Professional Statement

My research advances data-driven algorithms for transportation systems modeling, control, and optimization. My approach is *use-inspired* [1, 2]: I gravitate to problems with societal stakes for which existing tools fall short, and I build methods that generalize beyond the original problem. This orientation led me to transportation, a domain vital to economic and social well-being yet far from meeting its sustainability, safety, and access goals. Data and computing are already reshaping the field, primarily through monitoring and retrospective analysis. *Prospective counterfactual analysis* instead estimates how proposed strategies would affect outcomes before deployment. Making such analysis *scalable and trustworthy* remains difficult but would complement costly field and user studies. It could answer questions like "What would adding bike lanes in Boston do to traffic and safety?" or "Which US cities should implement congestion pricing after NYC?" My work develops computational tools needed for such analysis.

A key barrier to scalable, trustworthy prospective analysis is the difficulty of optimization: a trustworthy model must be *calibrated* to data (a notoriously difficult optimization problem), and decision-relevant analysis optimizes that model to weigh *alternatives*. This generates a **long tail of case-specific optimization problem variants**, each shaped by local networks, constraints, objectives, and data. Vehicle routing problems (VRPs) illustrate the pattern [3, 4]: each added feature, such as time windows, charging, or uncertainty, yields a distinct NP-hard variant, and deployments combine many such features. The *10,000 VRP papers published in 2025 alone* reflect the scale of this tail. This pattern pervades transportation, from signal control and transit planning to automated vehicle policy and urban air mobility.

Prior methods tackle this tail through intensive, per-problem specialization: hand-tuned heuristics for classical solvers, or the data and tuning that reinforcement learning (RL) needs (Fig. 1) [5–W8]. Because no method excels everywhere, this specialization is unavoidable, but it need not be handcrafted for each problem. My thesis is that specialization can instead be *learned* and amortized across a family of related problems by generalizing a learned solver across many scenarios or accelerating the solvers we already have, sharply lowering the burden on the analyst. The central question I pursue is **how far data-driven algorithms can reduce the expertise, data, and tuning needed to effectively solve families of related optimization problems**. My 2023 NSF CAREER Award recognizes and supports this agenda [W9].

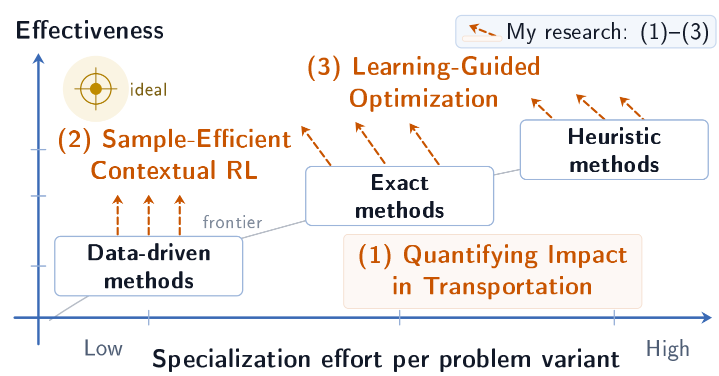

Figure 1: Optimization solver taxonomy and my research, pushing the frontier toward effective, low-effort long-tail optimization.

My group pursues this agenda in three directions (Fig. 1). **Thrust 1** quantifies impact at scale, stress-testing current optimization tools on consequential transportation questions both to surface methodological gaps and to demonstrate value. For instance, it estimates that eco-driving can cut intersection CO_2 by *11–22%* across three cities [**W10**].[1] **Thrusts 2 and 3** are complementary approaches to the long tail, differing in whether a capable solver already exists: when none does, Thrust 2 develops systematic methods to ease the use of deep RL, such as increasing sample efficiency by *30×* [**W11**]; when a solver does exist, Thrust 3 augments it with machine learning, making it *2–7× faster* at equal quality [**W12**, **W13**]. This work is grounded in partnerships with public agencies and industry (UDOT, JR East, Cintra, Amazon, and Symbotic [**W14**]) and is helping to define an emerging field at the intersection of transportation, optimization, and AI.

[1] References to my work are prefixed with a "W" and appear in blue. **Bold orange** highlights five representative articles.

Thrust 1: Methodologies for Quantifying Impact at Scale in Transportation Systems

There is tremendous interest in simulation-based modeling of transportation strategies because it can test many scenarios without costly field deployment. Its accuracy, however, depends on the quality of the data, assumptions, and algorithms. I identify methodological gaps by applying existing methods and then develop tools to fill them. Two problems I study are eco-driving and active traffic management.

My work quantifies the policy relevance of **eco-driving at signalized intersections** as a means to mitigate carbon emissions [**W10**]. Over 1,000 papers study *how* to drive efficiently, but not *whether* doing so would cut emissions enough to warrant large-scale deployment. We formalized the Regional Eco-driving Problem, which seeks driving behaviors that cooperatively minimize expected system-level emissions across region-wide traffic-intersection scenarios while preserving throughput and safety. A key challenge was solving one million scenarios and evaluating their emissions. Traditional model-based control techniques would be time-intensive to adapt to the wide range of scenarios, motivating the use of deep reinforcement learning (RL) as a workhorse for solving the scenarios [W15]. We also mitigated the risk of "reward hacking" (shifting emissions to other parts of the network) through careful handling of upstream and downstream queues. Furthermore, EPA's MOVES, the industry-standard emissions simulator, is too slow at scale, so we distilled it into NeuralMOVES, a set of neural surrogates that run $10^4\times$ faster [W16].

