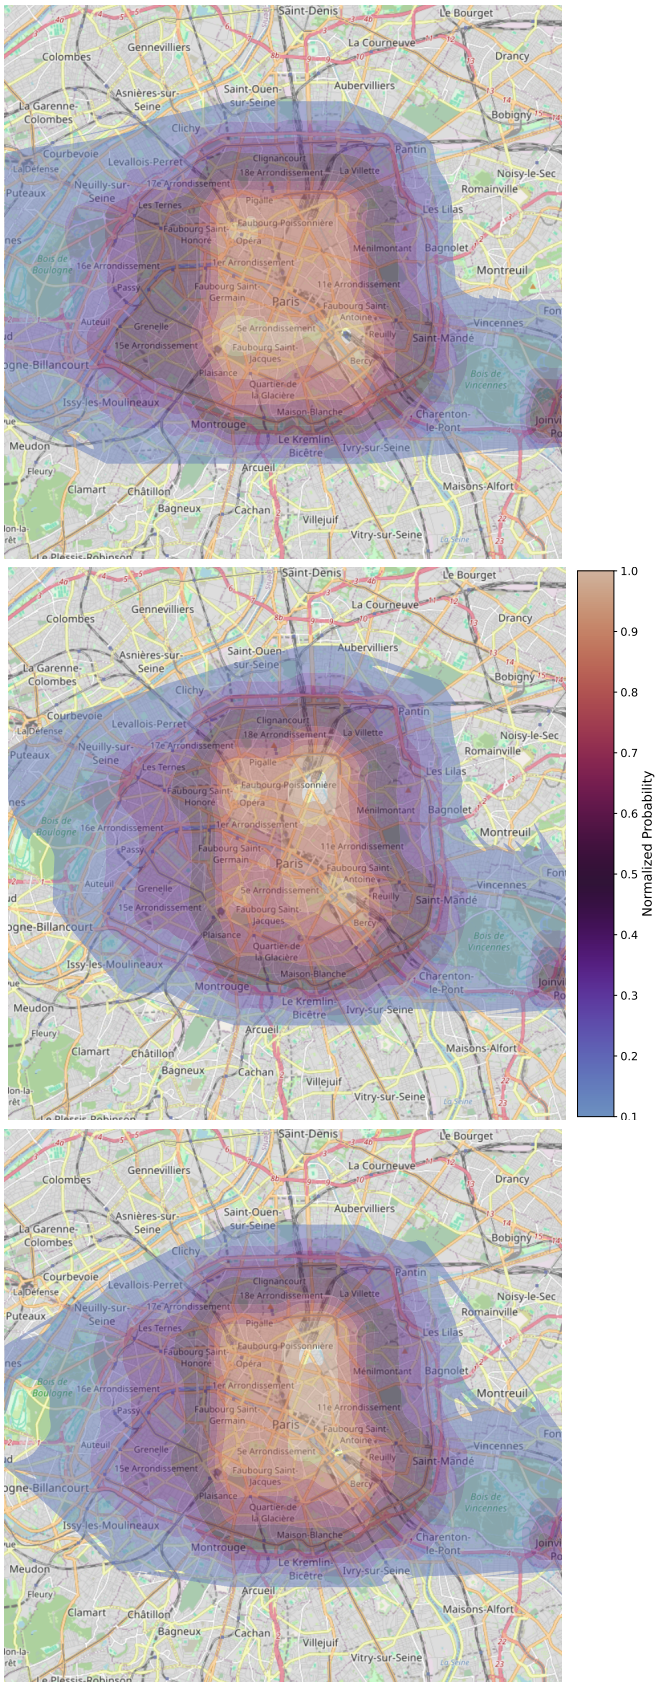

Fig. 1. Heatmaps for the three case studies: (Top) BMW, (Mid) Older adults, (Bottom) parents living with underage children.

REFERENCES

1. Hedges and Company, New BMW owner demographics: Who drives a BMW anyway? https://hedgescompany.com/blog/2019/03/new-bmw-owner-demographics/, Retrieved: 14 July 2025.

Measuring Individual Exploration in Human Mobility: A Novel Metric Based on the Returner-Explorer Dichotomy

AKIHIRO FUJIHARA[1,*]

[1] Graduate School of Engineering, Chiba Institute of Technology, 2-17-1 Tsudanuma, Narashino, Chiba 275-0016, JAPAN
[*] akihiro.fujihara@p.chibakoudai.jp

1. INTRODUCTION

Understanding human mobility is essential for various fields such as public health, urban planning, and economic forecasting. The increasing availability of real-world datasets (e.g., mobile phone records, GPS traces) has enabled researchers to uncover statistical regularities in movement patterns. While human mobility is generally predictable [1], low-frequency exploratory behavior remains difficult to model, yet crucial for improving prediction accuracy. This challenge has motivated theoretical studies using mathematical models and data analyses.

Song *et al.* demonstrated that human displacements and waiting times follow power-law distributions, leading them to propose the Individual Mobility (IM) model [2], which captures movement as a balance between exploration (visiting new places) and preferential return (revisiting familiar ones). Though empirically valid, this model lacks explicit trajectories. To address this, Fujihara and Miwa introduced the Homesick Lévy Walk (HLW) model [3], which combines Lévy flights with probabilistic returns to a home base. HLW reproduces key properties observed in human interaction networks: one-time encounters ("Ichi-go Ichi-ye") and scale-free revisit distributions.

Beyond modeling mechanisms, on the contrary, Pappalardo et al. [4] revealed that individuals tend to polarize into returners and explorers. This behavioral dichotomy is captured by the s_k index, which reflects the diversity and ranking of visited locations. However, no metric has been proposed to quantify the degree of exploration per individual grounded in the returner-explorer dichotomy, limiting comparative and demographic analyses.

This paper addresses this gap by introducing a novel indicator of exploratory behavior. Our main contributions are as follows:

- We analyze the k-dependence of s_k in the NetMob25 dataset[5] and confirm that the returner–explorer dichotomy holds. By extending previous studies that only considered $k \leq 10$ to $k \leq 100$, we find that while s_k generally increases with k, some individuals show significant decreases, revealing complex k-dependence of s_k.

- We also find that the distribution of k^*, *i.e.*, the smallest k such that $s_k > 0.5$, follows a power-law, indicating diverse variation in the transition point between behavioral types. We propose k^* as a new metric for quantifying individual exploratory tendencies.

2. RELATED WORKS

Several studies have examined the returner-explorer dichotomy using the s_i index. Cuttone *et al.* [6] proposed the *exploration rate*, that is, the ratio of new to total visited locations—and found that 20–25% of visits are typically to new places, though with substantial individual variation. Wang *et al.* [7] analyzed how temporal resolution affects classification, showing that returner-explorer labels can shift depending on the observation window.

He *et al.* [8] showed that even returner-type individuals may adopt explorer-like behavior in disaster scenarios, as reflected in changes in s_k and displacement patterns. Iwai and Fujihara [9] used reinforcement learning to simulate taxi mobility in San Francisco, demonstrating that optimal exploration behavior emerges at a specific exploration parameter ϵ, with learned trajectories exhibiting strong exploratory tendencies.

While these studies provide insights into individual mobility typologies and exploratory modeling, no prior work has proposed a systematic, empirical measure to quantify individual-level exploratory behavior from observed mobility trajectories based on the returner-explorer dichotomy.

3. METHOD

We use anonymized individual-level mobility traces collected from the NetMob2025 dataset [5]. Each user's trajectory consists of a time-ordered sequence of location visits, including geographic coordinates and timestamps.

Pappalardo *et al.* [4] introduced the index s_k to quantify the extent to which an individual's overall mobility is concentrated around their top k most frequently visited locations. This index is based on the concept of the radius of gyration, which characterizes the typical distance traveled by an individual:

$$r_g = \sqrt{\frac{1}{N}\sum_{i\in L} n_i \|\mathbf{r}_i - \mathbf{r}_{\mathrm{cm}}\|^2} \tag{1}$$

where L is the set of all visited locations, $\mathbf{r}_i$ is the position vector of location i, n_i is the number of visits at location i, $N = \sum_{i\in L} n_i$ is the total number of visits, and $\mathbf{r}_{\mathrm{cm}}$ is the center of mass of the trajectory. To focus on the top k locations, the k-radius of gyration is defined as:

$$r_g^{(k)} = \sqrt{\frac{1}{N_k}\sum_{i=1}^{k} n_i \|\mathbf{r}_i - \mathbf{r}_{\mathrm{cm}}^{(k)}\|^2} \tag{2}$$

where the locations are sorted in descending order of visitation frequency, $N_k = \sum_{i=1}^{k} n_i$, and $\mathbf{r}_{\mathrm{cm}}^{(k)}$ is the center of mass computed over the top k locations. Then, the exploration index s_k is defined as the ratio: $s_k = r_g^{(k)}/r_g$. In particular, individuals can be classified based on the threshold $s_k = 0.5$: those with $s_k > 0.5$ are considered *k-returners*, whose mobility is dominated by a small number of frequently visited places, whereas those with $s_k \leq 0.5$ are classified as *k-explorers*, who exhibit more diverse movement beyond the top k locations.

4. RESULTS

Figure 1 shows the distribution of s_k for $k = 2, 4, 6, 8$ in the NetMob25 dataset. Prominent peaks appear near $s_k = 0$ and $s_k = 1$, reflecting the returner-explorer dichotomy. For small k, explorer-like behavior ($s_k < 0.5$) dominates, but as k increases, returner-like behavior ($s_k \geq 0.5$) becomes more frequent, indicating a shift from exploratory to localized movement.

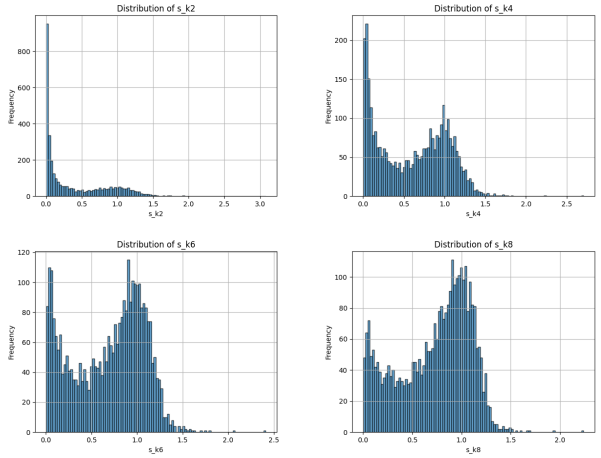

Fig. 1. The distribution of the ratio s_k for $k = 2, 4, 6, 8$.

Extending the analysis beyond $k \leq 10$, Figure 2 illustrates the evolution of s_k up to $k = 100$. While s_k generally increases with k, some individuals display non-monotonic trends, including local decreases or crossings of the threshold $s_k = 0.5$. These diverse patterns suggest that individual mobility cannot be fully characterized by the early behavior of $k \leq 10$ alone.

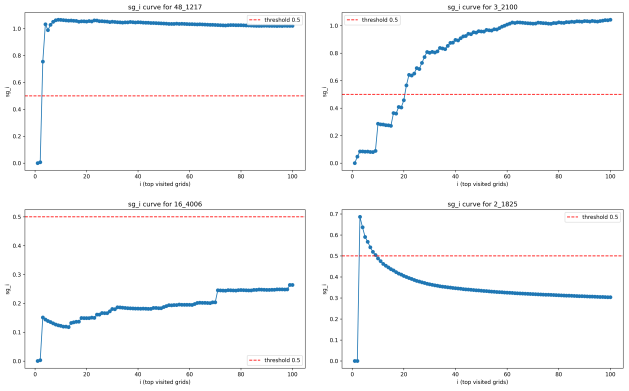

Fig. 2. Evolution of s_k as k varies from 1 to 100 for typical four individuals. The red line indicates the threshold $s_k = 0.5$.

Figure 3 shows the distribution of k^*, the smallest index at which $s_k > 0.5$. The distribution follows a power law with an exponent of 1.82, indicating substantial heterogeneity in individual exploratory behavior. This suggests that k^* reflects individual differences in the transition from exploration to return, and may serve as an effective indicator of exploratory tendency, consistent with prior findings on mobility heterogeneity [1].

Finally, we investigate whether various attribute values (`AGE`, `DIPLOMA`, `NB_CAR`, etc.) in the individual database provided by the NetMob 2025 Data Challenge are correlated with k^*. As a result, no clear correlation is observed with any of the attributes. We also compute entropy, r_g, and s_k to examine their correlations with k^*. Although some negatively correlated patterns could be vaguely discerned, no high correlation coefficients are obtained.

5. CONCLUSION AND FUTURE WORK

In summary, this study introduced a novel indicator to quantify individual exploratory behavior in human mobility, based on the cumulative visitation ratio s_k. By analyzing how s_k evolves

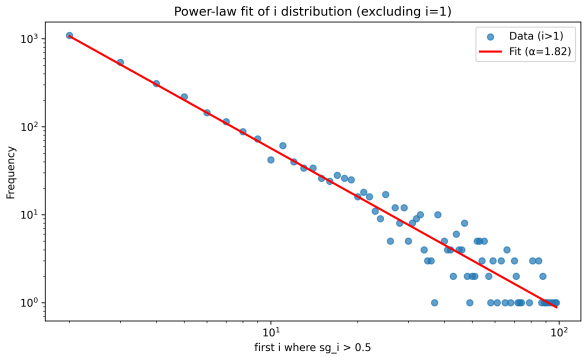

Fig. 3. Power-law distribution of k^*, the smallest index $k = k^*$ where $s_k > 0.5$.

with increasing location rank k, we identified the index k^*, defined as the smallest k such that $s_i > 0.5$, as an effective threshold distinguishing returner-like and explorer-like mobility patterns. Empirical analysis using the NetMob25 dataset showed that the distribution of k^* follows a power-law with exponent 1.82, revealing substantial heterogeneity in exploratory tendencies among individuals. These findings not only reinforce the returner-explorer dichotomy proposed in previous studies [4], but also offer new theoretical and empirical insights into the modeling and classification of human mobility.

We investigated whether the proposed new indicator k^*, which represents the degree of exploration, is correlated with various attribute values of the individual database from the NetMob 2025 Data Challenge, as well as with existing mobility metrics, such as entropy, r_g, and s_k. However, no clear correlations were observed. This suggests that k^* may constitute a new dimension of features distinct from existing attributes and metrics.

As future work, it is necessary to compute k^* using various human mobility datasets to verify whether power-law distributions can be universally observed. It would also be interesting to investigate, by means of k^* values, for which attribute values the disparities in explorer characteristics become more pronounced. Furthermore, extending existing mobility models, such as the IM model [2] and the HLW model [3], to theoretically explain the statistical properties of k^* through mathematical modeling is another important direction.

REFERENCES

1. C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, Science **327**, 1018 (2010).
2. C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, Nat. Phys. **6**, 818 (2010).
3. A. Fujihara and H. Miwa, "Homesick lévy walk: A mobility model having ichi-go ichi-e and scale-free properties of human encounters," in *2014 IEEE 38th Annual Computer Software and Applications Conference*, (2014), pp. 576–583.
4. L. Pappalardo, F. Simini, S. Rinzivillo, *et al.*, Nat. Commun. **6**, 8166 (2015).
5. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," (2025).
6. A. Cuttone, S. Lehmann, and M. C. González, EPJ Data Sci. **7**, 2 (2018).
7. R. Wang, N. Li, and Y. Wang, Sustain. Cities Soc. **69**, 102862 (2021).
8. Z. He, Y. Lv, S. Gao, *et al.*, Sci. Reports **14**, 13184 (2024).
9. Y. Iwai and A. Fujihara, "Considering a method for generating human mobility model by reinforcement learning," in *Advances in Intelligent Networking and Collaborative Systems*, L. Barolli, K. F. Li, and H. Miwa, eds. (Springer International Publishing, Cham, 2021), pp. 121–132.

The city as experienced: measuring individual differences in contextual exposure through mobility trajectories

LAURA SILVA[1], ANTONIO DESIDERIO[2,3,*], FRANCO BONOMI BEZZO[4], NANDINI IYER[5], AND RICCARDO DI CLEMENTE[5,2]

[1] Swiss Centre of Expertise in Life Course Research LIVES, University of Lausanne, Lausanne, 1015.
[2] ISI Foundation, Via Chisola 5, Turin, 10126.
[3] Department of Applied Mathematics and Computer Science, Technical University of Denmark, Copenhagen, 2800.
[4] Department of Social and Political Science, University of Milan, Milan, 20122.
[5] Complex Connections Lab, Network Science Institute, Northeastern University London, 58 St Katharine's Way, London, E1W 1LP.
[*] antonio.desiderio@isi.it

1. INTRODUCTION

Understanding how people experience the city requires moving beyond isolated measures of residential and work neighborhoods towards a more complex understanding of the diverse socioeconomic contexts individuals encounter daily. Individuals differ not only in the average levels of poverty or inequality they are exposed to, but also in the variability of these exposures across the places they traverse [1]. At the same time, mobility infrastructures and transport modes shape which neighborhoods people can access and the degree of socioeconomic diversity they encounter along the way. Finally, these patterns are structured by socio-demographic characteristics, such as education, which influence both the extent of exposure to deprivation and the breadth of contexts individuals move through [2, 3]. Examining these dimensions in combination allows us to uncover how urban mobility processes and social stratification intersect to generate unequal experiences of the city [4]. Because individuals often use their local environments as heuristic devices, the neighborhoods they pass through function as "inferential spaces" [5], providing interpretive cues about the social world that shape their attitudes and behaviors. This research can thereby inform urban policies designed to reduce socio-spatial inequalities.

2. DATA AND METHODS

We exploit data from the Île-de-France region. We use the NetMob25 dataset, which combines GPS-based mobility data collected from 3,337 volunteer residents between mid-October 2022 and mid-May 2023, covering one week per individual, with self-reported trip records and detailed demographic and socioeconomic information. We integrate these data with two contextual measures (poverty rate and Gini coefficient) available at the IRIS (which we define as our neighborhood unit) level from the 2021 Social and Fiscal Localized File System (FILOSOFI) [6] produced by INSEE (the French national institute for statistic and economic studies). For example, Figure 1 shows that the spatial distribution of poverty rates in Île-de-France is skewed towards values lower poverty, with most areas experiencing rates between approximately 10% and 40% , and a long tail extending toward higher values, especially in the north-east part of the region.

We employ a spatio-temporal matching algorithm to align each self-reported trip with its corresponding continuous GPS trace, thereby validating annotated journeys against objective location data. For each participant, our processed dataset comprises the full list of IRIS areas traversed with the precise duration in each, enabling detailed measurement of their distributed mobility exposures. For each individual, we calculate both

Fig. 1. Socioeconomic stratification of areas in Ile de France Spatial distribution of poverty rates for each IRIS area in the Paris region, which quantifies the socioeconomic context of locations.

the average and variability (standard deviation) of exposure to neighborhood poverty and inequality across all visited areas, excluding time spent in the home neighborhood. In the first part of the analysis, we focus on poverty: we construct a travel-to-stay poverty exposure matrix by comparing each trip's weighted average poverty exposure during travel with the poverty level at the destination. Both travel and destination poverty exposures are categorized into low, mid, and high bins. Transport modes were grouped into five categories based on shared mobility characteristics: Active Travel (walking, cycling, e-bikes, e-scooters), Private Vehicle (private cars as driver or passenger, motorcycles), Public Transit (trains, buses, subway, tramway), On-Demand/Taxi services, and Other modes (planes, commercial vehicles). We also measure within-trip poverty heterogeneity as the standard deviation of segment-level exposures, capturing the socioeconomic diversity encountered during each journey and allowing assessment of both exposure level and variability across transport modes. In the second part of the analysis, we extend the approach to both neighborhood poverty and inequality and use ordinary least squares (OLS) models to examine associations with individual characteristics, including gender, education, and socio-professional category.

3. RESULTS

Figure 2A illustrates individual differences in contextual exposure by comparing two users' week-long mobility trajectories. Each marker represents time spent traveling through or staying in a neighborhood (circle and square markers, respectively), with

size reflecting duration. The low-variability user (σ = 1.07) consistently encounters neighborhoods with similar poverty levels (8–12%), while the high-variability user (σ = 15.17) experiences a wide range (5–50%), moving between both low- and high-poverty areas.

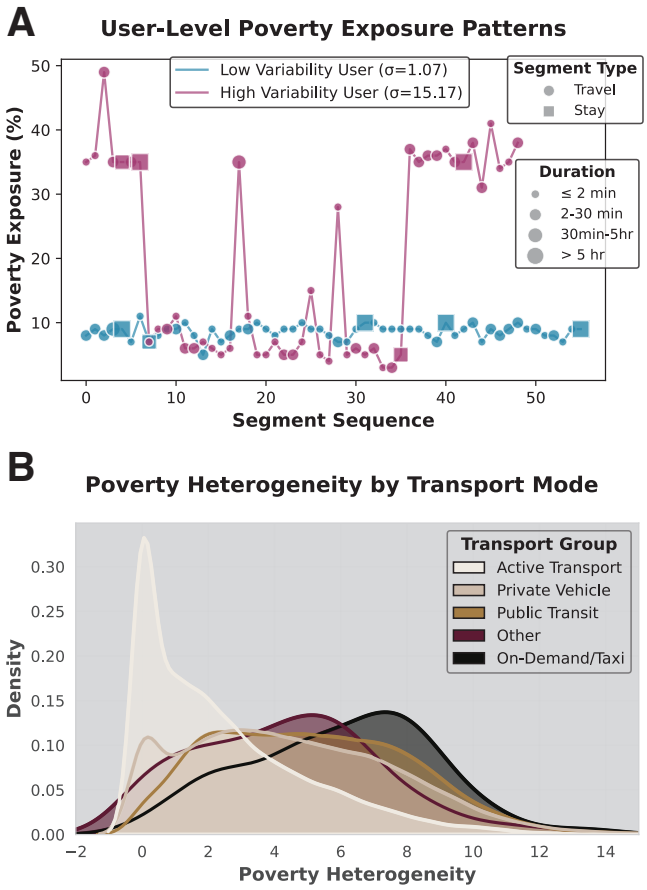

Fig. 2. A. Individual-Level Mobility Trajectories Illustrating Contextual Exposure. Weekly mobility trajectories for two users are plotted sequentially to show temporal variation in neighborhood poverty exposure. Circles denote travel segments through IRIS areas identified via the trip-matching process, and squares denote stays at specific locations. Marker size is proportional to the duration spent in each context – either traveling or stationary – highlighting differences in exposure profiles over the week. **B. Poverty Heterogeneity by Transport Modes.** Active travel trips (white) exhibit low heterogeneity while, while trips by private vehicles and public transit (beige and light brown, respectively) show broader distributions of heterogeneity in poverty exposure.

The travel-to-stay poverty exposure matrix reveals that transport mode strongly shapes average exposure to neighborhood poverty. We observe distinct mobility inequalities: public transit enables mobility across poverty levels, private vehicles mainly connect lower-poverty areas, active travel is spatially constrained to local, often higher-poverty neighborhoods, and on-demand services provide relatively balanced access with a slight bias toward higher-SES destinations. Moving beyond averages and examining within-trip heterogeneity in Figure 2B further reveals that active travel involves the least contextual diversity, while transit, private vehicles, and especially on-demand services expose users to more varied neighborhood conditions. Together, these results demonstrate how transport mode choice shapes both the range of socioeconomic contexts people encounter and the equity of urban mobility.

Secondly, we look into the effects of socio-demographic characteristics and find that gender, education, and individual socio-professional status all stratify not just the level of socioeconomic exposure but also its heterogeneity across daily mobility. Figure 3 shows that individuals with higher education levels are less exposed to poverty than those with lower education, yet they spend more time in unequal contexts (higher average Gini). Moreover, their trajectories are more heterogeneous: they move across neighborhoods with varying levels of poverty and inequality, unlike less-educated individuals whose mobility is more spatially confined. This suggests that education not only reduces exposure to deprivation but also broadens the range of socioeconomic environments people encounter in their daily lives.

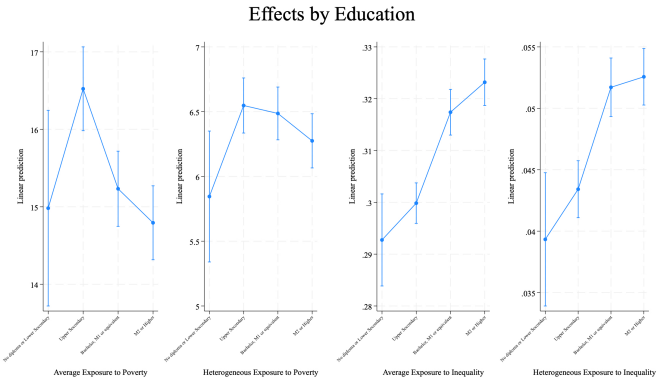

Fig. 3. Marginal effects of education on Average Exposure to Poverty, Heterogeneous Exposure to Poverty, Average Exposure to Inequality, Heterogeneous Exposure to Inequality

4. DISCUSSION

This study explores how urban experiences vary across social groups by analyzing the dynamic sequences of neighborhood contexts through which individuals travel. Our findings reveal that disadvantaged groups encounter neighborhoods with higher poverty ad less exposure variability. Contrast serves as a learning device [7, 8], limited variability in one's environmental exposures may reinforce existing perspectives and hinder awareness of alternative social realities. Limitations include not modeling trip purposes or other contextual factors like crime, highlighting the need for further research on individual mobility choices and urban infrastructure influences.

REFERENCES

1. C. R. Browning, J. Tarrence, C. A. Calder, *et al.*, Am. J. Sociol. **128**, 914 (2022). Publisher: University of Chicago Press.
2. Y. Liao, J. Gil, S. Yeh, *et al.*, Comput. Environ. Urban Syst. **117**, 102250 (2025). Publisher: Elsevier BV.
3. M. Macedo, L. Lotero, A. Cardillo, *et al.*, **17**, e0260874.
4. F. Xu, Q. Wang, E. Moro, *et al.*, Nat. Hum. Behav. **9**, 654 (2025). Publisher: Springer Science and Business Media LLC.
5. J. Mijs, Societies **8**, 64 (2018). Publisher: MDPI AG.
6. INSEE, "Social and fiscal localized income 2021 (filosofi)," https://www.insee.fr/en/metadonnees/source/serie/s1172 (2025). Accessed: 2025-09-18.
7. L. Festinger, Hum. Relations **7**, 117 (1954).
8. J. J. B. Mijs, Res. Soc. Stratif. Mobil. **85** (2023).

Zero-Shot Transportation Mode Classification in Île-de-France: Generalizing the PromptRide-LLM Pipeline

AMIR BADAWI[1], ANA-MARIA OLTEANU-RAIMOND[1], ARNAUD LE-GUILCHER[1], AND KARINE ZEITOUNI[2]

[1] Univ Gustave Eiffel, IGN-ENSG, LASTIG, France
[2] DAVID Lab, Université Paris-Saclay, Versailles, France
* amir.badawi@ign.fr

1. INTRODUCTION

Mobility data are central to urban analytics, and a key challenge is transportation mode detection (e.g., walk, bike, drive), which underpins policy and service design. Supervised approaches with hand-crafted features and city-specific labels [7] generalize poorly and require retraining. Critically, per-trip labels for active modes (i.e., walk, bike, run, scooter) are scarce or absent because surveys or operator metadata are intrusive and costly. This motivates label-efficient, zero-shot alternatives. Although LLMs are increasingly used in human mobility analysis [4, 5], prior work has not addressed transport mode detection.

To tackle these issues, we introduced **PromptRide-LLM**, a zero-shot pipeline that leverages LLMs to classify four transport modes (i.e., *bike, bus, car, and walk)* from natural-language trip summaries [2]. It was tested on the Geolife dataset (Beijing) as an interpretable, task-agnostic approach without labeled data.

Here, we extend PromptRide-LLM to NetMob25 (GPS trajectories from over 3,300 users in Île-de-France) and run four tests: (1) **E1**: the exact MDM PromptRide-LLM (unchanged) is applied to NetMob25; (2) **E2**: NetMob value-tuned numeric bands; (3) **E3**: human priors used only as soft tie-breakers (4-class); and (4) **E4**: the same as E3 but extended to *all* transport modes (10-class). We evaluate cross-city transfer and examine how *model family* (DeepSeek/32B vs. GPT-4o) and *answering style* (quick vs. step-by-step) interact with each setting, with emphasis on active mobility (walk, bike) versus motorized or transit classes.

Contributions. **LLM behavior & style:** First NetMob25 study of model family (DeepSeek/32B vs. GPT-4o) and answering style (quick vs. step-by-step) with explicit effects on *walk* and *bike* vs. *bus* and *car*. **Zero-shot transfer:** A generic hand-tuned prompt (E1) transfers to Île-de-France DeepSeek: Precision 0.559, macro-F1 0.554—establishing a practical baseline. **Prompting & priors:** Tight data-derived bands (E2) are brittle; soft p10–p90 envelopes and *relational cues* are more robust; user priors (E3) act as *soft* tie-breakers (quick favors bus and car, step-by-step restores bike and walk) without overriding strong kinematics. (E4) **Moving to 10 classes:** Rail modes remain difficult under GPS-only cues, motivating network-aware context (corridor alignment, station cadence).

2. PIPELINE

Common part. PromptRide-LLM follows a data→text→LLM flow with four stages: (1) *ingestion & cleaning* of GPS traces using plausibility filters and smoothing; (2) *trip segmentation* via small stops and speed guards into single-mode trips; (3) *feature enrichment* (distance, duration, avg and max speed, stop cadence, heading variability, etc); (4) *summary generation* into compact, human-readable trip descriptions consumed by an LLM for *zero-shot* mode inference (no training or fine-tuning). The summaries expose the same cues a human analyst would rely on, making decisions interpretable.

NetMob-specific additions. *Digital Elevation Model (DEM) augmentation.* Since NetMob25 lacks altitude, we sample a DEM along the GPS polyline to derive elevation gain and loss, grade (%), and vertical-speed proxies. These cues help separate active vs. motorized modes (e.g., bike and walk slowdowns on climbs) and stabilize stop detection on hilly segments.

Inference variants. We fix the primary model and decoding (**DeepSeek-R1-Distill-Qwen-32B**, T=0.1) and vary only prompt content; each variant is run with two answering styles: *quick* (one word) and *step-by-step* ($\leq$80 words then the label). In addition, we run a cross-model probe on **GPT-4o** with identical prompts and decoding to isolate model-family effects.

