

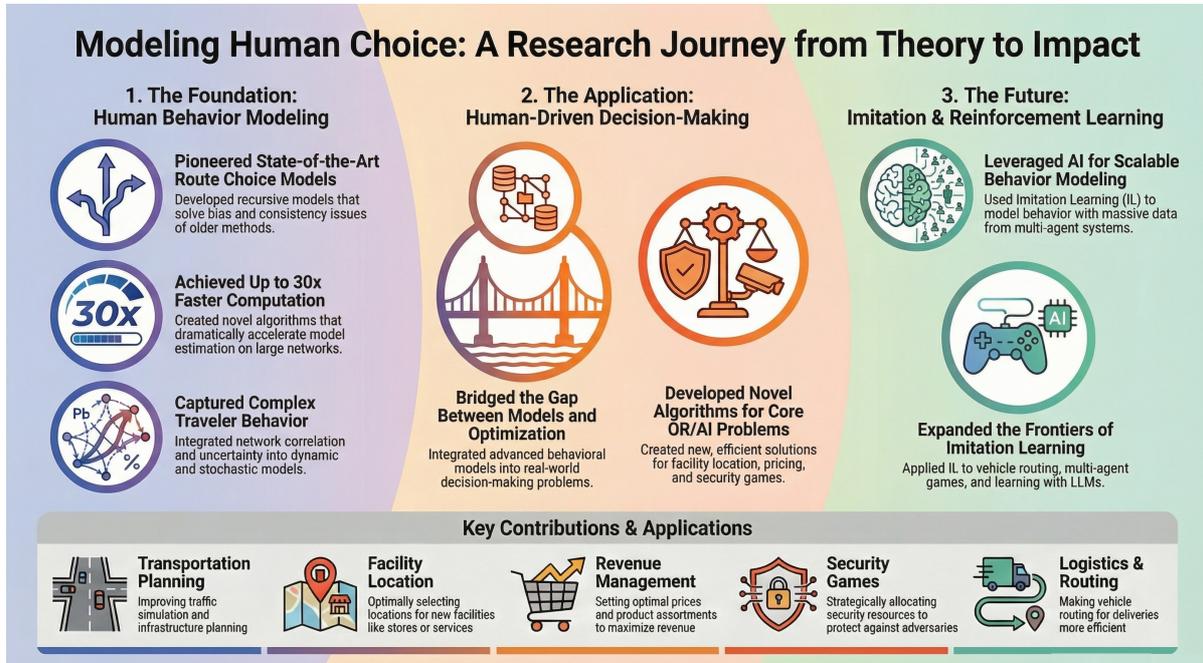
Research Statement

Mai Anh Tien

School of Computing and Information Systems, Singapore Management University

Email: atmai@smu.edu.sg

19/Dec/2025



My research seeks to understand how people make decisions in complex, everyday settings—such as choosing a travel route, selecting a product, or responding to risks—and to use this understanding to design better systems, services, and technologies. Such decisions are generally complex, influenced by many factors: personal preferences, changing environments, limited information, and increasingly, interactions with AI systems. My work addresses these complexities through three interconnected research areas: **human behavior modeling**, **human-driven decision-making**, and **imitation learning**.

First, I develop **behavioral models** that explain and predict how people choose among alternatives in large, dynamic environments such as transportation networks. These models help uncover why people behave the way they do and how their choices respond to congestion, uncertainty, or new information. My contributions include state-of-the-art techniques that improve both the accuracy and scalability of such predictions.

Second, I build **decision-making** tools that use these behavioral models to guide real-world planning tasks—for example, choosing where to build new facilities, how to design product assortments, or how to allocate security resources. These problems become substantially more realistic once human behavior is accounted for, but also far more difficult to solve. My work develops efficient optimization methods that make such human-aware decision-making practical, robust, and transparent.

