Transportation Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Special Issue on Machine Learning Methods for Urban Passenger Mobility

Xuan Di, Sean Qian, Carolina Osorio

informs
https://pubsonline.informs.org/journal/trsc
TRANSPORTATION SCIENCE
Vol. 59, No. 4, July–August 2025, pp. iii–vi
ISSN 0041-1655 (print), ISSN 1526-5447 (online)

Special Issue on Machine Learning Methods for Urban Passenger Mobility

Xuan Di,[a,b,*] **Sean Qian,**[c,d] **Carolina Osorio**[e]

[a] Department of Civil Engineering and Engineering Mechanics, Columbia University, New York, New York 10027; [b] Data Science Institute, Columbia University, New York, New York 10027; [c] Department of Civil and Environmental Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213; [d] Heinz College of Information Systems and Public Policy, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213; [e] Department of Decision Sciences, HEC Montréal, Montréal, Québec H3T 2A7, Canada

*Corresponding author

Contact: sharon.di@columbia.edu, https://orcid.org/0000-0003-2925-7697 (XD); seanqian@cmu.edu, https://orcid.org/0000-0001-8716-8989 (SQ); carolina.osorio@hec.ca (CO)

https://doi.org/10.1287/trsc.2025.intro.v59.n4

1. Background

In 2023, we announced this special issue (SI) focused on machine learning (ML), simulation, optimization, and game theory methods for the design of more efficient, sustainable, accessible, and equitable passenger transportation systems. This SI was in conjunction with the Google Research workshop on "Sustainable Urban Mobility: Simulation and Optimization" hosted in Mountainview, California, in June 2023. This workshop brought together experts in the fields of urban transportation, optimization, algorithmic computer science, and ML to discuss how integrating concepts and methods from these fields can help design the next generation of efficient, sustainable, and globally scalable mobility services. The cross-disciplinary conversation from this workshop was fruitful, which fostered the creation of this SI.[1]

We invited scholars to submit research papers that address emerging urban mobility challenges and opportunities with nonconventional data-driven methods. The SI features novel methodological approaches to complex transportation problems arising from the analysis of vehicular or passenger flows in large-scale metropolitan areas, such as city-scale demand modeling, large-scale traffic monitoring and management, and the underlying offline or real-time optimization problems that are high-dimensional and stochastic, and the need for optimization methods that can embed black box traffic models into optimal decision making to assist various stakeholders (e.g., simulation-based optimization, black box optimization, derivative-free optimization). The submissions to this SI consist of approaches including model-informed, data-driven, and hybrid methods that combine ML with domain knowledge, simulations, and decision making.

2. Motivation

Urban mobility has seen disruptive transformations in the past decades. Thanks to advances in sensing, commutation, networking, and computing, mobility services become increasingly on-demand, shared, multimodal, automated, and connected. Particularly, passenger mobility systems are the backbone of urban and regional infrastructure and mobility services, enabling people to access food, jobs, education, healthcare, and social opportunities. These systems encompass passenger trips in various classes (e.g., traditional human-driven vehicles in different classes, connected and automated vehicles) and in various modes (e.g., public transit, ridesharing, microtransit, active transportation). Passenger transportation is shaped by infrastructure, technology, policy, and human behavior. Emerging data from all elements of mobility systems provide unprecedented opportunities and challenges to design passenger mobility systems for economic vitality, environmental sustainability, and social equity. All those transformations call for innovations in the modeling and simulation of mobility systems, including but not limited to high-dimensional, continuous or discrete, offline or real-time, big or small data regime, stochasticity with uncertainties, and simulation-based optimization problems.

ML has been rapidly and widely applied in design, planning, and operation of urban passenger mobility systems. Research directly applying ML techniques includes (1) demand prediction (real-time forecasting of travel demand across modes, regions, and characteristics using multisource, historical, and real-time data, accounting for out-of-distribution generalization of test data); (2) passenger behavior modeling (learning individual or group preferences in driving or routing over various mobility choices to personalize mobility services and predict system-level impacts); and (3) multimodal integration (integrating multimodal passenger mobility from individual automated vehicle technologies to connecting passenger with multiple mobility options, leading to complex, versatile, and resilient systems).