- **E1 — Baseline (generic, hand-tuned).** We use the structure and generic trial-and-error value bands as the original PromptRide-LLM ([2]). Tests cross-city transfer "as is."
- **E2 — Value-tuned (NetMob-aware).** We use the same structure as E1, but numeric bands and verbal cues are *derived from NetMob analytics* (per-mode ranges, stop-cadence hints, segment-length guards) and applied as *fixed thresholds*. The purpose was to check whether dataset-calibrated values help without changing the reasoning steps.
- **E3 — Human priors (NetMob metadata).** After kinematic and geometric checks (deceleration–acceleration ratio, stop uniformity, max-speed cap), we append a short user-profile block {`licence`, `cars_binned`, `navigo_sub`, `bike_idx_binned`, `bike_sub`, `two_wheeler_idx_binned`, `e_scooter_idx_binned`, `age_band`, `sex`}. Priors act strictly as soft tie-breakers they never override strong sensor evidence and never flip a clear walk.
- **E4 — Full label set (10 modes).** E3 logic is extended to ten modes: {*car, walk, train, subway, tram, bus, bike, ebike, escooter, moto*}; plane and ferry are excluded due to small support number.

This experiments protocol preserves modularity, i.e., the DEM improves upstream evidence quality, and E1–E4 isolate how prompt values, user-level priors, and the size of the label set influence zero-shot inference downstream.

3. RESULTS

This section describes the results obtained for all the experiments E1-E4, with respect to the type of the answer, the model family, as well as using the metrics overall accuracy, precision, recall and m-F1 score. Moreover, for cross-dataset context, we also report the Geolife baseline from our prior as in [2]. Overall accuracy 0.665 on a balanced 4-class set (N=4,280).

Results for experiments E1–E3 in terms of answer choice, precision, and recall are shown in Table 1, and those for experiment E4 in Table 2. Precision and recall for class for the E4 are depicted in Table 3.

Setting (4-class)	Answer	Precision	Recall	m-F1
E1 (DeepSeek-32B)	step	0.559	0.559	0.554
E1 (DeepSeek-32B)	quick	0.458	0.458	0.396
E2 (DeepSeek-32B)	step	0.394	0.394	0.406
E3 (DeepSeek-32B)	quick	0.540	0.540	0.477
E3 (DeepSeek-32B)	step	0.470	0.470	0.421
Geolife (DeepSeek-32B)	quick	0.665	0.665	—

Table 1. Overall 4-class performance across prompts and answering styles. m-F1 denotes macro-F1. Geolife macro-F1 was not reported in our prior study.

Per-class highlights (recall). E1 DeepSeek—step: walk 0.724, bike 0.543, bus 0.324, car 0.647 (balanced). E1 DeepSeek—quick: walk 0.422, bike 0.000, bus 0.552, car 0.860. E2 DeepSeek—step: walk 0.309, bike 0.365, bus 0.555, car 0.348. E3 DeepSeek—quick: walk 0.760, bike 0.020, bus 0.640, car 0.740. E3 DeepSeek—step: walk 0.860, bike 0.680, bus 0.100, car 0.240. E4 (10-mode) — quick: bike 0.660, bus 0.160, ebike 0.160, escooter 0.460, moto 0.720, subway 0.000, train 0.300, tram 0.000, walk 0.860, car 0.430. E4 (10-mode) — step: bike 0.522, bus 0.087, ebike 0.438, escooter 0.356, moto 0.000, subway 0.000, train 0.118, tram 0.000, walk 0.936, car 0.528.

Regarding the 10-class extension (E4): precision is about 0.37 for both styles, with macro-F1 higher for quick (0.325) than for step-by-step (0.236).

Setting (10-class)	Answer	Overall Acc	Macro-F1
E4 (10-class)	quick	0.369	0.325
E4 (10-class)	step	0.372	0.236

Table 2. Overall 10-class results for E4 (*car, walk, train, subway, tram, bus, bike, ebike, escooter, moto*).

Per-class recalls (E4).

Ans.	bike	bus	car	ebike	escooter	moto	subway	train	tram	walk
quick	0.660	0.160	0.430	0.160	0.460	0.720	0.000	0.300	0.000	0.860
step	0.522	0.087	0.528	0.438	0.356	0.000	0.000	0.118	0.000	0.936

Table 3. Recall by class for E4 (*car, walk, train, subway, tram, bus, bike, ebike, escooter, moto*).

Across E1–E4, step-by-step helps the smaller model, quick suits the larger, soft relational cues outperform fixed thresholds, priors should only break ties, and rail needs network context.

What the experiments show. Across E1–E4, three patterns emerged. First, *answering style depends on model capacity*: smaller models (Deepseek-32B) benefit from step-by-step scaffolding to keep multiple weak cues in order, whereas larger models can compress rich summaries and often do better with a concise quick answer. Second, *prompt design matters*: fixed numeric cutoffs encourage brittle number matching, while soft ranges and relational cues (for example, stop regularity, deceleration–acceleration asymmetry, slope sensitivity) support more robust reasoning. Third, *priors help only at the margin*: user metadata should act as a late, low-weight nudge to resolve near-ties and must never override clear kinematic evidence. The ten-mode extension further revealed GPS-only blind spots for rail, which likely require network-aware context such as corridor alignment and station cadence.

Implications for zero-shot prompting. Answer style should match model capacity: step-by-step tends to suit smaller models, whereas quick is often effective when summaries are rich and the model is strong. Soft, relational cues are generally preferable to hard thresholds, which can be brittle. Priors function best as late tie-breakers after kinematic checks and the bus–car tie-breaker, and they should not overturn a confident walk decision. Elevation features can aid the separation of active modes, though their use is best kept judicious.

4. CONCLUSION

This work proposed an extended LLM approach for transportation mode classification on the NetMob2025 dataset. Although the results are not yet satisfactory, the experiments provided insights into how such models reason and raised new questions about handling active mobility modes without ground-truth data. In future work, we plan to redesign inference as a *multi-step* process: for each analysis facet (e.g., kinematics, stop/periodicity, DEM/grade sensitivity, human priors) the LLM will produce a sub-decision with confidence, and a final aggregator will combine these signals to predict the transport mode. We expect this factorized inference to yield higher accuracy and, crucially, much greater stability and change-resilience—the main shortcomings observed here.

Funding. This work is funded by the French National Research Agency as part of the Mob Sci-Dat Factory project (ANR-23-PEMO-0004) under the France 2030 program.

Acknowledgment. The authors thank the NetMob 2025 organizers for dataset access and support.

Disclosures. The authors declare no conflicts of interest.

Data availability. Data used in this study are part of the NetMob25 Challenge and subject to access conditions.

REFERENCES

1. Y. Zheng, H. Fu, X. Xie, W.-Y. Ma, and Q. Li, "GeoLife GPS trajectory dataset - User Guide," Microsoft Research Asia, Geolife GPS trajectories 1.1 edition, July 2011.
2. A. Badawi, A.-M. Raimond, A. Le Guilcher, and K. Zeitouni, "Harnessing Large Language Models for Predicting Mobility Modes," in 2025 26th IEEE International Conference on Mobile Data Management (MDM), 2025, pp. 222–227.
3. A. Chasse, A. J. Kouam, A. C. Viana, R. Stanica, W. V. Lobato, G. Ramos, G. Deperle, A. Bouroudi, S. Bussod, and F. Molano, "The NetMob25 Dataset: A High-resolution Multi-layered View of Individual Mobility in Greater Paris Region," arXiv preprint arXiv:2506.05903, 2025.
4. A. Haydari, D. Chen, Z. Lai, and C.-N. Chuah, "MobilityGPT: Enhanced Human Mobility Modeling with a GPT model," arXiv preprint arXiv:2402.03264, 2024.
5. T. Liu, M. Li, and Y. Yin, "Can Large Language Models Capture Human Travel Behavior? Evidence and Insights on Mode Choice," Available at SSRN 4937575, 2024.
6. X. Wang, M. Fang, Z. Zeng, and T. Cheng, "Where Would I Go Next? Large Language Models as Human Mobility Predictors," arXiv preprint arXiv:2308.15197, 2023.
7. S. Dabiri and K. Heaslip, "Inferring transportation modes from GPS trajectories using a convolutional neural network," Transportation Research Part C: Emerging Technologies, vol. 86, pp. 360–371, 2018.

Gendered mobility differences in the era of flexible work

MARIANA MACEDO[1, 2] AND ZSÓFIA ZÁDOR[1*]

[1] Northeastern University, London, UK
[2] Khoury College of Computer Sciences, Northeastern University, Boston, United States
[*] zsofia.zador@nulondon.ac.uk

1. INTRODUCTION

The adoption of remote work varies by income, gender and occupation. For instance, women show more interest in seeking out remote positions compared to men [1]. The general increase in the acceptance of remote working by employers alters urban mobility patterns and is expected to affect where people choose to reside [2], potentially revitalizing local businesses in residential areas [3]. There are insights into which occupations can plausibly transition to remote work [4], yet there is still limited empirical evidence on whether gender differences in preference for remote work manifest in practice. Women, especially mothers, may benefit from flexibility but risk greater economic disparity if excluded from these opportunities [5]. This disparity raises important questions about how to support vulnerable groups, who are unfortunately less likely to benefit from urban accessibility.

This study examines gender differences in the spatial landscape of urban mobility between remote workers and office-based employees. We estimate remote working days from the Netmob25 dataset [6] by counting the number of days each individual takes a work-related trip. We group workers into three categories; Office (more than three days in the office), Hybrid (three days in the office) and Remote worker (maximum of two days in the office). Our findings suggest that occupational characteristics, rather than gender or household composition, are the primary determinants of remote work adoption, although gender differences persist in mobility patterns.

2. MOBILITY DIVERSITY

We employ the mobility diversity indicator [7] measuring the Shannon entropy of the spatial coverage of the mobility of different groups (e.g. gender, socioeconomic status). Here, a set of travels made by group X can be defined by work type and gender for purpose d, the mobility diversity H_d^X is calculated as:

$$H_d^X = -\frac{1}{\log_2 N_Z} \sum_{i=1}^{N_Z} p_d^X(i) \log_2 p_d^X(i), \quad \textbf{(1)}$$

where $p_d^X(i) = \frac{N_d^X(i)}{N_d^X}$ represents the probability that group X visits zone i for purpose d. $N_d^X(i)$ denotes travels by group X to zone i, and N_d^X is the total travels by group X. The metric ranges from 0 (concentration in one area) to 1 (uniform spatial distribution).

We classify travellers by work type (office, hybrid, or remote) and gender, and analyse their mobility diversity across different trip purposes. To account for variations in sample size, we use bootstrapping: we repeatedly sample 80% of trips 100 times to estimate the distribution of H. Additional tests varying the sample proportion from 10% to 100% show that the results converge before reaching 60%, confirming that using 80% of trips provides robust and reliable estimates.

3. RESULTS AND DISCUSSION

We first assess whether gender and family circumstances shape the adoption (likelihood) of remote work arrangements. Despite documented female preference for remote work opportunities, our analysis of the Île-de-France sample using the Netmob25 dataset reveals that gender, the presence of children or dependents do not significantly influence remote work rates. To formally assess these relationships, we estimate an OLS regression of the form:

$$Y_i = \beta_0 + \beta_1 \,\text{HAS_CHILD}_i + \beta_2 \,\text{HAS_DEPENDENTS}_i + \beta_3 \,\text{SEX}_i + \sum_k \gamma_k \,\text{TYPE_HOUSE}_{ik} + \sum_m \delta_m \,\text{WORK_CAT}_{im} + \varepsilon_i, \quad \textbf{(2)}$$

where Y_i denotes the number of remote working days. The results confirm that occupation type emerges as the only consistent predictor: employees ($\beta = 0.99$, $p < 0.001$) and intermediate professions ($\beta = 0.89$, $p < 0.001$) work remotely more often than tradesmen. In contrast, having children, dependents, or identifying as a woman did not significantly predict remote work adoption. This suggests that while preferences are gendered, they do not translate into differences in remote work adoption, which appears to be shaped more by occupational characteristics. This may reflect a lag between preferences and labour market change, or due to biases in the sample.

Based on this, we analyse the spatial patterns of residential living and the mobility behaviours of different types of workers (office, hybrid, and remote) in Île-de-France, using the density of home locations and the indicator of mobility diversity (H) [7]. Fig. 1A illustrates residential segregation across work types; Office workers reside closer to central areas but widely spread on the urban fabric (blue zones), while remote workers cluster in peripheral and polycentric regions (orange zones). This is consistent with recent literature highlighting that remote work enables relocation away from city centres [2, 8]. This higher concentration of where remote workers live is supported by also calculating the diversity of residential locations [7] and in line with a more concentrated mobility in overall, to work, and not to work (Fig. 1C). Thus, remote workers show greater spatial concentration, whereas office workers are the ones that most evenly distributed their travels across the metropolitan area.

This unequal spatial distribution suggests that different worker types likely display distinct mobility patterns, including differences by gender. In Fig. 1B, we see that mobility diversity (H) [7] is lower for remote workers compared to all workers overall. Additionally, the gender differences in mobility diversity that are evident among office and hybrid workers decrease among remote workers. This indicates that when mobility is no longer driven by commuting, traditional gender-based differences in mobility tend to diminish.

Finally, we consider the purpose of the trip and observe a trend of equalisation between all categories of trips. Mobility diversity differences shrink between men and women considering overall mobility, work related and non-work related trips. There is a consistent gender gap between office workers and hybrid workers, but not among remote workers. This indicates that flexible work arrangements may reduce gendered mobility

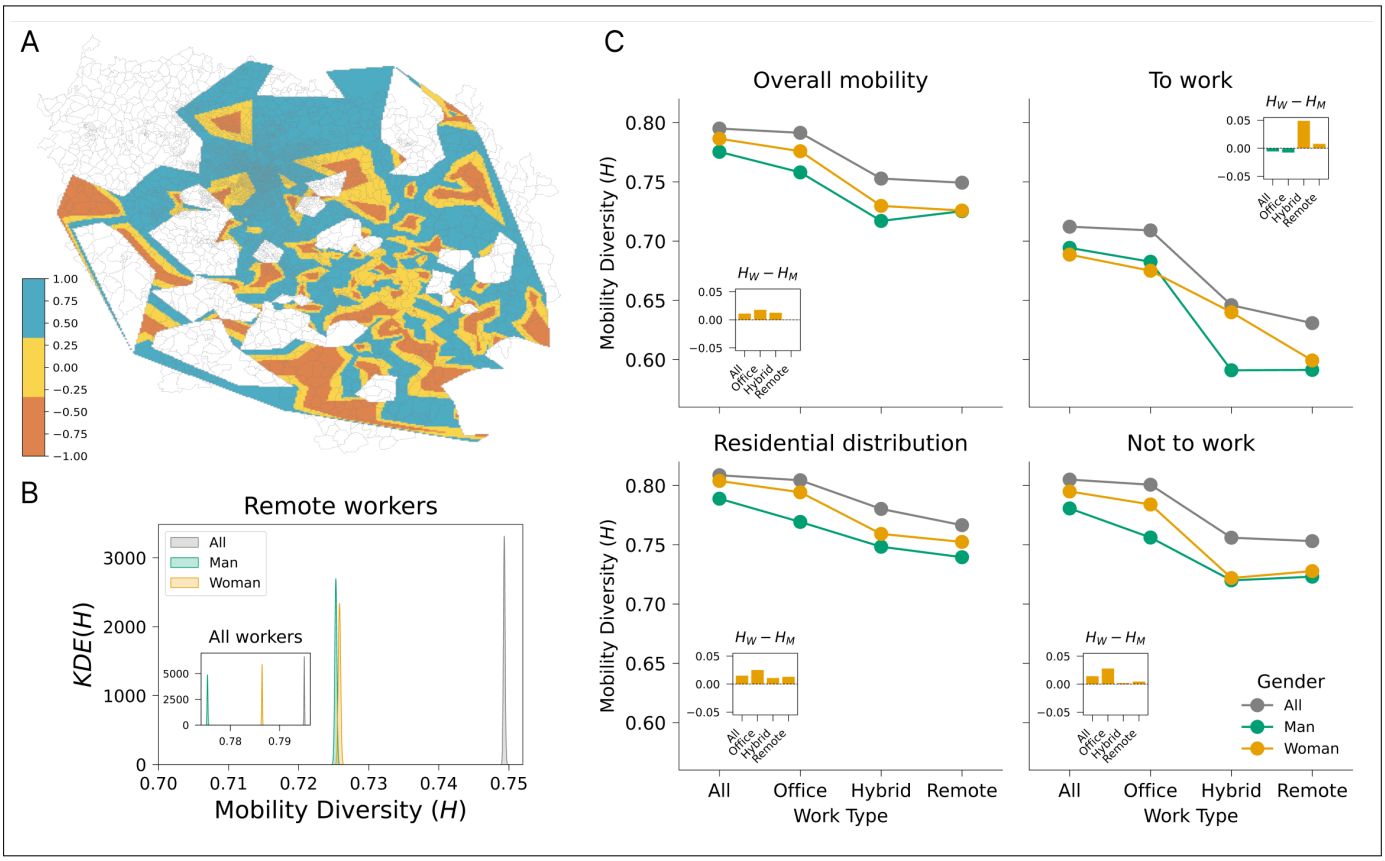

Fig. 1. A) Density map of home locations in Île-de-France. Colours denote whether, for a given zone, the majority of residents are office workers (positive values, blue), remote workers (negative values, orange), or whether proportions are similar (yellow). B) Kernel Density Estimation plots of mobility diversity (H) for all and remote workers, showing travel by all individuals (grey), men (green), and women (yellow). While men and women typically cover Greater Paris differently, this disparity disappears among remote workers. C) Average mobility diversity by work type and gender across different trip purposes. High H values above 0.70 indicate that workers are relatively evenly distributed across Île-de-France, with notable differences between worker types and gender. Work trips exhibit the lowest H, highlighting that work destinations are more spatially concentrated than non-work trips, though they are not clustered in a single main centre.

constraints. We identify one exception to the pattern, among hybrid workers, women exhibit the lowest mobility diversity on non-work trips, pointing to possible persistent constraints on their spatial behaviour even under more flexible arrangements.

Taken together, our work provides new empirical evidence showing that, despite documented female preferences for remote work, adoption rates might be primarily shaped by occupation type rather than gender or family circumstances. By leveraging large-scale mobility data from Île-de-France, we highlight how different work arrangements influence residential distribution and mobility diversity, with remote workers clustering in peripheral areas and exhibiting more spatially concentrated movements. In addition, we find that remote work is associated with a reduction in gendered mobility differences, particularly when mobility is no longer driven by commuting. These findings underscore the potential of flexible work arrangements to reduce gendered mobility constraints and inform urban and workplace policies aimed at promoting more equitable and inclusive cities. Finally, further analysis of individual mobility trajectories may reveal additional, subtler gendered patterns shaped by work type. We also acknowledge that our results are limited by the nature and composition of the available data.

REFERENCES

1. "LinkedIn Data: Women, Gen Z More Likely to Apply to Remote Jobs," .
2. A. Ramani, J. Alcedo, and N. Bloom, Proc. National Acad. Sci. **121**, e2408930121 (2024).
3. T. Yabe, B. G. B. Bueno, X. Dong, *et al.*, Nat. Commun. **14** (2023). Publisher: Springer Science and Business Media LLC.
4. J. I. Dingel and B. Neiman, J. Public Econ. **189**, 104235 (2020).
5. J. Tan, J. Lim-Soh, and P. L. Tan, Appl. Res. Qual. Life **19**, 2595 (2024). Publisher: Springer Science and Business Media LLC.
6. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The NetMob25 Dataset: A High-resolution Multi-layered View of Individual Mobility in Greater Paris Region," (2025). ArXiv:2506.05903 [cs].
7. M. Macedo, L. Lotero, A. Cardillo, *et al.*, PLOS ONE **17**, e0260874 (2022). Publisher: Public Library of Science (PLoS).
8. M. Akan, J. M. Barrero, N. Bloom, *et al.*, "The New Geography of Labor Markets," Tech. Rep. w33582, National Bureau of Economic Research, Cambridge, MA (2025).

Generation of Realistic Geolocated Agendas in a Digital Twin using GPS-based Mobility Dataset

MAXIME COLOMB[1,2,*]

[1] ASCII, Inria, CMAP, CNRS, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France.
[2] LASTIG, Univ Gustave Eiffel, ENSG IGN, F-94160 Saint-Mande, France
[*] maxime.colomb@inria.fr

1. INTRODUCTION

Digital twins of cities consist in very detailed representations of chosen intricated subsystems composing urban environment. The construction of such models depends on a deep understanding of each subsystem mechanisms. It also rely on the availability, quality, consistency and richness of diverse datasets. Realistic, representative, geo-localized and historicized sets of data are then required to ensure the quality and the precision of digital twins. Transportation is a key sub-system as it moves inhabitants within places and allows to simulate multiple externalities (transportation saturation, air pollution, etc.) However, this very dynamic sub-system has an expansive modelling cost in term of computation and size of input datasets. Human mobility modelling includes knowing the reason and the destinations of individual's trips. Many applications which need mobility simulation would only require to model those two components, for example analyses of the quality of access to various services, places saturation, epidemiology, and so on.

When possible, limiting human transportation models to geolocated agendas would save modelling computing time, data availability and reduce number of hypotheses and parameters. Focusing on improving this subsystem should increase digital twins accuracy and reflect on simulations which are applied within. Very detailed knowledge is required about why and where individuals are travelling, even more when their personal characteristics are taken into consideration. Individual socio-demographic classes would help build representative categorisations of individuals in order to extrapolate behaviours to similar individuals. Different contextual elements, when applied to single individuals (telework days vs on-site day of work) or to the whole population (holiday, strike, pandemic-related restrictions…), could help understand how individuals modify their schedules. Digital twins of cities could then adapt dynamically to specific conditions.

Multiple challenges shall be addressed. Datasets, however large they are, should be sufficiently representative to allow diverse extrapolation of behaviours to different classes of population. Such a representativity must apply to population socio-demographic characteristics, location and goal of trips. Probabilistic methods shall be used to transform such representative data into an exhaustive model. Rare events which are not captured by datasets shall also be modellised. Another challenge is to integrate such modelled information into a well described and feature-rich digital twin. Designed methods must be generic and optimized enough to apply to multiple contexts and on very large urban models.

This work aims to simulate this minimal aspect of mobility: generating individual agendas and locate each steps. Such information could then be injected in a digital twin of community and individuals in order to model how places are occupied and how individuals meet. A new model is designed and presented in this poster to perform such tasks in an efficient way. It uses the well detailed dataset provided by the *NetMob* 2025 Data Challenge [1]. Probabilistic tools are generated to provide a framework of agenda generation. Simulation of geo-localised agendas can then be applied to every individuals of the digital twin.

2. METHOD

Few distinct steps are needed to achieve such a goal. A probabilistic mobility framework shall first be generated (section A and B). The simulation of individual agendas can then be generated (section C). Adaptation of NetMob Data Challenge and results about the occupation are presented in a last sections (3 and 4)

A. Markov chain parametrization

We designed a Markov Chain, represented in Figure 1, which models the main types of activities available for any individuals. Probabilities to switch from an activity to another are set using calibration data. This chain is calibrated respectfully to a determined number of socio-demographic categories, for each time lapse of the simulation and for different types of days. A probabilistic list of activities can then be generated for every kind of individual. As specified in [1],

> Researchers should always consider (…) whether to disaggregate by day type or geographic unit.

A statistical verification must ensure that each discrete class contains enough values to be sufficiently representative, otherwise categories would be merged with neighbourhood categories.

As Markov Chains are let unconstrained, individuals can switch from an activity to any other one at every time of the simulation. Imposing events (lunch break, get children from school, etc.) could eventually be added, but the law of large numbers should reproduce such event at a satisfying level.

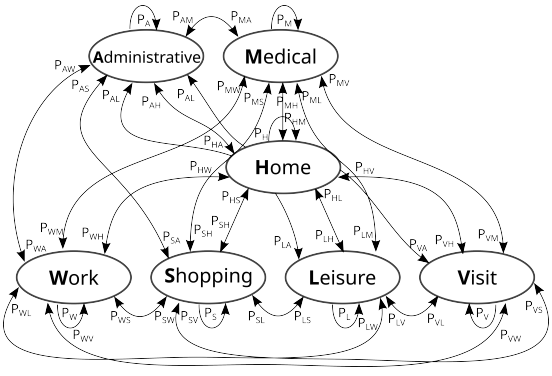

Fig. 1. Markov chain for activity selection

B. Micro-zoning with mobile data

The presence of different types of individuals at specific times and places of the city should be modelled using time-dependant origin and destination traces. A spatial aggregation is performed in order to determine a destination area rather than a precise place. Origin-Destination matrix are calculated, based upon mobility data. For each origin aggregate, the probability of going to any other aggregate is calculated based on all the weighted trips happening at that time and from that place. Such matrix are disaggregated to hours of the day and could also be disaggregated to few population classes if data are representative enough. Aggregates can be merged to meet statistical representativeness requirements. An example of OD Probability matrix can be found in Figure 2.

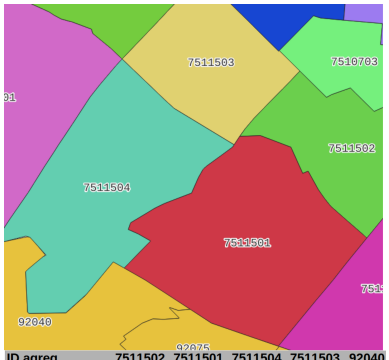

ID agreg	7511502	7511501	7511504	7511503	92040
7511502	0,341	0,091	0,024	0,041	0,045
7511501	0,102	0,368	0,046	0,032	0,014
7511504	0,022	0,061	0,353	0,039	0,009
7511503	0,032	0,038	0,037	0,407	0,143
92040	0,013	0,016	0,048	0,012	0,555

Fig. 2. Probability of going to different neighbourhood of the 15th arrondissement of Paris

C. Agenda generation

An agenda can be generated for every individual of the simulation by using the appropriate Markov Chain parameters. Location of some of the activities are linked to each individual by the digital twin (home and work location), some others are fluctuant and need to be spatialized. An IRIS is chosen using the OD probability matrix and within it, a random place with a type corresponding to the activity is attributed. The respect of the digital twin characteristics (type of activities open to the public, opening hours, level of frequentation, etc.) should reproduce places occupation in a realistic way.

3. INTEGRATION OF NETMOB DATA CHALLENGE

Data gathered for the NetMob Data Challenge [1] were the only mobility data used in this work. In order to improve speed of all requests on data that would happen latter on, the NetMob datasets have been loaded on a PostGis database. This model have initially been designed with survey for Markov chain parametrization and mobile data for OD matrix calibration. The latter dataset provides a high number of different traces, and the size of statistical aggregates can be small. Nevertheless, the fact that the same dataset calibrate every probabilistic model brings a unique consistency to the results.

4. RESULTS

Figure 3 shows the frequentation of a driving school in the 11th arrondissement of Paris. Frequentations of every activity can be found here: https://ici.saclay.inria.fr/simu/netmobresult.php. This digital twin, described in ([2] - submitted), cover the entire Île-de-France and includes 2 070 444 activities, 473 348 being open to the public. For now, frequentation for each arrondissement is presented separately.

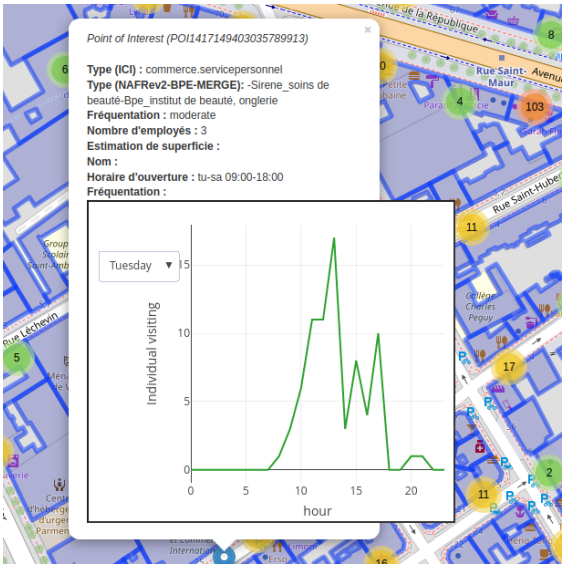

Fig. 3. Example of visits in an activity open to the public in an Île-de-France digital twin

5. CONCLUSION AND PERSPECTIVE

This model is actively used by the ICI project [4], which couples a digital twin with a contagion model. Sensitivity analysis should qualify uncertainty of this agenda generator and how its variability is affecting ICI epidemiological simulations. The model of the effect of particular events (strike, telework) on mobility behaviour necessitate the use of complementary datasets. For example, [5] could be used to model the influence of COVID-19 lockdowns on individual mobility.

6. BACK MATTER

Funding. This work will benefits from the Mob Sci-Dat Factory project [6].

Acknowledgment. This work have been developed as a part of the ICI project, involving the ASCII team of INRIA, LASTIG of IGN, CMAP of École Polytechnique, CRESS of INSERM.

Disclosures. The author declares no conflicts of interest.

REFERENCES

1. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, arXiv (2025).
2. M. Colomb and J. Perret, IJGIS (2025).
3. C. Lecomte, G. Le Roux, H. Commenges, and J. Vallée, "Le Mobiliscope. Un outil de géovisualisation de la composition sociale 24h/24 des territoires," in *Spatial Analysis and GEOmatics 2017*, (INSA de rouen, Rouen, France, 2017).
4. M. Colomb, Q. Cormier, J. Garnier, *et al.*, "ICI Project: Individual-based epidemic propagation simulator based on a Digital Twin," in *Health GIS: Spatial Thinking in Applied research 2024*, (O'Brien Institute for Public Health (OIPH) and University of Calgari, Calgari, Canada, 2024).
5. H. C. Fantecele, S. Rabenjamina, A. Carneiro Viana, *et al.*, "SafeCityMap (1st phase) – COVID INRIA mission: Investigating population mobility habits in metropolitan zones and the lockdown impact using mobile phone data," Research report, Inria (2021).
6. "Projet Ciblé Mob Sci-Dat Factory du PEPR MOBIDEC," (2025). [Online; accessed 15. Jul. 2025].