Finally, I explore **imitation learning**, a modern machine learning approach that enables computers to learn strategies directly from expert demonstrations. This offers a powerful complement to classical behavioral models: while the latter provide structure and interpretability, imitation learning provides scalability and flexibility when data are abundant or environments are interactive and multi-agent. My recent research shows how combining these ideas allows AI systems to better mimic human

expertise, operate safely under uncertainty, and support decision-making in complex multi-agent settings.

Across these three areas, a unifying theme emerges: *modeling how humans behave, then using those models to make better decisions—and in turn, using AI methods to learn more effectively from human expertise*. By connecting insights from econometrics, operations research, and machine learning, my work contributes new foundations for building intelligent systems that understand people, adapt to real-world complexity, and ultimately help societies make more informed and equitable decisions.

Human Behavior Modeling and Prediction

Motivation. Understanding how individuals make decisions in complex environments is fundamental to many fields, including transportation, economics, marketing, and public policy. Discrete choice models play a central role because they provide a principled, interpretable way to explain and predict how people choose among competing alternatives—whether selecting a travel route, a mode of transport, or a product. These models enable researchers and policymakers to quantify behavioral responses, evaluate new interventions, and design systems that align with human preferences.

In large transportation networks, however, accurately modeling route choice remains challenging. Traditional path-based estimation relies on sampling a subset of possible routes, which leads to biased and inconsistent behavioral predictions. Dynamic programming-based discrete choice models offer a theoretically sound alternative, but their use introduces new difficulties: representing correlation among overlapping paths, ensuring computational efficiency at city scale, and extending models to dynamic, uncertain, and information-rich travel environments. Addressing these challenges is essential for producing reliable behavioral predictions that can inform real-world planning, operations, and policy decisions.

Main Contributions. My collaborators and I developed a family of state-of-the-art recursive route choice models that address these methodological challenges. Our key contributions include:

1. *Modeling correlation structures* in large networks through an efficient and interpretable integration mechanism [31, 30, 29, 28, 27, 24].
2. *Fast estimation algorithms* that reduce computation time by up to 30× while maintaining statistical rigor [27, 20, 15].
3. *New modeling and solution approaches* for dynamic and stochastic networks, enabling realistic prediction of traveler responses to uncertainty and evolving information [17, 19].
4. *Generalization to large network-based demand systems*, allowing rapid estimation of high-dimensional behavioral models [26, 24]

These advances have been published in leading transportation journals (*Transportation Science*, *Transportation Research Part B & C*) and recognized with the **INFORMS-TSL Dissertation Prize**.

Future Work. My future work will extend these models to more general mobility environments, including multi-modal and highly dynamic urban systems where individuals switch across transport modes and react to complex real-time information. In parallel, I aim to integrate insights from imitation learning to further enhance scalability and generalizability, making behavioral prediction applicable to modern, data-rich smart city contexts such as Singapore’s transportation ecosystem.

Human-driven Decision-Making

Motivation. Many important operational and policy problems—such as deciding where to locate new facilities, how to design product assortments, how to price goods, or how to allocate security resources—depend critically on how humans will behave in response to these decisions. In practice, ignoring human behavioral responses can lead to suboptimal or even counterproductive outcomes: new facilities may fail to attract users, product offerings may not match customer preferences, and security resources may be deployed ineffectively against strategic adversaries.

Discrete choice models offer a powerful way to predict human behavior, but integrating these models into downstream decision-making creates difficult optimization problems. The resulting objective functions are often highly nonlinear, nonconvex, or stochastic, and classical solution techniques break down—especially when employing more expressive behavioral models that better reflect real-world decision processes. As organizations increasingly seek decision tools that are behaviorally realistic, robust, and scalable, developing new methods for optimization under advanced behavioral models has become a timely and important challenge.

