Inverse Learning and Intervention of Transportation Network Equilibrium

Zhichen Liu, PhD, University of Michigan, 2025

Dissertation Committee: Yafeng Yin (Chair), Yu Nie, Neda Masoud, Ruiwei Jiang, and Lei Ying

Introduction

Network equilibrium analysis has been a cornerstone of transportation science and has been used by planning agencies worldwide to forecast congestion, evaluate policies, and guide long-term network planning for more than six decades. Traditionally, equilibrium models are constructed as part of the four-step or activity-based modeling process, i.e., traffic assignment. Network equilibrium models usually rely on survey data, pre-specified behavioral assumptions, and labor-intensive calibration, making them expensive to build, slow to update, and prone to compressing daily variability into an average traffic pattern.

This dissertation establishes *inverse learning of user equilibrium* as a novel framework for constructing network equilibrium models directly from multiday traffic observations. The central premise is that equilibrium models need not be hand-specified through infrequently updated surveys and fixed behavioral assumptions. Rather, they can be learned directly from crowdsourced mobility data as context-dependent models that discover traveler behavior and congestion dynamics (see Figure 1).

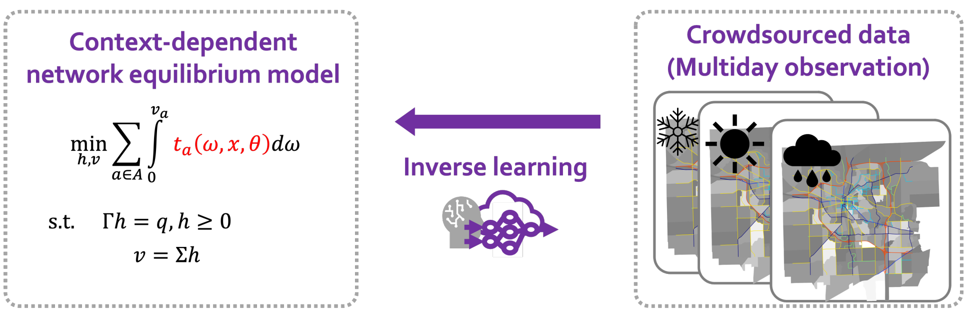

Figure 1: Inverse learning of transportation network equilibrium.

Building on this premise, the dissertation first develops inverse learning as a unified framework, then proposes a family of nonparametric and semi-parametric approaches for learning equilibria from multiday observations, and finally embeds learned contextual uncertainty into robust network design.

1. Inverse Learning as a Unified Framework

Constructing a network equilibrium model requires estimating three key components that are central to planning but not directly observable: travel demand, route-choice cost or utility functions, and link performance functions. For decades, researchers across transportation network modeling, data-driven optimization, inverse game theory, and physics-informed deep learning have studied how to infer such hidden components from data. Relevant efforts include origin–destination matrix estimation, route-choice calibration, and link-performance-function estimation. Yet these existing methods are often developed for specific data settings or modeling objectives, leaving the literature fragmented and making systematic comparison difficult.

This dissertation addresses this gap by establishing *inverse learning* as a unified perspective for constructing network equilibrium models directly from empirical data. The core idea is to formulate the equilibrium model-building process as an *inverse problem*: given a forward problem that describes an equilibrium flow distribution, such as the Beckmann–McGuire–Winsten formulation, an inverse problem is solved to find the parameters or functions of the forward problem that best reproduce observed traffic states (Figure 1). This framework[1] accommodates key modeling choices: (i) how to balance fitting observed data against enforcing equilibrium consistency; (ii) which metrics to use for measuring prediction accuracy and equilibrium deviation; and (iii) whether to use parametric or nonparametric functions to represent behavioral and physical mechanisms.

By organizing diverse calibration and learning methods around these three modeling choices, inverse learning offers a coherent perspective for comparing multi-disciplinary approaches and identifying gaps that isolated studies often obscure. One insight from this perspective is that data resolution changes the mathematical structure of the learning problem. Complete path-flow observations can reduce the otherwise bilevel inverse problem to a convex or even linear program. In practice, however, such observations are rarely available. Most empirical settings provide only aggregate link data, creating modeling and computational challenges that distinguish inverse learning of network equilibrium from standard inverse optimization and inverse game theory.

