Research Statement

Mai Anh Tien
School of Computing and Information Systems, Singapore Management University
Email: atmai@smu.edu.sg
19/Dec/2024

Background

Understanding human preferences and predicting human behavior in complex and interactive environments has been a longstanding focus across econometrics, social science, and machine learning, with successful applications in various domains such as resource planning and infrastructure development. Nowadays, the modeling problem becomes more and more complex given the availability of massive amounts of data collected from various channels (mobile phones, online platforms, social networks, etc.), especially when humans now can have access to various sources of information and even interact with AI agents when making decisions. My long-term aim is to address the following questions: *(i) How can we efficiently model and predict human behavior in complex, dynamic, and interactive environments*? (ii) *How can we make the modeling and prediction process adaptive and robust to handle different circumstances of data availability and human behavior constraints?* And (iii) *how can we make the behavioral models useful for various downstream decision-making tasks?* These questions pose several new challenges that necessity new methodologies, and my strategies are to address these challenges through methods rooted in Econometrics, Operations Research, and Machine Learning, focusing on applications in transportation behavioral modeling, location planning, revenue management, and game theory.

Research Areas

My research encompasses three main directions: human behavioral modeling, human-driven decision-making, and imitation learning. While the first direction was the focus of my PhD studies, the second addresses the question of how to efficiently solve decision-making problems that involve human behavioral models, and the last direction reflects my new explorations into a broader context of behavioral modeling where advanced machine learning techniques can be leveraged. Figure 1 below illustrates these three research directions and how they evolve over time.

Human Behavior Modeling and Prediction

This research direction is the main topic of my PhD dissertation, focusing on the use of discrete choice models to analyze and predict people's choice behavior. These models have been widely used in various research fields to describe, explain, and predict how decision-makers choose between two or more discrete alternatives. Examples include selecting a transport mode, buying a car, choosing a route, deciding whether to enter the labor market, or selecting a product from a set of options. Discrete choice models have found extensive application across domains like marketing, revenue management, game theory, transportation planning, environmental studies, energy forecasting, and policymaking.

Along this direction, my research primarily focuses on using discrete choice models to analyze and predict traveler behavior in large-scale transportation networks. This modeling is often referred to as the route choice problem, which has been a longstanding issue in transportation studies with applications in various domains such as traffic simulation, traffic signal control, and infrastructure planning. Prior to my Ph.D. studies, existing works typically relied on a sampling approach, often referred to as a *path-based* approach, where an algorithm is designed and used to generate a set of paths for modeling and estimation. The main issue with this approach is that the estimation results are biased and strongly depend on the algorithms used to generate the sample paths, leading to inconsistency issues. In a seminal work preceding my PhD studies, it has been shown that it is possible to achieve consistent estimates by employing techniques from dynamic programming. However, despite this promising result, utilizing these techniques in the route choice problem poses several challenging research questions. Namely, how to model network correlation—an important aspect in route choice modeling—within dynamic programming methods? How to speed up the computation—as employing dynamic programming is typically costly, especially with large-scale networks? Additionally, how to extend the approach to address the stochastic and dynamic nature of large urban networks, where travel times may change over the day and be uncertain to travelers?

My collaborators and I have been at the forefront of addressing these questions, leading to the development of several *recursive route choice models* that are now considered state-of-the-art (SOTA) behavioral models in the route choice modeling literature. Specifically, our contributions include: (i) an efficient and interpretable way to capture the correlation between path utilities in a complex network [23, 22, 20], (ii) novel algorithms that significantly accelerate the estimation process (up to 30 times faster) [19, 11], and (iii) new models and estimation techniques for stochastic and/or dynamic networks [17, 15]. Moreover, extending beyond the route choice problem, we have broadened the modeling framework to demonstrate a connection between route choice and a general context of network-based demand modeling. This has led to new methods for quickly estimating very large-scale demand models [22, 21].

My research in behavioral modeling has been published in leading transportation research journals (Transportation Science, Transportation Research Part B&C). My work in this field has received some of the most prestigious awards offered by the community, including the **INFORMS-TSL Dissertation Prize** (1st place), the oldest and most prestigious honor for doctoral dissertations in the transportation science and logistics area.

