Ashwattha Phatak

Raleigh, NC | ashwatthap@gmail.com | +1 919-971-3500 | linkedin.com/in/ashwatthaphatak | github.com/ashwatthaphatak

Education

North Carolina State University - Raleigh, NC Aug 2024 – May 2026
Masters in Computer Science **GPA: 3.91/4.0**
Coursework: Adv. Distributed Sys., Operating Sys., Machine Learning, Computer Networks, Autonomous Systems
Vishwakarma Institute of Technology - Pune, India Aug 2019 – May 2023
Bachelors in Electronics and Telecommunication; **GPA: 8.76/10.0**

Technical Skills

ML & Programming: Python, C++, CUDA, PyTorch, TensorFlow
Computer Vision & Perception: CNNs, LiDAR/3D Vision, Open3D, PCL, OpenPCDet
Data, Deployment & MLOps: SQL/NoSQL (MongoDB), Docker, Kubernetes, AWS, GCP
AI Systems & Retrieval: FAISS, HNSW, RAG, Vector Search, Prompt Engineering
Reinforcement & Optimization: Reinforcement Learning, Hyperparameter Tuning, Model Evaluation

Experience

Precision Sustainable Agriculture — Systems Software Intern Raleigh, NC
May 2025 – Present

- Developed a real-time diagnostics API that aggregated GPS, camera connectivity, and perception pipeline state, enabling centralized observability for distributed ML-driven robotic systems.
- Refactored ROS1 communication layers into a structured IPC-based middleware, improving reliability, determinism, and data throughput for real-time perception workloads.

Systems Lab (Dr. Yoon Man-ki) — Research Assistant Raleigh, NC
Jan 2025 – May 2025

- Extended OpenPCDet to support compressed LiDAR range-image representations (PBEA), enabling experimentation with alternative 3D data encodings for object detection.
- Benchmarked PointPillars and PV-RCNN models on the KITTI dataset across multiple range-image resolutions, rigorously measuring inference latency, recall, and average precision trade-offs.
- Built modular evaluation pipelines to quantify sensitivity of 3D detection performance to spatial resolution and metadata fidelity while analyzing end-to-end inference timing to derive efficiency–accuracy trade-offs for real-time autonomy systems.

State Street Corporation — Site Reliability Engineer Bangalore, India
Jul 2023 – Jan 2024

Projects

Vaxel — Deepfake-Resistant Image Sanitization Engine Feb 2026
PyTorch, OpenCV, scikit-image, Tauri (Rust), Next.js

- Designed a local-first ML defense pipeline that applies texture-focused perturbations to headshots to reduce vulnerability to face-attribute manipulation and identity misuse attacks.
- Implemented attack simulation modules to evaluate robustness against deepfake-style transformations such as age and attribute manipulation.
- Built a cross-platform desktop application using Tauri with a Python ML sidecar, enabling privacy-preserving, on-device inference without cloud dependency.

Semantic Concurrency Control for Shared Multi-Agent Memory Feb 2026 - April 2026
Python, FAISS, Vector Databases

- Proposed a Semantic Locking protocol to prevent logically conflicting writes in shared vector-database-backed memory systems for LLM-based agents.
- Designed a Semantic Lock Manager that uses FAISS-based similarity search to detect semantic conflicts via cosine similarity thresholds in high-dimensional embedding space.
- Implemented optimistic commit-time validation strategies for multi-agent RAG pipelines to mitigate semantic race conditions and contradictory memory updates.

Autonomous Reinforcement Learning Navigation (ROS2) Aug 2025 – Dec 2025
Python, PyTorch

- Implemented and compared on-policy (SARSA) and off-policy (Q-Learning) reinforcement learning algorithms for autonomous exploration and navigation.
- Integrated probabilistic localization using Kalman and Particle Filters into learning-based control loops for robust decision-making under sensor uncertainty.