

Swapnil Pande

804-869-8910 | swapnilpande.me | hello@swapnilpande.me | www.linkedin.com/in/SwapnilPande | www.github.com/SwapnilPande

PROFESSIONAL SUMMARY

Research engineer advancing learning-based control and planning for **embodied intelligence in the real world**. Conducts research **grounded in concrete application domains**, developing **reinforcement/imitation learning** methods to improve system performance in autonomous driving, manipulation, and mobile robotics. Brings deep expertise across embedded systems, mechanical design, and robotics software, enabling reasoning about learning across the **full hardware-software stack**. Experienced in independently structuring complex technical challenges, leading multi-disciplinary teams, and **driving research from concept to deployment**.

PROJECTS

Driving by Dreaming: Offline Reinforcement Learning for self-driving cars | *Master's Thesis, CMU* August 2022

- Developed a **model-based offline reinforcement learning** algorithm enabling policy learning entirely from logged driving data while being capable of outperforming the demonstrations.
- Proposed a neural dynamics model capable of long-horizon rollouts (500–1000 steps, compared to 1–5 in prior work), enabling policies to learn from simulated trajectories far out of the training distribution.
- Achieved **state-of-the-art benchmark performance** in the CARLA simulator with **25× less data and 10× faster training** than online RL baselines.
- Introduced a multi-agent self-play framework that improves policy **robustness in complex negotiation scenarios** (e.g., merging, right-of-way) and **increases data efficiency by over 10×** by collecting trajectories from all agents in parallel simulation.

Sim-to-Real Control of a Double Pendulum | *Every Flavor Robotics* Fall 2025 (In Progress)

- Configured **double pendulum system simulation** in Mujoco + MJX and trained an RL policy in Jax with PPO, successfully capable of dynamically transitioning between all stability modes.
- Engineered a **real-world counterpart** with sensing and motor controls running on a **real-time operating system**, streaming data to a computer executing the policy.
- Building a pipeline for **parameter identification and domain randomization** to achieve successful sim2real transfer.

Driver Route Optimization for Meal-Delivery Service | *Gourmet for Good* Spring 2025

- Formulated a route-optimization system for a meal-delivery service, framing delivery scheduling as a **constrained mixed-integer optimization problem** over up to 60 simultaneous orders with strict delivery-time and capacity constraints, with additional constraints on food temperature, order grouping, and food preparation time.
- Implemented, tuned, and benchmarked three optimizers to minimize total delivery cost: **simulated annealing, genetic algorithm, and cross-entropy method**.
- Validated performance using real delivery data from the Gourmet for Good ERP system, matching the performance of the human scheduler; **currently actively in use** for delivery optimization.

EtchBot: Ultra-High Performance Drawing System | *Every Flavor Robotics* Spring 2023

- Architected the full software stack for a **two-axis robotic drawing system** capable of rendering images on an Etch-a-Sketch in < 60 s, integrating computer vision, motion planning, and real-time motor control, with a dockerized server that powered **800+ consecutive live demos**.
- Devised a hybrid **trajectory optimizer** in Rust combining a **genetic algorithm** for global search with **graph-based refinement** to minimize travel distance and backtracking, processing **20k points drawings in under 10 s**.
- Implemented a closed-loop trapezoidal trajectory generator to synchronizes dual-axis motion with backlash compensation and **continuous replanning at 5 kHz**.

EXPERIENCE

Founder, CEO | *Every Flavor Robotics* February 2023 – August 2025

Robotics startup bridging research and industry by developing robust, accessible, and open-source hardware and software tools and sharing them through educational media.

- **Architected and maintained the firmware, Python API, and ROS2 interface** for MotorGo, an open-source robotic motor controller and sensor interface based on ESP32 and FreeRTOS; **deployed in Carnegie Mellon University classes**.
- Built and integrated **software stacks for 10+ robotic systems**, focused on **designing shared components**, such as motion-planning servers and reusable controller-tuning and teleoperation tools, for a **modular, scalable robotics framework**.