This research direction remains a primary focus of my work, and future endeavors will aim to address even more general and inclusive settings. This includes modeling multi-mode, dynamic, and stochastic transportation systems, where individuals may utilize multiple modes of transport and may be sensitive to uncertainties regarding future travel information. Furthermore, considering the relationship between route choice modeling and imitation learning (a major research area in the machine learning), we plan to explore advanced and recent machine learning techniques to further enhance generality and scalability, making these behavioral models applicable in modern, large, and dynamic urban networks such as the Singapore's transportation system.

Human-driven Decision-Making

I have been pursuing this direction since completing my Ph.D. studies. My primary motivation was to explore how human behavioral models can be efficiently leveraged in downstream decision-making problems, noting that this is an intriguing research topic in Operations Research (OR) and Artificial Intelligence (AI) over recent decades. Leveraging my advanced background and experience in behavioral modeling, I have identified several gaps in the literature, particularly regarding the integration of advanced behavioral models into optimization to achieve tractable, robust and trustworthy decision-making. While advanced behavioral models offer more accurate predictions, they often complicate the decision-making process, necessitating the development of new and innovative optimization methods. My research has been dedicated to bridging these gaps, focusing on applications in facility location, product assortment planning, and security games, which stand as some of the most important problems within the OR and AI communities.

Facility location is a major problem in both OR and AI communities, focusing on selecting an optimal subset of locations from a given pool of candidates to establish new facilities. This problem has many real-world applications and has been studied extensively for decades. However, most of existing studies assume that customer demand is fixed with certainty, which is not the case when using behavioral models (i.e., discrete choice models) to predict demand. Integrating choice models into facility location problems poses several challenges, primarily due to the resulting objective functions becoming highly nonlinear and difficult to solve using standard techniques. Additionally, existing methods often employ simple behavioral models (e.g., multinomial logit), which have limitations in capturing complex human behavior. Employing more advanced behavioral models further increases the problem's complexity, necessitating the development of new solution methods.

Given these challenges and research gaps, my students, collaborators, and I have focused on bringing generalized discrete choice models into the context. Our primary efforts have been concentrated on devising new algorithms to efficiently solve these novel facility location challenges. Some major contributions include [18] where we have developed a novel algorithm based on the outer-approximation scheme that significantly outperforms existing state-of-the-art algorithms. In [16], we made the pioneering effort to incorporate a general class of discrete choice models (i.e., the GEV family) into competitive facility location. To address the challenging problem, we have developed an efficient local search procedure with performance guarantees that surpasses prior methods. We have also explored solution methods to address uncertainty issues [12, 7], and developed new models and solution methods for a joint facility location and cost optimization problem under a new class of discrete choice models [8].

Product pricing and assortment planning is another major application of interest that I have been exploring. This key problem in revenue management involves choosing an optimal subset of products and their prices to offer customers, with the goal of maximizing expected revenue based on customer preferences. While the literature on assortment and price optimization is extensive, including approximation and exact methods developed for various choice models, it lacks a practical algorithm that works with multiple behavioral models simultaneously. This motivated us to develop a trust-region-based local search algorithm that efficiently solves the problem under various choice models [1]. Additionally, we have proposed new, general, and practical algorithms for solving a broader class of nonlinear decision-making problems. These algorithms are applicable to both assortment and pricing, as well as facility location problems [12].

The third topic of interest is security games, which involve a defender aims to allocate security resources to protect facilities from potential attacks by an adversary. I began exploring this application domain with a colleague upon joining SMU. Through this exploration, I discovered that discrete choice (or quantal response) models are widely utilized in Stackelberg security games (SSG) to capture adversary behavior, but the existing literature only focuses on simple discrete choice models. Consequently, the vast body of literature on discrete choice modeling would offer promising avenues for effectively and intuitively modeling adversary behavior. Motivated by this observation, we embarked on the first effort to introduce the nested logit model, from the family of discrete choice models, into SSG to address some limitations of previously employed quantal response models [13]. Additionally, we incorporated dynamic discrete choice concepts to model network-based adversary behavior [2]. Furthermore, we developed a model and solution algorithms that enable the defender to select operational locations, thereby significantly reducing the cost of potential attacks [9].