3. Paper Summary

The review of papers for this SI heavily focused on the innovation and methodological novelty of the work presented, as well as its real-world impacts on society and the transportation community. We eventually selected nine papers. The following details how ML is employed and the specific urban mobility challenges addressed.

Real-time forecasting of traffic states, such as counts, speed, and ridership, has gain tremendous attention. Zheng, Choi, and Sun (2025a) introduce an autoregressive framework to improve probabilistic multivariate traffic state forecasting by addressing correlated error structures often ignored in deep learning (DL)-based spatiotemporal models. Instead of assuming independent isotropic errors, the authors model the residual errors of a base DL forecaster using a matrix-variate autoregressive (AR) process and a nonisotropic Gaussian distribution with structured covariance. The method is model agnostic, compatible with any DL-based forecasting models, and improves both accuracy of point estimators and probabilistic traffic state estimators across a few tested generic DL models. The proposed framework learns from residual patterns rather than relying solely on tuning DL models with or on integrating external features fed to the DL model. Thus, it offers a statistically grounded, scalable solution for enhancing traffic state forecasting systems.

In a similar traffic state forecast context, Choi et al. (2025) improve probabilistic traffic state forecasting by capturing time-varying, multimodal, and spatiotemporally correlated errors through a scalable mixture model framework that integrates seamlessly with generic DL models. This paper proposes a dynamic mixture of matrix-variate Gaussian distributions to model forecasting errors over time of day. It addresses computational challenges of high-dimensional covariance by decomposing the full spatiotemporal covariance matrix into the Kronecker product of spatial and temporal components, thus enabling tractable training even for large sensor networks and long prediction horizons. The method acts as a plug-in loss module compatible with existing generic traffic state forecasting architectures based on DL approaches. The proposed ML framework leads to accurate, interpretable, and scalable forecasting for dynamic and complex urban traffic systems.

Zheng et al. (2025b) develop a method for detecting and correcting erratic (irregular, with unknown structure) measurement errors in traffic flow data collected from network-wide sensors. This paper introduces the concept of virtual sensors at intersections to detect violations of the flow conservation law (where inflow does not equal outflow). These binary indicators help identify potential sensor errors without knowing or processing exact error locations. With the information of those flow balance violations, an optimization model is employed to jointly estimate sensor error probabilities and traffic flow. This optimization model is embedded into the training loop of ML models, enabling joint estimation and decision making, avoiding bias, and improving performance over simple two-stage methods, especially when the assumptions of Poisson error processes and linearity do not hold. The paper combines ideas from the fields of ML, traffic flow theory, and optimization, enabling more robust and interpretable error estimation in traffic sensor systems. The approach significantly improves the reliability of large-scale traffic data for monitoring, planning, and forecasting purposes.

Guan, Huang, and Chen (2025) address the challenge of inferring missing trips of certain characteristics in travel demand data due to selection bias in both traditional surveys and modern mobile sensor data (e.g., location-based services). Selection bias arises because users self-select into data sources (e.g., only certain demographics use specific apps), making the data nonrepresentative. This paper proposes a behaviorally informed, likelihood-based framework to correct for self-selective biases in the process of selecting mobile sensors with large volumes of spatiotemporal data. Instead of treating each data set in isolation, the authors leverage multiple biased data sets whose biases are different to mitigate each other's weaknesses. The proposed ML model is tested using both simulated data and the 2017 National Household Travel Survey, which shows that the model robustly recovers 89%–106% of missing trips when behavioral drivers of data generation are properly modeled and recovers nearly 50% of missing trips even when key behavioral factors are omitted. The study bridges data science, ML, and trave behavior analysis and illustrates great potential to mitigate bias when constructing data from multiple sources.