Divided by design: how transport and amenities shape urban mobility and socioeconomic inequalities

PIERO BIRELLO[1], MARTINA CONTISCIANI[1], AND TIMUR NAUSHIRVANOV[1,*]

[1] Department of Network and Data Science, Central European University, Vienna, Austria
[*] naushirvanov_timur@phd.ceu.edu

1. INTRODUCTION

Spatial mobility patterns and transport infrastructure play a key role in shaping the socioeconomic and demographic structure of urban populations [1, 2]. These dynamics are further influenced by the inhomogeneous spatial distribution of amenities [3, 4]. In this study, using data from the NetMob25 Data Challenge for the Île-de-France region (Fig. 1A), we investigate the interplay between socioeconomic factors, demographic profiles, the spatial distribution of points-of-interest (POIs) and mobility behaviours. Our goal is to unveil potential inequalities in movement across different population groups and urban areas. To examine these inequalities, we analyse mobility patterns with respect to individuals' places of residence at the level of INSEE communes, and relate these movements to both the socioeconomic status (SES) and the spatial distribution of POIs. Together, these dimensions offer a more comprehensive characterisation of each commune, linking population characteristics with the functional structure of the built environment.

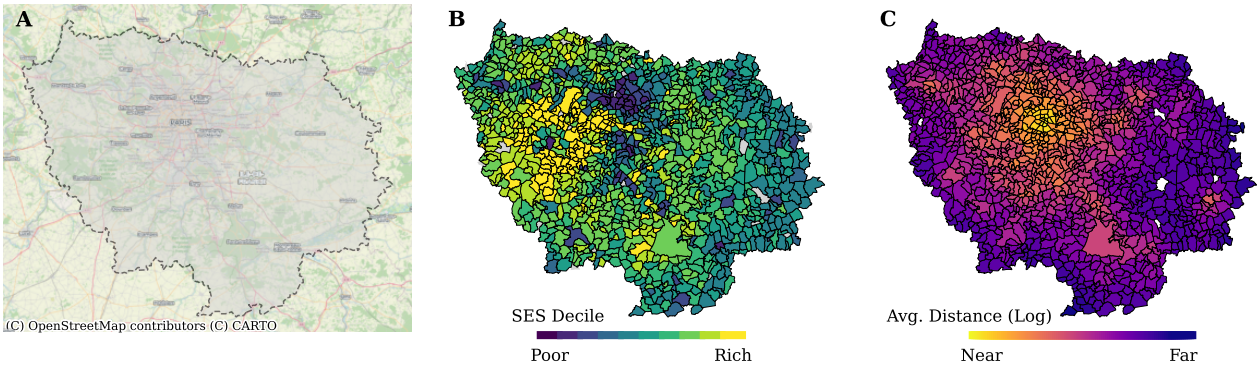

Fig. 1. Spatial context of socioeconomic conditions and accessibility in Île-de-France. (**A**) Administrative boundaries of the Île-de-France region. (**B**) Socioeconomic status (SES) deciles at the INSEE code level. (**C**) Average distance to POIs per INSEE unit.

2. METHODS

We estimate the SES of each commune using data on median income and population size. By computing population-weighted quantiles of median income, we ensure that each SES group—from lowest to highest—represents an approximately equal share of the regional population, rather than an equal number of communes. This approach yields a classification that divides the population into ten SES deciles (Fig. 1B). POIs are sourced from OpenStreetMap and classified into five functional categories: Purchase, Leisure, Education, Health, and Work or Business. To quantify POIs accessibility, we calculate the average Euclidean distance from potential residential locations—uniformly sampled over buildable land—to the nearest POI in each category. For each INSEE, we first compute category-specific average distances and then take the mean across categories to capture overall accessibility to commonly used POIs. The resulting values are shown in Fig. 1C. The Pearson correlation between SES and average POIs distance is $r = -0.10$, indicating a weak negative association. However, the relationship is non-linear, with both the poorest and richest areas exhibiting shorter average distances to POIs. This further highlights the importance of considering both SES and POIs accessibility, as they offer related yet complementary perspectives in characterising each commune.

To characterise the mobility patterns of different socioeconomic groups, we employ two complementary metrics: the assortativity coefficient and the socioeconomic preferential mobility index. The assortativity coefficient ρ is defined as the Pearson correlation between the income deciles of origin and destination communes, weighted by their mixing probability [5]. It ranges from -1 (trips occur exclusively between dissimilar socioeconomic areas) to 1 (trips are entirely within the same socioeconomic group), thus capturing the overall tendency of different SES groups to interact through mobility. In contrast, the index R quantifies the directionality of this mixing by measuring the probability of individuals to move towards higher or lower SES classes [6]. A value of $R = 1$ indicates that all trips are directed towards richer SES areas, while $R = -1$ reflects exclusive movement towards poorer communes. A value of $R = 0$ implies a balanced flow between higher and lower SES destinations. We use non-parametric bootstrapping for uncertainty quantification and permutation tests for statistical significance. Finally, we use Generalised Linear Models (GLM) to estimate the marginal effects of POI proximity on mobility outcomes, controlling for sociodemographic confounders (educational attainment, professional category, income, household size, dwelling type, and driving licence possession). Trip duration is modelled using a Gamma distribution with a log-link function, and the number of weekly trips is modelled using a Negative Binomial distribution with an identity link. The independent variable, average distance to POIs (x), is log-transformed. For the duration model, this specification yields an interpretation of the coefficient β as

an elasticity, representing the multiplicative percentage change in duration per percentage change in distance. For the frequency model, β represents the additive change in the number of weekly trips per unit change in $\log(x)$.

3. RESULTS

We compute the defined mobility metrics across various population characteristics and observe statistically significant patterns—relative to a null model—for variables such as age, profession, and Navigo subscription type. Notably, individuals aged 65 and over exhibit the highest assortativity coefficient (ρ), indicating a strong tendency to travel within their own socioeconomic group. In contrast, individuals aged 45–54 show both low assortativity and a positive R value, suggesting more frequent travel across SES boundaries, particularly toward higher-income areas.

When analysing trip-level annotations, we identify trip duration, trip distance, and transport mode as key variables influencing socioeconomic mixing. In general, assortativity is positive, indicating a preference for travel between communes of similar SES (Fig. 2A). Considering transport mode specifically, walking is associated with the highest assortativity—even for trips outside the residential area—followed by bus and tram trips, bike and scooter, private car, train, and subway. Interestingly, travelling by a private car has a negative preferential mobility index, implying that such travel also occurs towards areas with lower socioeconomic status. After controlling for trip duration and trip distance, some differences across modes persist, with active modes and bus/tram trips disproportionately confined within homogeneous areas, reinforcing within-area mobility, and private car, subway, and train facilitating broader cross-SES mixing. A key structural factor underlying these dynamics is the spatial distribution of wealth, which is strongly clustered. Spatial autocorrelation analysis confirms that communes with similar SES levels tend to be geographically close rather than randomly distributed (Moran's $I = 0.55$; Geary's $C = 0.45$; both $p = 0.001$). In particular, high- and low-SES areas are more spatially concentrated, while mid-SES communes are more diffuse across the region. While this spatial structuring helps explain some of the observed assortativity patterns, our analysis also suggests that structural inequalities may be influenced by how different transport networks are designed and accessed.

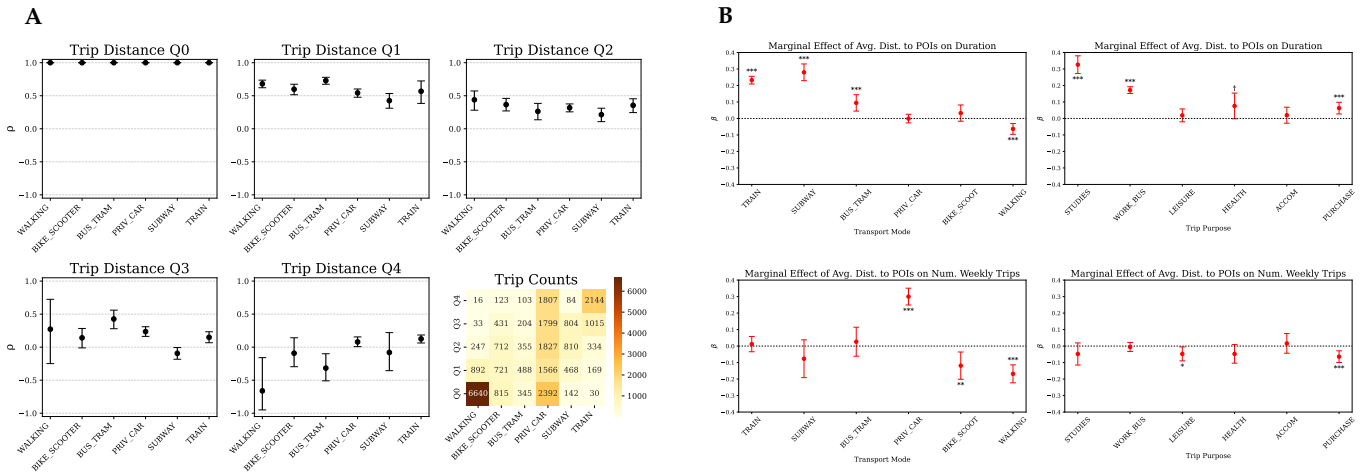

Fig. 2. Socioeconomic assortativity and marginal effects of POI proximity. A: Socioeconomic assortativity index ρ by main mode of transportation and trip distance quantile (Q1–Q4). Q0 corresponds to trips within one INSEE area, where no mixing occurs. Error bars represent 95% confidence intervals. The heatmap shows external trip counts by mode and distance quantile. **B:** Estimated marginal effects of the average distance to amenities on trip duration (top panels) and on the number of weekly trips per individual (bottom panels), across different transport modes (left panels) and trip purposes (right panels). Error bars show 95% confidence intervals, while stars indicate statistical significance (*** for $p < 0.001$, ** for $p < 0.01$, * for $p < 0.05$, † for $p < 0.1$).

Our findings suggest that mobility behaviour differs between central and peripheral areas of Île-de-France. We examine the potential role of Points of Interest (POIs) as a key determinant, and investigate in detail how travel behaviour is shaped by the local supply of POIs in an individual's area of residence, focusing on trips' duration and frequency. We find that greater distance to POIs increases the duration of trips by train, subway, and bus or tram, as expected (Fig. 2B). However, trip duration by car and by bike or scooter is unaffected. Interestingly, residents of POI-dense areas tend to walk longer. Analysing trip purposes provides further nuance: commuting time for study and work increases with distance to POIs, whereas activities such as leisure or accompanying children are unaffected. The duration of purchase trips is only marginally influenced, possibly because individuals visit multiple shops at once in POI-dense areas. Turning to trip frequency, distance from POIs pushes individuals towards car use, while proximity favours trips by bike, scooter, and walking. Furthermore, purchase trips are more frequent in POI-dense areas—note that a proxy for income was included among the control variables. Finally, we find evidence of a modest but statistically significant increase in the frequency of leisure trips in POI-dense areas ($p = 0.03$).

REFERENCES

1. N. Iyer, R. Menezes, and H. Barbosa, Environ. Plan. B: Urban Anal. City Sci. **51**, 1496 (2024).
2. H. Nilforoshan, W. Looi, E. Pierson, *et al.*, Nature **624**, 586 (2023).
3. C. Gaigné, H. R. Koster, F. Moizeau, and J.-F. Thisse, J. Urban Econ. **128**, 103394 (2022).
4. E. Graells-Garrido, F. Serra-Burriel, F. Rowe, *et al.*, PloS one **16**, e0250080 (2021).
5. E. Bokányi, S. Juhász, M. Karsai, and B. Lengyel, Sci. reports **11**, 20829 (2021).
6. M. Duran-Sala, A. K. Balachandran, M. Morandini, *et al.*, arXiv preprint arXiv:2407.01799 (2024).

Understanding Urban Trips: Mode and Purpose Prediction Using GPS Traces and Socio-Demographic Features

ASLIGÜL AKSAN ERAKIN[1, 2], İREM BETÜL KOÇAK[2, 3, *], AND OĞUZ YÜCEL[1, 4]

[1] Management Engineering Department, Istanbul Technical University, Reşitpaşa, 34467 Sarıyer/İstanbul
[2] Industrial Engineering Department, Bogazici University, Bebek, 34342 Beşiktaş/İstanbul
[3] Value Analytics Labs, 100 Cambridge Street, Suite 1400 Boston, MA, 02114, USA
[4] Management Department, Bogazici University, Bebek, 34342 Beşiktaş/İstanbul
[*] irem.kocak@std.bogazici.edu.tr

1. INTRODUCTION

Understanding why and how individuals travel is central to urban mobility research and policy-making. Accurate inference of trip purposes and transport modes enables analysis of commuting patterns, accessibility gaps, and the effects of sustainable transport policies. While major technology companies already provide such predictions within their platforms, their methods and data remain proprietary and inaccessible.

In this work, we address this gap by leveraging the NetMob25 Paris dataset, which links GPS trajectories, trip records, and socio-demographic information. Our goal is to benchmark transparent and reproducible models for predicting both trip purpose and travel mode. By combining trajectory-derived features with temporal, spatial, and demographic attributes, we aim to assess the feasibility and limitations of purpose and mode prediction on large-scale survey-weighted mobility data.

This approach contributes to methodological development and has direct policy relevance. Predictive models can help evaluate modal shifts, equity impacts, and accessibility improvements, offering decision-makers evidence-based tools when high-quality commercial data are unavailable.

2. RELATED WORKS

Transport mode detection has been widely studied using GPS-derived features such as speed, acceleration, and stop duration, with supervised learning approaches showing strong results [1, 2]. Deep learning methods have also been applied to sequential trajectory data, capturing richer temporal patterns [3].

Trip purpose inference has received less attention, traditionally relying on surveys and smart card data to characterize commuting and activity patterns [4]. Recent studies emphasize the role of socio-demographics and social context in shaping mobility choices [5, 6]. However, predicting minority purposes such as `HEALTH` or `BUSINESS` remains a challenge due to class imbalance.

Large-scale mobility challenges such as D4D [7] and D4R [8] demonstrated the potential of digital traces for mobility research. More recently, the NetMob conference series has provided rich datasets, with NetMob23 focusing on mobile service usage [9] and NetMob25 releasing GPS trajectories with socio-demographic attributes [10]. This dataset enables joint study of mode and purpose prediction with transparent, reproducible methods that bridge the gap between academic research and proprietary industry solutions.

3. DATASET AND FEATURE ENGINEERING

We use the NetMob25 dataset, which provides GPS trajectories for individuals in the Paris region, linked to validated trip records and socio-demographic attributes. Three components are available: (i) trip-level information such as origin, destination, time, duration, and annotated purpose; (ii) individual-level data including age, education, and profession; and (iii) per-user GPS traces enabling the extraction of trajectory features.

For **purpose prediction**, the target is the destination purpose (`Purpose_D`). Features include temporal indicators (hour, time of day, day type), spatial context (zones and areas), trip attributes (duration, previous purpose, grouped main mode), and individual socio-demographics. An exploratory analysis shows that `RETURN_HOME` and `WORK` dominate the distribution, while categories such as `HEALTH` and `BUSINESS` are relatively rare (Fig. 1). This imbalance motivates the use of weighted metrics and careful evaluation.

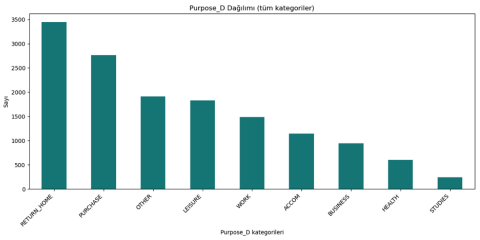

Fig. 1. Distribution of destination purposes (`Purpose_D`) in the dataset.

For **mode prediction**, the target is the grouped main transport mode. We derive trip-level features from GPS logs, including duration, mean and variance of speed, stop ratio, number of GPS points, and GPS quality flags. These are complemented with demographics and subscription indicators.

To account for survey design, sample weights are constructed as the product of trip-level (`Weight Day`) and individual-level (`WEIGHT INDIV`) coefficients, and are applied during training and evaluation. This ensures representativity and reduces bias toward overrepresented users or trips.

4. METHODOLOGY

We frame both tasks as multiclass classification problems. For **purpose prediction**, the label is the destination purpose (`Purpose_D`); for **mode prediction**, the label is the grouped main mode. To avoid information leakage, we use a leave-users-out split, holding out 20% of individuals for testing.

For purpose prediction, we benchmark:

- **Logistic Regression**: interpretable baseline with one-hot features.
- **XGBoost**: strong non-linear learner, robust to mixed data types.
- **Markov Chain**: sequence-aware baseline using first-order transitions of trip purposes.
- **Hybrid (XGB × Markov)**: combines predictive power of XGBoost with sequential context from Markov chains.

For mode prediction, we evaluate:

- **Decision Tree**: interpretable baseline.
- **XGBoost**: handles heterogeneous GPS-derived features and provides feature importance.
- **LSTM**: lightweight sequence model capturing temporal dependencies in tabular features.
- **Markov Chain**: transition baseline exploiting persistence in mode choices.

Evaluation. Performance is assessed using accuracy, weighted F1, and macro F1. Accuracy reflects overall correctness, weighted F1 accounts for class imbalance, and macro F1 emphasizes rare but policy-relevant categories. Survey expansion weights (`Weight Day` × `WEIGHT INDIV`) are applied in both training and evaluation.

5. RESULTS AND DISCUSSION

A. Trip Purpose Prediction

Table 1 reports the performance of different models. XGBoost achieved the best results (Accuracy = 0.60, F1-weighted = 0.56), outperforming logistic regression, Markov chains, and the hybrid. Markov alone performed poorly, while the hybrid did not improve over XGBoost, suggesting that sequential transitions are too noisy or imbalanced to add value. Frequent categories such as `WORK` and `RETURN_HOME` are predicted reliably, while rare purposes remain challenging.

Model	Accuracy	F1-weighted	F1-macro
Logistic Regression	0.57	0.50	0.40
XGBoost	**0.60**	**0.56**	**0.43**
Markov Chain	0.43	0.29	0.11
Hybrid	0.57	0.51	0.37

Table 1. Trip purpose prediction results.

B. Mode Prediction

Table 2 presents baseline performance. XGBoost again achieved the best weighted F1 (0.52), slightly ahead of Markov (0.51). Decision trees and LSTM performed worse, indicating limited generalization.

Model	Accuracy	F1-macro	F1-weighted
Markov Chain	0.52	0.45	0.51
XGBoost	**0.55**	0.31	**0.52**
Decision Tree	0.47	0.25	0.44
LSTM	0.50	0.23	0.45

Table 2. Mode prediction results before SHAP feature selection.

Table 3 shows performance after restricting to top SHAP features. Accuracy and F1 scores dropped slightly, indicating that lower-ranked features still add predictive value. SHAP confirmed that mean and variance of speed, trip duration, stop ratio, and GPS point count are the strongest predictors. This performance drop suggests that features ranked as less important by SHAP may still provide complementary information, and removing them reduces model diversity.

6. CONCLUSION AND FUTURE WORK

We showed that both trip purpose and mode prediction are feasible with the NetMob25 dataset. XGBoost achieved the best performance in both tasks, while Markov chains captured sequential persistence but struggled with rare classes. SHAP analysis confirmed that GPS-derived speed statistics, duration, and stop ratio are key predictors of mode.

Future work will focus on handling class imbalance, enriching spatial context, and applying advanced sequence models (e.g., transformers) to better capture temporal dependencies. These improvements can enhance the policy relevance of purpose and mode inference for sustainable and equitable transport planning.

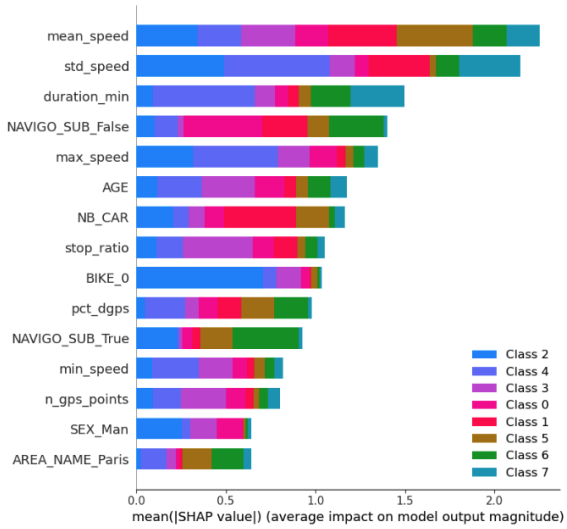

Fig. 2. SHAP summary plot for mode prediction (XGBoost).

Model	Accuracy	F1-macro	F1-weighted
Markov Chain	0.52	0.45	0.51
XGBoost	0.51	0.29	0.48
Decision Tree	0.48	0.26	0.45
LSTM	0.46	0.19	0.40

Table 3. Mode prediction results after SHAP-based feature selection.

REFERENCES

1. Y. Zheng, Q. Li, Y. Chen, *et al.*, "Understanding mobility based on gps data," in *Proceedings of the 10th international conference on Ubiquitous computing*, (2008), pp. 312–321.
2. L. Stenneth, O. Wolfson, P. Yu, and B. Xu, "Transportation mode detection using mobile phones and gis information," in *Proceedings of the 19th ACM SIGSPATIAL international conference on advances in geographic information systems*, (2011), pp. 54–63.
3. Y. Zheng, C. Liu, X. Xie, and X. Ma, IEEE Trans. on Knowl. Data Eng. **32**, 889 (2020).
4. X. Ma, C. Liu, H. Wen, *et al.*, J. Transp. Geogr. **58**, 135 (2017).
5. E. Moro, D. Calacci, X. Dong, and A. Pentland, Nat. Commun. **12**, 4633 (2021).
6. T. Yabe, N. K. Jones, P. S. C. Rao, *et al.*, Nat. Rev. Earth & Environ. **3**, 500 (2022).
7. V. D. Blondel, M. Esch, C. Chan, *et al.*, arXiv preprint arXiv:1210.0137 (2013).
8. A. A. Salah, A. Pentland, B. Lepri, *et al.*, "Data for refugees: the d4r challenge on mobility of syrian refugees in turkey," in *Proceedings of NetMob*, (2018).
9. O. E. Martínez-Durive, S. Mishra, C. Ziemlicki, *et al.*, Sci. Data (2023).
10. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, arXiv preprint arXiv:2506.05903 (2025).

Revisiting Algorithms for Activity Type Estimation from GPS Data in Diverse Built Environments

CHENCHEN SUN[1,*], CARL D. TORBJÖRNSSON[1], AND YUYA SHIBUYA[1]

[1] The University of Tokyo
[*] sun_chenchen@csis.u-tokyo.ac.jp

1. INTRODUCTION

GPS data has become increasingly useful for analyzing mobility-based urban issues due to its high sampling frequency [1–3]. Without costly post hoc travel diary surveys, GPS data often lacks information on individuals' socio-economic attributes, activity purposes, and travel behavior characteristics. Recent research has increasingly integrated complementary spatial datasets, such as land use maps, Points of Interest (POIs), and Areas of Interest (AOIs), to enrich GPS records and identify the purpose of visited locations [2–4]. Despite this, there has been limited systematic validation of activity type estimation algorithms [5–8]. Thus, this study aims to examine such an algorithm's robustness by leveraging the NetMob 2025 dataset[9], which includes both detailed GPS trajectories and corresponding travel survey data. We evaluate commonly used methods for estimating various activity purposes (e.g., work, hospital visits, commercial activities, accompanying others, etc.) and how their robustness varies per different built environment contexts (high-, mid-, and low-density areas) and different socioeconomic groups (diploma, age). The findings contribute to the development of more accurate and generalizable methods for GPS-based activity analysis, enhancing the applicability of GPS data in urban mobility and behavioral research.

2. METHODS

This article uses the Paris Attraction Area (Aire d'attraction de Paris) as a study area which is defined by INSEE to characterize the influence of Paris on the surrounding municipalities. It includes 1,926 municipalities, and 13,024,518 inhabitants in 2017, which accounting for 19.5 percent of the population at the national level.

A. Data

This paper mostly utilized three types of datasets. 1) High-resolution Multi-layered Datasets of Individual Mobility in the Greater Paris Region provided by NetMob2025 [9], which captures detailed daily mobility over a full week for 3,337 volunteers aged 16–80, collected between October 2022 and May 2023. It provides detailed GPS trajectories with a corresponding travel survey and individual socioeconomic attributes. 2) Point of Interest (POI) information from OMF *Places* dataset 2025-06-25 release provides a globally harmonized, structured set of POIs across major urban categories such as retail, education, healthcare, transport, and recreation. 3) Built-up environment information from the Global Human Settlement Layer (GHSL) dataset [10], including 100-meter grids of building height (GHS-BUILT-H) and 1km grids of land type (GHS-SMOD).

B. Comparison dataset building

B.1. Validation trip dataset

The additional trip table in NetMob 2025 data provides detailed start and end times of activities as well as their self-reported purposes of trips. However, this information is only available at broad spatial location levels (INSEE unit codes) and consists of 96,045 valid samples. To evaluate the accuracy of estimated activity types from GPS data, we matched each user's GPS-derived stay data with the trip table based on time, requiring the differences between start or end time be no more than 10 minutes. The GPS-derived stay point data were identified using the Scikit algorithm [11, 12] with 150m and 5 min duration parameters, considering the self-reported trip table's spatial unit level. This process yielded 67,023 matched records, which serve as the ground truth for validation. The sample characteristics before and after matching remain largely consistent, indicating that the procedure is reliable with limited error.

B.2. Activity type estimation algorithm

In this study, we adopt and evaluate one of the most widely used methods for estimating activity types from GPS data, as shown in Figure 1. This method uses POI data and employs common techniques found in GPS research to identify home, work, and non-routine activities.

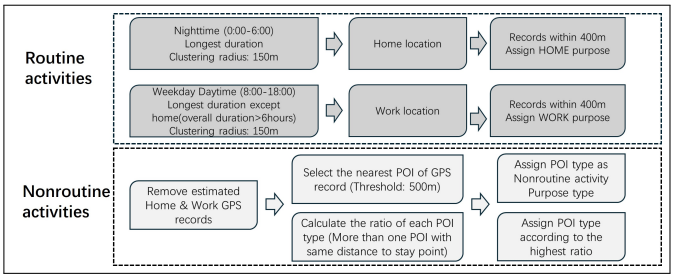

Fig. 1. Estimation rules of routine and nonroutine activities.

3. RESULTS

This study highlights the strengths and limitations of rule-based approaches for inferring activity types from GPS trajectories. Overall, Home (80. 07%)and Work (69. 96%) activities are identified with relatively high accuracy due to their temporal regularity and spatial stability, whereas Non-routine activities suffer from low recognition rates (Leisure:33.69%, Purchase:30.03%, Studies:29.16%), reflecting the limitations of distance-based POI matching. Figure 2 visualizes the actual types of ground-truth activity (validation dataset) corresponding to records that were inferred as Home, and various non-routine activity categories. Activities such as Purchase, Other, and Accompany were shown to be frequently misidentified as Home, while Home, Studies, and Purchase records tend to be incorrectly classified as Work. For non-routine activities, the most underestimated types, Purchase and Other, are frequently reassigned to Business and Leisure.

In addition, the accuracy comparison by demographic groups reveals systematic biases: children and elderly individuals are more accurately identified for home and non-routine activities, but perform poorly in work identification, possibly due to irregular or part-time work patterns (Figure 3). Education and socio-

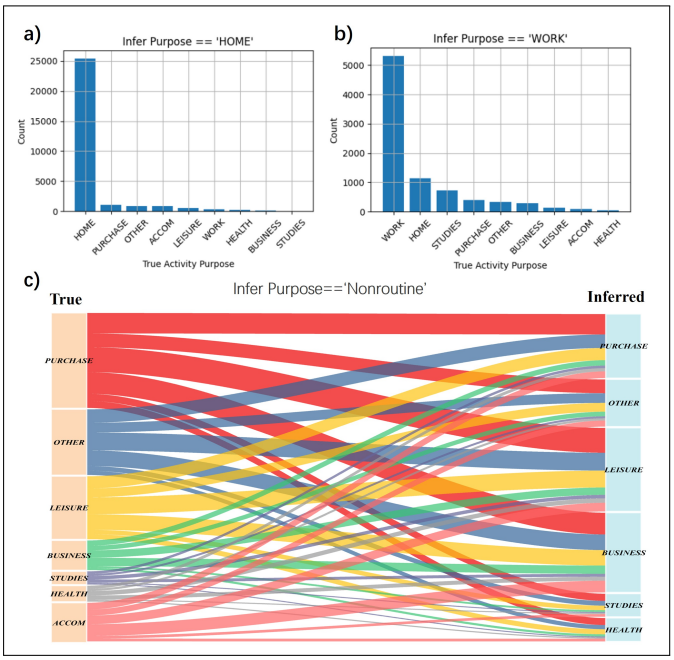

Fig. 2. Corresponding Relationship between Inferred Activity Type and Actual True Activity Type of Home(a), Work(b), and Nonroutine Activities(c).

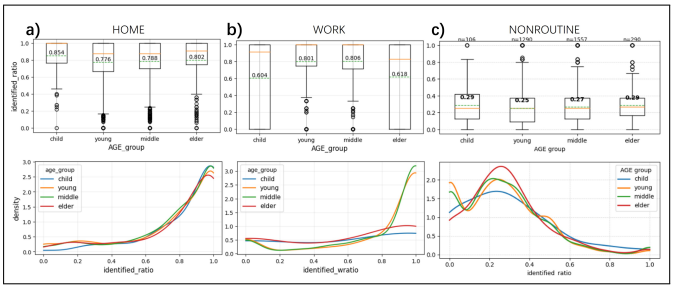

Fig. 3. Boxplot and Density Curves Plot of Identified ratio by Different Age Groups of Home (a), Work (b) and Nonroutine Activities (c).

professional status also shape inference performance: Higher education and middle social strata are associated with more stable work patterns and thus higher recognition accuracy, while disadvantaged groups are more prone to misclassification(Figure 4).