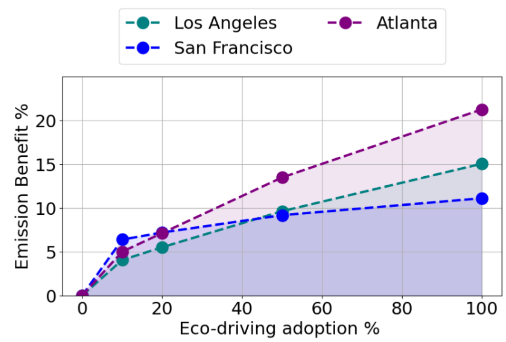

Figure 2: Eco-driving benefits are concave in adoption due to platooning effects, suggesting near-term deployment potential.

We found that, at full adoption, eco-driving can *cut intersection-level CO_2 by 11–22%* across three cities, chiefly through avoided acceleration (Fig. 2). Eco-driving has near- and long-term implications: 70% of its benefits are concentrated in 20% of drivers and intersections, and its gains accrue in addition to those from electric vehicle adoption. Eco-driving also yields co-benefits in air quality, throughput, and safety through smoother traffic. These benefits drew support from the Utah DOT, and we are now working with JR East on eco-driving for trains. The same RL-assisted approach could refine deployment guidance for professional fleets, gamification, or vehicle automation [W17–W24]. However, applying RL to each new problem required extensive engineering and trial and error, motivating dedicated research on more systematic, sample-efficient methods for deep RL (Thrust 2).

Controllable congestion on highways asks an analogous question about delay: how much congestion can active traffic management remove without resorting to more intensive interventions like tolling or lane expansion? We formalize *controllable congestion* as the maximum delay reduction from speed control, solved via nonlinear model predictive control subject to METANET macroscopic traffic dynamics [25] under varying demand and operational constraints [W26]. A key difficulty is calibration, despite decades of research, especially for freeways with stop-and-go congestion: we showed that two recent methods failed to generalize to a longer stretch of the same freeway [W27]. To capture rapid shifts in traffic regimes, we developed the first *dynamic* traffic calibration methods; our gradient- and RL-based methods improve METANET reconstruction accuracy by 48% and increase robustness to perturbations [W28–W30]. With global infrastructure firm Cintra, we used the calibrated models to estimate that on average *50% of rush-hour delay is controllable* in segments of I-24 (Nashville, TN) and I-35W (Fort Worth, TX).

To chart achievable gains measure by measure, we are now analyzing the traffic safety effects of road design with MassDOT, using behavior-prediction models adapted from autonomous driving [W31]; the modal shift effects of journey planning that accounts for parking-search time [W32]; and the accessibility effects of congestion pricing paired with transit reinvestment, as in NYC.

Thrust 2: Sample-Efficient Contextual Reinforcement Learning

Reinforcement learning (RL) is an appealing analysis tool: it directly optimizes strategies from complex objectives and simulators that encode case-specific variants, reducing reliance on specialized optimization expertise. But realizing this promise is hard: RL is tuning-intensive and unreliable across the scenario families that practical analysis needs. Analyzing eco-driving, for example, could span scenarios varying in traffic demand, adoption rates, signal plans, road geometry, and vehicle types. This thrust asks how to solve such scenario families with far less data, tuning, and engineering. We formalize them as contextual Markov decision processes (CMDPs): parameterized families of related MDPs (*tasks*) in which a context vector indexes the scenario [33–35]. We focus on fixed, observed contexts and continuous context spaces.

We evaluated *off-the-shelf* deep RL on CMDPs spanning 160 traffic-signal-control scenarios developed with the Utah DOT, as well as on benchmark control problems [W8]. We found that **RL methods are so sensitive to scenario variation** that, on average, they underperform a simple fixed-time traffic control baseline *even when evaluated on their own training scenario*. Whereas RL methods readily solve the canonical CartPole benchmark, they fail on its variations, suggesting overfitting of the *algorithm* (not the model) to specific scenarios exposed in public benchmarks. Further benchmarking showed trained models underperforming across the context space (Fig. 3, orange vs. gray) and multi-task learning failing to resolve it (green) [**W11**]. Because many RL applications are inherently contextual, these findings motivate sample-efficient contextual RL. To support research, the IntersectionZoo benchmark exposes one million data-informed traffic scenarios from ten US cities for cooperative eco-driving [W36].

Beyond transportation, this lens helped **Amazon** identify training brittleness in Nova Forge, one of the few commercial RL products, in this case for large language model (LLM) fine-tuning; we then co-developed and deployed a curriculum-based stabilization algorithm [W37, 38].

The gray curve in Fig. 3 shows that CMDP performance improves by selecting, for each context, the best transfer from independently trained RL policies. Counterintuitively, training on a target task may not yield its best performance, motivating **strategic task selection**. Model-Based Transfer Learning (MBTL) does this by using a Gaussian process to model per-context performance, representing transfer loss as a linear function of context distance, and applying Bayesian optimization to select source contexts under a limited budget; standard RL (e.g., PPO [39]) then trains only those policies [**W11**]. MBTL achieves up to *30× better sample efficiency* (brown vs. orange, Fig. 3), inherits GP-UCB's sublinear regret guarantees [41], and extends to many settings [W42, W43].