Forecasting multimodal transport demand in the presence of missing data requires models that can jointly handle imputation and prediction without error propagation. Li et al. (2025) propose a graph-guided generative-adversarial network (GIF) that tightly integrates generative adversarial network (GANs) transformer encoders, and a novel graph-based fusion mechanism for cross-modal learning. Unlike sequential approaches, GIF treats imputation and forecasting as a unified task, where the generator reconstructs both past and future demand, the discriminator enforces realism, and an encoder-decoder structure improves training stability. A graph-guided fusion layer enables knowledge sharing across transport modes like bike, taxi, bus, and rail by leveraging historical demand similarity. Experiments on six real-world data sets—including NYC Bike, NYC Taxi, BJ Taxi, and multimodal Australian transit—demonstrate that GIF outperform baselines such as Long short-term memory, Gated Recurrent Unit, Graph Convolutional Network, and GAN-based models under various missing patterns (random, segmental, consecutive) and rates (up to 70%). The study exemplified how integrating cross-task and cross-modal dependencies can significantly improve urban mobility forecasting robustness and accuracy.

Similar to the previous papers, Mo, Xiang, and Di (2025) also focus on spatiotemporal prediction over graphs, but from a different perspective: The focus is on improving the out-of-distribution (OOD) generalization capability of real-world data. This can occur when test data have a different distribution than training data. In this case, existing graph-based ML models (e.g., attention-based spatial-temporal graph convolutional network) might fail. Addressing this problem is critical for numerous real-world applications. An important example is that of human mobility prediction before and after COVID-19, given that mobility patterns experienced a significant distributional shift after the pandemic. There are two principles for OOD problems, namely, invariance existence (i.e., there exist invariant features that consistently relate to the labels across various environments) and environment diversity (i.e., diversifying training environments increases the likelihood that test environments align with training ones). To adhere to the above principles, this paper proposes a diffusion-augmented invariant risk minimization (diffIRM) framework that contains two modules: data augmentation and invariant learning. In the data augmentation process, a causal mask generator identifies causal features, and a graph-based diffusion model acts as an environment augmentor to generate augmented spatiotemporal graph data. In the invariant learning process, an invariance penalty is designed using the augmented data and then serves as a regularizer for training. The real-world validation experiments employ public human mobility data sets, namely, SafeGraph, PeMS04, and PeMS08. The proposed diffIRM outperforms baselines (including empirical risk minimization and invariant risk minimization) with interpretability by discerning invariant features while making predictions. This work combines knowledge in causal inference and ML.

Modeling metro passenger route choice behavior is challenging in large-scale transit systems, particularly when multiple routes between origin-destination (OD) pairs have similar travel times, making individual preferences difficult to infer from aggregate data. Du et al. (2025) present a fully differentiable, end-to-end simulation-based optimization framework for calibrating route choice ratios across the Hong Kong MTR network, which includes more than 20,000 potential route options. The proposed framework integrates a flow-based metro simulation model within a computational graph, allowing gradient-based optimization using iterative backpropagation and automatic differentiation. To reduce computational demands, a matrix-based optimization approach is introduced, which approximates simulator outputs using transition matrices; the two methods are combined in a hybrid algorithm that alternates between full simulation and matrix-based updates. The framework supports calibration under varying levels of model complexity, including a general C-logit model, a fully expressive OD-specific model, and an intermediate model with link-specific utilities. Experimental evaluations using synthetic OD flows and real-world smart card and train loading data show that the method produces consistent and interpretable results across a large parameter space. The framework is designed to be extensible to equilibrium formulations and can be applied to improve operational planning and behavioral inference in congested urban rail networks.

Liao et al. (2025) address the challenge of predicting human-like driving trajectories in complex, interactive traffic. It introduces HiT (human-like trajectory prediction), a behavior-centric model that moves beyond static graphs and rigid labels by incorporating dynamic geometric graphs and centrality-based behavioral descriptors—dynamic degree, closeness, and eigenvector centralities—to capture agent influence in real time. A fuzzy inference system with the q-rung orthopair fuzzy weighted Einstein Bonferroni mean operator computes continuous aggressiveness scores, which structure a hypergraph-based behavior encoder for higher-order behavioral grouping. Combined with an interaction-aware transformer and multimodal decoder, HiT predicts diverse future trajectories in polar coordinates, enhancing adaptability to irregular road layouts. Tested on five data sets (NGSIM, HighD, RounD, ApolloScape, MoCAD++), it outperforms prior methods, especially under aggressive and uncertain behavior. It also remains effective with only 25% training data and fewer parameters, demonstrating strong efficiency, robustness, and scalability.