Main Contributions. My research addresses these challenges by developing new optimization models and algorithms that incorporate sophisticated behavioral predictions while remaining tractable and scalable. Key contributions include:

1. **Facility Location under Discrete Choice Models:**
 - Developed efficient algorithms for facility location (and cost optimization) problems that incorporate general discrete choice models with demand uncertainty, achieving major improvements over existing approaches [25, 17, 16].
2. **Assortment and Price Optimization:**
 - Designed algorithms that works effectively across a variety of choice models, addressing the lack of practical methods that can handle multiple behavioral specifications [21, 8, 7, 2, 1].
3. **Security Games and Strategic Behavior:**
 - Introduced richer behavioral models—including nested logit and dynamic discrete choice—into Stackelberg security games to better capture adversary decision-making and overcome limitations of classical quantal response models [23, 22, 18, 9].

These works have appeared in leading *OR journals* (INFORMS Computing, EJOR, COR) and *major AI/ML conferences* (AAAI, NeurIPS, IJCAI), reflecting their broad relevance across both communities.

Future Work. Future research will extend these decision-making frameworks to dynamic, stochastic, and multi-agent settings, where human behavior interacts with evolving system states and strategic responses. A key direction is to incorporate advanced behavioral models, including those emerging from imitation and reinforcement learning, into optimization formulations. This integration will enable decision systems that are more adaptive, robust, and reflective of real-world complexity.

I also aim to explore large-scale, behaviorally informed optimization relevant to urban mobility, digital platforms, and public-sector planning—domains where decisions increasingly rely on massive data, uncertainty, and human–AI interaction. Ultimately, this line of work seeks to build principled methods that support trustworthy, human-centered decision-making in complex operational environments.

Imitation/Inverse Reinforcement Learning

Motivation. As modern systems increasingly generate massive amounts of data from interactive, multi-agent environments—such as transportation platforms, logistics systems, online marketplaces, and strategic games—classical behavioral models face limitations in scalability and expressiveness. Imitation learning (IL) and inverse reinforcement learning (IRL) offer powerful alternatives: they aim to learn policies directly from expert demonstrations, enabling models that can capture complex, sequential, and strategic behavior at scale.

These approaches are particularly relevant in practice because many real-world tasks—routing and scheduling, autonomous driving, cybersecurity, team coordination—require agents to act in environments where collecting interactive data is expensive or risky, and expert demonstrations may be noisy, incomplete, or of mixed quality. IL/IRL provide a data-driven way to learn from such expert signals, producing policies that can assist or augment human decision-makers. However, existing IL/IRL methods often struggle in multi-agent, offline, or imperfect-demonstration settings, which limits their applicability to the domains where they are most needed. Addressing these gaps is

essential for building AI systems that can replicate human expertise and support robust decision-making in complex environments.

Main Contributions. My research develops new IL and IRL algorithms tailored for multi-agent systems, offline learning scenarios, and realistic demonstration data. Key contributions include:

1. Learning from Noisy or Mixed-Quality Demonstrations:

- Pioneered approaches [5, 4, 3], which leverage LLMs to filter and extract high-quality information from unlabeled and heterogeneous expert demonstrations .
- Developed robust imitation learning methods that selectively learn from expert demonstrations while avoiding harmful behaviors, achieving state-of-the-art results in offline IL/RL benchmarks [12, 11, 10, 6].

2. Multi-Agent Imitation and Reinforcement Learning:

- Developed IL and preference-based methods for cooperative multi-agent reinforcement learning [13, 3–5].
- Proposed novel algorithms for competitive and cooperative multi-agent games where an imitator learns to model other agents’ behavior, allowing a target agent to achieve higher performance or win rates [19, 14].

Collectively, these works advance the theoretical and practical foundations of IL/IRL, enabling their application in challenging multi-agent and real-world contexts.

Future Work. My future research will focus on developing scalable and robust IL/IRL algorithms that operate effectively in multi-agent, human-in-the-loop, and high-stakes decision-making environments. Important directions include:

- **Unifying classical behavioral models and IL:** leveraging structural insights from discrete choice theory to improve interpretability, sample efficiency, and robustness of learning-based behavioral models.
- **Offline and limited-interaction learning:** creating algorithms that can learn safely from fixed datasets—a key requirement in domains such as transportation planning, healthcare decision support, and strategic security settings.
- **Human–AI collaboration:** designing IL/IRL methods that learn from experts but also support experts, enabling AI systems that provide guidance, critique, or strategic planning assistance.
- **Applications in complex multi-agent systems:** including mobility ecosystems, logistics networks, and competitive or adversarial environments where strategic behavior emerges.