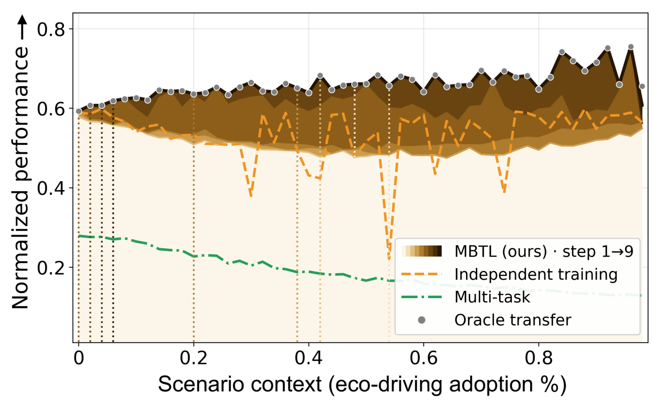

Figure 3: Standard and multi-task RL (orange, green) are sample-inefficient for solving CMDPs [39, 40]. MBTL iteratively selects training tasks (dotted) to exploit policy transfer (shades of brown) [**W11**].

Residual policy learning offers a complementary solution: start from a known nominal controller and learn a correction term [44, 45]. This approach fits transportation because controllers like the Intelligent Driver Model or Green Light Optimal Speed Advisory (GLOSA) encode useful operating rules, obviating imitation learning [46, 47]. We adapt this approach to contextual RL through Multi-Residual Task Learning (MRTL), which learns multi-task residual policies conditioned on the context vector atop these nominal controllers [W48]; a mixture-of-experts version combines multiple controllers [W49].

To make RL training more systematic, we are developing model-based approaches for other aspects of training, including fine-tuning dynamics [W50], hyperparameters, and training complexity [W51, W52].

Thrust 3: Accelerating Discrete Optimization with Machine Learning

Many transportation optimization problems, especially discrete ones like routing, scheduling, and assignment, are already served by strong solvers, from mixed-integer programming (MIP) to vehicle routing heuristics. For these problems, end-to-end RL may be unnecessary; this thrust instead keeps the solver and makes it faster. The key is that solvers run under a time budget: milliseconds for real-time control or hours to route millions of packages at Amazon. Within that budget, they rarely reach optimality, so steering the same solver toward productive decisions yields a better solution at the same computational cost. We call this *learning-guided optimization* [53, 54]: training small, often supervised predictors to steer a solver toward productive decisions, such as which parts of a solution to revisit or which routines to call, while preserving guarantees at modest cost (one GPU-day, 1,000 problem instances). Across a dozen studies, learning-guided solvers achieve **2–7× speedups** at equal quality, creating a repeatable workflow to specialize solvers to long-tail use cases.

One way to accelerate solvers is to **learn to reuse solutions across iterations**. Each solver pass uses substantial computation to adjust only a small portion of the incumbent (Fig. 4), so we learn which decisions will not change *in hindsight* and hand the solver a smaller residual problem. Motivated by rail-terminal dispatch in Boston, our learning-guided rolling-horizon optimization (L-RHO) for long-horizon flexible job-shop scheduling fixes machine assignments that stay consistent across planning windows, *cutting solve time 2×* [W12]. We also provide a probabilistic analysis of when L-RHO outperforms heuristic rolling-horizon optimization. In VRP solvers, we apply this idea spatially, contracting incumbent route segments predicted to remain stable into hypernodes for *7× speedups* [W13]. This prediction problem motivated the first joint non-autoregressive/autoregressive decoder in neural combinatorial optimization: a global pass flags unstable regions, and a local pass refines them; the global pass alone is too coarse, while the local pass alone is too myopic.

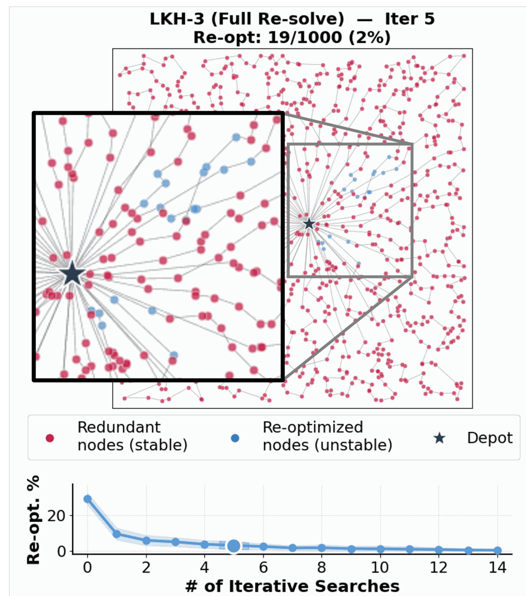

Figure 4: Solver redundancy in the capacitated VRP solver LKH-3 [55]: more than 98% of the solution is unchanged (red) across iterations.

A second way is to **replace hard-coded rules with data-driven ones**. My early work achieved 2–4× speedups in VRP, MIP, and multi-agent algorithms [W56–W58]. We then applied these ideas to *congestion in automated warehouses*. These facilities are transportation networks in miniature, where dozens to thousands of mobile robots form a dense traffic-assignment problem with subsecond replanning. Robots are packed densely to maximize throughput, but under a continuous task stream (lifelong multi-agent path finding, MAPF), congestion easily leads to deadlock. Rather than learning end to end, we hybridized RL with prioritized planning (PP) [59]: PP plans one robot at a time in a random order; we use RL to learn that order, leaving the planner untouched. With Symbotic, this improved package throughput by *20%* in simulated industrial warehouses [W14].

We are now creating a tutorial on learning-guided optimization and studying how agentic AI can discover acceleration techniques at scale. Other emerging interests include designing more sample-efficient neural architectures by encoding optimization inductive biases [W60] and extending these ideas to detailed simulators such as MATSim [61], where calibration difficulty and slow convergence inhibit adoption.

