

# Vineet Pasumarti

[✉ vineetpasumarti@gmail.com](mailto:vineetpasumarti@gmail.com) | [🏡 vineetpasumarti.github.io](https://vineetpasumarti.github.io) | [🔗 vineetpasumarti](https://vineetpasumarti.com) | [vineetpasumarti](https://vineetpasumarti.com)

## Summary

I'm a robotics researcher with a strong background in robot learning for agile high-speed autonomy. Experienced in reinforcement learning, multi-agent systems, and sim-to-real for agile quadrotors, with prior industry experience at Boeing Research on flight dynamics. At Penn, I host the Robot Learning Reading Group with my advisor to read and discuss state-of-the-art robot learning literature. **US citizen with Secret clearance.**

## Education

### University of Pennsylvania

M.S.E. IN ROBOTICS, GPA: 4.0/4.0

Philadelphia, PA

Aug. 2024 – May 2026

- **Thesis:** *Vision-Based Multi-Agent Drone Racing using Deep Reinforcement Learning* advised by Dr. Antonio Loquercio

### University of Wisconsin-Madison

B.S. IN MECHANICAL ENGINEERING WITH HONORS, GPA: 3.671/4.0

Madison, WI

Sep. 2019 – May 2023

- Awarded Faustin Prinz Undergraduate Research Fellowship due to research contributions and academic merit.

## Experience

### Janus Intelligent Robots Lab – GRASP Laboratory, University of Pennsylvania

Philadelphia, PA

GRADUATE STUDENT RESEARCHER (ADVISOR: ANTONIO LOQUERCIO)

Aug. 2024 – Present

- Led the execution of a research project (first-author submission to ICRA 2026) that formulates autonomous drone racing as a multi-agent reinforcement learning problem.
- Discovered the emergence of strategic drone racing behaviors, like overtaking, blocking, and attacking, as artifacts of long-horizon multi-agent policy optimization over sparse rewards.
- Built a custom Isaac Lab environment and bridged sim2real gap for agile quadrotor policies by enhancing simulation accuracy with physics-based motor dynamics, aerodynamic drag modeling, cascaded collective thrust and body rate control (CTBR), and domain randomization, then implementing a custom low-level controller in C within the brushless CrazyFlie 2.1 firmware for deployment.
- Developed recurrent RL methods that overcome the assumption of stationary transition dynamics in Markov Decision Processes, enabling smooth real-time decision-making for competitive multi-agent systems.
- Led the architectural design and ablation study of hybrid recurrent-feedforward policy networks with Feature-wise Linear Modulation layers for external conditioning.
- Designed a shared critic architecture in PyTorch with separate value heads to enable competitive agent-specific value estimates.

#### TEACHING ASSISTANT

Aug. 2025 – Present

- Fall 2025: ESE 651 Physical Intelligence (Deep RL, Imitation Learning, Behavior Models) with Dr. Antonio Loquercio

- Spring 2026: ESE 650 Learning in Robotics (Probabilistic Robotics, State Estimation, Optimal Control, RL) with Dr. Pratik Chaudhari

### Boeing Research and Technology

Huntington Beach, CA

RESEARCH ENGINEER II – LOADS AND DYNAMICS

Jul. 2023 – Jan. 2025

- Primary researcher on an Independent Research and Development (IRAD) to best represent the constituent material properties of carbon-matrix-ceramics (CMCs) for hypersonics.
- Responsible for loads envelope during flight-testing of the T7 Red Hawk trainer fighter jet at Edwards Air Force Base.
- Wrote code base for internal tools that allow structural analysts to parametrically build composite wing stringer co-bond configurations with wrinkles for Boeing 777X commercial airliner, reducing design+analysis iteration time by 60%.
- Designed and performed loads analysis for components on a hypersonic vehicle test bed, a satellite propulsion system, and a composite Mars habitat. All projects required rigorous GD&T and strict design standards.

### NASA Langley Research Center

Hampton, VA

VEHICLE ENGINEERING INTERN – ENTRY, DESCENT, AND LANDING

Jun. 2021 – Aug. 2021

- Designed collapsible drag systems for heat shields that deploy upon re-entry of the atmosphere to support sample-return missions. Validated via comprehensive simulations of loads at varying Mach numbers and angles-of-attack.

VEHICLE ENGINEERING INTERN – MECHANISMS

Jun. 2020 – Aug. 2020

- Collaborated in a team of 4 to develop a collapsible mobile lunar crane concept for the manipulation and offloading of large payloads to establish an outpost and permanent human presence on the lunar surface.
- Personally designed a hinge-rail hybrid structure to enable collapsible functionality without restricting motion of the crane trolley under load.

### Computational Flow Physics Modeling Lab, University of Wisconsin-Madison

Madison, WI

UNDERGRADUATE STUDENT RESEARCHER (ADVISOR: JENNIFER FRANCK)

Jan. 2020 – May. 2023

- Led Computational Fluid Dynamics (CFD) investigations of cross-flow turbine arrays for ARPA-E funded project on renewable energy capture where I discovered wake-recovery phenomenon in confined turbine arrays.
- Published and presented as first author at APS Division of Fluid Dynamics 2022 and AIAA Aviation 2022 on turbine interaction dynamics and confinement exploitation for hydrokinetic power generation.
- Wrote C++ code for our in-house CFD solvers, including custom OpenFOAM airfoil morphing library and 2D-to-3D blockage modeling techniques that reduce computational time while maintaining accuracy.