- **Led technical and research strategy**, prioritizing development across the MotorGo platform, robotics stack, and ML integration efforts to align engineering progress with community growth and long-term research goals.
- **Crafted the company's technical pitch** and **raised \$100K in friends-and-family funding** to expand the video production team and pursue larger, more ambitious robotics projects.
- **Spearheaded cross-functional collaboration** between engineering and creative teams, translating technical progress into accessible content and effective team coordination.

Robotics Contractor | Vanderbilt University

August 2022 – December 2022

Contract project with a Vanderbilt University research lab to prototype a robotic fleet for automated material handling in closed-loop experimental platforms.

- Jointly designed and **delivered five prototype mobile robots**, integrating custom motor controllers, 3D-printed mechanical assemblies, and commercial sensors to establish the foundation for future wet-lab automation research.
- Owned and engineered the **full software stack**, including navigation, **hybrid localization from intermittent sensor data**, and UI for control of multiple robots.
- Conceived, prototyped, and characterized an **ultra-low-cost IR localization and docking system (<\$1 in components)**, applying analytical modeling and experimental validation to achieve accurate pose resets for payload transfer.

Research Intern, Motion Planning | Argo AI

January 2022 – August 2022

Joined the motion planning and prediction research team exploring data-driven planning methods for autonomous vehicles.

- Applied methods from the Driving by Dreaming project to Argo's autonomous driving stack, **integrating model-based offline RL** pipelines trained and evaluated on **real-world driving data** within the company's motion planning framework.
- Collaborated with engineers across multiple teams as one of the first members of the learned motion planning team, evaluating the **feasibility of integrating learning-based approaches** within Argo's production autonomy stack.

Graduate Research Assistant | Auton Lab, Carnegie Mellon University

September 2020 – August 2022

Researched under Prof. Jeff Schneider at the CMU-Argo AI Center for Autonomous Vehicle Research.

- Explored **data-efficient reinforcement learning for autonomous driving**, contributing to publications on SAC variants and offline RL, and leading the Driving by Dreaming project on model-based offline learning from logged driving data.
- Maintained and extended the lab's shared CARLA reinforcement learning codebase, **supporting reproducible benchmarking** across three standardized benchmarks and experimentation across five concurrent research projects.
- **Mentored two undergraduate researchers** on reinforcement learning experiments, guiding experimental design and implementation.

Visiting Student | Carnegie Mellon University

May 2019 – August 2019

Worked under Prof. Howie Choset in the Biorobotics Lab.

- Conducted early research on multi-agent and model-based reinforcement learning for **cooperative decision-making**, contributing to methods for data-efficient policy learning in the Hanabi challenge.

TECHNICAL SKILLS

Languages : Python, C/C++, Rust, Bash

Learning & Simulation : PyTorch, JAX, Mujoco + MJX, Gymnasium, OpenCV

Robotics Stack : ROS/ROS2, Nav2, RViz, Foxglove

Embedded Systems : ESP32, FreeRTOS, UART/SPI/I2C, Real-Time Control Loops

Systems & Tooling : Docker, Linux, Git, Conda/uv, CI/CD (GitHub Actions)

SELECTED PUBLICATIONS

C. Igoe, **S. Pande**, S. Venkatraman, et al. "**Multi-Alpha Soft Actor-Critic: Overcoming Stochastic Biases in Maximum Entropy Reinforcement Learning**." *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.

A. Villafior, Z. Huang, **S. Pande**, et al. "**Addressing Optimism Bias in Sequence Modeling for Reinforcement Learning**." *Proceedings of the 39th International Conference on Machine Learning (ICML)*, 2022.

S. Pande. "**Driving by Dreaming: Offline Model-Based Reinforcement Learning for Autonomous Vehicles**." *Master's Thesis, Carnegie Mellon University*, 2022.

EDUCATION

Carnegie Mellon University

MS in Robotics | GPA: 4.14

August 2020 – August 2022

Vanderbilt University

BE in Mechanical Engineering | GPA: 3.98, *Summa Cum Laude*

August 2016 – May 2020