Together, these results lay a foundation for systematic, prospective analysis that addresses transportation's longstanding challenges and supports better outcomes in people's everyday lives.

Education. Three principles underlie how I prepare students to engineer increasingly complex systems: *systems thinking*, so students can match tools to problems; *experiential learning*, because understanding deepens when students put classroom concepts into practice and discover their limits; and *technical communication*, so they can explain ideas precisely and compellingly.[2]

6.7920 (Reinforcement Learning: Foundations and Methods), which I created, is MIT's first graduate course dedicated to RL and a modern successor to MIT's classic course on Dynamic Programming [62]. The central design question is what researchers must understand not only to know *how* RL works but also to judge *when to use it and when not to*. The course is proof-based and foundational, equipping students to exercise that judgment as RL enters high-stakes settings from LLM fine-tuning to robotic control.

1.041/1.200 (Transportation: Foundations and Methods) serves undergraduate and graduate students and is my platform for applying all three principles in a changing field in which new data sources, emerging technologies, and AI reshape what practitioners need to know. Hands-on labs make abstract topics tangible through exercises such as "Build your own traffic jam," "Build a queueing model for Seattle transit," and "Build an AI agent to help you find parking." They ask students to reason about when simulation-based, optimization-based, data-driven, or queueing-theoretic approaches are appropriate and how model assumptions affect outcomes. Technical communication culminates in student-produced educational videos, several of which have been published on the MIT CEE YouTube channel [63].

Outside the classroom, I have mentored 11 PhD, 16 master's, and 30 undergraduate students at MIT. My three PhD graduates are now at Anthropic, Toyota Research Institute, and Microsoft Research.

Service. My service aims to strengthen *scientific integrity* so that research translates reliably into knowledge and practice. I focus on improving the structures that help researchers verify, interpret, and build on one another's work.

I chair and co-founded the REproducible Research In Transportation Engineering (RERITE) Working Group to advance **reproducibility** [W64]. We launched this grassroots effort in 2024, motivated by concerns about limited artifact sharing. A subsequent LLM-assisted audit of 10,000 recent transportation studies quantified the scale of the problem: **only 3% shared data and code** (cf. 80% code sharing in ML), revealing substantial room to make the field's knowledge easier to verify and build upon [W65, 66]. RERITE supports and celebrates reproducible research by raising awareness, offering tutorials and awards, publishing special issues, curating resources, holding workshops, supporting research, and advising publishers on publication policies [W65, W67–W73]. In 2024, I led a tutorial at ITSC on reproducibility challenges and best practices, and Prof. Montasir Abbas (VT) called it "the most useful talk I have seen in all my years at ITSC" [W67]. We created a web app for discovering open artifacts [W72] and are launching a special issue to publish high-quality transportation benchmarks [W71].

Careful claim scoping presents a related opportunity: claims that extend beyond the available evidence can make it harder to assess what research can reliably deliver. As a program co-chair of the 2025 Reinforcement Learning Conference (RLC), I co-introduced a cover page requiring authors to state and scope their contributions explicitly; reviewers use these statements as a shared rubric for assessing claim–evidence alignment, helping authors communicate their contributions clearly and precisely [W74].

In **editorial service**, I hold associate-editor or equivalent roles at flagship venues spanning transportation, ML, and operations research: TR-C, ICML, NeurIPS, OPRE, and ICRA.[3] I also organize workshops and competitions to cross-pollinate these communities.

At MIT, I lead CEE Graduate Admissions (Systems & Transportation), serve on education committees for CEE and the College of Computing, and have organized LIDS seminars and CEE Research Days.

[2] Course materials are available at `web.mit.edu/1.041/www` and `web.mit.edu/6.7920/www`.

[3] Full names: Transportation Research Part C: Emerging Technologies, International Conference on Machine Learning, Conference on Neural Information Processing Systems, Operations Research, International Conference on Robotics & Automation.

References

[1] Stokes, Donald E. *Pasteur's Quadrant: Basic Science and Technological Innovation*. Brookings Institution Press, 2011.

[2] Rolnick, David, Aspuru-Guzik, Alan, Beery, Sara, Dilkina, Bistra, Donti, Priya L. Ghassemi, Marzyeh, Kerner, Hannah, Monteleoni, Claire, Rolf, Esther, Tambe, Milind, White, Adam. "Position: Application-Driven Innovation in Machine Learning". *Proceedings of the 41st International Conference on Machine Learning*. International Conference on Machine Learning. PMLR, July 8, 2024, pp. 42707–42718. URL: https://proceedings.mlr.press/v235/rolnick24a.html.

[3] Dantzig, G. B. Ramser, J. H. "The Truck Dispatching Problem". *Management Science* 6.1 (Oct. 1959), pp. 80–91. DOI: 10.1287/mnsc.6.1.80.

[4] Laporte, Gilbert. "Fifty Years of Vehicle Routing". *Transportation Science* 43.4 (Nov. 2009), pp. 408–416. DOI: 10.1287/trsc.1090.0301.

[5] Henderson, Peter, Islam, Riashat, Bachman, Philip, Pineau, Joelle, Precup, Doina, Meger, David. "Deep Reinforcement Learning That Matters". *Proceedings of the AAAI Conference on Artificial Intelligence* 32.1 (Apr. 2018). DOI: 10.1609/aaai.v32i1.11694.

