Google Scholar
GitHub

Manav Vora

(217) 205-7972
mkvora2@illinois.edu

Education

University of Illinois Urbana Champaign (UIUC) **Aug 2022 – Present**

- Ph.D. candidate in Aerospace Engineering (Robotics Focus). GPA: 3.8
- *Advisor:* Prof. Melkior Ornik
- Graduate Coursework: Reinforcement Learning; Advanced Computational Topics in Robotics; 3D Vision; Optimal Aerospace Systems; Estimation of Dynamical Systems.

Indian Institute of Technology Bombay **Aug 2018 – Aug 2022**

- B.Tech. (Honors) in Aerospace Engineering.
- Minor in Systems and Controls Engineering.
- Undergraduate Coursework: Stochastic Control; Foundations of Reinforcement Learning; Adaptive and Learning Control Systems; Navigation and Guidance; Motion Planning and Coordination of Autonomous Vehicles.

Objective

Scalable and Constrained Multi-Agent Planning and Robotics, using Meta-Learning, Deep Reinforcement Learning, Vision–Language Models, and Diffusion Models.

Publications and Preprints

1. **Vora, M.**, P. Thangeda, M. N. Grussing, and M. Ornik. 2023. "Welfare Maximization Algorithm for Solving Budget-Constrained Multi-Component POMDPs." *IEEE Control Systems Letters* 7: 1736–1741.

2. Puthumanaillam, G.*, **M. Vora***, and M. Ornik. 2024. "ComTraQ-MPC: Meta-Trained DQN-MPC Integration for Trajectory Tracking with Limited Active Localization Updates." In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

3. **Vora, M.**, M. N. Grussing, and M. Ornik. 2024. "Solving Truly Massive Budgeted Monotonic POMDPs with Oracle-Guided Meta-Reinforcement Learning." *arXiv preprint* arXiv:2408.07192. (Under review)

4. **Vora, M.**, I. Shomorony, and M. Ornik. 2024. "Capacity-Aware Planning and Scheduling in Budget-Constrained Monotonic MDPs: A Meta-RL Approach." *arXiv preprint* arXiv:2410.21249. (Under review)

5. Puthumanaillam, G.*, **M. Vora***, P. Thangeda, and M. Ornik. 2024. "A Moral Imperative: The Need for Continual Superalignment of Large Language Models." *arXiv preprint* arXiv:2403.14683.

6. Puthumanaillam, G.*, **M. K. Vora***, T. Shafa, Y. Li, M. Ornik, and S. Mitra. 2024. "Assured Collision Avoidance for Learned Controllers: A Case Study of ACAS Xu." In *AIAA SCITECH 2024 Forum*, 1168.

**Equal contribution.*

Research Experience

Research Assistant **LEADCAT Lab, UIUC** **Aug 2022 – Present**

Massive Multi-Agent Systems [1], [3], [4]

- Developed scalable algorithms for **massive multi-agent** systems under resource and capacity constraints, introducing budgeted-POMDPs and hierarchical planning approaches.
- Applied advanced machine learning techniques—including **meta-reinforcement learning and random forest regression**—to achieve **linear scalability** in problems with up to $\mathbf{10^{10000}}$ **states** and $\mathbf{3^{10000}}$ **actions**.
- Outperformed existing methods and industry baselines by over **70%** in infrastructure maintenance efficiency and over **90%** in large-scale robot swarm maintenance.

Robot Navigation in Uncertain Environments [2]

- Developed **ComTraQ-MPC**, a novel framework integrating **Deep Q-Networks (DQN) and Model Predictive Control (MPC)** to improve trajectory tracking in partially observable, stochastic and GPS-denied environments with limited active localization updates.
- Demonstrated significant improvements in operational efficiency and accuracy, validated through simulations and hardware experiments.
- Developing a universal algorithm for autonomous navigation, enabling robots to perform context-aware self-reasoning—understanding their strengths, limitations, and mission demands—to optimize real-time decision-making using **self-reasoning Vision-Language Models (VLMs) and diffusion models.**

Large Language Models and AI Alignment [5]

- Investigated the challenges of achieving life-long **superalignment** in **Large Language Models (LLMs)**, highlighting inherent limitations in adapting to the dynamic nature of human ethics and evolving global scenarios.
- Proposed **strategies** to mitigate alignment discrepancies between static AI models and dynamic human societies.

Verification of Neural Network Control Systems [6]

- Introduced an **integrated verification pipeline** by combining the **nnenum framework** with the **Verse** toolkit, closing the gap between neural network output analysis and system dynamics.
- Achieved robust safety guarantees for **ACAS Xu** by **iteratively verifying** network decisions and propagating system states, maintaining accurate system representations over time.
- Demonstrated **scalability and adaptability** across a range of **autonomous systems**, effectively handling non-linear dynamics and complex operational conditions.

Research Intern **GRASP Lab, UPenn** **Apr 2021 – Feb 2022**

Neural Network-Based ODEs and Lie Algebra for Robotic Control Systems

- Explored **Lie Algebra**-based Neural ODE architectures to enhance dynamics prediction.
- Demonstrated **comparable baseline performance** on RRT-star for Dubins' Car with Lie Algebra-generated dynamics.
- Investigated **Hamiltonian-based Neural ODEs** for pendulum and quadrotor systems, leveraging underlying system structure.

Honors and Awards

- IROS 2024: Selected among the top 10% of accepted papers for oral presentation.
- IEEE CDC Student Travel Award
- UIUC Aerospace Engg. Department Fellowship (2023) for research & academic achievements
- Chanakya Research Fellowship (2021) for research on autopilot design for delivery drone

Languages and Frameworks

- C++; C; Python; Julia; MATLAB; Bash
- ROS/ROS2; PyTorch; TensorFlow; OpenAI Baselines; OpenCV; Gazebo; SUMO; Simulink; Git;

Research Projects

Vision-based Spacecraft Trajectory Estimation

- Developed a methodology for spacecraft trajectory estimation using **2D images** captured by onboard cameras to reconstruct **3D spatial paths**, focusing on the OSIRIS-REx mission's Departure Flyby phase.
- Implemented **Structure from Motion (SfM)** techniques, integrating **SIFT** for feature detection and **RANSAC** for outlier mitigation, and evaluated trajectory estimation using global and incremental SfM via **OpenMVG**, achieving high precision compared to actual mission trajectory data.

RL-Framework for Intelligent Traffic Light Control

- Created a traffic light intersection with incoming and outgoing traffic using the **SUMO simulator**.
- Programmed and evaluated the performance of **DQN and PPO reinforcement learning** agents using metrics such as **average wait time and average queue length**.