2. Discovering Equilibria from Multiday Observations

Building on the unified inverse-learning framework, this dissertation next studies the new opportunity created by multiday traffic observations. Repeated observations across days reveal variation in demand, congestion, and context, allowing forward equilibrium components to be learned directly from data rather than being pre-specified based on modelers' beliefs. This section develops two complementary approaches to exploit this opportunity: an expressive nonparametric approach based on implicit deep learning and a scalable semi-parametric approach for limited-sample settings.

2.1 Prioritizing Expressivity: Nonparametric Inverse Learning via Implicit Layers

When expressivity is the priority, the dissertation develops a nonparametric inverse-learning approach that represents unknown equilibrium components using neural networks. These learnable neural components are embedded in a variational inequality that defines the user equilibrium conditions. The variational inequality is then encoded as an *implicit layer* within a deep learning pipeline that maps context features to equilibrium flows (see Figure 2). This nonparametric approach allows the model to learn complex behavioral and congestion relationships directly from observed traffic states without requiring the modeler to commit in advance to a specific route-choice rule or congestion function.

This modeling flexibility comes with computational challenges. Each observed context induces its own equilibrium constraints, so training requires repeatedly solving and differentiating through hundreds or thousands of user equilibria. The dissertation addresses this challenge with two main contributions[2]. First, it develops a topology-aware neural-network architecture that guarantees the existence of an equilibrium solution while accommodating future network-topology changes in what-if analyses. Second, it develops scalable algorithms for both forward and backward propagation (see Figure 2). In forward propagation, operator-splitting and mirror-descent methods are used to solve batches of link- and path-based variational inequalities efficiently by coding the solution trajectory as computational graphs. In backward propagation, iterative and inexact automatic-differentiation-based methods differentiate through the implicit layer. Together, these contributions make large-scale nonparametric equilibrium learning computationally feasible. Extensive numerical experiments on the Braess, Sioux Falls, and

[1] Liu, Zhichen and Yafeng Yin (2025), *"Constructing Transportation Network Equilibrium Models from Empirical Data: An Inverse Learning Perspective"*, *Under revision with Transportation Research Part B*

[2] Liu, Zhichen, Yafeng Yin, Fan Bai, et al. (2023), *"End-to-end learning of user equilibrium with implicit neural networks"*, *Transportation Research Part C*

Chicago networks demonstrate that the nonparametric framework can accurately recover equilibrium models and support what-if analysis.

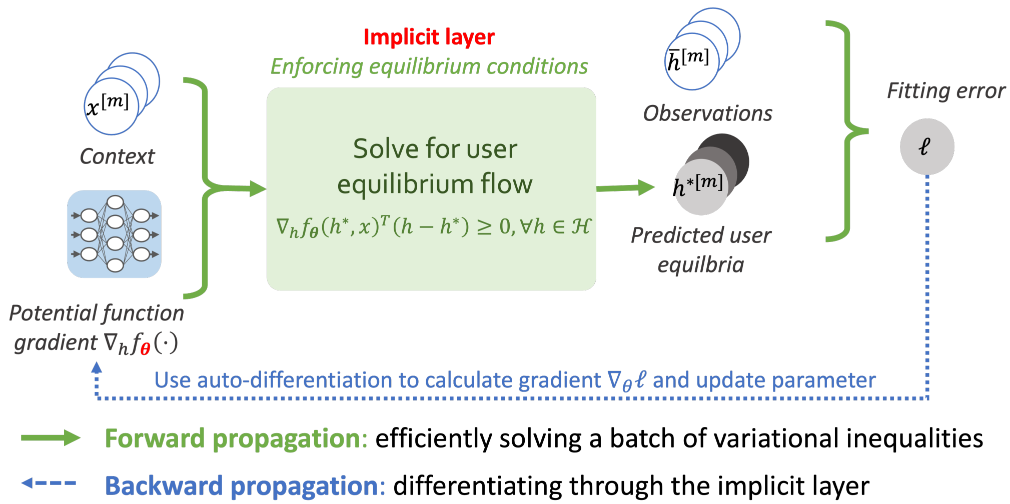

Figure 2: Nonparametric inverse learning via implicit layers. Forward propagation solves context-dependent equilibrium problems for batched samples, and backward propagation differentiates through the implicit layer to update neural network parameters during training.

2.2 Prioritizing Scalability: Semi-parametric Inverse Learning via Multiconvex Optimization

Although the nonparametric approach offers high goodness-of-fit, it is also data- and computation-intensive (see Figure 3). In many practical settings, available datasets contain only a few hundred samples, which may be insufficient to train neural networks reliably. To address this limitation, the dissertation develops a semi-parametric inverse learning framework that prioritizes computational tractability.