Furthermore, built environment characteristics also influence inference outcomes. Population density and building height both exhibit non-linear effects: Home and Work activities' inference accuracy improves up to medium building density levels but declines at higher building density intensities, whereas Non-routine activities generally perform worse in denser, more vertical environments, though their accuracy shows a modest rebound at the very highest density levels (Figure 5 a and b). Detailed zoning type analysis (Figure 5 c) shows stronger alignment: activities are more accurately inferred when they occur in zones functionally associated with their purpose.

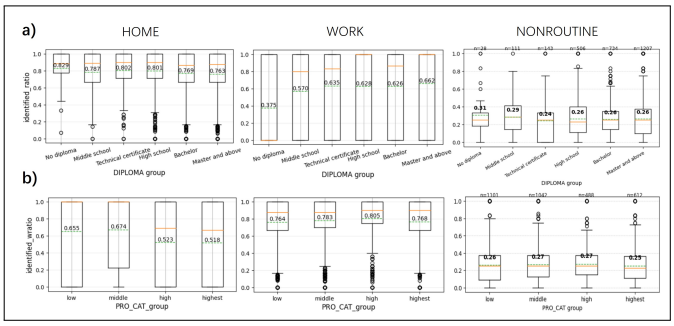

Fig. 4. Boxplot and Density Curves Plot of Identified ratio by Different Diploma(a) and Social Professional Level(b) Groups of Home, Work, and Nonroutine Activities.

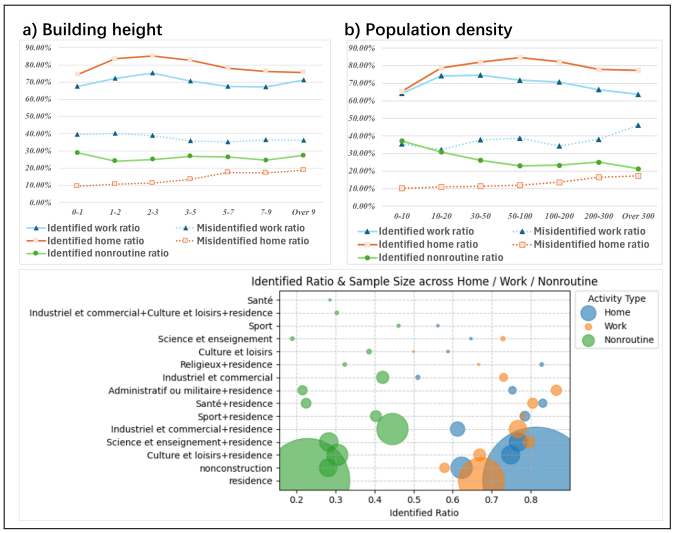

Fig. 5. Inference Accuracy of Home, Work, Nonroutine Activities by Area's Average Building Height (a), Population Density (b), and Zoning Type(c).

REFERENCES

1. S. Lu, Z. Fang, X. Zhang, *et al.*, ISPRS Int. J. Geo-Information **6**, 7 (2017).
2. E. Moro, D. Calacci, X. Dong, and A. Pentland, Nat. Commun. **12**, 4633 (2021).
3. Z. Fan, T. Su, M. Sun, *et al.*, PNAS Nexus **2**, pgad077 (2023).
4. Y. Yang, A. S. Pentland, and E. Moro, EPJ Data Sci. **12**, 15 (2023).
5. J. Wu, C. Jiang, D. Houston, *et al.*, Environ. Health **10**, 101 (2011).
6. B. Furletti, C. Cintia, C. Renso, and L. Spinsanti, "Inferring human activities from gps tracks," in *Proceedings of the 2nd ACM SIGKDD International Workshop on Urban Computing,* (Chicago, IL, USA, 2013), pp. 1–8.
7. T. Feng and H. J. Timmermans, Eur. J. Transp. Infrastructure Res. **15**, 662 (2015).
8. L. Gong, T. Yamamoto, and T. Morikawa, Asian Transp. Stud. **4**, 114 (2016).
9. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," (2025).
10. M. Pesaresi, M. Schiavina, P. Politis, *et al.*, Int. J. Digit. Earth **17**, 2390454 (2024). Open Access.
11. H. Ramaswamy and K. Toyama, "Project lachesis: parsing and modeling location histories," in *International Conference on Geographic Information Science,* (2004), pp. 106–124.
12. Y. Zheng, ACM Trans. on Intell. Syst. Technol. **6** (2015).

Improving Activity Classification Using Highly Detailed Demographic Profiles

SANTIAGO GARCIA-GABILONDO[1], KETIAN CHEN[1], ZIYUE CHEN[2], JUE MA[1,*], AND YOSHIHIDE SEKIMOTO[1]

[1] The University of Tokyo
[2] Technical University of Munich
[*] majue@csis.u-tokyo.ac.jp

1. INTRODUCTION

Accurately classifying human activities from location data is a central task in mobility analysis, with applications in transportation planning and behavioral research. While prior research has explored the integration of ancillary data such as activity surveys and point of interest (POIs) [1–3], limited attention has been given to the influence of specific attributes in improving activity classification accuracy. This work examines whether and to what extent fine-grained demographic profiles (e.g., age, gender, education level, professional category, and household type) combined with POI attributes improve activity classification models. Building on this motivation, we develop a multi-step approach that leverages commonly available data (i.e. Point of Interest (POI) data and time-use surveys) to enhanced stay point classification performance across diverse demographic subgroups.

2. METHODOLOGY

A. Preliminary Analysis of Behavioral Heterogeneity by Demographic Attributes

Users are grouped into subpopulations based on five demographic variables: gender, age, education level, professional category, and household type. Each subgroup contains over 100 users, ensuring a sufficient sample size. We analyze two behavioral dimensions: (1) activity probability by time of day and (2) activity duration. To compute activity probability, we divide the 24-hour day into 96 fixed 15-minute intervals. Each activity episode is mapped to intervals based on its start time and duration. We calculate the proportion of individuals in a given subgroup who engaged in that activity during the interval. This produces a time-series probability distribution for a specific activity type. Figure 1 shows the probability distribution of leisure activity by subgroup, demonstrating variation by time of day. The probability distribution based on professional category reveals more distinct patterns than that based on gender.

To measure subgroup heterogeneity, we use the two-sample Kolmogorov–Smirnov (KS) test to quantify the difference in activity probability distribution between each subgroup and the overall population. Table 1 shows the subgroups with the highest KS statistics, indicating significant deviations compared to the overall population. We hypothesize that attributes with the most variance would have the largest impact on activity classification accuracy.

Table 1. All activities KS test relative to the overall population

Subgroup		Exp. Pop.	KS Stat.
Age Group:	over 65	120,839	0.920
	16-24	207,231	0.863
	55-64	208,598	0.783
	25-34	284,815	0.675
	45-54	346,931	0.554
Household Type:	Living with one or both parents	139,455	0.679
	Another family member in the household	30,331	0.628
	Not related to other household members	21,128	0.524
Professional Category:	Retired	150,935	0.838
	Unemployed	200,328	0.744

B. Stay Point Detection

To test the robustness of our stay point detection and activity classification methodology under varying data quality conditions, we artificially reduced the quality of the dataset by randomly sampling the total GPS, preserving 1%, 5%, 10%, 25%, 50%, and 100% of the original trajectories of each user. The resulting datasets were processed using the trajectory stop detector algorithm from the MovingPandas library. We also analyzed the impact of different algorithm parameters in the detection of activities and their duration [4]. We tested the maximum diameter of a stop using 125, 200, and 300 meters; and the minimum stay duration using 2, 5, 10, 15, and 20 minutes.

C. Activity Classification

We estimate joint probabilities at any given stay point for all activity types described by the dataset (home, work, studies, business, leisure, health, purchase, accompany, and other). These probabilities combine: (1) POI data to identify the feasibility of a given activity type occurring at a given location and time of the day; and (2) survey-based tables that capture activity likelihoods based on time-of-day and duration for each stay point.

To identify the feasibility of a given activity, we could not rely on POI proximity as it has been done in prior literature [2, 4, 5] because the data was spatially blurred to the H3 level 10. To overcome this limitation, we focus on identifying the probability of a particular activity by time of day given the open businesses within a specific H3 grid. To do so, we manually mapped specific business types to the nine activity types and weighted the probability of each activity using the average opening hours by POI category (obtained from OSM) to approximate the behavioral logic of businesses. To identify the likelihood for specific activities by time of day, we followed the procedure described on subsection A. Finally, we derived the probability distribution for each activity type by its duration from survey tables without considering our defined subpopulations.

After the joint probability estimation for each stay, we follow a conventional classification workflow [4, 6, 7], initially extracting a person's most significant anchor points (i.e. home and work/study locations) and subsequently classifying all other stay points. We use DBSCAN to cluster spatially proximate stay points, obtaining the cumulative probability (i.e. summing all individual stay point probabilities) by activity type for each cluster. The home and then work locations are identified as the clusters with the highest cumulative probability for those activity types.

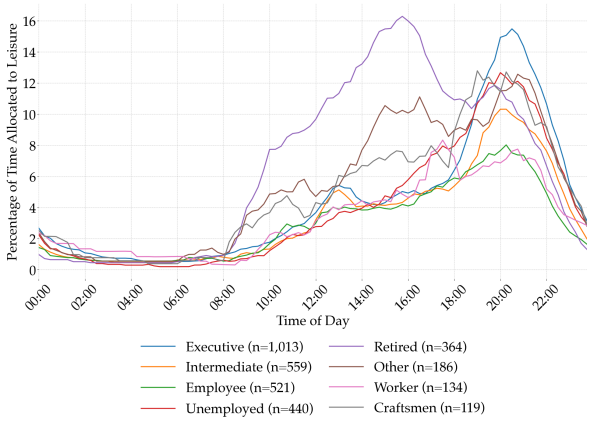

(a) Probability of Leisure Activity by Professional Category

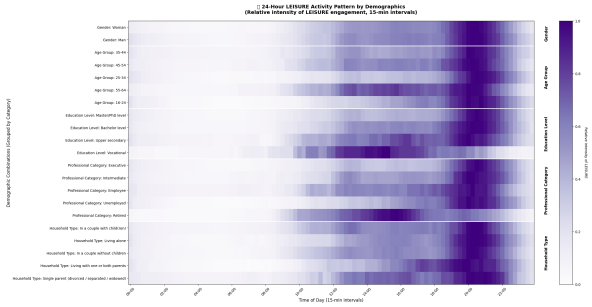

(b) Probability of Leisure Activity for each Demographic Group by Time of the Day

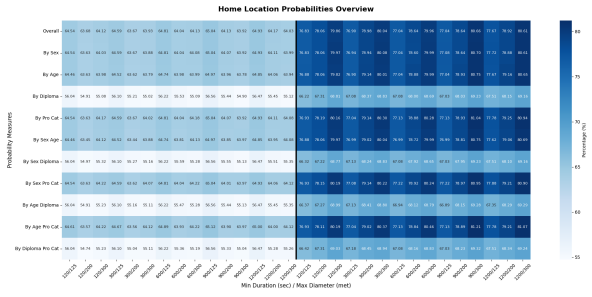

(c) Home Location Probabilities Comparison: left-1% sampling rate, right-100% sampling rate

Fig. 1. Preliminary Data Analysis and Classification Results

3. RESULT AND CONTRIBUTIONS

Validation is performed by checking whether the highest-probability prediction matches the ground-truth activity, scoring 1 for correct matches and 0 otherwise. Table 2 summarizes the accuracy of classification evaluation. We found that using time-use and duration statistics is more effective for activity classification than solely using POI data. We confirmed our hypothesis that more detailed demographic summary statistics improve classification, but the increase is less than expected. Furthermore, we confirmed the effectiveness of combining POI and time-use surveys derived statistics to improve activity classification across all demographic groups and sampling rates (Table 3).

This study contributes to the literature by addressing four key questions: 1. **What demographic attributes matter most for activity classification?** We analyze subgroups deviations from general behavior and assess how incorporating demographic attributes improves classification accuracy. In doing so, we also answer how well can activities be classified without demographic data (using the overall population statistics). 2. **How can we leverage POI data with spatially blurred mobility data?** We propose a novel method to aid activity classification using POI opening schedules by category, which works even on spatially blurred data. 3. **How does data sparsity affect classification?** By simulating degraded GPS sampling, we assess robustness under typical app-based mobility data constraints.

Table 2. Classification Accuracy Results(Excluding Home and Work Activities)

Model Type	Accuracy
POI Data Only	0.1553
Time/Duration Profiles (Overall)	0.1728
Time/Duration Profiles (by Sex)	0.1726
Time/Duration Profiles (by Age)	0.1864
Time/Duration Profiles (by Diploma)	0.1714
Time/Duration Profiles (by Professional Category)	0.1862
Combined Models (with POI)	
Overall w/ POI	0.2573
Sex w/ POI	0.2584
Age w/ POI	0.2713
Diploma w/ POI	0.2352
Professional Category w/ POI	0.2813
Sex & Age w/ POI	0.2741
Sex & Diploma w/ POI	0.2387
Sex & Professional Category w/ POI	0.2845
Age & Diploma w/ POI	0.2512
Age & Professional Category w/ POI	0.2854
Diploma & Professional Category w/ POI	0.2609

Table 3. Impact of GPS Data Sampling Rate on Classification Accuracy (Excluding Home and Work Activities)

GPS Sample Rate	POI Data Only	Overall	Overall w/ POI	Age & Pro. Cat.	Age & Pro. Cat. w/ POI
1%	0.132	0.110	0.217	0.139	0.247
5%	0.151	0.166	0.253	0.194	0.283
10%	0.157	0.182	0.263	0.208	0.293
25%	0.162	0.190	0.268	0.216	0.296
50%	0.167	0.197	0.274	0.220	0.299
100%	0.162	0.192	0.269	0.215	0.295

4. BACK MATTER

Funding. This study was supported by JST SPRING, Grant Number JPMJSP2108.

Disclosures. The authors declare no conflict of interest.

REFERENCES

1. B. Yin and F. Leurent, Transp. Res. Rec. **2675**, 771 (2021). Publisher: SAGE Publications Inc.
2. L. Huang, Q. Li, and Y. Yue, "Activity identification from GPS trajectories using spatial temporal POIs' attractiveness," in *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Location Based Social Networks,* (ACM, San Jose California, 2010), pp. 27–30.
3. N. Sari Aslam, M. R. Ibrahim, T. Cheng, *et al.*, Geo-spatial Inf. Sci. **24**, 711 (2021). Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/10095020.2021.1985943.
4. C. Zhong, S. Aslam, Y. Wang, *et al.*, Work. Pap. Ser. (2024). Publisher: University College London.
5. L. Yin, N. Lin, and Z. Zhao, Cities **109**, 103013 (2021).
6. M. Allahviranloo and W. Recker, Transportation **42**, 561 (2015).
7. A. Chakirov and A. Erath, p. 23 p. (2012). Artwork Size: 23 p. Medium: application/pdf Publisher: ETH Zurich.

MULTI-MODAL MOBILITY FOR GENERATING ELECTRICITY: WALKING AS CASE-STUDY

HOSSAMALDEEN ELGHAZALY[2], YASMIN RADWAN[2], YAHYA ZAKARIA[3], NADINE TAMISH[4], NADA SALEH[5], MOHAMED H. ABDELLATIF[6], AND NOHA GAMALELDIN[1]

[1] School of Information Technology and Computer Science (ITCS), Center for Informatics Science (CIS), Nile University, Giza, Egypt
[2] School of Information Technology and Computer Science (ITCS), Nile University, Giza, Egypt
[3] Faculty of Computers and Data Science, Alexandria University, Alexandria, Egypt
[4] Faculty of Engineering, Alexandria University, Alexandria, Egypt
[5] Faculty of Computers and Information, Tanta University, Gharbia, Egypt
[6] Faculty of Systems and Biomedical Engineering, Cairo University, Giza, Egypt

CORRESPONDING AUTHORS: H. ELGHAZALY (e-mail: h.rabie2272@nu.edu.eg) AND N. GAMALELDIN (e-mail: Ngamal@nu.edu.eg).

I. INTRODUCTION

FROM human activity to human utility, current global energy scenarios indicate a substantial increase in future electricity demand, driven not only by the pursuit of comfort and improved quality of life but also by the growing reliance on promising technologies in this era of rapid technological advancement.

Consequently, alternative energy sources are becoming increasingly vital for meeting these demands sustainably[1]. France's domestic coal, oil, and natural gas resources are rather limited, which is why the state was concerned with the issue of energy independence[2].

In this context, piezoelectric energy harvesting offers a promising method for converting mechanical energy, which is often lost in urban roadway traffic, into usable electricity. This approach has gained attention from both academia and industry[3]. Our study aims to examine the NetMob 2025 Data Challenge dataset by leveraging mobility traces to identify locations in the Île-de-France region with the highest energy generation potential.

This analysis will inform the strategic placement of piezoelectric modules, support the simulation and optimization of energy collection of unused energy, and contribute to a more informed assessment of the economic viability of the technology for large-scale deployment. Furthermore, the study explores the potential of piezoelectric energy harvesting based on spatial mobility data by developing a machine learning model that simulates the energy conversion process and classifies zones into low, medium, or high energy generation categories.

II. METHODOLOGY

The methodology adopted a sequential, data-driven approach to estimate the potential piezoelectric energy generation from pedestrian footsteps in the Île-de-France region. The process began with the acquisition and integration of multi-source data, including a large-scale GPS mobility dataset (NetMob 2025), weather records, and OpenStreetMap road network data. A critical preprocessing phase was undertaken to clean the data, handle missing values, and apply spatial deblurring techniques to recover accurate pedestrian trajectories from anonymized GPS data. The cleaned data were then enriched through feature engineering, creating variables that captured temporal, weather, and road-type characteristics.

Subsequently, the processed data were aggregated to the road segment level to calculate the daily pedestrian flow. These aggregated values were used to classify each road into categories of low, medium, or high pedestrian density using quantile binning. Machine learning models were then developed to predict these density classifications and estimate continuous pedestrian flow values. An ensemble of algorithms, including XGBoost and LightGBM, was trained and validated on a stratified subset of the data to achieve high predictive accuracy.

Following model development, the methodology was applied to a selection of case study streets chosen for their high pedestrian activities, as shown in Figure 1. Machine learning models were used to forecast the daily footstep counts on these streets. These predictions were fed into a standardized mathematical framework that converted footsteps into estimated energy output and economic savings based on the specifications of piezoelectric tiles and local electricity costs over a projected 20-year lifespan. Finally, a simulation analysis was conducted to evaluate how weather conditions affect pedestrian footsteps and, consequently, electricity generation potential. This study systematically perturbed key meteorological input features, specifically the normalized average temperature and relative humidity, to assess their impact on pedestrian mobility predictions under varying scenarios.

III. RESULTS

The analysis of the NetMob 2025 dataset revealed that walking constituted 28.7% of all transport modes in Île-

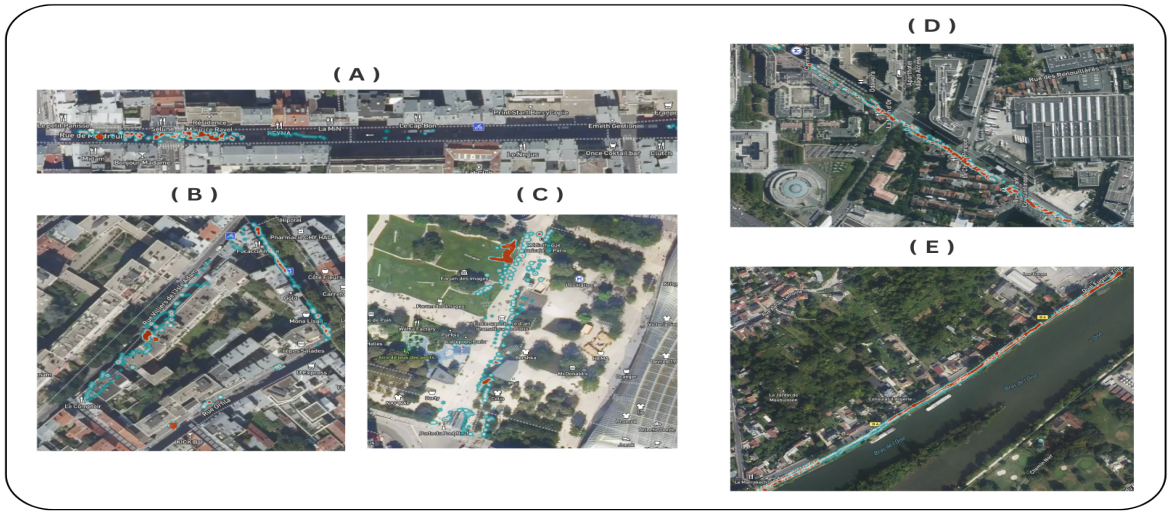

FIGURE 1. Illustration of the selected roads: (A) Rue de Montreuil [Paris], (B) Rue des Pyrénées [Paris], (C) Allée André Breton [paris], (D) Boulevard Anatole France [IS-(Inside Suburbs)], and (E) Quai Eugene Turpin [OS-(Outside Suburbs)].

de-France, highlighting its significance for energy harvesting. Machine learning models demonstrated exceptional performance in predicting pedestrian activity, with XGBoost achieving 98.9% accuracy in classifying road segments into three balanced density categories. The regression model also showed strong predictive capability (R^2 = 0.87) for estimating pedestrian flow.

TABLE 1. Daily Energy Output from High-Density Streets

Street Name	Steps	Energy Generated (kW h)
Rue de Montreuil	56,325	4,506.00
Rue des Pyrénées	76,299	6,103.92
Allée André Breton	49,855	3,988.40
Boulevard Anatole France	79,797	6,383.76
Quai Eugene Turpin	93,164	7,453.13
Total	**355,440**	**28,435.21**

The energy harvesting potential across five high-density streets yielded a collective daily output of 28,435 kWh (Table 1). Economic analysis projected substantial returns, with yearly savings ranging from $4.37–$8.16 Million and 20-year returns between $87.4–$163.3 Million, achieving 99.68% operational payback.

Simulation revealed non-linear relationships between meteorological conditions and pedestrian mobility, with colder temperatures paradoxically increasing mobility in specific contexts. These findings demonstrate robust scalability potential for urban energy harvesting applications.

IV. DISCUSSION

This study moves beyond theoretical estimates by using real GPS mobility data to identify high-density pedestrian streets, enabling more accurate energy forecasts, shorter payback periods, and more efficient piezoelectric tile deployments. By combining mobility data and machine learning, this study offers a replicable framework for sustainable infrastructure planning. The limitations of this study include the reliance on volunteer GPS data and idealized assumptions about long-term efficiency without considering maintenance or degradation. The harvested energy can support street lighting, IoT networks, public buildings, and emergency backups and serve as a symbol of renewable innovation. Overall, piezoelectric tiles represent a scalable urban energy source with environmental, economic, and social benefits that contribute to the vision of smarter, safer, and more sustainable cities.

References

[1] C. Jettanasen, P. Songsukthawan, and A. Ngaopitakkul, "Conversion of mechanical energy to electrical energy using piezoelectric materials for bicycle lane lighting systems," *Applied Sciences*, vol. 12, no. 14, p. 1, 2022. DOI: 10.3390/app12147237.

[2] U.S. Energy Information Administration, *Nuclear power plants generated 68% of france's electricity in 2021*, https://www.eia.gov/todayinenergy/detail.php?id=55259, Accessed: Jul. 9, 2025, 2021.

[3] R. M. A. D. Rajapaksha, P. K. V. M. Senevirathne, and M. N. V. Fernando, "Design and optimization of piezoelectric energy harvesting systems," *J. Res. Technol. Eng.*, vol. 6, no. 1, 2025.

Beyond Willpower: Structural Constraints and Inequality in Mobility-Related CO_2 Emissions

YUXI ZHANG[1], KANATA TAKAHASHI[2], SIJIAN TIAN[1], HIBIKI SUMIOKU[3], AND YUYA SHIBUYA[1,*]

[1] The University of Tokyo
[2] Hitotsubashi University
[3] Nagoya University
[*] yuya-shibuya@iii.u-tokyo.ac.jp

1. INTRODUCTION

Anthropogenic greenhouse-gas (GHG) emissions—especially carbon dioxide (CO_2)—are the main driver of global warming [1]. The Paris Agreement recognizes the urgent need to reduce CO_2 emissions, yet contributions to climate change are not evenly shared [2]. Because mitigation policies act on people and places with very different capacities and constraints, understanding who drives emission growth and who delivers mitigation is essential for effective and fair climate action; otherwise, well-meant measures may miss their targets [2].

Transport is a major source of GHGs in Europe and remains difficult to decarbonize [3]. Despite broad progress in economy-wide emissions, transport-related emissions have risen in many member states, underscoring the need to identify the drivers behind mobility emissions if long-term goals are to be met [4]. Recent work also shows that mobility emissions are highly unequal across individuals and are concentrated in the upper tail of the distribution [4]. Same mobility-related carbon inequality also being observed in Paris (Figure 2). This brings the distributional question to the foreground: sector trends translate into uneven burdens and benefits at the individual level.

Although previous studies have documented substantial disparities in mobility-related emissions, the relative contributions of structural constraints and individual behaviors and their potential interactions remain insufficiently understood. In this study, we ask: Which structural conditions and individual attributes are significantly associated with inequality in mobility-related carbon emissions?

We use Hägerstrand's [5] time-geography "constraint" framework as an analytical lens to categorize potential determinants into capability, coupling, and authority constraints, and test their associations with individual weekly travel emissions.

2. RESEARCH QUESTION AND HYPOTHESES

This study addresses the core question: **Which structural conditions and individual attributes are associated with inequality in mobility-related carbon emissions?** Following Hägerstrand's time-geography framework [5], we conceptualize three types of constraints: capability, coupling, and authority, as potential structural determinants of individual travel-related emissions. Accordingly, we test two hypotheses:

H1: A combination of capability, coupling, and authority constraints is systematically associated with the baseline level of an individual's travel-related carbon emissions.

H2: Constraints are not independent but interactively shape individual emissions.

3. DATA & METHODS

The analysis is based on a large-scale travel diary dataset from the Greater Paris region [6]. After preprocessing, the final modelling sample comprises $N = 3{,}179$ individuals.

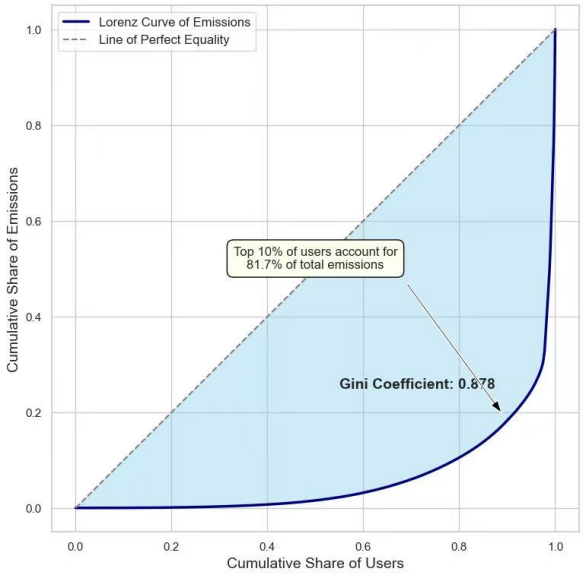

Fig. 1. Lorenz curve of weekly carbon emissions, showing strong inequality (Gini coefficient = 0.878; top 10% emit 81.7% of total emissions).
Source: Authors' calculation based on the NetMob25 dataset [6] (additional data as described in Methods [7] [8]).

Dependent variable (Y). The outcome of interest is each user's *weekly travel-related CO_2 emissions*. We used the origin–destination area codes from the trip data [6] and matched them with INSEE geospatial data [7] to estimate trip distances. These distances were combined with transport-mode-specific emission factors from Île-de-France Mobilités [9] and aggregated at the user level into weekly totals.

Independent variables (X). Predictors were structured according to Hägerstrand's time-geography "constraint" framework [5]:

- *Capability constraints*: demographic and resource-based limits on action, including age, car ownership, driving license, and public transport subscriptions.
- *Coupling constraints*: interdependencies with household and occupational structures, including family composition (e.g. presence of children, single parent households) and employment status.
- *Authority constraints*: spatial and institutional restrictions, operationalized as residential location categories (city center, inner suburbs, outer suburbs).

To test whether the effect of capability constraints was moderated by authority constraints, we also pre-specified interaction terms, including Age × Suburbs and Car ownership × Suburbs.

Analytical procedure. The dependent variable was right-skewed; we applied a common 3×IQR filter to remove 158 extreme outliers. We first performed correlation pruning, removed highly correlated predictors($|r| > 0.7$), thus excluded were `OTHER_SUB_PT` (indicating whether the individual holds a public transport subscription other than Navigo or Imagine R)and `RATIO_CHILDREN` (the proportion of children in the household), the latter likely due to collinearity with household structure dummies. The final model was estimated using Ordinary Least Squares (OLS) with HC3 robust standard errors, after confirming that all Variance Inflation Factors (VIFs) were below 5.

4. RESULT

The final Ordinary Least Squares (OLS) regression model was statistically significant (Prob(F-statistic) < 0.001), explaining approximately 21.1% of the variance in weekly mobility-related emissions (Adjusted $R^2 = 0.211$) for the sample of 3,179 individuals. The following sections detail the statistically significant associations ($p < 0.05$) organized by the theoretical constraint framework.

A. Capability Constraints

As hypothesized, several capability constraints were significantly associated with emissions. The strongest positive correlate was car ownership. In contrast, age showed a significant negative association, as did having no diploma.

B. Coupling Constraints

Two coupling constraints were found to be significant. Contrary to some expectations, being a single-parent household was associated with significantly higher emissions. In contrast, being in a specific occupational category 'unemployed or inactive'—indicating individuals without paid employment, was linked to lower emissions.

C. Authority Constraints

The most powerful authority constraint was geographical. Residence in the outer suburbs (`LOC_CAT_Outer_Suburbs`) was strongly and positively associated with higher weekly emissions. However, pre-specified interaction terms (Age × Suburbs, Car ownership × Suburbs), did not reach statistical significance ($p > 0.05$).

5. DISCUSSION & LIMITATION

This study provides an associational map of mobility-related carbon inequality. First, our results show that individual emissions are systematically linked to structural constraints rather than being mere outcomes of personal choice. Capability constraints (notably car ownership) and authority constraints (residence in the outer suburbs) emerge as the strongest drivers of high-carbon mobility. This pattern is consistent with other European contexts [4, 10], reinforcing the view that suburban, car-oriented lifestyles are a central driver of carbon inequality.

