

Hrishikesh Pawar

📍 Worcester, MA 📞 +1(774)-519-2900 ✉️ hpawar@wpi.edu 🌐 [linkedin.com/in/hp39](https://www.linkedin.com/in/hp39) 🔗 hrishikesh-st.github.io

Education

Worcester Polytechnic Institute

August 2023 - May 2025

Master of Science in Robotics Engineering; GPA: 4.0/4.0

Worcester, MA

Courses: CS541-Deep Learning, RBE549-Computer Vision, RBE595-Vision based Robot Manipulation, REB550-Motion Planning

Teaching Assistant: RBE595-Vision based Robot Manipulation, RBEBL01-Robots for Recycling

Experience

PeARlab, WPI | Research Collaborator advised by Prof. Nitin Sanket

August 2024 – Present

- Developed **adaptive optics** with **event cameras** for navigation, adjusting focal length and aperture to obtain depth cues.
- Implementing RL policy for adaptive control of optical parameters, improving navigation efficiency and obstacle avoidance.
- Built a custom simulator using **Gaussian Splats** generating high frame-rate, realistic frames and events, enabling HITL testing.

Nobi | AI Software Intern | Remote

May 2024 - August 2024

- Led R&D efforts for smart ceiling lamps improving real-time fall detection and emergency response in elderly care.
- Engineered a **rotation-aware** detection model by integrating **Swin Transformers** with **Hausdorff distance matching** and **Adaptive Query Denoising**, improving accuracy from **85 mAP** to **92 mAP**.
- Experimented with **LLaVa** and **CLIP**, applying **LoRA** for **PEFT** fine-tuning for vision-language detection task generalization.
- Reduced deployment time by **30%** by automating end-to-end deployment pipeline using **Jenkins**, **Kubernetes** and **Docker**.

Adagrad AI | Computer Vision Engineer | Pune, India

November 2020 - July 2023

- Worked on R&D for **Gate-Guard**, an edge-based Boom Barrier system using **Automatic License Plate Recognition (ALPR)**.
- Developed data collection, training and deployment pipelines for lightweight object detection models: **Yolo-X** and **Yolo-v5**.
- Achieved accuracy of **97%** for four-wheelers and **95%** for two-wheelers with **50 fps** throughput on **Nvidia Jetson-TX2**.
- Designed interactive analytics and monitoring services using **Django**, **Azure**, **WebSockets**, **Kafka**, **Celery** and **Redis**.
- Deployed across **200+** private sites, **30+** government-based BRT (Bus Rapid Transport) and **5** smart city sites.
- Featured links: [Indian Express](#) | [Times of India](#) | [YouTube](#)

Projects

Deep Visual-Inertial Odometry | [GitHub](#) | [Link to project](#)

- Won the **best project** award for the **Computer Vision course (RBE-549)** at **WPI** amongst **40** graduate students.
- Developed a VIO stack, coupling **CNNs** for image processing and **LSTMs** for sequential inertial data enhancing pose estimation.
- Generated synthetic datasets in **Blender** simulating trajectories with realistic sensor noise to improve model robustness.
- Trained Visual-Inertial models combining **MSE** and **Geodesic Loss** reducing the RMSE **Absolute Trajectory Error** by **28%**.

Classical Structure from Motion Pipeline | [GitHub](#)

- Developed SfM pipeline starting with foundational two-view geometry and scaling to handle multiple views (total **five** views)
- Implemented **PnP** with **Lavenberg-Marquardt** optimization, reducing the reprojection errors across multiple views.
- Reduced reprojection error by **42%** using **bundle adjustment** with **sparse Jacobians** and **Trust Region Reflective optimization**.

Zero-Shot Semantic Neural Style Transfer for Images | [GitHub](#)

- Implemented **AdaAttN** module enabling per-point attentive normalization for **zero-shot semantic neural style transfer**.
- Led ablation studies showcasing the removal of deeper attention layer decreases training time by **13%** without loss of quality.
- Integrated **AdaAttN** with **CLIPSeg**'s prompt-based segmentation, enabling user-defined, localized neural style transfer.

Skills

- **Languages:** Python, C, C++, SQL
- **Frameworks:** Pytorch, OpenCV, TensorFlow, ONNX, Numpy, Pandas, Django, Flask, Celery, Kafka, Matlab, CMake, Docker, TensorRT, Jenkins, Kubernetes, Git, ROS, ROS2