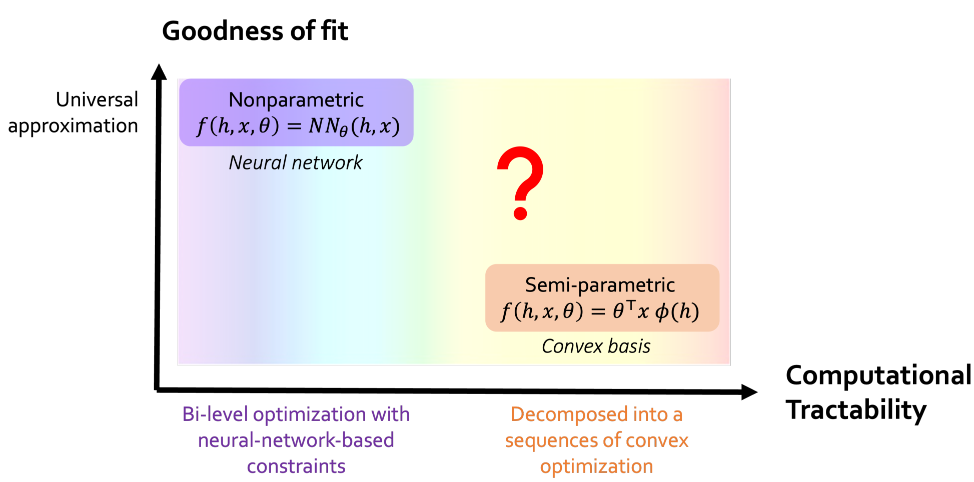

Figure 3: Trade-offs among alternative parameterizations in inverse learning. Nonparametric approaches provide greater expressivity but require solving bilevel optimizations with neural-network-based constraints, whereas semi-parametric models impose structure that improves scalability under limited data.

The core idea is to approximate the unknown equilibrium potential function using a nonnegative combination of *convex bases*. With a sufficiently rich basis family, this structure can capture context-dependent travel behavior while keeping learning tractable: it avoids implicit gradients and decomposes the inverse problem into smaller convex subproblems that can be solved in parallel. The dissertation further establishes statistical consistency under noisy link-flow observations, showing that both prediction quality and parameter recovery improve as more data become available.

3. Choosing Among Parametrizations: Trade-offs Among Expressivity, Generalization, and Optimization

The nonparametric and semi-parametric approaches above are two instances of a broader family of equilibrium models enabled by the inverse learning framework (see Figure 3). This flexibility in parametrization raises a fundamental methodological question: *how should modelers choose among alternative parameterizations?* This dissertation addresses this question by analyzing the feasibility of learning equilibrium states from multiday observations and decomposing learning errors into three sources: expressivity risk, generalization risk, and optimization risk[3] (see Figure 4).

Expressivity risk arises when the chosen equilibrium model class is too restrictive to represent the true underlying travel process. *Generalization risk* arises when a model fits the observed days well but performs poorly on unseen contexts and network topologies. *Optimization risk* arises from the computational challenges of repeatedly solving and differentiating through equilibrium conditions. The dissertation theoretically shows that the nonparametric equilibrium model is *universally expressive*: when parameterized by sufficiently rich neural networks, it can replicate any unique, differentiable equilibrium state induced by a well-posed variational inequality. This expressivity, however, increases data requirements for generalization and computational burden for optimization. Semi-parametric models reduce the data and computational requirements by imposing domain-informed structure, often improving robustness when data are limited.

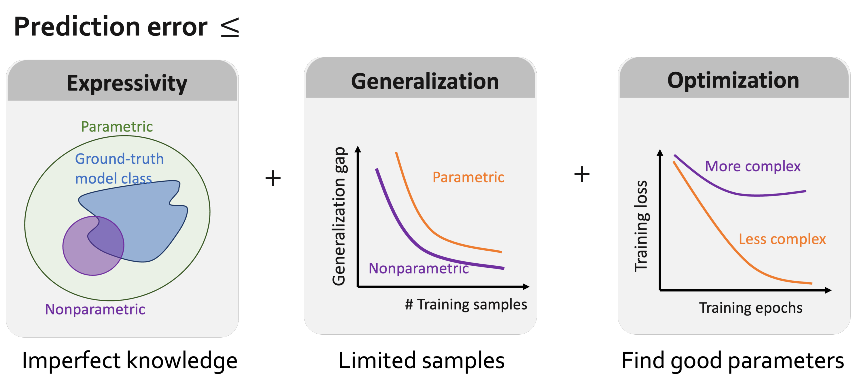

Figure 4: Learning feasibility of the inverse learning framework. Overall performance is decomposed into expressivity, generalization, and optimization risk.