[6] Engstrom, Logan, Ilyas, Andrew, Santurkar, Shibani, Tsipras, Dimitris, Janoos, Firdaus, Rudolph, Larry, Madry, Aleksander. "Implementation Matters in Deep Policy Gradients: A Case Study on PPO and TRPO". *International Conference on Learning Representations*. 2020. URL: https://openreview.net/forum?id=r1etN1rtPB.

[7] Andrychowicz, Marcin, Raichuk, Anton, Stańczyk, Piotr, Orsini, Manu, Girgin, Sertan, Marinier, Raphael, Hussenot, Léonard, Geist, Matthieu, Pietquin, Olivier, Michalski, Marcin, Gelly, Sylvain, Bachem, Olivier. "What Matters for On-Policy Deep Actor-Critic Methods? A Large-Scale Study". *International Conference on Learning Representations*. 2021. URL: https://openreview.net/forum?id=nIAxjsniDzg.

[W8] Jayawardana, Vindula, Tang, Catherine H, Li, Sirui, Suo, Dajiang, **Wu, Cathy**. "The Impact of Task Underspecification in Evaluating Deep Reinforcement Learning". *Advances in Neural Information Processing Systems (NeurIPS)*. Oct. 2022. DOI: 10.48550/arxiv.2210.08607. [26% acceptance rate].

[W9] National Science Foundation, *NSF Award Search: Award # 2239566 - CAREER: Learning for Generalization in Large-Scale Cyber-Physical Systems*. https://www.nsf.gov/awardsearch/showAward?AWD_ID=2239566. Mar. 2023.

[W10] Jayawardana, Vindula, Freydt, Baptiste, Qu, Ao, Hickert, Cameron, Sanchez, Edgar R. Tang, Catherine, Taylor, Mark, Leonard, Blaine, **Wu, Cathy**. "Mitigating Metropolitan Carbon Emissions with Dynamic Eco-driving at Scale". *Transportation Research Part C: Emerging Technologies (TR-C)* (Apr. 2025). DOI: 10.1016/j.trc.2025.105146.

[W11] Cho, Jung-Hoon, Jayawardana, Vindula, Li, Sirui, **Wu, Cathy**. "Model-Based Transfer Learning for Contextual Reinforcement Learning". *Advances in Neural Information Processing Systems (NeurIPS)*. Vol. 37. Dec. 2024, pp. 88279–88319. URL: https://proceedings.neurips.cc/paper_files/paper/2024/hash/a10c3d85879c29331ba4d73ede56c7d3-Abstract-Conference.html. [26% acceptance rate].

[W12] Li, Sirui, Ouyang, Wenbin, Ma, Yining, **Wu, Cathy**. "Learning-Guided Rolling Horizon Optimization for Long-Horizon Flexible Job Shop Scheduling". *International Conference on Learning Representations (ICLR)*. Apr. 2025. URL: https://openreview.net/forum?id=Aly68Y5Es0. [32% acceptance rate] **MIT Home Page Feature**.

[W13] Ouyang, Wenbin, Li, Sirui, Ma, Yining, **Wu, Cathy**. "Learning to Segment for Vehicle Routing Problems". *International Conference on Learning Representations (ICLR)*. Apr. 2026. URL: `https://openreview.net/forum?id=pN261iTKvr`. [28% acceptance rate] **Oral (1%)**.

[W14] Zheng, Han, Ma, Yining, Araki, Brandon, Chen, Jingkai, **Wu, Cathy**. "Learning-Guided Prioritized Planning for Lifelong Multi-Agent Path Finding in Warehouse Automation". *Journal of Artificial Intelligence Research (JAIR)* (Mar. 2026). DOI: `10.1613/jair.1.20611`.

[W15] Jayawardana, Vindula, **Wu, Cathy**. "Learning Eco-Driving Strategies at Signalized Intersections". *2022 European Control Conference (ECC)*. London, United Kingdom: IEEE, July 2022, pp. 383–390. DOI: `10.23919/ECC55457.2022.9838000`. **MIT Home Page Feature**.

[W16] Sanchez, Edgar Ramirez, Tang, Catherine, Xu, Yaosheng, Renganathan, Nrithya, Jayawardana, Vindula, He, Zhengbing, **Wu, Cathy**. "NeuralMOVES: Extracting and Learning Surrogates for Diverse Vehicle Emission Models". *Transportation Research Part C: Emerging Technologies (TR-C)* (Mar. 2026). DOI: `10.1016/j.trc.2026.105530`.

[W17] Niazi, M. Umar B. Cho, Jung-Hoon, Dahleh, Munther A. Dong, Roy, **Wu, Cathy**. "Eco-Driving Incentive Mechanisms for Mitigating Emissions in Urban Transportation". *IEEE Transactions on Control of Network Systems (T-CNS)* (Oct. 2025). DOI: `10.1109/TCNS.2025.3623943`.

[W18] Sridhar, Mayuri, **Wu, Cathy**. "Piecewise Constant Policies for Human-Compatible Congestion Mitigation". *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, Sept. 2021, pp. 2499–2505. DOI: `10.1109/ITSC48978.2021.9564789`.

[W19] Kim, Jeongyun, Cho, Jung-Hoon, **Wu, Cathy**. "Reinforcement Learning for Robust Advisories under Driving Compliance Errors". *IEEE Transactions on Intelligent Transportation Systems (T-ITS)* (2025). DOI: `10.1109/TITS.2025.3550418`.

[W20] Hasan, Aamir, Chakraborty, Neeloy, Chen, Haonan, **Wu, Cathy**, Driggs-Campbell, Katherine. "Cooperative Advisory Residual Policies for Congestion Mitigation". *ACM Journal on Autonomous Transportation Systems (JATS)* 2.2 (Dec. 2024), pp. 1–31. DOI: `10.1145/3699519`.

