Chuyang (Vera) Xiao

(+1) 412-608-5597 | vxiao@andrew.cmu.edu | LinkedIn | GitHub | Personal Website

Education

Carnegie Mellon University — **Pittsburgh, PA**
Master of Science in Robotics, GPA: 4.09/4.0 — *May. 2027*

ShanghaiTech University — **Shanghai, China**
Bachelor of Science in Computer Science, GPA: 3.74/4.0 (Top 10%) — National Scholarship (Top 0.4% nationwide) — *Jun. 2025*

University of California, Berkeley — **Berkeley, CA**
Exchange Student in Electrical Engineering and Computer Sciences, GPA: 4.0/4.0 — *Sep 2023 – Aug 2024*

Selected Publications

[1] **Xiao**, et al. "JAMB: Joint Action–Motion Diffusion for Bimanual Manipulation." ICRA 2027 (under review).

[2] **Xiao**, Zhan, et al. "H2RBench: A Real-to-Sim Benchmark for Evaluating Human-to-Robot Transfer." CoRL 2026.

[3] Zhong, Liu, **Xiao**, et al. "FreqPolicy: Frequency Autoregressive Visuomotor Policy with Continuous Tokens." NeurIPS 2025.

[4] Wang, Ye, **Xiao**, et al. "DexH2R: A Benchmark for Dynamic Dexterous Grasping in Human-to-Robot Handover." ICCV 2025.

[5] **Xiao**, et al. "Optimizing Efficiency of Mixed Traffic through Reinforcement Learning: A Topology-Independent Approach and Benchmark." ICRA 2025.

Research Experience

RPAD Lab at Carnegie Mellon University — **Aug 2025 – Present**
Robotics Manipulation Research Assistant — *Pittsburgh, PA*

- Built the first real-to-sim benchmark for evaluating Human-to-Robot transfer methods across four manipulation tasks using Isaac Lab, Marble, and Trellis, with simulated and real-world demos collected via cuRobo motion planning and GELLO teleoperation.
- Evaluated four representative H2R transfer methods under a unified protocol and validated simulation as a reliable proxy for real-world evaluation across 16 method-task pairs, achieving Pearson $r = 0.89$, Spearman $\rho = 0.85$, and MMRV $= 0.06$.
- Made a frozen Diffusion Policy failure-aware using 50 success and 50 failure rollouts per task, implementing candidate selection and denoising-gradient guidance that raised FurnitureBench success from 19.5% to 23.2% and 27.6%, respectively, at 1.17× latency.
- Developed a joint action–3D track diffusion policy with spatiotemporal track representations, achieving a 23.9% higher task success rate than the baseline on the RoboTwin bimanual manipulation benchmark.

4DV Lab at ShanghaiTech University — **Sep 2024 – May 2025**
Computer Vision and Robotics Research Assistant — *Shanghai, China*

- Co-designed DynamicGrasp, a three-stage pipeline combining grasp preparation, learned approach generation, and goal-pose alignment to adapt ShadowHand trajectories to moving objects during Human-to-Robot handovers.
- Adapted Diffusion Policy and 3D Diffusion Policy for approach generation and benchmarked them on 803 held-out DexH2R trajectories, with DP3 outperforming DP by 26.9 and 18.8 percentage points in Easy and Hard settings.
- Designed a six-metric evaluation suite covering success, safety, path length, penetration depth, and inference speed and used it to show that DP3 matched MotionNet in close-range success (27.1% vs. 26.6%) while doubling safety (33.7% vs. 15.4%).
- Implemented DCT-based hierarchical action generation within Diffusion Policy and 3D Diffusion Policy, contributing to 67.9% average success across 48 simulation tasks and enabling 10× faster inference than Diffusion Policy on Robomimic.

The University of Hong Kong — **Apr 2024 – Sep 2024**
Autonomous Driving Research Assistant — *Hong Kong, China*

- Developed a topology-independent mixed-traffic controller using Soft Actor-Critic and centralized training with decentralized execution, enabling autonomous vehicles to coordinate traffic using only local observations across intersections and roundabouts.
- Built the first benchmark for mixed-traffic control in complex real-world scenarios by converting OpenStreetMap data into SUMO, covering 444 scenarios, 111 road topologies, and 20 countries with traffic demands up to 5,000 vehicles/hour.
- Evaluated the policy against traffic-light, unsignalized, and state-of-the-art controllers across 72 held-out scenarios, reducing average wait time by 74% and increasing throughput by 64% over traffic-light control at full autonomous-vehicle penetration.

Work Experience

Spider Robot Company — **May 2025 – Aug 2025**
Robotics Perception Intern — *Shanghai, China*

- Calibrated camera systems and configured NVIDIA Jetson AGX Orin devices for reliable on-robot perception deployment.
- Annotated robotic-scene images and trained YOLO models for target-object detection.
- Evaluated VLA models and manipulation platforms, including π_0, RoboTwin, and ALOHA, for technical feasibility.

ShanghaiTech University — **Feb 2025 – Jun 2025**
Teaching Assistant, Introduction to Machine Learning — *Shanghai, China*

- Prepared lecture slides, recitation materials, homework assignments, and exam questions, and graded coursework while providing feedback and support to students.

Technical Skills

Programming & ML: Python, C/C++, Bash, PyTorch, TensorFlow, Weights & Biases, TensorBoard, YOLO, SAM2, DepthAnything
Robotics & Systems: Isaac Lab, MuJoCo, SAPIEN, SUMO, cuRobo, OpenGL, Linux, Git, NVIDIA Jetson AGX Orin, Google Cloud Platform, SolidWorks, MeshLab