Second, we identify a notable social dimension: single-parent households display higher emissions, a result that may reflect unobserved coupling constraints such as 'time poverty'. This mechanism remains speculative but highlights the importance of incorporating household structures into future analyses.

Third, our findings suggests a policy implication: interventions should focus on altering the macroscopic structures that limit low-carbon choices, such as improving suburban public transport, and on supporting vulnerable households—particularly single-parent households, which in our analysis were associated with significantly higher emissions—rather than focusing solely on individual behavior choice.

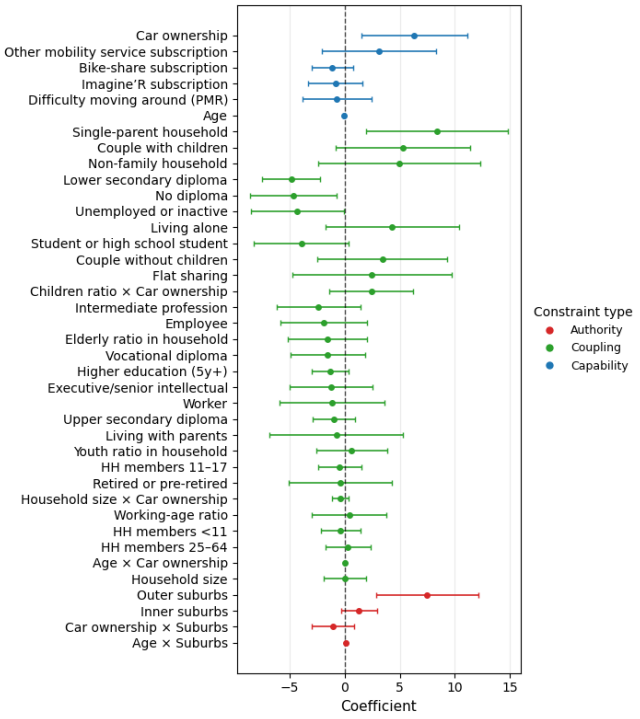

Fig. 2. OLS coefficient estimates with 95% confidence intervals, colored by constraint type.

Finally, the study has limitations that inform our future work. Our findings are associational and not causal; issues such as residential self-selection remain unresolved. In addition, the model's modest explanatory power (Adj. $R^2 \approx 0.211$) suggests that unmeasured factors—including accessibility conditions and individual-level difference (e.g., environmental attitudes, habitual practices)—play an important role. Future research could strengthen robustness by integrating richer structural indicators and applying alternative modeling strategies, like generalized linear models.

Disclosures. The authors declare no conflicts of interest.

Data Availability Statement. The authors do not have permission to share the data used in this analysis.

REFERENCES

1. S. J. Davis and K. Caldeira, Proc. National Acad. Sci. **107**, 5687 (2010).
2. L. Chancel, Nat. Sustain. **5**, 931 (2022).
3. World Bank, "Enhancing transport decarbonization in the european union," https://www.worldbank.org/en/region/eca/publication/eu-transport-decarbonization (2025).
4. F. Klein and N. Taconet, Energy Econ. **136**, 107630 (2024).
5. T. Hägerstrand, Pap. Reg. Sci. Assoc. **24**, 6 (1970).
6. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," (2025).
7. INSEE, "Base officielle des codes postaux," https://www.data.gouv.fr/datasets/base-officielle-des-codes-postaux/ (2025).
8. Expert Group for Urban Mobility [EGUM], "Social climate fund: Public transport and shared mobility. european commission," (2024).
9. Île-de-France Mobilités, "Emission de coe par voyageur-kilomètre sur le réseau," https://prim.iledefrance-mobilites.fr/en/jeux-de-donnees/emission-de-co2e-par-voyageur-kilometre-sur-le-reseau (2024).
10. C. Brand and B. Boardman, Energy Policy **36**, 224 (2008).

Higher-order Network Analysis of Human Mobility Data

CHEN ZHANG[1], ESTÉBAN NOCET-BINOIS[1], TIMOTHY LAROCK[1,*], AND JÜRGEN HACKL[1]

[1] *Complex Infrastructure Systems Group, Department of Civil and Environmental Engineering, Princeton University, Princeton, New Jersey 08544*
[*] *larock@princeton.edu*

Understanding and predicting mobility dynamics in transportation networks is critical for many engineering and policymaking tasks such as infrastructure planning, resilience assessment, and traffic management. However, analyzing human mobility data presents unique challenges, from the difficulty of collecting representative datasets with high spatial and temporal resolution to the methodological and computational demands of their analysis, compounded by the need to rigorously protect individual privacy. These challenges have led to research on how to model human behavior without collecting large-scale data through the use of synthetic populations and mobility simulations [1–5]. Such data can be leveraged to build detailed models of mobility, however their validity and correspondence to real-world patterns must be systematically evaluated to ensure that insights derived from synthetic data generalize to real systems.

In this work, we develop a framework for analyzing human mobility data using *higher-order network science* and explore how this framework can be used to evaluate the validity of simulated mobility datasets [6, 7]. The key difference between traditional network-based models and the higher-order approach we take here is the treatment of *memory*. Traditional network models assume that movement between nodes is memoryless, meaning that the next step is based only on the current position. This assumption limits the ability of these methods to capture sequential dependencies inherent to real-world mobility data, which arise from the intentionality of human movement from origins (*e.g.*, home) toward destinations (*e.g.*, work or school) [8–10]. It is important to understand whether simulated mobility datasets accurately reproduce these higher-order dependencies. Towards this end we make three primary contributions in this work:

1. We build a pipeline to convert raw GPS traces representing movement through geographic space into sequences of nodes visited in a network (*i.e.*, network *walks* or *paths*).
2. Taking paths as our unit of analysis, we apply higher-order network analysis techniques to understand mobility patterns in the 2025 NetMob Data Challenge Dataset [11].
3. We use the Open Île-de-France MATSim scenario [5] to generate a state-of-the-art synthetic population mobility dataset and compare the observed and synthetic datasets from a path-based higher-order perspective.

We analyze two datasets representing human mobility in the Île-de-France region. In both datasets we focus our analysis on the use of private cars mainly due to the fact that the underlying network data better covers road travel than other modes.

Data. The 2025 NetMob Data Challenge Dataset includes over 80,000 mobility trajectories of 3,337 participants living in the Île-de-France regio over a period of 7 days between October 2022 to May 2023 and is described in detail in [11]. Each participant volunteered to carry a dedicated GPS device and log the purpose of each of their trips. The dataset also includes surveyed demographic data about individual participants as well as metadata for trips and geographic regions. The dataset represents the movements of participants around the Île-de-France region as sequences of GPS coordinates representing participant position over time. These sequences are logically split into trips based on participant logs. Each trip also includes the date, start and end time, primary mode of travel (e.g. private car, taxi, or bus), and the purpose of the trip. We focus our analysis on trips taken using private cars, which includes 21,972 individual GPS trajectories (more than 27% of all trips).

We generate a second source of mobility data using the simulation software MATSim, a state-of-the-art iterative, agent-based methodology for dynamic traffic assignment [1–5]. We use the open-source Ile-de-France MATSim synthetic population generation and MATSim simulation code made available by Horl et al. [5]. Constructing the synthetic population and running the MATSim simulation takes advantage of a variety of public regulatory data as documented in the paper describing the simulation [5]. The output of the MATSim simulation includes the paths agents took through the network using each possible mode of travel. We generated 17,816 synthetic paths for agents using private automobiles for comparison with the NetMob 2025 dataset.

Methods. We transform the raw GPS traces into paths through an underlying network by performing map matching. We adopt a method based on Markov Decision Proccesses, which has been benchmarked against hand-labeled data as well as other map matching methodologies and implemented in the Map Matching 2 package [12]. We match the NetMob participant trajectories to the OpenStreetMaps (OSM) [13] street network of the Île-de-France region [14]. Figure 1a shows a partial map of the region with orange lines corresponding to aggregated map-matched trajectories as well as stylized examples of paths through the underlying street network. Once each GPS coordinate in a trip is assigned to an edge in the road network, the resulting sequence of edges (duplicates removed) defines the path of the participant through the network, resulting in a total of 21,119 paths.

We adopt the framework for higher-order network analysis introduced by Scholtes [6]. In this framework, sequential patterns are modeled with higher-order Markov chains, represented formally as directed and weighted *de Bruijn graphs* [15]. In a de Bruijn graph of order k, each node represents a sequence of k consecutive nodes, and directed edges connect overlapping subpaths, encoding valid transitions observed in empirical trajectories. A de Bruijn graph can be directly constructed from observed trajectories by sliding a window of k nodes across them and counting the frequency of each subpath. Given a set of paths through a network, we construct the observed de Bruijn graph for any order from $k = 1$ (a traditional weighted and directed network) to $k = \ell$, where ℓ is the maximum observed path length in the input dataset. We can also combine de Bruijn graph models of increasing orders into multi-order models and apply statistical tests to determine the optimal order that explains the observed data [6]. Such models can also be used to compute the likelihood of a path given the data, which can in turn be used to predict the most likely next step given a partial sequence.

Results. We plot path length distributions for the NetMob (1b) and MATSim (1c) datasets, both in terms of the number of nodes visited (horizontal axis) and the length measured in meters (vertical axis). Paths in the observed NetMob data are notably longer than those from the MATSim simulation.

We want to understand how well the path datasets cover

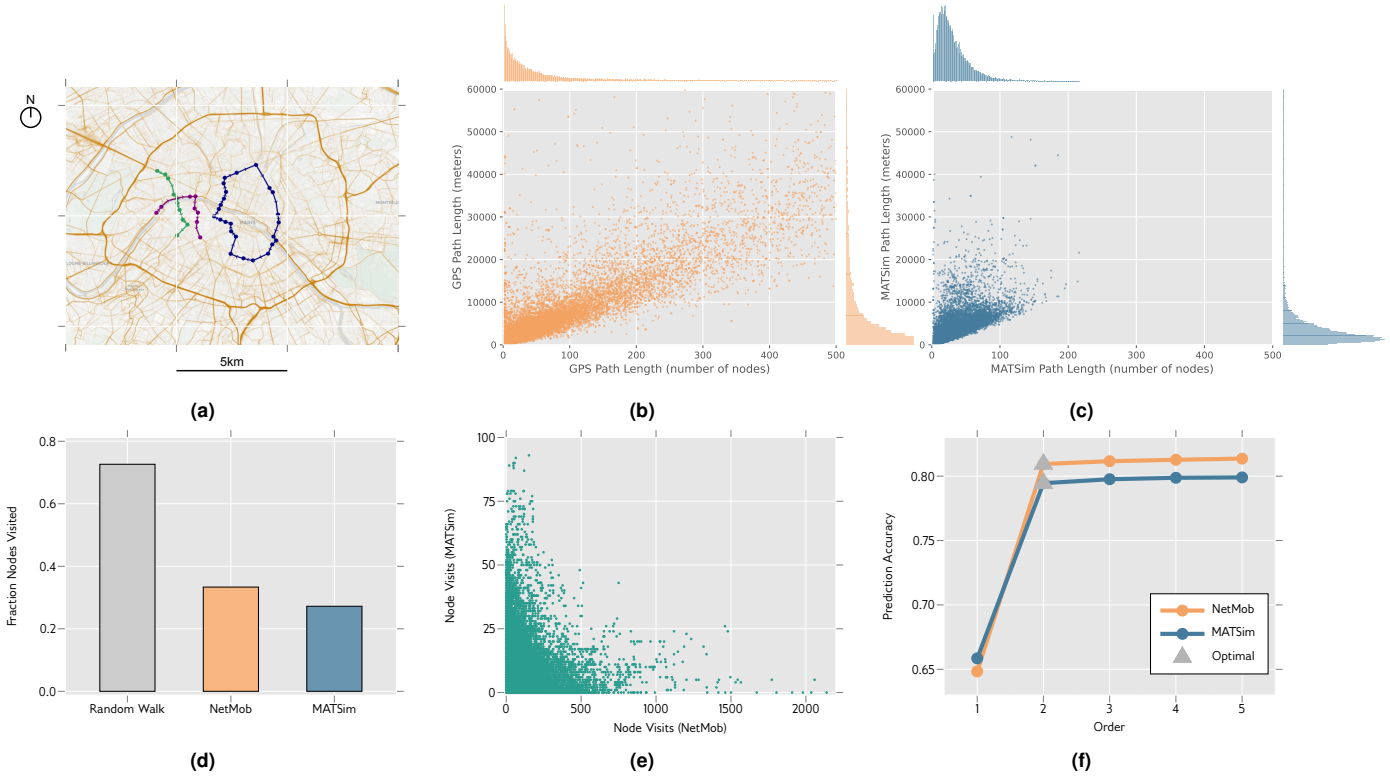

Fig. 1. (a): Map matched GPS trajectory data and (stylized) paths overlayed on the map of Île-de-France. (b-c) Distribution of path lengths in nodes (horizontal axis) and meters (vertical axis) corresponding to private car driving trips matched to the road network for NetMob (b) and MATSim (c). (d) Number of visits by participants/agents for each node in the road network. (e) Fraction of nodes in the network visited by NetMob participants, MATSim agents, and a random walker following the same path length and starting positions distributions as the two datasets combined. (f)): Next step prediction accuracy of higher-order network models.

the nodes in the underlying physical network. We count the number of times each node is visited by a participant or agent in the datasets by constructing a vector with an entry for each node and recording the number of visits the node recieved from NetMob participants and MATSim agents.

In Figure 1d, we report the fraction of nodes in the network that were visited by at least one agent, As a baseline, we also show results for a set of random walks (RW) over the undirected road network that has the same path length distribution and starting node distribution as the combination of the two datasets. The fraction of nodes visited by the NetMob participants is larger than the fraction visited by MATSim agents and both fractions are smaller than the random walker.

We also compare the number of times each node was visited in the two datasets in Figure 1e. The visitation counts are not directly correlated and in fact there are many nodes with 0 visitations in one dataset and nearly the maximum observed visitation in the other, indicating that the participants and agents visit nodes at different frequencies.

Finally, we test the ability of higher-order models with increasing maximum order to predict the next node in mobility paths by splitting the paths into training and testing sets, then using the training set to compute the conditional transition probability $P(v_i \mid v_{i-1}, \ldots, v_{i-K})$ for the next node in the sequence [6]. We then select the node with the highest probability as our next-step prediction. Figure 1f shows the maximum order of a multi-order model against the next step prediction accuracy. We find that the prediction accuracy improves dramatically when incorporating paths of length 2, which is also the optimal order according to the likelihood test proposed in [6], while increasing the order beyond 2 only marginally improves prediction accuracy while increasing the number of estimated model parameters.

Conclusion. In this work we have compared two human mobility datasets from the perspective of path-based higher-order network science, with the goal of assessing to what extent patterns observed in the NetMob 2025 dataset can be reproduced within the synthetic MATSim model. We have found that the datasets show similar overall trends in terms of the basic higher-order network measures, as well as next-step prediction, suggesting they share some similar characteristics. We also found that study participants and MATSim agents visited different nodes in the underlying network and visited common nodes at different rates. These results, while preliminary, suggest that there may be some gaps between the real-world mobility data and state-of-the-art simulations from the perspective of paths through the network.

There are many open questions and challenges to overcome based on the analysis we have shown here. For example, our analysis did not take into account the statistical weighting associated with the NetMob dataset, which should be accounted for in the future, as using these weights to adjust the importance of different trips may lead to different results for some of our metrics. Similarly, future work could align the synthetic population in the MATSim simulation with that of the NetMob 2025 Data Challenge, potentially yielding a closer correspondence between observed and simulated mobility patterns. Looking ahead, applying higher-order network science in this way can support the development of mobility models that not only better reflect empirical movement patterns but also provide more reliable tools for transportation planning and policy.

REFERENCES

1. K. Nagel and G. Flötteröd, “Agent-based traffic assignment: Going from trips to behavioral travelers,” (2009).
2. A. Chakirov and P. J. Fourie, “Enriched Sioux Falls scenario with dynamic and disaggregate demand,” (2014).
3. A. Horni, K. Nagel, and K. W. Axhausen, *The Multi-Agent Transport Simulation MATSim* (Ubiquity Press, 2016).
4. J. Hackl and T. Dubernet, Future Internet **11**, 92 (2019).
5. S. Hörl and M. Balac, Transp. Res. Part C: Emerg. Technol. **130**, 103291 (2021).
6. I. Scholtes, “When is a Network a Network? Multi-Order Graphical Model Selection in Pathways and Temporal Networks,” in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,* (Association for Computing Machinery, New York, NY, USA, 2017), KDD ’17, pp. 1037–1046.
7. C. Zhang and J. Hackl, “Beyond Connectivity: Higher-Order Network Framework for Capturing Memory-Driven Mobility Dynamics,” (2025).
8. S. Zhu and D. Levinson, PLOS ONE **10**, e0134322 (2015).
9. A. Lima, R. Stanojevic, D. Papagiannaki, *et al.*, J. The Royal Soc. Interface **13**, 20160021 (2016).
10. E. J. Manley, J. D. Addison, and T. Cheng, J. Transp. Geogr. **43**, 123 (2015).
11. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, “The NetMob25 Dataset: A High-resolution Multi-layered View of Individual Mobility in Greater Paris Region,” (2025).
12. A. Wöltche, Trans. GIS **27**, 1959 (2023).
13. OpenStreetMap contributors, “Planet dump retrieved from https://planet.osm.org,” (2017).
14. Geofabrik, “Ile-de-France,” https://download.geofabrik.de/europe/france/ile-de-france.html (2021).
15. N. G. Bruijn, de, Proc. Koninklijke Nederlandse Akademie van Wetenschappen **49**, 758 (1946).

How Daily Mobility Practices Reshape Greenspace Exposure Inequalities: Individual-level GPS Tracking Evidence from Greater Paris

HONG DENG[1,*], SHITIAN ZHANG[1], HEZHISHI JIANG[1], AND MIAO ZENG[1]

[1] The Bartlett Centre for Advanced Spatial Analysis, University College London (UCL), W1T 4TJ, UK
[*] deng.hong.22@ucl.ac.uk

1. INTRODUCTION

Exposure to greenspaces provides numerous health benefits, including mitigating air and noise pollution, reducing stress, and offering venues for physical activity and social interactions[1]. Accurately evaluating individual greenspace exposures is essential for understanding exposure-health relationships and addressing public health challenges through effective interventions.

Extensive research has typically employed a static approach to estimate greenspace exposure based on fixed locations, such as people's residences or workplaces [2]. However, mobility-based exposure estimates often differ substantially from residential assessments since individuals spend considerable time outside their residential neighbourhoods[3][4]. People with limited residential greenspace access may achieve higher exposure through daily mobility (commuting, shopping), whilst the reverse may also occur. Kwan [5] termed this convergence toward mean exposure levels the 'neighbourhood effect averaging problem' (NEAP), which stems from mobility constraints and complex travel patterns shaped by residential environment, income, and location preferences.

Subsequent research has applied NEAP theory to air pollution and greenspace studies[6], revealing that individuals strategically navigate environmental risks through daily mobility. For greenspace exposure, evidence suggests a 'polarisation effect' where mobility either amplifies residential advantages or compensates for disadvantages, rather than simply averaging exposure levels[7].

Despite recent preliminary findings on greenspace exposure and NEAP revealing temporal, spatial, and sociodemographic mechanisms[8][9][10], a critical knowledge gap exists regarding how mobility patterns are associated with NEAP, and how these relationships are moderated by transportation access, household composition and mobility impairment. Emerging micromobility services like e-scooters may enable residents in greenspace-poor areas to seek higher exposure through enhanced mobility opportunities. Conversely, declining mobility in older adults may constrain movement and exacerbate green space deficits for elderly residents in underserved areas. Households with intensive caring responsibilities may rely more heavily on private vehicles rather than active transportation modes, potentially reducing greenspace contact[11]. Meanwhile, low-income workers without cars may walk and use public transit from inner-city areas to suburban employment locations[12], achieving higher mobility-based greenspace exposure than residential assessments suggest. This study innovatively considers three evolving contexts: later-stage mobility decline, emerging transportation technologies, and family caring responsibilities to inform interventions promoting equitable greenspace access through targeted mobility solutions. The NetMob dataset, combining socioeconomic and GPS data, enables investigation of how inequalities shape greenspace exposure through daily mobility.This study addresses three research questions:

RQ1: How do different mobility pattern clusters relate to differences between trajectory-based and residence-based greenspace exposure?

RQ2: How do transportation resources moderate the mobility cluster-exposure relationships?

2. METHODOLOGY

We constructed a comprehensive variable framework capturing greenspace exposure, daily weather, multidimensional mobility patterns, and individual- and household-level characteristics.

Dependent Variable

Greenspace Exposure Difference, defined as the relative difference between static (residential-based) and dynamic (mobility-based) greenspace exposures for each participant-day. Greenspace exposure is operationalized using the Normalized Difference Vegetation Index (NDVI), derived from MODIS satellite imagery, which ranges from -1 to 1 (0 = no vegetation, 1 = dense vegetation). *Static exposure* is measured as the time-weighted mean NDVI within 300 m buffers around participants' home locations, while *dynamic exposure* is calculated analogously for all non-home locations visited during trips, weighted by dwell time. The difference (*static* $-$ *dynamic*) captures whether participants are exposed to greener or less green environments outside the home: positive values indicate greener residential surroundings compared to activity locations, negative values indicate greater greenness in out-of-home environments, and values close to zero suggest comparable levels across home and activity settings.

Weather Data

Daily meteorological variables were obtained from the Meteostat database. For each trip origin, the geographically nearest weather station ($\leq$ 100 km) was identified, and station-level daily data—including temperature, precipitation, wind speed and direction, air pressure, sunshine duration, and snow cover—were retrieved. These records were merged with trip data by date and station, thereby assigning each trip with its corresponding daily weather exposures.

Mobility Patterns

This study develops a comprehensive framework for characterizing individual daily mobility patterns using GPS-corrected travel diary data, extracting key indicators that capture multidimensional aspects of human movement behavior. Specifically, we calculate **Purpose Diversity Entropy** using Shannon entropy (base-2) of destination purpose frequencies to measure activity variety, **Transportation Mode Ratios** categorizing trips into active (walking, cycling), public transport, and motorized modes to assess sustainability preferences, and **Activity Radius** computed as the maximum distance from the centroid of all visited

locations to quantify spatial movement range. Additional indicators include trip chain length, total travel duration, distance, and number of unique locations visited. Collectively, these indicators provide a holistic characterization of individual mobility behavior, enabling the identification of distinct mobility profiles ranging from routine single-purpose commuters with limited spatial range to diverse multi-modal travelers with extensive urban exploration patterns.

Demographic and Socioeconomic Variables

Individual-level demographic variables include age (centered at the sample mean) and gender (male/female). Socioeconomic factors include educational attainment, occupational category, and mobility impairment status (PMR). These characteristics capture heterogeneity in resources, constraints, and accessibility that may systematically influence both mobility and environmental exposures.

Transport Access

Transport access variables include driving license ownership, household car ownership (capped at a maximum of three vehicles), two-wheeler ownership, bicycle ownership, and electric scooter access. Public transport accessibility is further proxied by Navigo pass (Île-de-France monthly transit subscription) usage. These indicators account for differential transport opportunities that may shape mobility behaviors and greenspace exposure.

Household Composition

Household characteristics include total household size, household type, and the number of minor children (under 18). In addition, the age structure of co-residents is recorded by counts of working-age adults (25–64 years) and older adults (65+ years). These variables capture family structure and dependency ratios, which may constrain travel behaviors and influence exposure opportunities.

Statistical Approach

We applied Bayesian multilevel regression models using the `brms` package, which are well-suited to our nested data structure where daily mobility observations are clustered within individuals across geographic regions. The general model specification regressed the daily *Greenspace Exposure Difference* on mobility clusters, demographic covariates, weather conditions, and household transport access, while including individual-level random intercepts to account for repeated measures. Formally, the outcome for individual i on day t was modeled as:

$$\begin{aligned} \text{GED}_{it} \sim\ & \beta_0 + \beta_1 \text{Cluster}_i + \beta_2 \text{Age}_i \\ & + \beta_3 \text{Gender}_i + \beta_4 \text{Temperature}_t \\ & + \beta_5 \text{Precipitation}_t + \beta_6 \text{CarOwnership}_i \\ & + u_i + \epsilon_{it} \end{aligned} \tag{1}$$

where GED_{it} denotes the Greenspace Exposure Difference for individual i on day t, $u_i \sim \mathcal{N}(0, \sigma_u^2)$ represents individual-level random effects, and $\epsilon_{it} \sim \mathcal{N}(0, \sigma^2)$ is the residual error term. Weakly informative priors were applied: $\mathcal{N}(0, 200)$ for the intercept, $\mathcal{N}(0, 100)$ for fixed effects, and Exponential(0.01) for variance components. Models were estimated using four Markov chains with 2,000 iterations each (1,000 warmup).

Our analytical strategy consisted of three phases: systematic variable selection, weight sensitivity analysis, and spatial effects modeling.

Phase 1: Testing Which Variables to Include. We started with a basic model that included mobility patterns (grouped using clustering analysis), demographic factors (age, gender), weather data, and individual differences. The mobility groups represented different travel behaviors like Active Commuters, Public Transit Users, Long-Distance Travelers, and Mixed-Mode Commuters. We then systematically tested additional variables (household transport access, income factors, family structure) one at a time using three criteria: (1) the model had to improve significantly, (2) the new variable had to show a meaningful statistical effect, and (3) our core results had to remain stable. We continued testing until we either evaluated all possible variables or three variables in a row failed our criteria. Phase 2: Accounting for Survey Weights. Since our main statistical software couldn't directly handle survey weights (which ensure our sample represents the broader population), we used three different approaches: (1) divided participants into weight groups and compared them, (2) adjusted our results after analysis, and (3) validated findings using specialized survey analysis software. This allowed us to check whether survey weights affected our conclusions. Phase 3: Handling Geographic Differences. We used statistical models that account for regional variations across 597 administrative areas. We measured how much of the total variation came from individual differences versus regional differences versus random variation. We compared different ways of modeling these regional effects to find the best approach.

3. DISCUSSION

Across all model specifications, two primary factors showed consistent and significant associations with differential greenspace exposure through daily mobility. Public transit users systematically encountered greener environments during their trips compared to their residential areas, with effect sizes ranging from -202.6 to -299.1 units depending on model specification. Car ownership consistently facilitated access to greener environments, with effects ranging from -110.8 to -193.9 units per additional vehicle.

Mixed-mode mobility patterns, long-distance travel patterns, and demographic variables (age, gender) showed no statistically significant associations across any model specification. Weather conditions (temperature and precipitation) also had no detectable influence on mobility-environment relationships. The substantial individual-level random variation (ICC = 0.372) and modest regional clustering (4.9% of variance) indicate considerable heterogeneity in how individuals experience environmental differences through their daily travel patterns, with most variation occurring within rather than between spatial units.

REFERENCES

1. S. de Vries, S. M. E. van Dillen, P. P. Groenewegen, and P. Spreeuwenberg, **94**, 26.
2. E. D. Ekkel and S. de Vries, **157**, 214.
3. M.-P. Kwan, **69**, 1311. Publisher: Elsevier BV.
4. M.-P. Kwan, **102**, 958. Publisher: Routledge _eprint: https://doi.org/10.1080/00045608.2012.687349.
5. M.-P. Kwan, **15**, 1841. Number: 9 Publisher: Multidisciplinary Digital Publishing Institute.
6. S. Gyanwali, S. Karki, K. M. Jang, *et al.*, **27**, 7.
7. J. Wu, B. Wang, N. Ta, and Y. Chai, **113**, 346. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/24694452.2022.2098086.
8. L. Zheng, M.-P. Kwan, Y. Liu, *et al.*, **258**, 119499.
9. J. Wang, M.-P. Kwan, G. Xiu, *et al.*, **243**, 104970.
10. Y. Lu, J. Kim, X. Shu, *et al.*, **253**, 105222.
11. C. D'Agostino, E. Piva, P. Pucci, and C. Rossi, **44**, 1016. Publisher: Routledge _eprint: https://doi.org/10.1080/01441647.2024.2350613.
12. Y. Xu, L. E. Olmos, D. Mateo, *et al.*, **3**, 611. Publisher: Nature Publishing Group.

Graphing the Jam: GNSS-Powered Modeling and Forecasting of Urban Traffic Congestion

NASRAH HUSSEIN , AHMED ZAKI , YOUSSEF OTHMAN , FADY DESOKY S. ABDELAZIZ , DINA OMAR , AND NOHA GAMAL EL-DIN[1*]

[1] Center of Informatics Science, School of Information Technology and Computer Science, Nile University, El Sheikh Zayed City, Egypt
[*] ngamal@nu.edu.eg

1. INTRODUCTION

The traffic congestion in cities has become an urgent issue of the health conditions, economic efficiency, and sustainability of the environment. It causes more time in travels, more consumption of fuel and more emissions. The Global Traffic Scorecard stated that vehicles on the street in big cities risk loss of thousands of dollars to traffic delays and associated expenses every year. [1]. Moreover, the world health organization associates vehicle related air pollution with severe respiratory and cardiovascular diseases. [2].

Today, when high-resolution GNSS mobility data is available, it becomes possible to study the individual mobility patterns instead of relying on aggregate models. In this paper, a graphical approach is put forward by which GNSS traces are transformed into a dynamic, directed graph over time of flows within and among zones in the urban area. It assists in congestion modeling and short-range traffic prediction and can provide better planning tools regarding mobility management and in the actual decision-making to manage congestion.

2. METHODOLOGY

The objective of this work is to create a graph representation to describe and forecast congestion based on GNSS mobility data in cities. To simulate the velocity of formation and propagation of congestion in different zones, it represents traffic as a dynamic, directed graph. The methodology aims at finding the main parameters affecting congestion, predicting future congestion levels using learning schemes, assisting practical mobility studies and parameter sensitivity studies, and allowing extension to multi-modal transport systems.