[W21] Hasan, Aamir, Chakraborty, Neeloy, Chen, Haonan, **Wu, Cathy**, Driggs-Campbell, Katherine. "Lessons in Cooperation: Driver Sentiments towards Real-Time Advisory Systems". *IEEE Intelligent Transportation Systems Magazine (ITSM)* (2025). DOI: `10.1109/MITS.2025.3555564`.

[W22] Yan, Zhongxia, **Wu, Cathy**. "Scalability of Platoon-based Coordination for Mixed Autonomy Intersections". *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2024. DOI: `10.1109/IROS58592.2024.10801701`.

[W23] Yan, Zhongxia, **Wu, Cathy**. "Reinforcement Learning for Mixed Autonomy Intersections". *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. Nov. 2021, pp. 2089–2094. DOI: `10.1109/ITSC48978.2021.9565000`.

[W24] Suo*, Dajiang, Jayawardana*, Vindula, **Wu, Cathy**. "Model-Free Learning of Corridor Clearance: A Near-term Deployment Perspective". *IEEE Transactions on Intelligent Transportation Systems (T-ITS)* (2024). DOI: `10.1109/TITS.2023.3344473`.

[25] Messmer, A. Papageorgiou, M. "METANET: A Macroscopic Simulation Program for Motorway Networks". *Traffic Engineering & Control* 31.9 (Sept. 1990). URL: `https://trid.trb.org/View/348807`.

[W26] Raghavan, Shreyaa, Ramirez-Sanchez, Edgar, **Wu, Cathy**. "Controllable Congestion: Toward Actionable and Model-based Performance Measures for Highways". *IEEE Intelligent Vehicles Symposium (IV)*. Detroit, MI, USA: IEEE, June 2026, pp. 22–25. DOI: `10.1109/IV66570.2026.11623991`.

[W27] Hickert, Cameron, Wang, Athena, Samaei, Maryam, Zhang, Chengyuan, Sun, Lijun, Ameli, Mostafa, **Wu, Cathy**. "AutoTune: A Unified Benchmark for Highway Traffic Microsimulation

Calibration". *IEEE Intelligent Vehicles Symposium (IV)*. Detroit, MI, USA: IEEE, June 2026. DOI: 10.1109/IV66570.2026.11624093.

[W28] Raghavan, Shreyaa, Hickert, Cameron, Chan, Monica, **Wu, Cathy**. "Dynamic Gradient-Based Calibration for Robust and Accurate Traffic Macrosimulation". *International Conference on Intelligent Transportation Systems (ITSC)*. Sept. 2026. DOI: 10.48550/arXiv.2605.19056.

[W29] Chan, Monica S. H. "Reinforcement Learning for Robust and Dynamic Second-Order Macroscopic Traffic Flow Calibration". MA thesis. Cambridge, MA: Massachusetts Institute of Technology, May 2026.

[W30] Chan, Monica, Raghavan, Shreyaa, **Wu, Cathy**. "Open-Source METANET Calibration for Reproducible Freeway Traffic Macroscopic Simulation". *International Conference on Intelligent Transportation Systems (ITSC)*. Sept. 2026. DOI: 10.48550/arXiv.2605.23042.

[W31] Lewis, Ryan Hardesty, Hickert, Cameron, Xie, Yuanchang, Vinitsky, Eugene, **Wu, Cathy**. "CrossRoads: A Robustness Benchmark for Driving Foundation Models on Roadway Counterfactuals". In review.

[W32] Hickert, Cameron, Li, Sirui, He, Zhengbing, **Wu, Cathy**. "Probability-Aware Parking Selection". *IEEE Transactions on Intelligent Transportation Systems (T-ITS)* (Feb. 2026). DOI: 10.1109/TITS.2025.3649170. **MIT Home Page Feature**.

[33] Benjamins, Carolin, Eimer, Theresa, Schubert, Frederik, Mohan, Aditya, Döhler, Sebastian, Biedenkapp, André, Rosenhahn, Bodo, Hutter, Frank, Lindauer, Marius. "Contextualize Me – The Case for Context in Reinforcement Learning". *Transactions on Machine Learning Research* (June 2, 2023). DOI: 10.48550/arXiv.2202.04500.

[34] Modi, Aditya, Jiang, Nan, Singh, Satinder, Tewari, Ambuj. "Markov Decision Processes with Continuous Side Information". *Proceedings of Algorithmic Learning Theory*. PMLR, Apr. 2018, pp. 597–618. URL: https://proceedings.mlr.press/v83/modi18a.html.

[35] Hallak, Assaf, Di Castro, Dotan, Mannor, Shie. *Contextual Markov Decision Processes*. Feb. 2015. DOI: 10.48550/arXiv.1502.02259.

[W36] Jayawardana, Vindula, Freydt, Baptiste, Qu, Ao, Hickert, Cameron, Yan, Zhongxia, **Wu, Cathy**. "IntersectionZoo: Eco-driving for Benchmarking Multi-Agent Contextual Reinforcement Learning". *International Conference on Learning Representations (ICLR)*. Apr. 2025. URL: https://openreview.net/forum?id=XoulHHQGFi. [32% acceptance rate].

[W37] Zheng, Han, Ma, Yining, Gunasekaran, Karthick, Balaji, Bharathan, Du, Zheng, Vitaladevuni, Shiv, **Wu, Cathy**. "Internalizing Curriculum Judgment for LLM Reinforcement Fine-Tuning". DOI: 10.48550/arXiv.2605.11235. In review.

