Introduction. The deployment of AI in autonomous systems—such as self-driving cars, drones, and robotic manipulators—remains challenging due to safety risks, the cost of real-world data, and the difficulty of specifying precise task objectives. As a result, development increasingly relies on digital twins and high-fidelity simulations for safe and efficient training of reinforcement learning (RL) agents and planning. Yet, deployment still faces challenges like the sim-to-real gap, sensitivity to misspecifications, and misaligned objectives. These issues motivate the following research objectives:

- **RO 1** Using abstractions and multifidelity simulations to accelerate learning while preserving real-world relevance through alleviating sim-to-real gap in RL
- **RO 2** Ensuring robustness to uncertainty, noise, and model error across the learning pipeline
- **RO 3** Leveraging structured objectives to align with desired behavior and support learning

My Research. My research contributes to the theoretical and algorithmic foundations necessary to address the challenges outlined above. While the three research questions may appear distinct, I have found that they are deeply interconnected and benefit from being addressed jointly. As such, most of my work aligns with a primary research objective (RO) while also supporting a secondary one.

RO 1: In [1], I study how state abstractions with formal performance guarantees can be used to efficiently compute robust, near-optimal policies. I demonstrate that seemingly similar approaches to leveraging low-fidelity simulations can yield vastly different results due to fundamental phenomena emerging in autonomous systems.

RO 2: My work on constrained decision-making [2] develops efficient methods to characterize and bound the set of initial distributions and constraint thresholds under which a policy is robust, i.e. remains feasible and near-optimal. In [3], we show that incorporating structured reward priors enhances robustness to model misspecification. Similarly, in [5], I show that using dynamics learning as an auxiliary task improves robustness to observation noise in reinforcement learning.

RO 3: In [3], we study inverse reinforcement learning under known structural reward priors. By incorporating structural properties common in autonomous systems, we improve reward identifiability and develop efficient algorithms to enforce these constraints. Furthermore, since many autonomous tasks are not well-modeled by single-objective formulations, I address gaps in multiobjective and constrained systems in [2] and [4]. Finally, in [6], I propose a method for synthesizing subtask progressions, which improves both interpretability and efficiency in solving complex tasks.

Application to Autonomous Driving: In parallel with my theoretical work, I am actively involved in applied research on autonomous driving. I am developing a full pipeline for designing, testing, and deploying control algorithms in collaboration with Mcity—the University of Michigan's advanced testbed for connected and automated vehicles. Mcity provides vehicles, infrastructure, high-fidelity simulations (running on HPC clusters), and a digital twin environment for realistic validation. Our team recently won an NSF-supported competition, earning enhanced access to Mcity's remote testing capabilities. This applied work grounds my research in real-world challenges and informs my efforts on robustness, abstraction, and reward learning.

Research Plan. I plan to continue my computational and algorithmic research towards the three research objectives. Our foundational preliminary work summarized above has focused on proof-of-concept experimentation. In the remaining two years of my PhD studies, I will focus on developing novel algorithms for and running experiments with high-dimensional sensory data and larger continuous state spaces for which both building multifidelity abstractions algorithmically and designing policies are computationally intensive.

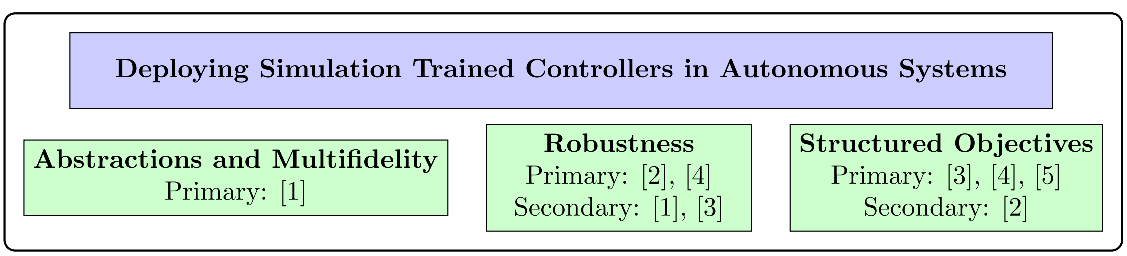

Figure 1: Summary of My Research and Proposed Work

References: [1] **Tercan**, Ozay, "On the relation of bisimulation, model irrelevance, and corresponding regret bounds" NeurIPS 2025 Workshop on Aligning Reinforcement Learning Experimentalists and Theorists. [2] **Tercan**, Ozay, "Initial Distribution Sensitivity of Constrained Markov Decision Processes" CDC 2025 . [3] Shehab*, **Tercan***, Ozay, "Efficient Reward Identification In Max Entropy Reinforcement Learning with Sparsity and Rank Priors" CDC 2025. [4] **Tercan**, Anderson "Increased Reinforcement Learning Performance through Transfer of Representation Learned by State Prediction Model" IJCNN 2021. [5] **Tercan**, Prabhu, "Thresholded Lexicographic Ordered Multiobjective Reinforcement Learning" ECAI 2024. [6] **Tercan** et al. "Synthesizing a Progression of Subtasks for Block-Based Visual Programming Tasks" AAAI 2024 Workshop on AI for Education.