# Projects

---

## Learning to Overtake in Drone Racing via Recurrent RL

PERSONAL PROJECT - ESE 650: LEARNING IN ROBOTICS

- Proved that temporal memory is critical for end-to-end quadrotor control for non-stationary multi-agent environments.
- Designed a two-stage PPO-based training pipeline for overtaking in head-to-head drone racing on a life-size 3D Figure-8 track in simulation.
- Demonstrated that feedforward MLP policies with full opponent state fail due to lack of temporal reasoning, even with increased network capacity.
- Proposed a recurrent policy architecture (2-layer LSTM + 4-layer MLP) that encodes opponent dynamics over time, improving strategic overtaking behavior.
- Augmented the observation space with full opponent state and evaluated policies under dense and sparse competition rewards to promote tactical behavior.
- Achieved 85.7% gate success and 50% sustained overtake rate using the recurrent policy, outperforming all MLP baselines across metrics including crash rate and lap completion.

## Autonomous VIO-Based Quadrotor Navigation and Control

TEAM PROJECT - MEAM 620: ADVANCED ROBOTICS

- Developed complete classical autonomy stack for quadrotor navigation using Visual-Inertial Odometry (VIO) with an Error State Extended Kalman Filter (ES-EKF) for onboard state estimation.
- Implemented stereo vision fusion with IMU data to achieve optimal state estimation by strategically tuning process noise covariance to weight visual measurements over inertial predictions and prevent drift-induced hover errors.
- Designed geometric nonlinear PD controller with fine-tuned gains achieving 2s settling time and 2% steady-state error on the CrazyFlie 2.0 platform.
- Integrated A\* path planning with Ramer-Douglas-Peucker sparsification and minimum-jerk trajectory generation, as well as optimized voxel resolution and safety margins for obstacle avoidance.
- Extended framework with an online replanning pipeline using a limited 5m sensor range and 7.5m planning horizon by implementing collision detection and dynamic replanning every 20 time steps for unknown environment exploration.

## 'Beat-the-Expert': An IL-RL Framework to Outperform Expert Racing Policies

TEAM PROJECT - ESE 680: REAL-WORLD ROBOT LEARNING

- Developed an imitation learning-reinforcement learning (IL-RL) framework combining HG-DAGGER and PPO to surpass expert racing performance on the F1TENTH platform
- Implemented Human-Gated DAGGER imitation learning bootstrap using 54-dimensional LiDAR observations and multi-layer perceptron policy networks to achieve 3x faster convergence than end-to-end RL
- Fine-tuned pre-trained IL policies with Proximal Policy Optimization, achieving 28.2% lap time improvement over easy expert demonstrations while maintaining sample efficiency
- Demonstrated policy generalization across multiple real-world F1TENTH tracks (Spielberg, Nürburgring, Spa-Francorchamps) with up to 100% lap completion rates
- Conducted comprehensive performance analysis comparing bootstrapped RL, end-to-end RL, and expert policies, contributing insights to autonomous racing control strategies

## Img2GPS: GPS Coordinate Prediction from Single Images via ResNet

TEAM PROJECT - CIS 519: MACHINE LEARNING

- Framed geolocation as a supervised learning task to regress 2D GPS coordinates from RGB input using deep residual networks.
- Collected a dataset of 741 images across diverse lighting, weather, and viewpoints in the Penn Engineering Quad, then extracted GPS labels from EXIF metadata and normalized coordinates for regression.
- Tuned batch size and learning rate through ablation: identified batch size 16 and learning rate 0.001 with learning rate scheduler (step size 5,  $\gamma = 0.1$ ) as optimal.
- Conducted in-depth analysis of learning rate–batch size interactions and architecture bias-variance behavior.

# Publications

---

**V. Pasumarti, L. Bianchi, A. Loquercio.** "Agile Flight Emerges from Multi-Agent Competitive Racing." **in-review at IEEE International Conference on Robotics and Automation (ICRA) 2026.**

**V. Pasumarti, M. Dave, J. Franck.** "Simulation of two cross-flow turbines under confinement." **AIAA AVIATION Forum 2022.**

**V. Pasumarti, M. Dave, J. Franck.** "Simulation of dual cross-flow turbines under confinement." **APS Division of Fluid Dynamics (APSDFD) 2022.**

# Awards & Honors

---

2022 **Faustin Prinz Undergraduate Research Fellowship**, \$5,000 tuition credit and \$2,000 research funding

University of Wisconsin–Madison

2020 **NASA Academy Award Recipient**, \$7,000 from Wisconsin Space Grant Consortium (WSGC)

NASA Langley Research Center

2022 **Graduated with Honors in Research**, Recognized for excellence in undergraduate research thesis

University of Wisconsin–Madison