

Haixi Zhang

[LinkedIn](#) - haixizhang02@gmail.com - (607) 339-9816 - San Jose, CA 95113 - [Portfolio](#) - [GitHub](#)

EDUCATION

Cornell University

M.Eng in Electrical & Computer Engineering (GPA: 3.81/4.0)

Ithaca, NY

Aug. 2024 – May. 2025

University of Rochester

B.S. in Electrical & Computer Engineering; Minor in Statistics (GPA: 3.86/4.0)

Rochester, NY

Sep. 2020 – May. 2024

TECHNICAL SKILLS

Languages & Systems: C++, Python, CUDA, Bash, Linux, ROS/ROS2, Git, Bazel, CMake, Docker, Kubernetes

Computer Vision & Robotics: Camera models, multi-sensor extrinsic calibration, $SE(3)/SO(3)$, LiDAR-camera fusion, point-cloud processing, OpenCV, Eigen, PCL/Open3D, state estimation (EKF), SLAM, occupancy-grid mapping, A* planning

Machine Learning & AI: PyTorch, CNNs, Vision Transformers, multimodal learning, cross-attention, multi-task learning, representation learning, PyTorch DistributedDataParallel (DDP), NCCL, distributed data loading/evaluation, LLM APIs, AI agents

Production ML & Data: ONNX, TensorRT, C++ inference, model export and optimization, numerical-parity testing, Airflow, ETL pipelines, dataset curation, data-quality validation, production monitoring

PROFESSIONAL EXPERIENCE

Machine Learning Engineer, Tensor Auto

San Jose, CA

Production Computer Vision, 3D Geometry, Multimodal Learning, Onboard Inference

Sep. 2025 – Present

- Own three company-level ML/CV systems as sole technical owner, taking two from initial requirements to onboard production deployment and leading the third from requirements through active validation and integration.
- Designed a residual extrinsic-calibration framework spanning camera-LiDAR 6-DoF alignment and LiDAR-LiDAR rotational alignment, combining differentiable projection and depth rasterization, $SE(3)/SO(3)$ correction modeling, multimodal feature fusion, and geometry-aware objectives.
- Co-developed Airflow/Kubernetes data pipelines and curated 50K+ training pairs from 500+ GB of multimodal sensor logs; both calibration systems recovered controlled rotational perturbations up to 10° to approximately 0.5° residual, while LiDAR-LiDAR evaluation improved rotation by 62.5% across 95.9% of held-out frames.
- Re-engineered CPU-heavy camera-LiDAR training via geometry precomputation and TAR-sharded streaming, reducing single-accelerator runs from ~ 60 to ~ 20 hours; built world-size-configurable PyTorch DistributedDataParallel/NCCL with deterministic sharding and globally correct evaluation to reach < 5 hours ($\sim 13\times$ overall), while profiling LiDAR-LiDAR scale-out limits to CPU-side projection and point correspondence.
- Lead development of an in-cabin seat-occupancy and occupant-posture classification system from fisheye imagery, establishing session-isolated evaluation, multi-seed benchmarking, class-imbalance diagnostics, deterministic output gating, and an ONNX-compatible production contract.

Undergraduate Research Assistant

University of Rochester, ECE Department

Python, PyTorch, C, Graph Learning, HPC, Docker/Kubernetes

Jan. 2023 – May 2024

- Built reproducible HPC pipelines for graph learning and nature-inspired models in financial and power-system applications, improving benchmark performance by up to 25% and containerizing distributed experiments with Docker/Kubernetes.

SELECTED RESEARCH & PROJECTS

Robotic Tool Handler (Senior Design)

University of Rochester, Lead Developer

C++, ROS, NVIDIA Jetson, Real-time Perception [GitHub](#)

Jan. 2024 – Jun. 2024

- Led perception and system integration for a seven-member robotics team; deployed a containerized ROS vision pipeline on NVIDIA Jetson, improved throughput by 20%, and built GitLab CI with software- and hardware-in-the-loop validation.
- Developed real-time 3D object detection and classification by fusing camera and LiDAR evidence, achieving sub-20 ms inference latency and improving mean average precision by 25%.

Robot Mapping, Estimation, and Navigation

University of Rochester, Solo Developer

ROS, C++, LiDAR, Kalman Filter, TurtleBot2 [GitHub](#)

Jan. 2024 – May. 2024

- Implemented and deployed a C++/ROS navigation stack on TurtleBot2, integrating RGB-D, LiDAR, and wheel odometry for EKF-based state estimation, occupancy-grid mapping, A* planning, and sensor fault detection.
- Reduced localization error by 10% and planning time by 15% through EKF tuning, occupancy-map parameterization, and unit-tested sensor integration and failure handling.

Monocular Real-time Perception for Autonomous Driving

University of Rochester

Python, PyTorch, Computer Vision, Depth Estimation [GitHub](#)

Apr. 2023 – Aug. 2023

- Designed a 30+ FPS multi-task perception pipeline combining object tracking, trajectory prediction, road segmentation, and monocular depth estimation through shared visual features and lightweight task decoders.