Our work along this direction have been published in leading OR venues (EJOR, COR, Math Prog) and major Machine Learning conferences (AAAI, NeurIPS, IJCAI, AAMAS). This has been and remains one of my long-term research focuses. There are several unanswered questions I aim to explore, particularly considering the increasing complexity of behavioral models, which necessitates the development of new algorithms to address resulting decision-making challenges. One major challenge is on how to efficiently integrate dynamic and stochastic behavioral models, such as advanced route choice models and those derived from the imitation learning literature (as will be discussed in the next section), into decision-making. Such integration would lead to highly complex optimization problems, but it would also offer valuable opportunities for real-world applications involving intricate human behavior.

Imitation/Inverse Reinforcement Learning

Since completing my Ph.D. studies, I have been exploring ways to move beyond classical behavioral models (i.e., discrete choice models) to find more scalable and efficient modeling techniques capable of handling massive amounts of data collected from interactive multi-agent environments, for which classical behavioral models are limited to handle. Imitation learning (IL) emerged as an appealing direction since it shares a similar objective with classical behavioral models: replicating a policy from demonstration data. Moreover, IL techniques have been known for their scalability and accuracy. Particularly, the integration of modern deep learning techniques into IL algorithms further enhances their efficiency in handling large parameter sets, offering the potential to better predict human behavior within complex systems. In my research, delving into IL not only brings forth novel techniques for behavioral modeling and decision-making, but also fosters opportunities for collaboration with colleagues who specialize in reinforcement learning and deep learning.

The literature on IL is extensive and profound, yet it still has limitations in addressing my specific topics of interest, particularly about whether IL can be leveraged to enhance decision-making processes. This inquiry has led to some notable findings where I demonstrate that IL can indeed assist agents in better learning and making actions. For example, in [14], we explore a class of capacitated vehicle routing problems (VRP) – a major challenge in OR – and propose a method to employ imitation learning to mimic the behavior of heuristic operators. This marks the first exploration of IL in the VRP context, opening up several new promising research avenues. In [10, 4, 3], we investigate multi-agent games where we introduce novel IL-enhanced algorithms where an imitator agent is trained to learn and predict the behavior of the players, thereby aiding the victim agent in achieving higher win rates. In [6, 5], we demonstrate how IL can be utilized to solve reinforcement learning problems, resulting in SOTA results.

My current research in this domain focuses on the development of novel IL algorithms tailored to handle multi-agent settings and scenarios where learning needs to be done without interaction with environments (offline learning) or with low-quality demonstrations. I believe I am still in an early stage of exploring this direction. My long-term goal is not only to advance state-of-the-art IL algorithms to accurately mimic human (or expert system) behavior, but also to connect IL techniques with classical behavioral modeling and decision-making, enabling applications in much more complex and inclusive settings.

Selected Publications and Outputs

[1] T. Mai and A. Lodi, “An Algorithm for Assortment Optimization Under Parametric Discrete Choice Models,” in *Fields Institute Communications Series on Data Science and Optimization* , 2024. doi: 10.2139/ssrn.3370776.

[2] T. Mai, A. Bose, A. Sinha, T. Nguyen, and A. K. Singh, “Tackling Stackelberg Network Interdiction against a Boundedly Rational Adversary,” in *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, 2024, pp. 2913–2921. doi: 10.24963/ijcai.2024/323.

[3] T. V. Bui, T. Mai, and T. H. Nguyen, “Inverse Factorized Q-Learning for Cooperative Multi-agent Imitation Learning,” in *Proceedings of the Thirty-Eighth Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

[4] T. V. Bui, T. Mai, and T. H. Nguyen, “Mimicking To Dominate: Imitation Learning Strategies for Success in Multiagent Competitive Games,” in *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

[5] M. H. Hoang, T. Mai, and P. Varakantham, “Imitate the Good and Avoid the Bad: An Incremental Approach to Safe Reinforcement Learning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.