These learning feasibility analyses yield a practical message for transportation modelers: *there is no universally best equilibrium model*. Instead, parametrization choice should depend jointly on the desired goodness-of-fit, available data, interpretability requirements, and the available computational resources. The nonparametric approach may be preferable when both data and computing resources are sufficient. Multi-resolution data can further reduce the dimensionality of the inverse problem by pre-calibrating selected components. Thus, the dissertation provides both a theoretical foundation and a practical guide for choosing among equilibrium-learning parametrizations.

[3] Liu, Zhichen and Yafeng Yin (2025), *"End-to-end learning of user equilibrium: Expressivity, generalization, and optimization"*, *Transportation Science*

4. Distributionally Robust Transportation Network Design with Contextual Uncertainty

Beyond enabling nonparametric learning, multiday observations create a second opportunity: they reveal how equilibrium states vary across recurring contexts such as day of week, weather, and other conditions. Although transportation network analysis has long recognized traffic uncertainty, planners have often lacked data to observe or quantify it. Multiday data makes this uncertainty observable, yet traditional equilibrium models compress daily variation into a single average traffic state and cannot take full advantage of these new observations. Instead, this dissertation learns *context-dependent* equilibrium models in which equilibrium flows vary with observable context, capturing systematic variation that average-value models miss and enabling uncertainty-aware planning.

Learning a context-dependent equilibrium model is not an end in itself. The ultimate purpose of network equilibrium models is to support network management and long-term what-if analysis. A context-dependent model captures recurring variation more accurately than an average-value model, but it may still misrepresent behavior over a decade-long planning horizon because traveler behavior may shift over time. To address this issue, the dissertation develops a *context-dependent distributionally robust network design* framework under behavioral shift [4]. This formulation constructs an ambiguity set around the learned uncertainty distribution and optimizes against the worst case distribution within that set. This allows the planner to account for endogenous uncertainty in travel demand and for distributional mismatch between predicted and realized behavior. The dissertation shows that, under suitable assumptions, this robust design reduces underestimation of social cost relative to a non-robust counterpart.

To solve the resulting large-scale min-max problem with equilibrium constraints, the dissertation develops an automatic-differentiation-based algorithm that solves and differentiates through batches of equilibrium problems in parallel, with local convergence guarantees. Although motivated by continuous network design, the framework applies more broadly to transportation interventions with continuous controls, including signal timing and second-best tolling. In doing so, it closes the loop from learning to intervention by turning learned contextual uncertainty into robust planning decisions.

City-Scale Validation and Broader Impact

The proposed inverse learning framework is demonstrated in a city-scale Ann Arbor deployment using one year of crowdsourced trajectory data from General Motors, in collaboration with the Michigan Department of Transportation[5]. In this deployment, the inverse-learning-based model reduced link travel-time prediction error from 83.6% under the benchmark four-step travel demand forecast model to 34.3%, while capturing interpretable behavioral patterns such as lower travel demand on weekends and snow days. The inverse learning procedure required less than ten minutes, demonstrating that the framework can deliver planning-relevant accuracy on a practical timescale while supporting frequent updates.

Overall, this dissertation establishes inverse learning of user equilibrium as a novel paradigm for transportation network equilibrium analysis. By building equilibrium models from readily available crowdsourced mobility data, the framework transforms transportation network modeling from a survey-based, labor-intensive calibration process into an accessible and scalable service. This paradigm shift can help cities reduce congestion and emissions without additional investments in data collection, while expanding access to advanced modeling tools for resource-constrained communities.

[4] Liu, Zhichen, Yafeng Yin, and Xi Lin (2026), *"Distributionally Robust Transportation Network Design with Contextual Uncertainty"*, *In preparation for submission to Transportation Science*

[5] Liu, Zhichen, Yafeng Yin, Xi Lin, and Zihao Wang (2026), *"Large-Scale Inverse Learning of User Equilibrium via Multiconvex Optimization"*, *Transportation Research Part B*