A. Congestion Modeling

We can represent urban mobility as a directed graph and, to structure it, the geographic areas, such as IRIS units, are encoded as nodes, and the movement of people between zones is encoded as edges. This framework will best fit to record the direction as well as the strength of flows over time. One reason is that the graphical representation will help in any change that occurs on a given node, like a change in flow or congestion, which will affect the entire network through its connections. These inter-dependencies can be modeled explicitly in the graph formalism, such that an observer can witness the way in which local conditions influence the wider mobility system. In our graph, in contrast to the origin-destination models that emphasize the start and end points of trips, each of the trajectories has zones between the starts and ends of the movements, allowing a fuller and more natural sense of the dynamics of movements.

To estimate local congestion, we define an activity measure a_s^t for each zone s in a time interval t, combining internal and external trip flows, as proposed in previous work using the NetMob25 dataset [3]:

$$a_s^t = \text{MoveInside}_s^t + \frac{\text{InComing}_s^t + \text{OutGoing}_s^t}{2}$$

- **MoveInside**: counts trips that begin and end within the same zone.
- **InComing**: counts trips arriving into the zone.
- **OutGoing**: counts trips departing from the zone.

We then compute a congestion proxy D_s^t as:

$$D_s^t = \frac{a_s^t}{A_s}$$

where A_s is the physical area of zone s. This density-based metric serves as a proxy for congestion, reflecting how concentrated the activity is within a given area [3].

In our simulation model, we vary key input parameters such as flow volume, directionality, and temporal distribution to analyze their impact on D_s^t over time. This setup allows us to identify which factors contribute most significantly to congestion buildup and how changes in mobility patterns affect the broader network.

We use a dynamic thresholding method to label zones with high congestion: at each time step, zones in the top 10% of D_s^t values are flagged as congested. This percentile-based approach adapts to natural variations and avoids reliance on fixed thresholds.

This type of formulation is designed as a modular system that can grow. Such future developments might include the incorporation of data on infrastructure, the differences in transport modes, or duration of stay, to allow capturing the dynamics of congestion in a particular urban setting.

The simulation is also run at discrete time steps, producing a timeline record containing the graph snapshots of the evolution of congestion as it builds and clears out. Measurements that come out of them are congestion heatmaps, node-level density profiles, and flow bottleneck diagnostics that serve both analysis and forecasting purposes.

B. Traffic Forecasting

To predict congestion in the short-term, traffic dynamics is modeled as a spatio-temporal learning task on a time varying graph. For each zone, the objective is to predict the next time step's congestion level D_s^{t+1} using historical trends in activity, flows, and local congestion. We incorporate features from both the focal zone and its connected neighbors to account for flow propagation across the network.

These temporal and structural dependencies are first realized by machine learning models. Inspired by recent advances in spatio-temporal modeling [4, 5], we consider extending to graph-based deep learning models such as spatial-temporal graph neural networks (STGNNs) and attention-based architectures. These models allow for learning from irregular graph topologies while

effectively integrating spatial and temporal patterns in traffic data [6].

Graph-based temporal attention mechanisms [4] have shown improved accuracy by leveraging both topology and periodic behavior in traffic flows. Such designs align with our data, which captures repeated urban movement behavior over time. By modeling traffic patterns as learned transformations over graph embeddings, these approaches can support more robust forecasting, even under noisy or partially missing observations.

This forecasting process allows for early congestion detection, facilitates real-time control policy on the traffic approach, and demonstrates the enablement of adaptive policy decisions on the urban mobility system.

3. PRELIMINARY OUTPUT

In order to illustrate the suggested method, we have built a directed mobility graph out of a subset of GNSS traces composed by the NetMob25 collection. The figure is constructed where 100 areas of the IRIS are selected; directed edges are provided where trip flows between two areas are observed. The nodes symbolize zones, and the edges cipher direction and power of the move.

Figures 1 and 2 illustrate two snapshots of the mobility graph, captured during rush hour and non-rush hour respectively. In each graph, node color reflects the computed congestion level D_s^t: dark-red shades indicate higher congestion based on the density metric introduced in Section 2. Congestion clearly concentrates in specific zones during rush hour, while the network shows more distributed and lower-density activity during off-peak periods.

To complement these visuals, Tables 1 present the normalized congestion parameter values for the five zones with the highest congestion scores in rush hours during the same time intervals. The values include the computed attendance a_s^t and the resulting congestion score D_s^t, both normalized using min-max normalization. These metrics help quantify how congestion evolves under changing flow conditions in the most congested areas.

This exploratory result justifies the practicality of producing explainable, time-dependent congestion maps of GNSS mobility data. It also establishes the grounds for further analyses, like the comparison of the pattern of congestion at times of day or the regression of the effect of partitioning of transport modes in future work.

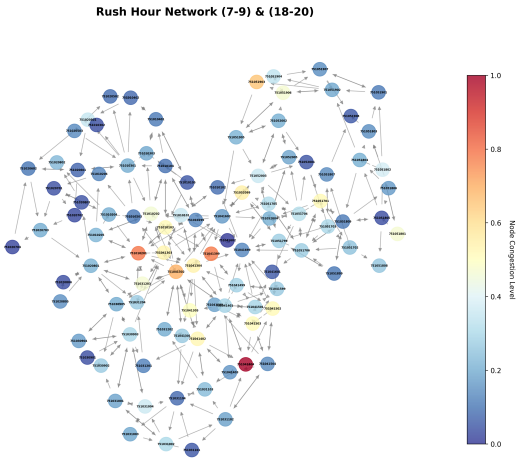

Fig. 1. Mobility graph during rush hour with nodes colored by congestion level D_s^t.

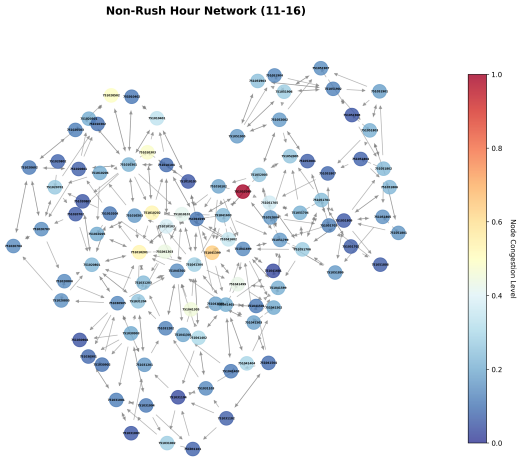

Fig. 2. Mobility graph during non-rush hour with nodes colored by congestion level D_s^t.

Table 1. Normalized Congestion Scores (D_s^t) for Congested Zones During Rush and Non-Rush Hours

Zone	Rush Hour D_s^t	Non-Rush Hour D_s^t
751041404	1.000	0.294
751010201	0.787	0.533
751041399	0.781	0.634
751041302	0.692	0.147
751051903	0.655	0.244

4. CONCLUSION AND FUTURE WORK

In this work, we propose a graph-based modeling, simulation, and prediction framework based on GNSS mobility data to study a key problem of urban congestion. We model mobility as a directed graph of dynamic changes that allows us to display the flow patterns between zones and to produce measurable metrics of congestion characterized by activity and density. Preliminary tests indicate that the approach is capable of creating a structured flow network out of real data. In the future, it is anticipated to provide support to several forms of transport and their characteristic congestion, and also work on graph-based deep learning to provide more precise and helpful input in relation to urban planning.

REFERENCES

1. INRIX, Inc., “Inrix global traffic scorecard,” https://inrix.com/scorecard/ (2022). Accessed: 2025-07-17.
2. World Health Organization, “Ambient (outdoor) air quality and health,” https://www.who.int/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health (2023). Accessed: 2025-07-17.
3. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, “The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region,” https://arxiv.org/abs/2506.05903 (2025). ArXiv:2506.05903v1.
4. S. Zhang, Y. Guo, P. Zhao, *et al.*, IEEE Trans. on Intell. Transp. Syst. **23**, 7743 (2022).
5. Y. Li, R. Yu, C. Shahabi, and Y. Liu, “Diffusion convolutional recurrent neural network: Data-driven traffic forecasting,” in *International Conference on Learning Representations (ICLR)*, (2018).
6. Z. Wu, S. Pan, F. Chen, *et al.*, “Graph wavenet for deep spatial-temporal graph modeling,” in *IJCAI*, (2019), pp. 1907–1913.

Optimizing Store Location Based on Urban Accessibility, Travel Behavior, and Customer Experience Insights

Alaa Atef[1], Rghda Salah[2], and Noha Gamal[3]

[1] *Faculty of Computers and Data Science, Alexandria University*
[*] *cds.Alaaatef41563@alexu.edu.eg*
[2] *Faculty of Engineering, Zewail City of Science and Technology, Giza, Egypt*
[*] *rghdasalah2@gmail.com*
[3] *School of Information Technology and Computer Science, Nile University, Giza, Egypt.*
[*] *ngamal@nu.edu.eg*

1. INTRODUCTION

Rapid urbanization and the availability of large-scale mobility data have opened new opportunities for data-driven planning and retail site selection. Traditional methods relying on surveys or static GIS often miss the complexity of urban dynamics [1, 2]. Recent studies leverage mobile phone and GPS data to analyze accessibility, trip complexity, and behavioral diversity [3], with accessibility measures [4] and centrality metrics widely applied. However, prior work has typically reduced trip complexity to mode counts [5, 6] overlooking structural choice behavior and temporal efficiency.

This study addresses these gaps by introducing a multidimensional trip complexity index and a *Reachability Score* that integrates both trip frequency and travel duration. Building on these, we define a composite *Opportunity Score* that combines behavioral, spatial, and demographic factors to rank urban zones. Using the NetMob25 dataset for the Paris metropolitan region [7], enriched with IRIS geospatial data [8], OpenStreetMap [9], and GTFS transit schedules [10], we classify zones into functional types (e.g., leisure magnets, commuter hubs, low-flow areas). Preliminary results reveal central high-activity communes (e.g., PARIS-1E, PARIS-8E) as commercial hotspots, while peripheral areas show untapped potential. The framework thus provides scalable, interpretable, and actionable tools for urban policy and retail investment.

2. METHODOLOGY

We constructed a data-driven pipeline to evaluate urban accessibility and visitation outcomes using mobility and sociodemographic indicators at the commune level. Outcomes of interest were (i) Visit Intensity (VI_{raw}), capturing overall visitation frequency, and (ii) Repeat Visitor Index (RVI_{raw}), reflecting normalized repeat visitation. A composite proxy $Y_{composite}$ was initially tested but excluded after correlation analysis due to lack of robustness.

2.1 Independent variables

Trip-level parameters were aggregated to destination communes and normalized so that higher values represented more favorable conditions. Indicators included origin diversity, age diversity, weekend–weekday ratio, trip distance metrics, transfers, travel time, and proximity to transit. Cost-related features were inverted for interpretability, and all predictors were clipped at extreme percentiles and min–max scaled to $[0, 1]$.

2.2 Correlation analysis

Spearman rank correlations were computed between all predictors and outcomes to ensure estimability. Both VI_{raw} and RVI_{raw} showed consistent significant associations with all predictors, while $Y_{composite}$ failed the criterion and was excluded.

2.3 Regression modeling

Ordinary Least Squares (OLS) regression was used to validate explanatory power. Results confirmed strong behavioral meaning: for example, origin diversity positively predicted repeat visitation, while age diversity and distance metrics showed contrasting effects on volume vs. loyalty. Transit-related indicators further highlighted the role of accessibility.

2.4 Machine learning extensions

To improve predictive performance and capture nonlinearities, we trained Ridge, Lasso, and Elastic Net regressions, along with Random Forests, Gradient Boosting, Histogram-based Gradient Boosting, XGBoost, and LightGBM. Hyperparameters were optimized via 5-fold cross-validation, and performance assessed with R^2, adjusted R^2, and RMSE. Ensemble stacking was also tested for RVI, combining multiple learners with Ridge as a meta-estimator.

Models were applied across 1,030 communes (94.8% coverage). Predicted scores were discretized into five quantile-based categories (Very Low–Very High) to facilitate comparative mapping and policy interpretation.

3. RESULTS

3.1 Visit Intensity (VI_{raw})

Regression analysis confirmed strong explanatory power of the predictors. OLS achieved $R^2 = 0.858$ (Adj. $R^2 = 0.855$, RMSE=0.852), while Ridge regression performed similarly ($R^2 = 0.857$, RMSE=0.855, $\alpha = 1.15$). Ridge reduced extreme coefficients from OLS, improving stability. Key positive drivers included Origin Diversity, Transit Proximity, Access Transit, and Travel Efficiency, while transfer-related indicators showed weaker effects. Overall, models explained over 85% of VI variance.

3.2 Repeat Visitor Index (RVI_{raw})

Linear models performed modestly (Elastic Net $R^2 = 0.509$, RMSE=0.253), while tree-based ensembles outperformed. Random Forest provided the strongest generalization ($R^2 = 0.942$, CV-$R^2 \approx 0.55$, RMSE=0.087). HistGradientBoosting offered balanced fit ($R^2 = 0.743$, CV-$R^2 \approx 0.58$). XGBoost and LightGBM achieved near-perfect in-sample fits ($R^2 \geq 0.98$) but overfit in validation (CV-$R^2 \approx 0.48$–0.51). Thus, ensemble trees, especially RF and HGB, provided the most reliable performance.

3.3 Composite Outcome ($Y_{Composite}$)

Linear penalized regressions underperformed (Ridge CV-R^2 = 0.23, Lasso=0.57, Elastic Net=0.60). Tree-based methods improved results (RF CV-R^2 = 0.66, GB=0.63, HGB=0.62), with strong in-sample accuracy (train $R^2 > 0.9$). Stacking ensembles further boosted generalization (RF+GB CV-R^2 = 0.68; full stack CV-R^2 = 0.66), striking the best trade-off between accuracy and stability.

3.4 Visualization & Mapping

Predicted scores were mapped using quantile-based classification into five categories (Very Low–Very High). Choropleth maps revealed clear spatial patterns: high-performing communes clustered in central Paris (Île-de-France), while lower scores appeared in peripheral areas. An interactive version enabled users to explore commune-level predictions (e.g., score, class, INSEE code), making results actionable for planners and stakeholders.

Fig. 1. Quantile-based choropleth map of predicted visitation intensities across Paris communes. Darker colors indicate higher accessibility and attractiveness.

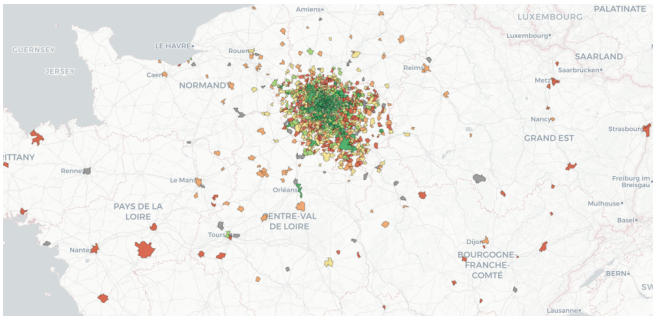

Overall, tree-based ensemble models consistently outperformed linear regressions, and visualization translated predictions into directly usable decision-support tools for urban policy and retail planning.

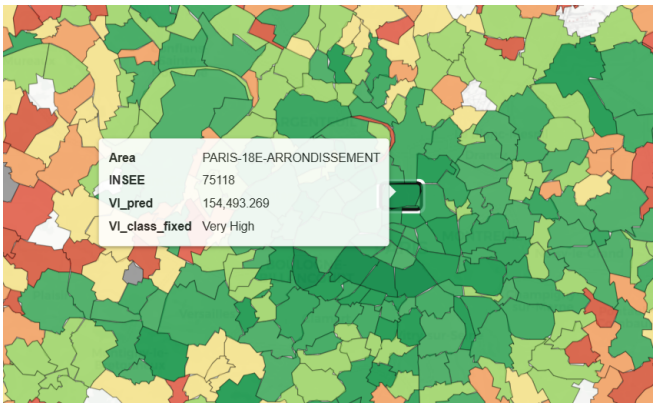

Fig. 2. Interactive classification map (snapshot). Hovering over a commune reveals its name, INSEE code, predicted value, and assigned class (Very Low–Very High), enabling planners to explore local accessibility and visitation patterns.

4. CONCLUSION

This work presented a data-driven framework for optimizing store locations by combining diversity indicators, accessibility measures, and composite indices. Results show that diversity and accessibility are strong predictors of visitation: origin diversity enhances loyalty, age diversity increases footfall but lowers repeat visits, and better accessibility consistently raises visitation.

The framework further introduces practical tools, including an interpretable *Opportunity Score* and an improved reachability measure that accounts for travel time and trip volume. Together, these provide actionable guidance for planners and retailers in evidence-based site selection.

5. FUTURE WORK

Future work will extend the framework by incorporating mobility network analysis and a Reachability Score to capture structural properties of trip flows[11, 12] and spatial accessibility.

4.1 Mobility Network Construction

We plan to build a directed network with zones (INSEE codes) as nodes and trips between zones as weighted edges, including self-loops. Key metrics for each zone will include in-degree (total trips ending at the zone), out-degree (trips originating from the zone), and self-loop ratio (share of internal trips). High self-loop ratios indicate self-contained or residential zones, while low ratios signal well-connected areas, such as transit hubs or business districts. These network features will summarize both local and inter-zone mobility patterns.

4.2 Reachability Score

Edges will be weighted by trip volume and median duration. Shortest-path lengths will be defined as travel duration divided by trip volume, allowing computation of out-closeness centrality. The resulting Reachability Score quantifies how efficiently a zone can access all others, integrating both travel frequency and efficiency. This metric will complement existing indicators, providing a richer understanding of urban accessibility for planning and retail decisions.

REFERENCES

1. W. G. Hansen, “How accessibility shapes land use,” J. Am. Inst. Planners, vol. 25, no. 2, pp. 73–76, 1959.
2. G. Clarke, “Changing methods of location planning for retail companies,” GeoJournal, vol. 45, no. 4, pp. 289–298, 1998.
3. X. Chen, S. Wang, and Y. Li, “Uncovering urban mobility patterns for retail site selection using mobile phone data,” *Cities*, vol. 108, p. 102960, 2021.
4. O. Mumm, M. Murad, V. M. Carlow, “Accessibility Score – Data analytics for urban mobility networks: Braunschweig case,” Environ. Plan. B, vol. 52, no. 3, pp. 594–613, 2024.
5. X. Ye, R. M. Pendyala, and G. Gottardi, “An exploration of the relationship between mode choice and complexity of trip chaining patterns,” *Transportation Research Part B: Methodological*, vol. 41, no. 1, pp. 96–113, 2007.
6. X. Chen, S. Wang, Y. Li, “Uncovering urban mobility patterns for retail site selection,” Cities, vol. 108, p. 102960, 2021.
7. Chasse et al., “The NetMob25 Dataset: Multi-layered view of individual mobility in Greater Paris,” arXiv:2506.05903, 2025.
8. *IRIS GeoPackage Dataset*, retrieved from https://www.data.gouv.fr/fr/datasets/iris-geopackage/.
9. OpenStreetMap. https://www.openstreetmap.org/#map=6/26.81/30.25
10. *GTFS Public Transit Data, Île-de-France Mobilités*, retrieved from https://prim.iledefrance-mobilites.fr/en/jeux-de-donnees/offre-horaires-tc-gtfs-idfm
11. M. Flores-Garrido et al., “Mobility networks in Greater Mexico City,” Sci. Data, 2023. https://www.nature.com/articles/s41597-023-02880-y.pdf
12. Nanni M, Tortosa L, Vicent JF, Yeghikyan G(2020), “Ranking places in attributed temporal urban mobility networks.” *PLOS ONE*, 2020. https://journals.plos.org/plosone/article/file?id=10.1371/journal.pone.0239319&type=printable

When Movement Reveals Poverty: Spatial Mobility as a Marker of Socioeconomic Exclusion

Rghda Salah*[1], Abdullah Allama[2], Islam Aref[3], Ahmed Makboul[4], Abdallah Wael[4], and Noha Gamal†[5]

[1]Faculty of Engineering, Zewail City of Science and Technology, Giza, Egypt
[2]Faculty of Computers and Data Science, Alexandria University, Egypt
[3]Faculty of Computer and Informatics, Zagazig University, Egypt
[4]Faculty of Computer Science, Helwan University, Egypt
[5]School of Information Technology and Computer Science, Nile University, Giza, Egypt

Abstract

This study investigates how realized daily mobility encodes socioeconomic disadvantage by combining high-resolution GPS trajectories, trip records, Points of Interest (POIs), transport-mode signals and neighborhood statistics for the Île-de-France region. We derive compact behavioral indicators — Daily Path Area (DPA), Average Places Visited (APV), a duration-weighted transport mode score, and a Purpose Discretionary Score (Purpose DPS) based on POI-linked purposes — and evaluate their explanatory power for median household income at fine spatial scales (IRIS). Using regression and classification models, we show that mobility extent (DPA) and trip-purpose diversity (Purpose DPS) are consistently the strongest predictors of income, while ensemble learners (Random Forest, XGBoost) outperform linear baselines. Our results demonstrate that mobility-based measures provide a scalable, complementary lens on urban poverty that can inform targeted transport and social policies.

Litrature Review

Mobility is increasingly recognized not merely as physical displacement but as a dimension of social participation and access to opportunities. Prior GPS studies introduced indicators such as Daily Path Area (DPA) and Average Places Visited (APV) to quantify mobility extent and diversity; these measures correlate with employment status, educational attainment and transport access. Complementary work linking POIs to trajectories has shown that the distribution of place categories visited (necessity vs. discretionary) predicts socioeconomic status with high fidelity. Collectively, the literature suggests mobility captures both spatial reach and activity purpose, providing behavioral signals of structural disadvantage that are not fully revealed by income statistics alone.

Methodology

Data and linkage. We use the NetMob25 mobility dataset for the Île-de-France region containing high-frequency GPS traces, self-reported trip-purpose fields, and individual demographic attributes. Trips and stops are spatially aggregated to IRIS zones to link mobility traces with official INSEE socioeconomic indicators (disposable median income, 2021).

Feature construction. We extract four compact indicators:

- **Daily Path Area (DPA)**: for user u on day d, compute the convex hull of origin/destination coordinates and take its area (projected to UTM). User DPA is the day-average:

$$\mathrm{DPA}_u = \frac{1}{D}\sum_{d=1}^{D} \mathrm{Area}(\mathrm{ConvexHull}(\{p_1, \ldots, p_n\}_{u,d})).$$

- **Average Places Visited (APV)**: number of distinct purpose categories per day, averaged across days.

- **Weighted Mode Score**: duration-weighted mean of transport-mode affordability weights (1=walking/bike up to 4=premium modes):

$$\mathrm{ModeScore}_u = \frac{\sum_{t=1}^{T_u} w_{m_t}\, d_t}{\sum_{t=1}^{T_u} d_t}.$$

- **Purpose Discretionary Score (Purpose DPS)**: trips are mapped to discretionary weights in $[0, 1]$ (necessity ≈ 0, discretionary ≈ 1) and averaged per user:

$$\mathrm{PurposeDPS}_u = \frac{1}{T_u}\sum_{t=1}^{T_u} w_{p_t}.$$

Modeling and evaluation. We link user-level features to IRIS median income (continuous) for regression and to income quintiles for classification. Models compared include Linear Regression, Logistic Regression, Random Forest, Gradient Boosting, and XGBoost. Preprocessing includes imputation, scaling,

and quantile discretization for classification tasks. Performance is evaluated with R^2 (regression) and accuracy/precision/recall (classification).

Results

Regression. Linear models explain very little variance. Ensemble regressors, particularly Random Forest and XGBoost, yield substantially higher R^2 values, indicating non-linear relationships between mobility features and income. Feature importance from Random Forest identifies Purpose DPS and Avg DPA (km^2) as the primary predictors.

Classification. In binary classification (High vs Low income IRIS), ensemble classifiers outperform linear baselines. Table 1 summarizes benchmark accuracies.

Table 1: Benchmark classification accuracies

Model	Accuracy
Random Forest	0.783
XGBoost	0.775
Gradient Boosting	0.622
Logistic Regression	0.554

Interpretation. Higher DPA and higher Purpose DPS are associated with higher median incomes, suggesting wealthier respondents exhibit broader spatial horizons and visit a greater share of discretionary destinations. Weighted Mode Score and APV contribute but with smaller effect sizes and more heterogeneity across IRIS.

The following figure 1 illustrate spatial and individual-level patterns from the mobility and socioeconomic data. Figure 4 shows the spatial distribution of users across the Île-de-France region using commune-level representative locations and marker clustering to reduce overplotting. Colored by income category, the map highlights that higher-income respondents are more common in suburban and peri-urban areas, while lower-income groups are concentrated in dense central urban zones. As locations are approximated by commune centroids, dense districts may appear visually dominant despite moderate per-commune populations.

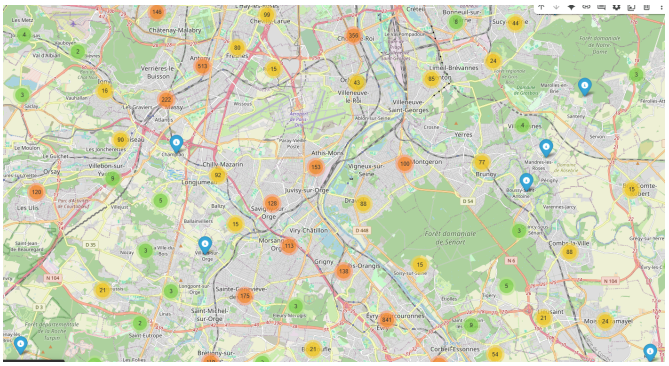

Figure 1: Spatial distribution of users across the Île-de-France region (marker clustering by income category).

Discution

This work shows that compact mobility descriptors derived from GPS and POI linkages are effective behavioral markers of socioeconomic status. Two contributions stand out: (1) combining spatial extent (DPA) with qualitative purpose information (Purpose DPS) yields richer predictive power than either alone; (2) ensemble, non-linear models are necessary to capture complex interactions among mobility features and neighborhood income. Policy implications include the potential for near-real-time, low-cost monitoring of transport-related exclusion and for guiding interventions (e.g., targeted service provision, transit subsidies, or localized infrastructure changes).

Limitations & future work. The study draws on a single regional dataset (Ile-de-France), which limits generalizability. POI-linking provides an operational proxy for purpose but cannot fully replace subjective motivations. Future work should evaluate transferability across diverse urban contexts, incorporate qualitative validation, and explore causal designs to link mobility constraints with outcomes.

Acknowledgements. The authors thank contributors to the NetMob dataset and acknowledge funding/support as detailed in the full report.

References (selected)

1. C. van Dülmen, M. Sžimon, and A. Klärner, "Transport poverty meets car dependency: A GPS tracking study of socially disadvantaged groups in European rural peripheries," *Journal of Transport Geography*, vol. 101, p. 103351, 2022.
2. D. Li and J. Liu, "Uncovering the relationship between point-of-interests-related human mobility and socioeconomic status," *Telematics and Informatics*, vol. 39, pp. 49–63, 2019.

Evaluating Socio-spatial Inequalities in Paris's 15-Minute City Model Using Individual Mobility Data

BOWEN ZHANG,[1*] ZHENGZI ZHOU,[1] YIKANG WANG[1], CHEN ZHONG[1]

[1]*Centre for Advanced Spatial Analysis, University College London, London, WC1E 6BT, United Kingdom*

bowen.zhang.18@ucl.ac.uk

Received XX Month XXXX; revised XX Month, XXXX; accepted XX Month XXXX; posted XX Month XXXX (Doc. ID XXXXX); published XX Month XXXX

1. INTRODUCTION

The "15-Minute City" (15MC) has emerged as a powerful planning vision to restructure urban life around proximity, sustainability, and social equity. It proposes that most daily needs (e.g. grocery shopping, education, healthcare, leisure, and green space) should be reachable within a 15-minute walk, cycle, or public transport trip [1]. Paris has become the flagship case under Mayor Anne Hidalgo, where car-light streets, cycling infrastructure, and local service investment have been promoted as part of the "chrono-urbanism" agenda. While the concept has gained remarkable political and academic traction, it has also attracted scepticism. Critics argue that focusing on proximity risks obscuring inequalities in service quality and distribution, benefitting already advantaged central districts while peripheral areas remain underserved. The concept has even been entangled in conspiracy narratives, complicating public discourse. Despite the attention, rigorous behaviour-based empirical evaluations remain sparse. Most studies rely on buffer-based accessibility to points of interest (POIs) [2], which assume residents use their nearest option and therefore fail to capture actual travel behaviour.

2. Data and Method

This paper addresses that gap by conducting one of the first large-scale behavioural assessments of the 15MC vision using the newly released NetMob25 dataset. The dataset, collected during the EMG 2023 GNSS-based mobility survey in Île-de-France, records ~500 million GPS points for 3,337 participants over seven days, yielding 80,697 validated trips with detailed purposes and modes. Combined with socio-demographic attributes and 254,966 POIs from OpenStreetMap, the dataset provides a solid basis to examine whether and how residents achieve the 15MC ideal in practice.

We evaluate travel behaviour towards the concept of x minutes city by three indicators: (i) the 15-Minute Access Rate, defined as the proportion of service trips completed within 15 minutes; (ii) the Sustainable Mode Share, measuring the proportion of those trips made by walking, cycling, or public transport; and (iii) the Localisation Equity Index (LEI), which normalises accessible POIs by access rate to evaluate whether local amenities offer real choices. We analysed outcomes at both the trip and individual level, disaggregated by service type, and examined social and spatial inequalities. To uncover determinants, we applied SHAP analysis for a Random Forest regression model, linking behavioural outcomes with attributes such as age, education, household composition, car ownership, and mobility impairment. Our findings reveal a nuanced picture.

3. Results

At the trip level, 51.4% of service trips across Île-de-France are completed within 15 minutes of home, suggesting that proximity is partly realised. Yet, at the individual level, achievement is far more uneven (Fig 1): the mean share of local trips per resident is 42.9%, with a standard deviation of 20.4%. One quarter of participants achieve less than 28.6% localisation, while another quarter surpass 57.9%. At the extremes, some residents report virtually no local access, while others achieve near-total localisation (up to 96.6%). These results show that aggregate achievement masks deep disparities between individuals.

