

Junhua Huang

Los Angeles, CA | 1-585-286-7948 | junhuahuang1218@gmail.com

EDUCATION

University of California, Los Angeles

M.S. in Electrical and Computer Engineering

Los Angeles, CA

Expected 2027

University of Rochester

B.S. in Computer Science and B.S. in Business

Rochester, NY

May 2025

- GPA: **3.82**/4.00; Highest Distinction, Cum Laude; Dean's List (8 semesters); Schwartz Discover Grant; TA for CSC171/172

SELECTED EXPERIENCE

Vision and Autonomy Intelligence Lab (VAIL), UCLA

Los Angeles, CA

Student Researcher | Prof. Bolei Zhou

Jan 2026 – Present

- ChairNav (first author):** built a data-efficient cross-embodiment training-data pipeline adapting 170 h of source teleoperation data to **Permobil M3** using synthetic trajectories; used imitation learning on 3.19 h of human demonstrations and human-in-the-loop post-training, achieving **84.9% autonomous driving time across 7.8 km** of outdoor navigation.
- Identified nonlinear wheelchair dynamics for physics simulation; used Isaac Sim and UrbanSim; built semi-automated RGB-LiDAR calibration, conditional auto-exposure, 3D-printed mounts, and automated robot-data reconstruction, visualization, and filtering pipelines.
- Accelerated ChairNav model inference from **0.6 to 20 Hz** through model quantization and TensorRT optimization on Jetson AGX Orin; reverse-engineered R-net CAN and implemented feedforward and residual compensation. Broader engineering included GPS/IMU/odometry localization and multi-IMU/high-precision GPS collection.
- Cue the Flow (CoRL 2026):** deployed VLM grounding and a frozen $\pi_{0.5}$ VLA on AgileX PiPER with Jetson image processing and remote inference in a synchronous control loop; integrated the arm, power, communications, and streaming on M20 Pro, and built semi-automated demonstration collection and cleaning.
- Built shared perception between **Aria Gen 1 and the M20 Pro** for online mapping and target navigation; combined Aria Gen 2 streaming, gaze, head pose, image depth, and hand gestures to infer target direction, distance, and user intent.

UCLA NLP Group

Los Angeles, CA

Student Researcher | Prof. Kai-Wei Chang

Dec 2025 – Mar 2026

- Co-developed **ToolSmith**: generated 264 verified tools from a 9,174-query pool through tool synthesis and verification, improving Qwen3-VL-8B MMSearch accuracy from **51.46% to 55.56%**.

Seizure Lab, ECE Department, UCLA

Los Angeles, CA

Student Researcher

Aug 2025 – Oct 2025

- Built a clinical seizure-semiology dataset (10,000 records, 8 features) contributing to **S3**; deployed 3 VLMs with vLLM asynchronous inference, increasing throughput from **0.3 to 5 requests/s (16.7 $\times$)**.

Farsee2 Technology

Hybrid

Machine Learning Intern

Jul 2024 – Sep 2024

- Reduced floating artifacts by **approximately 40%** in city-scale Gaussian Splatting with incidence-angle weighting and depth/normal supervision; built multi-machine preprocessing pipelines for reconstruction prototypes.

SELECTED PUBLICATIONS

- J. Huang***, Z. Liu*, et al. *ChairNav: Cross-Embodiment Pre-training and Personalization for Long-Horizon Wheelchair Navigation*. Under submission, 2026. *Equal contribution.
- H. Wang et al. (incl. **J. Huang**). *Cue the Flow: Steering Flow-Matching Policies for Open-World Delivery Manipulation*. **CoRL 2026**.
- R. Sun et al. (incl. **J. Huang**). *ToolSmith: An Automatic Tool-Generation Framework for Multimodal Search Agents*. **EMNLP 2026 (Findings)**.
- L. Zhang et al. (incl. **J. Huang**). *Seizure-Semiology-Suite (S3): A Clinically Multimodal Dataset, Benchmark, and Models for Seizure Semiology Understanding*. **ICML 2026, Spotlight**.
- P. Liu et al. (incl. **J. Huang**). *GestureLSM: Latent Shortcut based Co-Speech Gesture Generation with Spatial-Temporal Modeling*. **ICCV 2025**.
- J. Huang**, C. Huang, C. Xu. *Semantic Visually-Guided Acoustic Highlighting with Large Vision-Language Models*. arXiv 2026.
- P. Liu et al. (incl. **J. Huang**). *GestureFAR: Streaming Co-Speech Gesture Generation with Flow Autoregression*. arXiv 2026.

PLATFORMS

Robots: Permobil M3, AgileX PiPER, DEEP Robotics M20 Pro, Booster T1, Unitree Go2

Wearable Platforms: Meta Aria Gen 1 and Gen 2

Edge Computing: NVIDIA Jetson Xavier NX, NVIDIA Jetson AGX Orin

TECHNICAL SKILLS

Robot Learning: imitation/preference/reinforcement learning, flow matching, VLA/VLM, cross-embodiment transfer

Robotics Systems: system identification, Isaac Sim, UrbanSim, ROS/ROS2, CAN, TensorRT, SLAM, RGB-LiDAR calibration