[6] H. Hoang, T. Mai, and P. Varakantham, “SPRINQL: Sub-optimal Demonstrations driven Offline Imitation Learning,” in *Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

[7] T. T. Dam, T. A. Ta, and T. Mai, “Robust maximum capture facility location under random utility maximization models,” *Eur J Oper Res*, vol. 310, no. 3, pp. 1128–1150, 2023.

[8] N. H. Duong, T. T. Dam, T. A. Ta, and T. Mai, “Joint location and cost planning in maximum capture facility location under random utilities,” *Comput Oper Res*, p. 106336, 2023.

[9] T. Mai and A. Sinha, “Securing Lifelines: Safe Delivery of Critical Services in Areas with Volatile Security Situation via a Stackelberg Game Approach,” in *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023, pp. 5805–5813. doi: 10.1609/aaai.v37i5.25720.

[10] T. V. Bui, T. Mai, and T. H. Nguyen, “Imitating Opponent to Win: Adversarial Policy Imitation Learning in Two-player Competitive Games,” in *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2023, pp. 1285–1293.

[11] T. Mai and E. Frejinger, “Estimation of undiscounted recursive path choice models: convergence properties and algorithms,” *Transportation Science*, vol. 56, no. 6, pp. 1469–1482, 2022.

[12] A. Bose, A. Sinha, and T. Mai, “Scalable Distributional Robustness in a Class of Non-Convex Optimization with Guarantees,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

[13] T. Mai and A. Shina, “Choices Are Not Independent: Stackelberg Security Games with Nested Quantal Response Models,” in *Proceedings - 36th AAAI Conference on Artificial Intelligence (AAAI)*, 2022.

[14] T. V. Bui and T. Mai, “Imitation Improvement Learning for Large-scale Capacitated Vehicle Routing Problems,” in *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS))*, in working paper. 2022, pp. 551–559.

[15] T. Mai, X. Yu, S. Gao, and E. Frejinger, “Route choice in a stochastic time-dependent network: the recursive model and solution algorithm,” *Transportation Research Part B: Methodological*, vol. 151, p. 42, 2021.

[16] T. T. Dam, T. A. Ta, and T. Mai, “Submodularity and local search approaches for maximum capture problems under generalized extreme value models,” *Eur J Oper Res*, 2021.

[17] G. de Moraes Ramos, T. Mai, W. Daamen, E. Frejinger, and S. P. Hoogendoorn, “Route choice behaviour and travel information in a congested network: Static and dynamic recursive models,” *Transp Res Part C Emerg Technol*, vol. 114, pp. 681–693, 2020, doi: 10.1016/j.trc.2020.02.014.

[18] T. Mai and A. Lodi, “A multicut outer-approximation approach for competitive facility location under random utilities,” *Eur J Oper Res*, vol. 284, no. 3, pp. 874–881, 2020, doi: 10.1016/j.ejor.2020.01.020.

[19] T. Mai, F. Bastin, and E. Frejinger, “A decomposition method for estimating recursive logit based route choice models,” *EURO Journal on Transportation and Logistics*, vol. 7, no. 3, pp. 253–275, 2018, doi: 10.1007/s13676-016-0102-3.

[20] M. Zimmermann, T. Mai, and E. Frejinger, “Bike route choice modeling using GPS data without choice sets of paths,” *Transp Res Part C Emerg Technol*, vol. 75, pp. 183–196, 2017, doi: 10.1016/j.trc.2016.12.009.

[21] T. Mai, E. Frejinger, M. Fosgerau, and F. Bastin, “A dynamic programming approach for quickly estimating large network-based MEV models,” *Transportation Research Part B: Methodological*, vol. 98, pp. 179–197, 2017, doi: 10.1016/j.trb.2016.12.017.

[22] T. Mai, “A method of integrating correlation structures for a generalized recursive route choice model,” *Transportation Research Part B: Methodological*, vol. 93, pp. 146–161, 2016, doi: 10.1016/j.trb.2016.07.016.

[23] T. Mai, M. Fosgerau, and E. Frejinger, “A nested recursive logit model for route choice analysis,” *Transportation Research Part B: Methodological*, vol. 75, pp. 100–112, 2015, doi: 10.1016/j.trb.2015.03.015.