[38] *Reinforcement Fine-Tuning (RFT) with Amazon Nova models - Amazon Nova*. June 2026. URL: https://docs.aws.amazon.com/nova/latest/nova2-userguide/nova-reinforcement-fine-tuning.html#nova-rft-adaptive-curriculum.

[39] Schulman, John, Wolski, Filip, Dhariwal, Prafulla, Radford, Alec, Klimov, Oleg. *Proximal Policy Optimization Algorithms*. 2017. DOI: 10.48550/arXiv.1707.06347.

[40] McLean, Reginald, Chatzaroulas, Evangelos, Terry, J. K. Woungang, Isaac, Farsad, Nariman, Castro, Pablo Samuel. "Multi-Task Reinforcement Learning Enables Parameter Scaling". *Reinforcement Learning Conference*. May 2025. URL: https://openreview.net/forum?id=eBWwBIFV7T.

[41] Srinivas, Niranjan, Krause, Andreas, Kakade, Sham M. Seeger, Matthias W. "Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting". *IEEE Transactions on Information Theory* 58.5 (May 2012), pp. 3250–3265. DOI: 10.1109/TIT.2011.2182033.

[W42] Cho, Jung-Hoon, Li, Sirui, Kim, Jeongyun, **Wu, Cathy**. "Temporal Transfer Learning for Traffic Optimization with Coarse-grained Advisory Autonomy". *IEEE Transactions on Robotics (T-RO)* (2025). DOI: `10.1109/TRO.2025.3636819`.

[W43] Zhou, Tianyue, Cho, Jung-Hoon, **Wu, Cathy**. "Structure Detection for Contextual Reinforcement Learning". *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. Jan. 2026. URL: `https://ojs.aaai.org/index.php/AAAI/article/view/40137`. [17.6% acceptance rate].

[44] Silver, Tom, Allen, Kelsey, Tenenbaum, Josh, Kaelbling, Leslie. *Residual Policy Learning*. Jan. 3, 2019. DOI: `10.48550/arXiv.1812.06298`. Pre-published.

[45] Johannink, Tobias, Bahl, Shikhar, Nair, Ashvin, Luo, Jianlan, Kumar, Avinash, Loskyll, Matthias, Ojea, Juan Aparicio, Solowjow, Eugen, Levine, Sergey. "Residual Reinforcement Learning for Robot Control". *2019 International Conference on Robotics and Automation (ICRA)*. 2019 International Conference on Robotics and Automation (ICRA). May 2019, pp. 6023–6029. DOI: `10.1109/ICRA.2019.8794127`.

[46] Treiber, Martin, Hennecke, Ansgar, Helbing, Dirk. "Congested Traffic States in Empirical Observations and Microscopic Simulations". *Physical Review E: Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics* 62.2 (Aug. 2000), pp. 1805–1824. DOI: `10.1103/PhysRevE.62.1805`.

[47] Eckhoff, David, Halmos, Bastian, German, Reinhard. "Potentials and Limitations of Green Light Optimal Speed Advisory Systems". *2013 IEEE Vehicular Networking Conference*. IEEE. IEEE, 2013, pp. 16–18. DOI: `10.1109/VNC.2013.6737596`.

[W48] Jayawardana, Vindula, Li, Sirui, **Wu, Cathy**, Farid, Yashar, Oguchi, Kentaro. "Generalizing Cooperative Eco-driving via Multi-residual Task Learning". *2024 IEEE International Conference on Robotics and Automation (ICRA)*. Yokohama, Japan: IEEE, May 2024, pp. 6836–6842. DOI: `10.1109/ICRA57147.2024.10610586`.

[W49] Jayawardana, Vindula, Li, Sirui, Farid, Yashar, **Wu, Cathy**. "Multi-Residual Mixture of Experts Learning for Cooperative Control in Multi-vehicle Systems". DOI: `10.48550/arXiv.2507.09836`. In review.

[W50] Zhou, Jianan, Cho, Jung-Hoon, Zhou, Tianyue, Zheng, Han, Zhang, Jie, Dong, Roy, Ma, Yining, **Wu, Cathy**. "Task Specialization Fine-Tuning for Contextual Reinforcement Learning". DOI: `10.48550/arXiv.2608.17180`. In review.

[W51] Cho, Jung-Hoon, Zhang, Heling, Du, Siqi, Dong, Roy, **Wu, Cathy**. "Formalizing Task-Space Complexity for Zero-Shot Generalization". *Learning for Dynamics & Control Conference (L4DC)*. June 2026. URL: `https://arxiv.org/abs/2606.20967`.

[W52] Cho, Jung-Hoon, Zhang, Heling, Du, Siqi, Dong, Roy, **Wu, Cathy**. "Formalizing and Estimating Task-Space Complexity for Zero-shot Generalization". In review [Revise and Resubmit as Technical Note, IEEE Transactions on Automatic Control (TAC)].

[53] Bengio, Yoshua, Lodi, Andrea, Prouvost, Antoine. "Machine Learning for Combinatorial Optimization: A Methodological Tour d'Horizon". *European Journal of Operational Research* 290.2 (Apr. 2021), pp. 405–421. DOI: `10.1016/j.ejor.2020.07.063`.

[54] Amos, Brandon. "Tutorial on Amortized Optimization". *Foundations and Trends in Machine Learning* 16.5 (June 2023), pp. 592–732. DOI: `10.1561/2200000102`.