Sustainability adds another layer of inequality. Only 33.2% of 15-minute trips are made by walking, cycling, or public transport, revealing a persistent reliance on cars even for short journeys. Person-level results confirm the disparity: the mean sustainable share is 27.4%, but one quarter of residents achieve less than 10%. Many suburban residents remain locked into car dependency, not because distances are long, but because infrastructure and service provision make sustainable choices unviable.

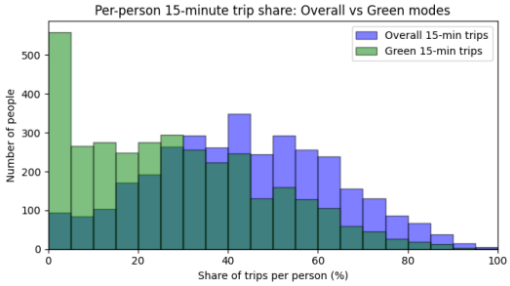

Figure 1. Per-person achievement rate for 15-minute trips

Socio-demographic analysis confirms the stratification of local accessibility by the Random Forest model. Households with children under 10 are more likely to achieve localisation, reflecting the proximity demands of school and childcare. Older adults also report higher localisation, likely due to reduced mobility ranges. Conversely, young adults (18–24) and higher socio-professional groups travel further, bypassing local options for variety or quality. Car ownership is paradoxical: while it correlates positively with achievement (as households with cars can access nearby services more flexibly), it also sustains car dependency. Permutation importance, applied on the test set to mitigate overfitting, refines this picture (Fig. 2). The most influential attributes include mobility impairment status (PMR), sex, and educational attainment (diploma level), followed by household composition variables such as living alone, flat-sharing, or being in a couple with or without children. These results suggest that both resource-based factors (cars, subscriptions) and socio-demographic characteristics (age, education, household type) play important roles in shaping local accessibility outcomes.

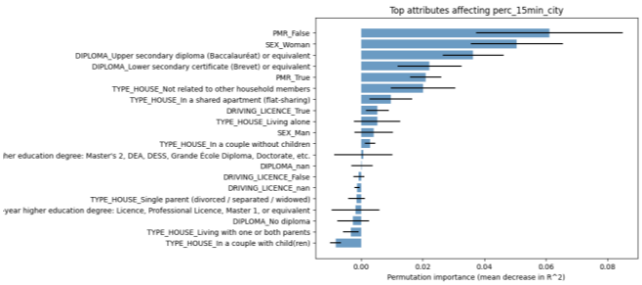

Figure 2. Top attributes influencing the 15-minute city achievement rate

Spatial inequalities are perhaps the most striking. Figure 3 provides a spatial representation of sustainable mode usage for 15-minute trips across the Île-de-France region. It clearly depicts higher percentages of sustainable mode usage in central Paris, decreasing significantly towards peripheral and suburban areas. This visualisation underscores the spatial inequities in sustainable mobility options and reinforces the need for targeted infrastructure development and policy interventions to extend the 15-minute city benefits equitably across the entire metropolitan region.

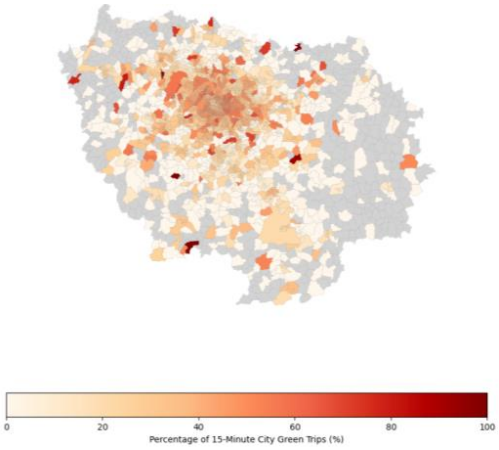

Figure. 3 the percentage distribution of 15-minute trips by area in the Île-de-France region

Meanwhile, linking the POI dataset enables us to compare whether areas with similar levels of 15-minute access rate share diversities of service. The Localisation Equity Index (LEI) shows the Central Paris (département 75) offers 20.9 times more POIs than the outer suburbs (Grande Couronne). The inner ring (Petite Couronne) has ~4 times more POIs than the outer ring. Yet, surprisingly, behavioural achievement rates in some suburban zones are comparable to central ones. This situation, a similar achievement with radically different underlying opportunities, is revealed by the Service Diversity Index. Residents in outer suburbs may complete many service trips within 15 minutes, but often to a narrow set of local providers. This limits choice, flexibility, and resilience, whereas central Parisians achieve localisation in environments rich with alternatives.

4. Conclusion

This study demonstrates that while the 15-Minute City vision is partly realised in Île-de-France, its benefits are unevenly distributed. Proximity alone does not guarantee sustainability or equity: Socio-demographic differences further shape outcomes, with age, household structure, and resources strongly influencing local accessibility. Spatially, many residents complete short trips by car, and suburban areas achieve localisation with far fewer service options than central Paris. These results suggest that achievement rates must be interpreted alongside transport mode and service diversity to avoid masking inequality.

References

1. Moreno, C., et al., *Introducing the "15-Minute City": Sustainability, Resilience and Place Identity in Future Post-Pandemic Cities.* Smart Cities, 2021. **4**(1): p. 93-111.
2. Papadopoulos, E., A. Sdoukopoulos, and I. Politis, *Measuring compliance with the 15-minute city concept: State-of-the-art, major components and further requirements.* Sustainable Cities and Society, 2023. **99**: p. 104875.
3. Chasse, A., et al., *The NetMob25 Dataset: A High-resolution Multi-layered View of Individual Mobility in Greater Paris Region.* arXiv preprint arXiv:2506.05903, 2025.

Paris as a 15-Minute City: An Explainable AI Perspective

ANDRÁS J. MOLNÁR[1], CSABA I. SIDLÓ[1,*], RITA RÓNAI[1], AND DOMONKOS RÓZSAY[1]

[1] *Institute for Computer Science and Control, Hungarian Research Network, Hungary*
[*] *sidlo@sztaki.hun-ren.hu*

1. INTRODUCTION AND RELATED WORK

The need for climate action, better urban livability, and sustainable mobility has revived interest in concepts reimagining modern cities. The 15-minute city, a concept emphasizing human-scale urbanism where essential services are within a 15-minute walk or bike ride, promotes decentralization, strong communities, and reduced reliance on private motorized transport [1–4] while advancing equity, sustainability, and well-being [5].

Diverse approaches operationalize the 15-minute city: Knap et al. introduce a composite cycling accessibility index in Utrecht [6], while Abbiasov et al. use mobility datasets to measure disparities in global cities [7]. Paris is a flagship case, with Thaury et al. analyzing its accessibility patterns [8].

The 15-minute city faces criticism and implementation challenges, including social surveillance, restrictions on individual mobility, and feasibility doubts in car-oriented or segregated areas [9, 10]. Without inclusive planning, the model could reinforce existing inequalities [11–13].

We join the NetMob dataset with scoidemographic and map data, and extend analyzing the 15-minute city concept with advanced explainable AI (XAI) methods, targeting causal relationships. We provide examples of extracted patterns, and establish a robust framework for continuous evaluation.

2. DATA AND ANALYTICAL METHODOLOGY

Our analysis integrates the NetMob 2025 trajectories and **demographic data** with external sources, including 2020 INSEE demographics [14] and October 2022 OpenStreetMap (OSM) **Point of Interest data** (POI) provided by the Ohsome API [15].

We refined trajectories by applying stop detection [16], defining a stop as remaining within a 100 m radius for at least 15 minutes. This allowed us to segment trips into meaningful sub-trajectories for more accurate analysis. To quantify '15-minuteness,' we created a POI availability indicator by counting key POI categories (education, healthcare, municipal, finance/telecom, commercial, and accessibility) within a 1.2 km walking and a 4 km biking radius of start locations. To mitigate the issue of multiple OSM tags representing a single facility, we applied a logarithmic transformation to the raw POI counts.

We applied **correlation analysis** first to identify relationships among POI availability, trip characteristics, and demographics. Categorical demographic features were binarized and all features were min-max normalized. We then employed **XAI**, training gradient-boosted trees, for their strong predictive performance and interpretability via frameworks like SHAP [17]. The logic of these models provides actionable insights for urban policy. Finally, we extended our analysis by calculating asymmetric contributions of features to explore **causal relationships**, similarly to Kelen et al. [18].

3. RESULTS AND DISCUSSION

By encompassing approaches from statistical analysis to advanced explainable AI, we uncovered several noteworthy patterns, which we showcase through selected examples.

The **X-minute mobility** concept flexibly defines a "15-minute city" across modes and time frames. Behavior shifted around 5 km/h (walking) and 15 km/h (biking), supporting these as suitable thresholds. The common 15-minute frame lies in the most variable region of travel, justifying its use as a definition.

Correlations among POI categories (Figure 1) were strong, supporting a unified "15-minuteness" indicator. Mobility patterns showed a car–public transport divide, with leisure travel increasing with POI proximity mainly for walkers and bikers. Single and higher-educated residents correlated positively with POI access, while families relied more on cars.

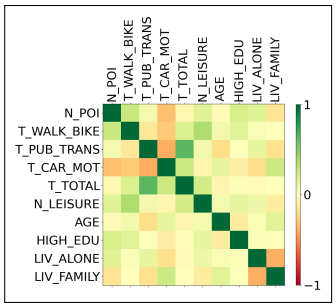

Fig. 1. Correlations of POIs, travel times, number of leisure-aimed trips, age, higher education and living arrangement

Trip duration is a key factor of urban livability: XGBoost models predicting trip segment duration were not highly accurate (R^2 0.21–0.35), but reveal significant factors, with the *purpose of the trip* as the most important. Longer trips were also predicted for individuals with greater home-to-work distances and for those with public transport subscriptions. *15-minute POI features* were influential. *Sociodemographic features* also played an important role; e.g. households with children were associated with shorter trip durations. Or, the disposable income ratio suggests that greater income inequality associate with shorter trip durations.

What Makes a Trip "15 Minutes"? Models predicting if a trip segment will be "15-minutes" were quite accurate (AUC 0.75 - 0.85), with a consistent set of influential features across configurations. The *purpose of a trip* is a critical, as is the distance between home and workplace. Interestingly, the number of young children in a household correlates with an increase in 15-minute segments. For bicycle-based trips, a user's total number of trips is a key factor. The overall number of POIs relevant to the 15-minute city concept is also a major predictor of a trip's duration.

Predicting **transport modes of 15-minute trips** also has high accuracy (AUC 0.7–0.85). We highlight some observations: For trips within a 15-minute walking distance, the *purpose of the trip* is the most significant factor. The presence of *educational POIs* at the starting point strongly increases the likelihood of walking. Additionally, *socio-demographic factors* like having a higher intellectual profession or income make walking more likely, while trips for work or business are less often on foot, in contrast to trips for purchasing goods like lunch.

Private car use for short trips - which is a widely debated

topic with environmental, economic, and social impacts - serves as our primary example to demonstrate XAI methods. Models predicting car use were stable and accurate (AUC 0.8–0.85). Figure 2 summarizes which factors correlate with the use of cars.

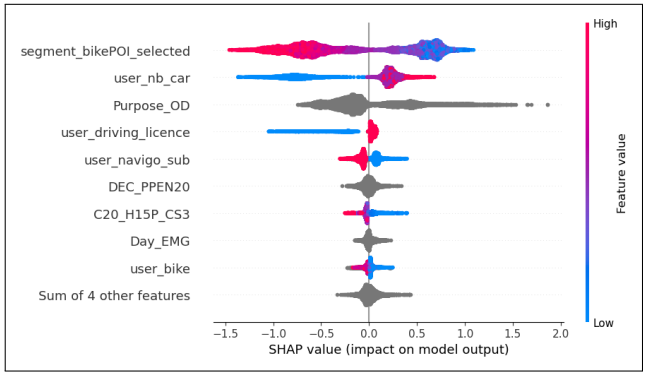

Fig. 2. Impact of the most important features of a model predicting if cars were used for a potential 15-minute bike trip

POI proximity is a crucial factor, decreasing the likelihood of car use. Owning a car or a driver's license strongly predict car use for these short trips. Trip purpose is also highly influential: business and purchase-related trips are especially noteworthy.

Trees within the XGBoost models can provide actionable insights. Figure 3 shows an example where high density of POIs is associated with reduced car use. However, in areas with few POIs, having a public transport subscription significantly decreases car use, suggesting that it can serve as a substitute.

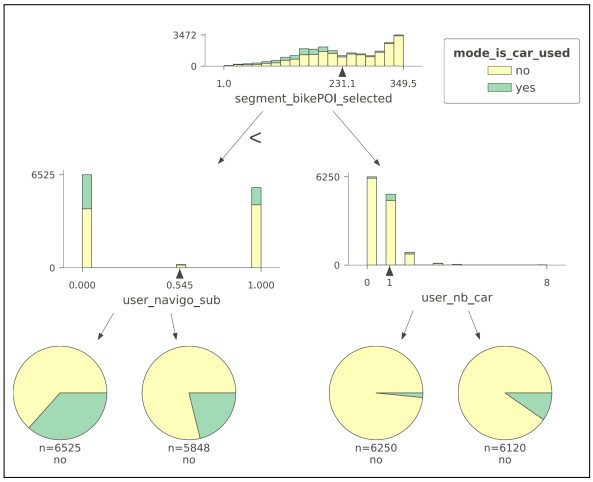

Fig. 3. An influential decision tree when predicting car use for trip segments not longer and faster than a 15-minute bike trip.

Asymmetric SHAP allows us to define and evaluate chains of causal relationships between features, determining if the effect of one feature on the target is, in fact, mediated by another. Figure 4 shows the most affected features when car ownership-related features precede socio-demographics of users: "Ordering A" corresponds to this case, and "Ordering B" is the opposite. The role of owning a car is masked by the socio-demographic properties, it can be considered as an indirect effect. Interestingly the purpose of the trip, and the number of POIs are affected the most, which indicates that sociodemographic features are a key factor, indirectly influencing transport mode selection through these features.

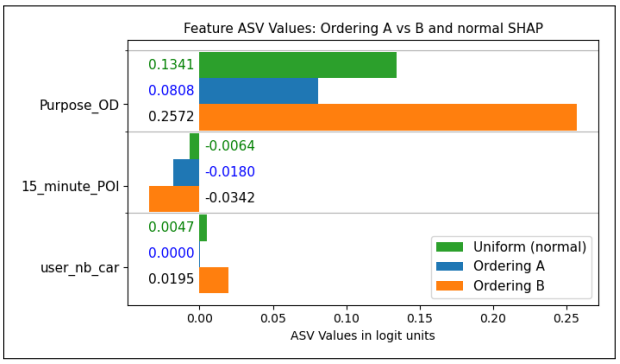

Fig. 4. Evaluating two asymmetric SHAP orderings of the input features for a model predicting private car use for less than 15-minute slow trips

4. CONCLUSION AND FUTURE WORK

Our study uses the NetMob dataset, enriched with geographic and sociodemographic data, to research the 15-minute city concept using advanced analytics. We demonstrate that empirical evidence supports key assumptions of the 15-minute concept, revealing a strong correlation between POI proximity and trip characteristics. Our methodology provides actionable insights that can be directly translated into urban policies. Our current findings already highlight significant insights related to the 15-minute concept, which we plan to extend as a future work. To achieve this, we seek interdisciplinary collaborations with experts in urban planning and sociology.

REFERENCES

1. C. Moreno, Z. Allam, D. Chabaud, *et al.*, Smart Cities **4**, 93 (2021).
2. M. Weng, N. Ding, J. Li, *et al.*, J. Transp. & Health **13**, 259 (2019).
3. G. Pozoukidou and Z. Chatziyiannaki, Sustainability **13** (2021).
4. L. Abdelfattah, D. Deponte, and G. Fossa, Transp. Res. Procedia **60**, 330 (2022).
5. Driving Urban Transitions Partnership, "15-minute city mapping report," https://dutpartnership.eu/wp-content/uploads/2024/04/DUT_15-minute-City-Mapping_04-2024.pdf (2024). Accessed: 2025-07-14.
6. E. Knap, M. B. Ulak, K. T. Geurs, *et al.*, J. Urban Mobil. **3**, 100043 (2023).
7. T. Abbiasov, C. Heine, E. L. Glaeser, *et al.*, "The 15-minute city quantified using mobility data," Working Paper 30752, National Bureau of Economic Research (2022).
8. M.-O. Thaury, S. Genet, L. Maurice, *et al.*, Front. big Data **7**, 1354007 (2024).
9. F. Caprotti, C. Duarte, and S. Joss, Cities **155**, 105497 (2024).
10. University of the Built Environment, (2023). Accessed: 2025-07-14.
11. D. Szymańska, E. Szafrańska, and M. Korolko, Bull. Geogr. Socio-economic Ser. **66**, 137 (2024).
12. Z. Allam and S. E. Bibri, Smart Cities **7**, 83 (2023).
13. K. Mouratidis, Cities **153**, 105274 (2024).
14. Insee, "Les naissances en 2021 - tableaux de séries longues," (2022).
15. M. Raifer, R. Troilo, F.-B. Mocnik, and M. Schott, "Oshdb - openstreetmap history data analysis (0.7.2)," (2021).
16. A. Graser, GIForum **1**, 54 (2019).
17. S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems 30,* I. Guyon, U. V. Luxburg, S. Bengio, *et al.*, eds. (Curran Associates, Inc., 2017), pp. 4765–4774.
18. D. M. Kelen, M. Petreczky, P. Kersch, and A. A. Benczúr, "Theoretical evaluation of asymmetric shapley values for root-cause analysis," in *2023 IEEE International Conference on Data Mining (ICDM),* (2023), pp. 210–219.

The privacy bounds of fine-grained human mobility

Sachit Mishra, Abhishek K. Mishra
Inria, Lyon, France
{sachit.mishra, abhishek.mishra}@inria.fr

***Abstract*—This paper revisits the concept of user uniqueness in human mobility, extending foundational work based on sparse CDR datasets to modern, high-resolution GNSS traces. Using a rich dataset collected in Île-de-France for NetMob 2025, we evaluate how user re-identifiability evolves as a function of spatial and temporal sampling. Our analysis reveals that while spatial-only knowledge leads to a gradual increase in uniqueness, combining even a small number of spatio-temporal observations results in near-perfect re-identification. These results highlight the increasing privacy risks posed by high-frequency mobility data and call for renewed attention to anonymization strategies and privacy-aware data releases.**

***Index Terms*—Mobility data, re-identification, uniqueness**

I. Introduction

The foundational work "Unique in the Crowd" by de Montjoye *et al.* [1] demonstrated that human mobility traces, even when sampled coarsely, can lead to high re-identification risks. Using Call Detail Records (CDRs), the authors showed that knowing just four spatio-temporal points was sufficient to identify 95% of users in a mobile phone dataset uniquely. While impactful, this analysis was conducted on CDRs, which offer limited spatial and temporal resolution.

Several follow-up studies have echoed these concerns, extending the analysis to other datasets [2], [3], [4]. Collectively, this line of research suggests that even limited adversarial knowledge may compromise user anonymity in large-scale mobility data. However, these prior works mostly focus on data with limited spatial resolution or simulated user behavior, leaving open the question of how uniqueness manifests in dense, real-world GNSS trajectories.

In this paper, we revisit the question of user uniqueness using a fundamentally different data source: high-resolution GNSS traces. Unlike CDRs, which capture discrete and sparse communication events, GNSS data records continuous and fine-grained geolocation trajectories. This raises important questions about how user identifiability behaves under modern sensing conditions, particularly when adversaries have access to a small subset of spatial and temporal observations.

We use the `NetMob 2025` [1] dataset, a rich and multimodal collection of GPS trajectories, validated trips, and demographic metadata from over 3,300 individuals in the Île-de-France region. We analyze how user uniqueness varies with the number of observed locations, the addition of temporal information, and the level of granularity at which space and time are represented. Our study provides a detailed view of how quickly anonymity erodes when only a handful of data points are known, offering new insights into privacy risks in real-world high-resolution mobility systems.

[1]https://netmob.org/www25/datachallenge

II. Methodology

Spatial Preprocessing. Each GPS coordinate is first mapped to its corresponding IRIS unit, the smallest administrative division in France used for population statistics. This mapping is performed using shapefiles from INSEE and spatial joins via GeoPandas. Only GNSS points located within Île-de-France are retained. This results in a trajectory per user, represented as a time-ordered sequence of IRIS codes.

Temporal Alignment. To incorporate temporal granularity, each GPS point is also assigned to a fixed time bin. We experiment with intervals varying from minutes, hourly, and daily granularity throughout the dataset duration. Each observation is therefore encoded as a (zone, time-bin) pair. This representation enables us to analyze both spatial-only and spatio-temporal uniqueness.

Uniqueness Estimation. We simulate adversarial knowledge by randomly sampling k mobility observations from each user's trajectory. Each observation is either a spatial location (i.e., IRIS zone) or a spatio-temporal pair consisting of a zone and a time bin. Let $\mathcal{D} = \{u_1, u_2, \ldots, u_N\}$ denote the set of all users in the dataset. For each user u_i, we define their sampled signature of size k as:

$$S_i^{(k)} = \{(z_{i,1}, t_{i,1}), \ldots, (z_{i,k}, t_{i,k})\}$$

In the spatial-only case, we omit the temporal component: $S_i^{(k)} = \{z_{i,1}, \ldots, z_{i,k}\}$. A user u_i is considered *unique* if there is no other user $u_j \in \mathcal{D}$, $j \neq i$, such that $S_j^{(k)} = S_i^{(k)}$.

The uniqueness score $U(k)$ for a given k is defined as the proportion of users whose signature is unique:

$$U(k) = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}\left[\forall j \neq i, S_j^{(k)} \neq S_i^{(k)}\right]$$

where $\mathbb{1}[\cdot]$ is the indicator function, returning 1 if the condition is true and 0 otherwise.

To ensure robustness, the experiment is repeated across multiple random draws of $S_i^{(k)}$ for each user, and we report the average uniqueness across trials. We vary k from 2 to 10 to assess how the amount of adversarial knowledge impacts re-identifiability under both spatial and spatio-temporal sampling assumptions.

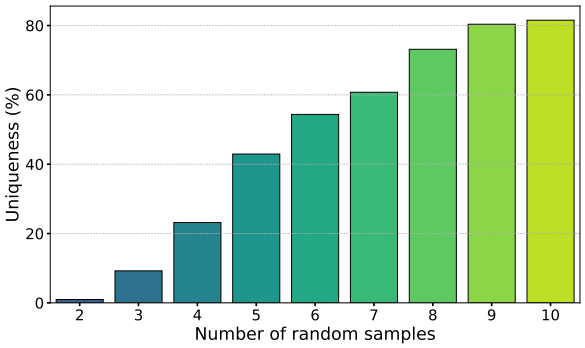

Fig. 1: User uniqueness as a function of number of randomly selected spatial zones (k).

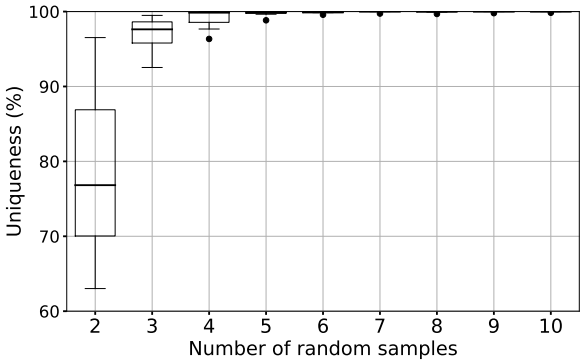

Fig. 2: Distribution of user uniqueness across random draws of $(\text{zone}, \text{time-bin})$ pairs. Each box summarizes varying time bins for each value of k.

III. Results

In the spatial-only setting shown in Figure 1, we observe a progressive but expectedly slow increase in user uniqueness as the number of sampled zones k increases. For $k = 5$, around 42% of users are uniquely identifiable. This value grows steadily with each additional known location, surpassing 80% for $k = 9$. These results indicate that even a modest number of location points, when known with IRIS-level granularity, is sufficient to uniquely re-identify the majority of users, though the knowledge of the temporal dimension is likely to significantly aid the adversary.

In the complete spatio-temporal setting (see Figure 2), where both location and time-bin are considered, uniqueness is very high even for small k values and exhibits variability when changing time bins. The median uniqueness for $k = 2$ is just above 75%, and it rises to over 97% by $k = 3$. Even with two random samples, with a time-bin of 5 minutes, more than 95% of users are unique. For greater values of k, i.e., $k \geq 5$, each user is perfectly identifiable even at the IRIS-level granularity.

Looking forward, we will extend this work along both spatial and temporal axes. On the spatial side, we plan to recompute uniqueness over a hierarchy of grid resolutions—moving from IRIS polygons down to regular cells of 1 km × 1 km, 100 m × 100 m, 10 m × 10 m, and as fine as 5 m × 5 m—to quantify how anonymity erodes as granularity sharpens. In parallel, we will test increasingly fine and regular temporal bins. Finally, leveraging the dataset's rich metadata, we will stratify users by age group, residential context, and mobility intensity to identify cohorts that are disproportionately vulnerable to re-identification. These extensions will provide a more nuanced map of the privacy–utility trade-off and guide the design of context-aware mitigation strategies.

IV. Discussion

Our analysis reveals that the privacy assumptions often held in earlier mobility studies, particularly those based on coarse-grained CDRs, do not hold when applied to high-resolution GNSS traces. When time is taken into account, even just two spatio-temporal samples are sufficient to uniquely identify the vast majority of users. While spatial-only knowledge leads to slower growth in uniqueness, the addition of even moderately fine-grained time bins (e.g., 5 minutes) accelerates re-identification dramatically.

These results underscore the growing fragility of anonymity in modern mobility datasets. As location data becomes more granular and frequent, traditional de-identification techniques, such as spatial aggregation or hourly binning, offer limited protection. More rigorous mechanisms are needed, including formal privacy frameworks (e.g., local differential privacy), synthetic data generation, or controlled access policies that weigh utility against individual risk. Our findings also reinforce the need to embed uniqueness assessments into the design and evaluation process of any mobility data-sharing initiative.

V. Conclusion

This paper revisits user uniqueness in the context of high-resolution GNSS mobility traces, leveraging one of the most detailed real-world datasets collected at urban scale. By varying both spatial and temporal sampling parameters, we quantify how adversarial knowledge impacts re-identifiability. Our findings demonstrate that while spatial-only uniqueness increases gradually with the number of known locations, incorporating even a small number of time-stamped observations leads to near-complete user identifiability, often with just two or three samples.

References

[1] Y.-A. De Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel, "Unique in the crowd: The privacy bounds of human mobility," *Scientific reports*, vol. 3, no. 1, p. 1376, 2013.

[2] D. Malandrino, A. Petta, V. Scarano, L. Serra, R. Spinelli, and B. Krishnamurthy, "Privacy awareness about information leakage: Who knows what about me?," in *Proceedings of the 12th ACM workshop on Workshop on privacy in the electronic society*, pp. 279–284, 2013.

[3] Y. Cao, Y. Xiao, L. Xiong, L. Bai, and M. Yoshikawa, "Protecting spatiotemporal event privacy in continuous location-based services," *IEEE Transactions on Knowledge and Data Engineering*, vol. 33, no. 8, pp. 3141–3154, 2019.

[4] B. Lubarsky, "Re-identification of "anonymized data"," *Georgetown Law Technology Review. Available online: https://www. georgetownlawtechreview. org/re-identification-of-anonymized-data/GLTR-04-2017 (accessed on 10 September 2021)*, 2010.

Advanced Crime-Mobility Relationship Modeling: A Multi-Task Learning Approach for Urban Safety Prediction

AHMED YAHIA[1], REEM ALI[1], SARAH AHMED[1], AND NOHA GAMAL ELDIN SAAD[2,]

[1] Department of Computer Science, Nile University, Giza, Egypt, Emails: {a.ashraf2136, r.khaled2365, s.ahmed2397}@nu.edu.eg
[2] School of Information Technology and Computer Science (ITCS), Center of Informatics Science (CIS), Nile University, Egypt
[1] NGamal@nu.edu.eg

1. INTRODUCTION

Urban transportation decisions reflect complex interactions between mobility needs, temporal constraints, and safety concerns. While traditional mode choice models emphasize travel time, cost, and convenience [1], they often underestimate how crime risk shapes travel behavior, particularly for vulnerable populations such as women and caregivers [2, 3]. Crime and mobility exhibit bidirectional relationships through routine activity theory and crime pattern theory, where offender–target encounters and activity nodes produce spatio-temporal crime concentrations that both influence and are influenced by mobility flows [4, 5].

Recent empirical work demonstrates that integrating human mobility flows enhances short-term crime forecasting and that jointly modeling related urban tasks improves predictive performance [6]. To capture these interdependencies, we propose a unified multi-task learning (MTL) framework that simultaneously models five interconnected tasks: crime count prediction, crime risk classification, route safety assessment, mode choice prediction, and individual vulnerability estimation [7, 8].

Our framework is validated using the NetMob Challenge 2025 dataset (80,697 trips across 600 French areas) and French national crime statistics, harmonized via official postal/INSEE codes [9–11]. By jointly leveraging geographic, temporal, demographic, and transportation features, the model captures crime–mobility interdependencies and yields interpretable, policy-relevant insights for urban safety and equitable transportation planning.

2. DATA SOURCES AND METHODOLOGY

We combined large-scale mobility and crime datasets to analyze the interaction between urban travel behavior and safety outcomes. The NetMob Challenge 2025 dataset [9] includes 80,697 trips recorded across France, capturing origin–destination coordinates, transportation modes (seven categories), trip purposes, and detailed temporal attributes. Each trip is linked to demographic profiles containing information on age, gender, occupation, and household composition, thereby enabling the integration of mobility and socio-demographic characteristics.

To incorporate safety information, we merged official French national crime statistics for the years 2023–2024, which report monthly incident counts for 600 administrative areas [10]. Spatial harmonization was achieved using the official postal code database [11], resulting in a 92.4% alignment between trip origins and destinations and their corresponding crime reporting areas.

The preprocessing pipeline involved multiple steps to ensure data integrity. Missing demographic values were imputed, implausible trips such as those with speeds exceeding 200 km/h were excluded, and mobility records were temporally aligned with crime data by month and year. After cleaning, the dataset was reduced to 74,564 validated trips with complete feature coverage suitable for modeling.