Traditional transportation network equilibrium models are difficult to calibrate due to unobservable congestion dynamics, demand endogeneity, and behavioral heterogeneity. Liu and Yin (2025) propose an end-to-end learning framework that replaces the bottom-up model assembly with a neural network that jointly learns supply, demand, and equilibrium states directly from multiday traffic data. User equilibrium conditions are enforced via an implicit layer formulated as a variational inequality, and the model is trained using auto-differentiation with either iterated differentiation or inexact implicit differentiation methods. Theoretical contributions include formal bounds on expressivity, generalization, and convergence within a mathematical program with equilibrium constraints formulation. Empirical validation on three synthetic networks—Braess, Sioux Falls, and Chicago Sketch—demonstrates the framework's ability to recover equilibrium states accurately, outperforming ablated variants and emphasizing the importance of forward equilibration. This work bridges ML and network modeling by enabling theory-grounded, data-driven equilibrium inference.

4. Needs for ML-Based Research in Transportation

Because of the widespread and broad application of urban passenger transportation research, ML has been applied in numerous ways. Similar to various ML fields, for example, object detection, in order to move the broad transportation field forward with continuous performance improvement, we urgently call for standard application scenarios, benchmark data sets and methods, and unified test environments, as well as consensus on open-sourcing data, codes, algorithms, and results replicability to all.

Endnote

[1] See https://rsvp.withgoogle.com/events/urban-mobility-workshop/home.

References

Choi S, Saunier N, Zheng VZ, Trépanier M, Sun L (2025) Scalable dynamic mixture model with full covariance for probabilistic traffic forecasting. *Transportation Sci.*, ePub ahead of print March 14, https://doi.org/10.1287/trsc.2024.0547.

Du K, Lee E, Ma Q, Su Z, Zhang S, Lo HK (2025) Modeling metro passenger routing choices with a fully differentiable end-to-end simulation-based optimization (SBO) approach. *Transportation Sci.*, ePub ahead of print February 7, https://doi.org/10.1287/trsc.2024.0557.

Guan X, Huang S, Chen C (2025) Using multiple biased data sets to recover missing trips with a behaviorally informed model. *Transportation Sci.*, ePub ahead of print May 16, https://doi.org/10.1287/trsc.2024.0550.

Li C, Liu W, Ma W, Yang H (2025) Simultaneous multimodal demand imputation and forecasting via graph-guided generative and adversarial network. *Transportation Sci.*, ePub ahead of print May 13, https://doi.org/10.1287/trsc.2023.0326.

Liao H, Li Z, Zhang G, Li K, Xu C (2025) Toward human-like trajectory prediction for autonomous driving: A behavior-centric approach. *Transportation Sci.*, ePub ahead of print May 7, https://doi.org/10.1287/trsc.2023.0366.

Liu Z, Yin Y (2025) End-to-end learning of user equilibrium: Expressivity, generalization, and optimization. *Transportation Sci.*, ePub ahead of print March 13, https://doi.org/10.1287/trsc.2023.0489.

Mo Z, Xiang H, Di X (2025) diffIRM: A diffusion-augmented invariant risk minimization framework for spatiotemporal prediction over graphs. *Transportation Sci.*, ePub ahead of print April 8, https://doi.org/10.1287/trsc.2024.0562.

Zheng VZ, Choi S, Sun L (2025a) Probabilistic traffic forecasting with dynamic regression. *Transportation Sci.*, ePub ahead of print April 8, https://doi.org/10.1287/trsc.2024.0560.

Zheng Z, Wang Z, Fu H, Ma W (2025b) Estimating erratic measurement errors in network-wide traffic flow via virtual balance sensors. *Transportation Sci.*, ePub ahead of print February 24, https://doi.org/10.1287/trsc.2023.0493.