[55] Helsgaun, Keld. "An Extension of the Lin-Kernighan-Helsgaun TSP Solver for Constrained Traveling Salesman and Vehicle Routing Problems". *Roskilde: Roskilde University* 12 (2017), pp. 966–980. URL: `http://webhotel4.ruc.dk/~keld/research/LKH-3/`.

[W56] Li*, Sirui, Yan*, Zhongxia, **Wu, Cathy**. "Learning to Delegate for Large-scale Vehicle Routing". *Advances in Neural Information Processing Systems (NeurIPS)*. Vol. 34. Dec. 2021, pp. 26198–26211. URL: https://openreview.net/forum?id=rm0I5y2zkG8. [26% acceptance rate] **Spotlight (3%)**.

[W57] Yan, Zhongxia, **Wu, Cathy**. "Neural Neighborhood Search for Multi-agent Path Finding". *International Conference on Learning Representations (ICLR)*. 2024. URL: https://openreview.net/forum?id=2NpAw2QJBY. [31% acceptance rate] **MIT Home Page Feature**.

[W58] Li*, Sirui, Ouyang*, Wenbin, Paulus, Max, **Wu, Cathy**. "Learning to Configure Separators in Branch-and-Cut". *Advances in Neural Information Processing Systems (NeurIPS)*. Dec. 2023. URL: https://openreview.net/forum?id=gf5xJVQS5p. [26% acceptance rate] **MIT Home Page Feature**.

[59] Erdmann, Michael, Lozano-Pérez, Tomás. "On Multiple Moving Objects". *Algorithmica* 2.1 (Nov. 1, 1987), pp. 477–521. DOI: 10.1007/BF01840371.

[W60] Huang, Peixin, Wu, Yaoxin, Ma, Yining, **Wu, Cathy**, Song, Wen, Zhang, Wei. "A General Neural Backbone for Mixed-Integer Linear Optimization via Dual Attention". *International Conference on Machine Learning (ICML)*. July 2026. URL: https://openreview.net/forum?id=fBAMvSNH8N.

[61] Andreas Horni, Kai Nagel, and Kay W. Axhausen, eds. *The Multi-Agent Transport Simulation MATSim*. London: Ubiquity Press, 2016.

[62] Bertsekas, Dimitri P. "Dynamic Programming and Optimal Control". *Belmont, MA: Athena Scientific* 1 (2005).

[63] MIT Civil and Environmental Engineering, *CEE Classes*. May 2026. URL: https://www.youtube.com/playlist?list=PL2oXDPa_CEVHjtvaSlzXU5QvgXkXcY6G3.

[W64] RERITE, *REproducible Research In Transportation Engineering – RERITE Working Group*. URL: https://www.rerite.org/.

[W65] Ji*, Junyi, Lu*, Ruth, Belkessa, Linda, Wang, Liming, Varotto, Silvia, Dong, Yongqi, Saunier, Nicolas, Ameli, Mostafa, Macfarlane, Gregory S. Madadi, Bahman, **Wu, Cathy**. "Measuring the State of Open Science in Transportation Using Large Language Models". DOI: 10.48550/arXiv.2601.14429. In review [Major Revision, Transportation Research Part C: Emerging Technologies (TR-C)].

[66] Pineau, Joelle, Vincent-Lamarre, Philippe, Sinha, Koustuv, Lariviere, Vincent, Beygelzimer, Alina, Buc, Florence, Fox, Emily, Larochelle, Hugo. "Improving Reproducibility in Machine Learning Research (A Report from the NeurIPS 2019 Reproducibility Program)". *Journal of Machine Learning Research* 22.164 (2021), pp. 1–20. URL: http://jmlr.org/papers/v22/20-303.html.

[W67] **Wu, Cathy**, Ghosh, Bidisha, Zheng, Zuduo, Martínez, Irene. *Reproducibility in Transportation Research: A Hands-on Tutorial*. 2024. URL: https://www.rerite.org/itsc24-rr-tutorial/.

[W68] Zheng, Zuduo, Levinson, David, Zhou, Xuesong, Mohammadian, Saeed, **Wu, Cathy**, Ghosh, Bidisha, Martínez, Irene. *Reproducibility in Transportation Research: A Hands-on Tutorial 2.0 – Reproducible Manuscripts*. 2025. URL: https://zuduozheng.github.io/RERITE_TUTORIAL_2.0/.

[W69] *Reproducible Research Prize at hEART 2026 – RERITE Working Group*. June 2026. URL: https://www.rerite.org/heart-2026-prize.html.

[W70] *Reproducible Research in Transportation*. June 2026. URL: https://link.springer.com/collections/dgachijbgi.

[W71] *Call for Interest: Transportation Evaluations & Datasets (Dual Special Issue with TR-C and T-ITS)*. June 2026. URL: https://forms.gle/7ggEkAPZe6QKrNoG7.

[W72] Ji*, Junyi, Lu*, Ruth, Belkessa, Linda, Wang, Liming, Varotto, Silvia, Dong, Yongqi, Saunier, Nicolas, Ameli, Mostafa, Madadi, Bahman, **Wu, Cathy**. *Open Science Explorer for Transportation Research*. Mar. 2026. URL: https://www.rerite.org/MOST/explorer.html.

[W73] *Workshop on Open Science for Intelligent Vehicles | IEEE IV 2026*. June 2026. URL: https://www.rerite.org/ivs-2026.

[W74] *RLJ | RLC 2025 Submission Instructions*. June 2026. URL: https://rl-conference.cc/2025/submissionInstructions.html.