Feature engineering expanded the 29 raw variables into 83 derived predictors organized into four main domains. Geographic features captured zone identifiers, spatial clustering patterns, and neighborhood crime densities. Temporal features included cyclical encodings of hours, days, and months, as well as holiday and strike indicators. Individual-level predictors represented demographics, household structure, and vulnerability factors. Transportation-related features reflected mode type, trip complexity, transit availability, and mode-specific exposure to crime.

To reduce dimensionality and ensure generalizability across tasks, recursive feature elimination was applied, retaining only predictors that contributed consistently to performance in the multi-task learning framework.

3. MACHINE LEARNING INTEGRATION

To capture the multidimensional relationship between mobility and crime, we developed a multi-task learning (MTL) framework that simultaneously models five outcomes: crime count regression, crime risk classification, route safety assessment, mode choice prediction, and individual vulnerability estimation. This design reflects the interdependence of urban safety and mobility behavior, ensuring that insights from one task inform the predictions of others [7, 8].

Each task was implemented using Random Forest ensembles [12], selected for their robustness with mixed feature types, ability to capture nonlinear effects, and interpretability through feature importance analysis. Shared feature representations were learned across tasks to improve generalization, while task-specific optimization preserved the nuances of each predictive objective. Hyperparameters, including the number of estimators, maximum tree depth, and minimum sample split, were tuned using stratified 5-fold cross-validation to balance accuracy and overfitting.

For benchmarking, we compared the MTL framework against single-task baselines, including independent Random Forests, logistic regression models, and gradient boosting classifiers. Regression tasks were evaluated using R^2, mean absolute error (MAE), and root mean squared error (RMSE), while classification tasks were assessed using accuracy and macro-averaged F1 scores. These evaluations confirmed that multi-task learning consistently outperformed single-task approaches across both predictive accuracy and interpretability, validating the advantage of a unified framework in modeling urban safety and mobility.

4. FINDINGS

Key findings from our analysis are consistent with criminological theories linking mobility and crime [4, 5] and with recent

evidence showing that incorporating human mobility flows enhances crime prediction [6]. Our results further demonstrate that multi-task learning provides additional gains in accuracy and interpretability across related urban safety outcomes [7, 8].

2. **Spatial Crime Distribution Analysis:** Spatial analysis reveals a highly concentrated geographic distribution of crime risk across French administrative areas. High-risk zones (>20,000 incidents annually) are focused in major urban centers, representing just 8.3% of areas but accounting for 34.7% of total incidents. Medium-risk areas (5,000–20,000 incidents) account for 23.4% of areas and 41.2% of incidents, while low-risk rural areas comprise the majority of geographic space (68.3% of areas) but only 24.1% of incidents.

This extreme spatial concentration validates our emphasis on geographic features in the predictive model and strongly suggests that targeted interventions in high-risk zones could achieve substantial overall crime reduction.

3. **Gendered Safety Perception:** Women show significantly greater safety concerns (58% higher) and avoidance behaviors (27/% more likely) than men, particularly during evening hours. These gender disparities in mobility patterns suggest that improved nighttime infrastructure could substantially enhance women's freedom of movement.

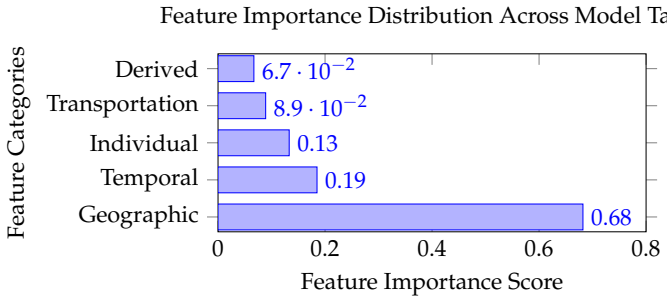

Fig. 1. Feature importance by category. (Representative sample; not all features depicted.)

Table 1. Multi-Task Model Performance with Cross-Validation Results

Task	Type	Metric	Score ± SD	Other Metrics
Crime Count	Reg	R^2	0.847 ± 0.032	MAE: 1,247 ± 89
Crime Risk	Clf	Macro F1	0.906 ± 0.018	Acc: 90.79% ± 1.2%
Mode Choice	Clf	Macro F1	0.984 ± 0.008	Acc: 99.94% ± 0.3%
Route Safety	Clf	Macro F1	0.963 ± 0.015	Acc: 96.29% ± 0.8%
Vulnerability	Reg	R^2	0.798 ± 0.041	MAE: 0.034 ± 0.007

Our multi-task framework achieved strong predictive performance across all tasks: crime count prediction (R^2=0.847), risk classification (90.8% accuracy), mode choice (99.9% accuracy), route safety assessment (96.3% accuracy), and individual vulnerability modeling (R^2=0.798), demonstrating the model's comprehensive effectiveness.

5. CONCLUSION

This research introduces a unified multi-task learning framework for modeling the complex relationships between urban mobility and crime. By jointly predicting crime counts, risk levels, route safety, transportation mode choice, and individual vulnerability, the framework achieves high predictive performance while maintaining interpretability for policy applications.

Key findings include the dominant role of geographic features (68.2% of predictive power), the elevated crime exposure associated with public transportation (52% higher than active modes), and systematic gender differences in safety behavior, with women avoiding high-risk areas 27% more often than men.

These insights translate into actionable strategies such as dynamic allocation of security resources, targeted safety improvements for transit systems, and gender-sensitive infrastructure design. The framework bridges advanced machine learning with evidence-based planning, supporting safer and more equitable urban transportation systems.

Future work should extend this analysis through causal inference, cross-cultural validation, and longitudinal tracking of policy impacts to strengthen both scientific understanding and practical effectiveness.

ACKNOWLEDGMENTS

This research utilizes the NetMob Challenge 2025 dataset [9] and French national crime statistics [10]. We thank the organizers for providing access to the data resources.

REFERENCES

1. M. Ben-Akiva and S. R. Lerman, *Discrete choice analysis: Theory and application to travel demand*, MIT Press, 1985.
2. A. Loukaitou-Sideris, "How to ease women's fear of transportation environments: Case studies and best practices," *Research on Women's Issues in Transportation*, vol. 2, pp. 1–8, 2009.
3. S. Hanson, "Gender and mobility: new approaches for informing sustainability," *Gender, Place & Culture*, vol. 17, no. 1, pp. 5–23, 2010.
4. L. E. Cohen and M. Felson, "Social change and crime rate trends: A routine activity approach," *American Sociological Review*, vol. 44, no. 4, pp. 588–608, 1979.
5. P. J. Brantingham and P. L. Brantingham, *Patterns in crime*, New York: Macmillan, 1984.
6. C. Wu, F. Li, and H. Chen, "Enhancing short-term crime prediction with human mobility flows: Evidence from multiple cities," *EPJ Data Science*, vol. 11, no. 1, pp. 1–22, 2022.
7. R. Caruana, "Multitask learning," *Machine Learning*, vol. 28, no. 1, pp. 41–75, 1997.
8. S. Ruder, "An overview of multi-task learning in deep neural networks," arXiv:1706.05098, 2017.
9. NetMob Challenge Organizers, "NetMob Challenge 2025 Dataset," 2025. Available online: https://netmob.org/www25/datachallenge#participation
10. Ministère de l'Intérieur, "Bases statistiques communale, départementale et régionale de la délinquance enregistrée," 2025. Available online: https://www.data.gouv.fr/datasets/bases-statistiques-communale-departementale-et-regionale-de-la-delinquance-enregistree-par-la-police-et-la-gendarmerie-nationales/
11. La Poste, "Base officielle des codes postaux," 2025. Available online: https://www.data.gouv.fr/datasets/base-officielle-des-codes-postaux/
12. L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

Discovering Functional Urban Zones via Heterogeneous Graph Convolution on High-Resolution Mobility Data

CHUAN LI[1,2,3], VINCENT GAUTHIER[2*], HASSINE MOUNGLA[1,2], AND HU ZHAOBO[2]

[1] LIPADE, Université Paris Cité, France.
[2] SAMOVAR, Telecom SudParis, Institut Polytechnique de Paris, Palaiseau, France.
[3] Sorbonne Université, Paris, France.
[*] vincent.gauthier@telecom-sudparis.eu

1. INTRODUCTION

The complex web of human mobility defines the functional structure and vitality of modern cities. Traditional urban analysis, often relying on static census data or coarse-grained surveys, struggles to capture the dynamic, multi-modal nature of daily travel. The recent availability of high-resolution, longitudinal mobility data from mobile devices presents a new frontier for urban science and computing [1–3].

Most graph-based approaches simplify urban mobility into homogeneous networks, losing crucial information about different modes of transport. However, a trip by subway and a trip by private car represent fundamentally different types of connections between two locations. Ignoring this heterogeneity limits our understanding of the urban fabric.

In this paper, we embrace this complexity. We introduce a novel approach that models the multi-modal transport system of the Greater Paris region as a **heterogeneous graph**, a structure well-suited for capturing multi-modal relationships [4, 5]. Using the rich NetMob25 dataset [3], we construct a graph where census blocks (IRIS) are nodes, and each transport mode (walking, subway, car, etc.) defines a unique type of relationship (edge) between them.

We then deploy a sophisticated Graph Neural Network (GNN) directly on this rich, multi-relational structure. The goal is not merely to predict a single variable, but to learn a powerful, low-dimensional embedding for each IRIS zone that encapsulates its functional "DNA"—a signature derived from its residents' demographics and its connectivity to the rest of the city. This approach builds on prior work that has successfully used human mobility to discover urban functional zones and learn zone embeddings [6–8]. We demonstrate that our method successfully uncovers non-obvious functional similarities between neighborhoods and automatically partitions the city into coherent zones, offering a powerful new tool for urban analytics.

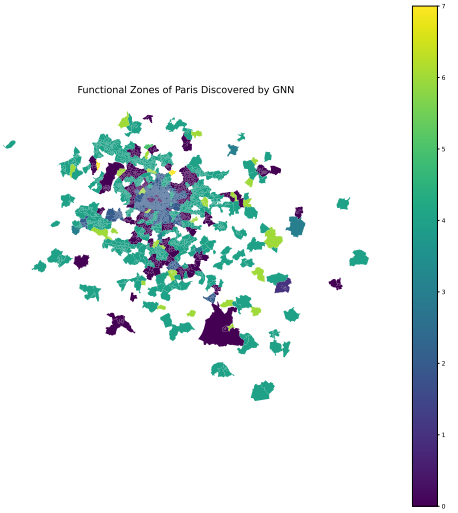

Fig. 1. Map of discovered functional zones in Paris. K-Means clustering applied to the GNN-generated node embeddings automatically partitions the IRIS census blocks into 8 distinct functional communities. Each color represents a different community, revealing coherent spatial structures corresponding to business, residential, and mixed-use areas.

2. METHODS

Dataset and Graph Construction

Our study leverages the NetMob25 dataset, which chronicles the mobility of 3,337 residents in the Île-de-France region over one week [3]. We utilize the *individuals* and *trips* databases to construct our graph, fusing data from multiple domains to create a richer representation [9].

Graph Representation. We model the region as a heterogeneous graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ using the PyTorch Geometric library.

- **Nodes ($\mathcal{V}$):** Each of the **384** unique IRIS census zones observed in our weekly data sample is treated as a node of type 'iris'.
- **Edges ($\mathcal{E}$):** A directed edge exists from node i to node j if at least one trip occurred between them. Each of the **15** observed transport modes ('WALKING', 'BIKE', 'SUBWAY', etc.) is modeled as a distinct edge type, forming relational tuples such as $(\text{iris}, \text{walking}, \text{iris})$.

Feature Engineering. To enrich the graph, we engineered features for both nodes and edges.

- **Node Features:** For each IRIS node, we aggregated demographic data of its residents from the 'individuals' dataset. This resulted in a feature vector for each node containing its residents' average age, average car ownership, and the proportion holding a driving license.
- **Edge Features:** To capture the characteristics of travel, we calculated aggregated statistics for trips. For example, features like average trip duration were added to their corresponding edges, providing the model with more context about the connections between zones.

Heterogeneous GNN Model

Graph-based deep learning has emerged as a powerful paradigm for traffic and mobility analysis [10, 11]. We extend this approach to a multi-modal, heterogeneous context.

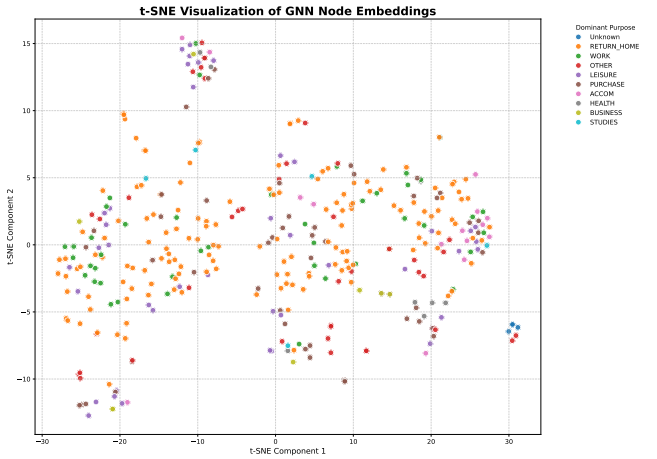

Fig. 2. t-SNE visualization of the learned node embeddings. Each point is an IRIS zone, colored by its dominant purpose. While clusters are not perfectly separated—reflecting the mixed-use nature of urban zones—the model begins to group functionally similar nodes in the embedding space.

Learning Task. We frame our problem as a node classification task. The goal is to predict the dominant function of each IRIS zone, a classic problem in urban computing [12]. We define the zone's function as the most frequent destination purpose ('Purpose_D') of trips ending in that zone (e.g., 'WORK', 'RETURN_HOME', 'LEISURE'). This labeling process resulted in **10** distinct functional classes.

Model Architecture. We implement a GNN using the 'HeteroConv' wrapper from PyTorch Geometric, designed for heterogeneous graphs. The model consists of two 'HeteroConv' layers, each using 'SAGEConv' as the base graph convolution operation. This architecture allows the model to learn distinct message-passing functions for each transport mode before aggregating the results for each node. A dropout layer with a rate of $p = 0.3$ is applied between the convolutional layers for regularization.

Training. The model is trained end-to-end using an Adam optimizer with a learning rate of **0.01** and a standard cross-entropy loss function. To ensure a fair evaluation, we partition the nodes into training (60%), validation (20%), and test (20%) sets.

3. RESULTS AND DISCUSSION

The model's ability to identify urban functional zones was evaluated on an unseen test set, achieving an impressive final accuracy of 37.33% and a macro F1-score of 32.22%(Table 1). These metrics are substantially better than a 10% random baseline, and the small performance gap between the training (79.28% accuracy) and test sets highlights the model's strong generalization capabilities.

Beyond its predictive accuracy, the primary contribution of this work lies in applying the learned GNN embeddings for urban analysis. We used K-Means clustering ($k = 8$) to automatically partition the IRIS zones into functional communities, as mapped in Fig. 1. Post-hoc analysis allows us to assign a functional identity to these discovered zones: (1) Residential Zones , including both strongly residential areas and hybrid residential-business districts; (2) Commercial & Leisure Hubs , separating dedicated retail zones from tourist and entertainment centers; (3) Specialized Areas , such as dedicated recreational parks and unique special-purpose zones (e.g., airports); and (4) Mixed-Use Centers , which encompass both polyfunctional community hubs and balanced suburban centers. This automated partitioning demonstrates the model's ability to generate high-level, structural insights about the city's functional fabric.

In addition to partitioning the city, the embeddings also serve as a powerful tool for similarity queries. For instance, the model identified IRIS zone '75114' (in Paris's 14th arr.) as functionally similar to key inner-suburb hubs like Clichy ('92026', similarity: 0.95) and central Paris districts like the 1st arrondissement ('75101', similarity: 0.94), highlighting its ability to identify relationships that transcend mere geographic proximity.

While this analysis is powerful, it is based on a graph constructed from a one-week time window. Future work could extend this framework to a spatio-temporal GNN to capture dynamic urban evolution. Furthermore, enriching node and edge features with more detailed census or GPS data would likely yield further performance gains.

Table 1. GNN Model Performance on Node Classification.

Data Split	Accuracy	F1-Score (Macro)
Training	0.7928	0.7523
Validation	3514	0.2114
Test	**0.3733**	**0.3222**

CONCLUSION

In this work, we presented an end-to-end pipeline for transforming raw, multi-modal mobility data into actionable urban insights. By representing the city as a feature-rich heterogeneous graph and applying a GNN, we learned powerful functional embeddings for urban zones. These embeddings enabled the automatic discovery of coherent functional communities and the identification of similar neighborhoods, demonstrating a significant step forward in data-driven urban analytics.

REFERENCES

1. Y. Zheng, L. Capra, O. Wolfson, and H. Yang, ACM Trans. on Intell. Syst. Technol. **5**, 38:1 (2014).
2. Y. Zheng, ACM Trans. on Intell. Syst. Technol. (TIST) **6**, 1 (2015).
3. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," (2025).
4. X. Wang, H. Ji, C. Shi, *et al.*, "Heterogeneous graph attention network," in *The World Wide Web Conference, WWW '19,* (ACM, 2019), pp. 2022–2032.
5. C. Zhang, D. jin Song, C. Huang, *et al.*, "Heterogeneous graph neural network," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19,* (ACM, 2019), pp. 793–803.
6. J. Yuan, Y. Zheng, and X. Xie, "Discovering regions of different functions in a city using human mobility and pois," in *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12,* (ACM, 2012), pp. 186–194.
7. Z. Yao, Y. Fu, B. Liu, *et al.*, "Representing urban functions through zone embedding with human mobility patterns," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18,* (ijcai.org, 2018), pp. 3919–3925.
8. X. Zhai, J. Jiang, A. Dejl, *et al.*, Inf. Fusion **120**, 103057 (2025).
9. Y. Zheng, IEEE transactions on big data **1**, 16 (2015).
10. R. Jiang and X. Wu, CoRR **abs/2203.01138** (2022).
11. M. N. Tygesen, J. D. Nielsen, F. M. Rodrigues, and F. C. Pereira, Transp. Res. Part C: Emerg. Technol. **145**, 103946 (2022).
12. N. J. Yuan, Y. Zheng, X. Xie, *et al.*, IEEE Trans. on Knowl. Data Eng. **27**, 712 (2015).

Gendered Mobility Patterns in Île-de-France: Insights from High-Resolution GNSS Traces

DANIELA PERROTTA[1,*], EGOR KOTOV[1], TOM THEILE[1], AND EMILIO ZAGHENI[1]

[1] Max Planck Institute for Demographic Research, Rostock, Germany
[*] corresponding author: perrotta@demogr.mpg.de

1. INTRODUCTION

Mobility is not gender-neutral. Across cities, travel behaviors—how far, when, and where people move—are shaped by social roles, resources, and safety concerns that differ systematically between men and women. Women's mobility often reflects caregiving responsibilities, constrained time use, and security concerns, resulting in shorter, more complex, and less flexible trips [1, 2]. These inequalities are especially pronounced among women in lower-income or informal employment, who face additional barriers to affordable and accessible transportation [3]. Despite its importance as a lens on gender inequality, research on gendered mobility has often been marginal or limited in scale. Traditional surveys suffer from small samples, recall bias, and low resolution [4], while digital traces such as mobile phone or transaction data offer larger coverage but often lack spatial detail or demographic representativeness [5–8].

To address these gaps, we leverage high-resolution GPS traces from the NetMob 2025 Data Challenge to analyze gendered mobility patterns in Île-de-France. We compute individual-level mobility indicators—including unique and core locations, radius of gyration, Shannon entropy, and immobility rates—disaggregated by gender and weighted to ensure population representativeness. This approach provides a fine-grained characterization of spatial mobility inequalities, advancing the study of how gender and sociodemographic factors intersect to shape daily movement.

2. METHODS

We analyzed GNSS traces from over 3,000 participants in the NetMob 2025 dataset [9] to study gendered mobility patterns in Île-de-France at high spatial and temporal resolution. The dataset includes individual- and trip-level weights calibrated to regional sociodemographic profiles, ensuring representativeness. We computed five individual mobility indicators: (i) **unique locations visited** (distinct communes), (ii) **core locations visited** (communes accounting for 80% of activity), (iii) **radius of gyration**, measuring spatial dispersion around the user's center of mass, (iv) **Shannon entropy**, quantifying the diversity of visited locations [10], and (v) **immobility rate**, defined as the share of days with no trips or daily dispersion below 200m. Metrics were computed per user, disaggregated by sex and age group, and weighted to ensure representativeness. We excluded out-of-region traces and duplicates, retaining 308 communes. To control for local composition and context, gender comparisons were made at the commune level. Distributions were summarized with weighted statistics and visualized with bar and box plots. Paired t-tests compared men's and women's mobility within the same commune, accounting for local variation.

Commune-level income data were obtained from the INSEE Filosofi 2019 dataset [11]. Grid-level statistics (200m precision in dense areas) were aggregated to communes via point-in-polygon assignment and averaged to obtain a representative median income per commune.

3. RESULTS

Our findings reveal consistent gender differences in mobility behaviors across communes in Île-de-France.

Unique and Core Locations Women visited fewer distinct locations than men. On average, women visited 23.64 communes (95% CI [22.29, 25.00]) compared to 28.30 for men (95% CI [26.32, 30.28]), a difference of 4.66 (95% CI [2.41, 6.91]; $p < 0.001$), or about 20% fewer. The gap was largest among those aged 25–44, with men visiting 7.87 more communes on average (95% CI [4.25, 11.48]). A similar pattern appeared for core locations, defined as those covering 80% of daily activity. Men averaged 8.46 core locations, versus 6.77 for women—a difference of 1.69 (95% CI [0.95, 2.44]; $p < 0.001$), or a 25% increase. Again, the largest gap was among participants aged 25–44, where men averaged 2.60 more core locations (95% CI [1.40, 3.79]).

Spatial Extent and Diversity of Mobility Women also showed more localized mobility, with a radius of gyration 668 meters shorter on average than men (95% CI [100, 1,237]; $p < 0.05$). Among those aged 25–44, men's radius exceeded women's by 1,258 meters (95% CI [361, 2,154]; Fig. 1A). Mobility diversity was likewise lower among women, who averaged 0.22 points lower Shannon entropy than men (95% CI [0.10, 0.34]; $p < 0.001$), indicating more regular and concentrated routines (Fig.1B). Consistently, women spent more time at their top-ranked locations: frequency rank plots (Fig.2) show higher visit probabilities for their first and second locations, with parity from the third onward.

Immobility Rates Women also experienced higher rates of immobility. Across the study period, the average immobility rate was 1.11 for women (95% CI [1.01, 1.20]) versus 0.86 for men (95% CI [0.76, 0.96]), with a difference of -0.25 (95% CI [-0.38, -0.11]; $p < 0.001$). Notably, the youngest participants (aged 16–24) exhibited the highest immobility rate overall (1.26, 95% CI [1.09, 1.42]), even compared to older adults aged 65 and above (1.09, 95% CI [0.89, 1.29]).

Income and Gender Disparities in Mobility To assess socioeconomic influences, communes were stratified into tertiles of median disposable income (Q1=lowest to Q3=highest). As shown in Fig. 3, average radius of gyration decreased with income: individuals in Q1 had the widest spatial reach, while those in Q3 had the shortest. The gender gap was also largest in Q1, with men traveling 1,410 meters further on average than women (95% CI [87, 2,732]). This gap diminished in higher-income tertiles, suggesting socioeconomic context both structures mobility and amplifies gender inequalities.

4. DISCUSSION

Our findings reveal systematic gendered patterns in spatial behavior, with women consistently showing more localized mobility. They visited fewer unique and core locations, had shorter radii of gyration, lower entropy, and higher immobility than men. These differences were strongest in early adulthood (25–44), a

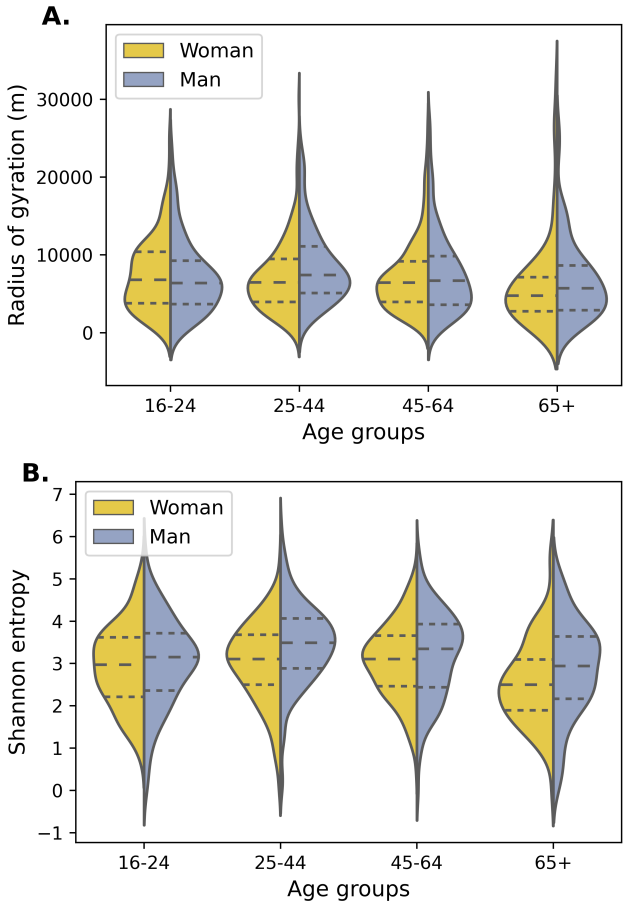

Fig. 1. Distributions of the radius of gyration (A) and Shannon entropy (B), by gender and age.

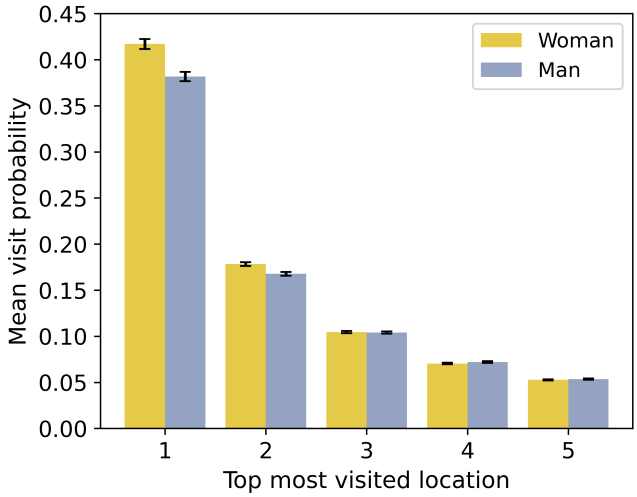

Fig. 2. Distributions of the mean probability of visiting the 5 most frequented locations of each user, by gender. Error bars correspond to the standard deviation of the mean.

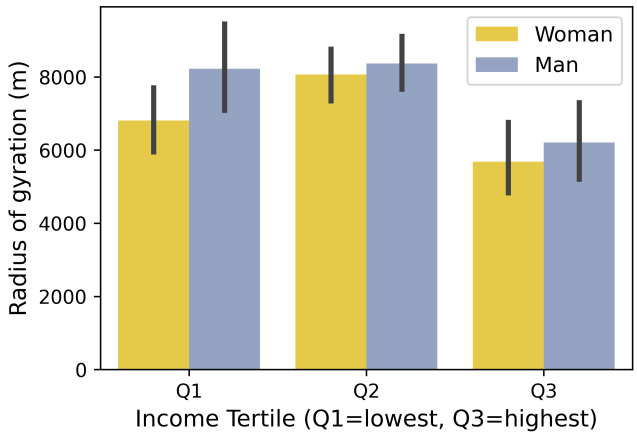

Fig. 3. Distributions of the radius of gyration, by income tertiles (Q1=lowest, Q3=highest) and gender. Error bars correspond to the standard deviation of the mean.

period often linked to caregiving. Overall, results highlight persistent structural inequalities in access, autonomy, and daily mobility.

Income-stratified analysis showed that socioeconomic context shapes gender disparities. Lower-income communes had the widest mobility overall, but the gender gap was largest, with men traveling much farther than women. In contrast, wealthier groups had shorter mobility ranges, and gender differences were smaller. Socioeconomic disadvantage thus both influences overall mobility and amplifies gender inequality.

These findings build on Gauvin et al. [5], who documented gendered mobility in Santiago, Chile. We find similarly lower spatial mobility and more constrained activity spaces among women, reinforcing broader evidence of structural constraints on women's movement. Using high-resolution GPS data, we add finer spatial and contextual detail, showing how gender and socioeconomic factors jointly structure daily movement. The dataset's unique resolution and representativity provide rare insights into gendered mobility and contribute to understanding structural inequalities in everyday life.

REFERENCES

1. A. Loukaitou-Sideris, Secur. journal **27**, 242 (2014).
2. D. Brown, G. McGranahan, and D. Dodman, *Urban informality and building a more inclusive, resilient and green economy* (JSTOR, 2014).
3. C. Tacoli, *Urbanization, gender and urban poverty: paid work and unpaid carework in the city* (Human Settlements Group, International Institute for Environment and . . . , 2012).
4. R. M. Groves, Int. J. Public Opin. Q. **70**, 646 (2006).
5. L. Gauvin, M. Tizzoni, S. Piaggesi, *et al.*, Humanit. Soc. Sci. Commun. **7**, 1 (2020).
6. D. Garcia, Y. Mitike Kassa, A. Cuevas, *et al.*, Proc. National Acad. Sci. **115**, 6958 (2018).
7. M. Lenormand, T. Louail, O. G. Cantú-Ros, *et al.*, Sci. reports **5**, 10075 (2015).
8. E. Jacobs, T. Theile, D. Perrotta, *et al.*, Popul. Dev. Rev. (2024).
9. A. Chasse, A. J. Kouam, A. C. Viana, *et al.*, "The netmob25 dataset: A high-resolution multi-layered view of individual mobility in greater paris region," (2025).
10. C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, Science **327**, 1018 (2010).
11. INSEE, "Revenus, pauvreté et niveau de vie en 2019 – données carroyées (au carreau niveau naturel) – dispositif fichier localisé social et fiscal (filosofi)," (2023). Institut National de la Statistique et des Études Économiques (INSEE), Ministère de l'Économie et des Finances, France.