Through these directions, my long-term goal is to bridge the gap between interpretable behavioral modeling and scalable machine learning, enabling AI systems that capture human expertise and support effective decision-making in complex, interactive environments.

Selected Publications and Outputs

- [1] Q. Shao, T. Mai, and S. Cheng, “Constrained Pricing in Logit-based Revenue Management,” *INFORMS J Comput*, 2026 (forthcoming).
- [2] H. G. Pham and T. Mai, “Constrained Assortment Optimization under the Mixed-Logit Model: Approximation Schemes and Outer Approximation Approaches,” *Eur J Oper Res*, 2025.
- [3] T. V. Bui, T. Mai, and H. T. Nguyen, “MisoDICE: Multi-Agent Imitation from Unlabeled Mixed-Quality Demonstrations,” *Proceedings of the Thirty-Eighth Conference on Neural Information Processing Systems (NeurIPS)*, 2025, [Online]. Available: <https://arxiv.org/abs/2505.18595>
- [4] T. V. Bui, T. H. Nguyen, and T. Mai, “ComaDICE: Offline Cooperative Multi-Agent Reinforcement Learning with Stationary Distribution Shift Regularization,” in *Proceedings of the 13th International Conference on Learning Representations (ICLR)*, 2025. [Online]. Available: <https://arxiv.org/abs/2410.01954>

- [5] T. V. Bui, T. Mai, and T. H. Nguyen, "O-MAPL: Offline Multi-agent Preference Learning," in *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025. [Online]. Available: <https://arxiv.org/abs/2501.18944>
- [6] H. Hoang, T. Mai, and P. Varakantham, "No Experts, No Problem: Avoidance Learning from Bad Demonstrations," in *Proceedings of the 39th Conference on Neural Information Processing Systems (NeurIPS 2025)*, 2025.
- [7] 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.
- [8] C. Le and T. Mai, "Constrained Assortment Optimization under the Cross-Nested Logit Model," *Prod Oper Manag*, 2024.
- [9] 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.
- [10] 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.
- [11] H. Jiang, T. Mai, P. Varakantham, and H. Hoang, "Reward Penalties on Augmented States for Solving Richly Constrained RL Effectively," in *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 2024, pp. 19867–19875. doi: 10.1609/aaai.v38i18.29962.
- [12] 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. [Online]. Available: <https://arxiv.org/abs/2402.13147>
- [13] 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. [Online]. Available: <https://arxiv.org/abs/2310.06801>
- [14] 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. [Online]. Available: <https://arxiv.org/abs/2308.10188>
- [15] T. Mai, T. V. Bui, Q. P. Nguyen, and T. V. Le, "Estimation of Recursive Route Choice Models with Incomplete Trip Observations," *Transportation Research Part B: Methodological*, vol. 173, no. 2023, pp. 313–331, 2023.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] T. Mai, N. H. Duong, and T. A. Ta, "Binary-Continuous Sum-of-ratios Optimization: Discretization, Approximations, and Convex Reformulations," 2022.

- [22] 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.
- [23] 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. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/hash/59e02a1440e6667e01628ed4c325255c-Abstract-Conference.html
- [24] 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.
- [25] 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.
- [26] 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.
- [27] 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.
- [28] T. Mai, F. Bastin, and E. Frejinger, "On the similarities between random regret minimization and mother logit: The case of recursive route choice models," *Journal of Choice Modelling*, vol. 23, pp. 21–33, 2017, doi: 10.1016/j.jocm.2017.03.002.
- [29] 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.
- [30] T. Mai, "Modeling Route Choice with Real-Time Information : Comparing the Recursive and Non-Recursive Approaches," *arXiv preprint arXiv:2005.05783*, 2015.
- [31] 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.