

Jinghao Liu

jliu63@uw.edu | (206) 495-7111 | github.com/JasonHistoria | linkedin.com/in/jinghao-liu

Education

University of California, San Diego, La Jolla, CA

Incoming Sep 2026

M.S. in Computer Science (CS75) — joining Prof. Xiaolong Wang's lab

University of Washington, Seattle, WA

Expected Jun 2026

B.S. in Computer Science

GPA: 3.92

Technical Skills

Programming: Python, Java, C, C++, C#, R, JavaScript, HTML/CSS

ML/AI: PyTorch, Hugging Face, LeRobot, Accelerate/DeepSpeed, Diffusion & Flow models, vLLM, LoRA

Tools: Git, SLURM, VSCode

Research Experience

University of Washington, Robotics and State Estimation (RSE) Lab

Oct 2025 – Jun 2026

Undergraduate Research Assistant, Advised by Prof. Dieter Fox; mentored by Ge Yan

Flex- π : A Multi-Stream World-Action Model with Compute Flexibility

Co-first author | arXiv:2608.10860 | flex-pi.github.io

- Built and maintain the end-to-end training and closed-loop evaluation pipeline for Flex- π , a compute-flexible world-action model that embeds RGB, 3D point-maps and DINO semantics in a shared latent space, letting one checkpoint run any combination of input and output streams; reached **94.6%** on RoboTwin (50 tasks) and **99.2%** on LIBERO.
- Designed a unified 32-DoF end-effector action space (rot6d + per-dataset normalization) bridging heterogeneous embodiments (AgiBot, YAM, DROID, LIBERO), with large-scale DROID/LeRobot data pipelines (metric-depth + 3D point-maps) on multi-GPU HPC clusters.
- Validated on a real bimanual YAM robot across five tasks, averaging **83.0%** in distribution and **76.1%** out of distribution — including **98.8%** on long-horizon kitchen organization and **76.0%** on a dexterous screw-driving repair — against $\pi_{0.5}$ at 52.1%, ManiFlow at 58.0% and Fast-WAM at 31.7%.
- Earlier, built a 3D pipeline recovering camera trajectories and hand poses from egocentric video (Ego100K) for human-demonstration pretraining; completed, shelved pending compute.

Mohamed bin Zayed University of Artificial Intelligence (MBZUAI)

Jun 2025 – Oct 2025

Research Assistant

AI Security Beyond Core Domains: Resume Screening as a Case Study of Adversarial Vulnerabilities in Specialized LLM Applications

*Manuscript under review at the **International Journal of Machine Learning and Cybernetics***

- Built the first benchmark (**1,000 authentic LinkedIn profiles**, 14 sectors) and a systematic framework to study adversarial robustness of LLM-based resume screening—an unstudied domain in AI safety.
- Defined an attack taxonomy (4 types $\times$ 4 injection positions) and measured cross-model vulnerability via an automated vLLM + SLURM pipeline (**150k+ evaluations**), exposing >80% attack success.
- Built hybrid prompt-engineering + **LoRA SFT** defenses, achieving **2.6 $\times$ robustness gain** and **58.5% overall mitigation**.

Academic Project Experience

Sentra: AI-Powered Roleplay Chatbot Platform

Sep 2024 – Jun 2025

Developer | Advised by Prof. René Just | github.com/Yuzhifur/Sentra-Chatbot

- Built and led a full-stack roleplay chatbot (**sentra-4114a.web.app**; React + Firebase) integrating multiple AI models (Claude, Gemini) for real-time streaming responses, with a Cross-Friends Memory system, authentication, AI avatar generation, and 85%+ Jest coverage via Git/CI-CD.

Synthetic Image Detectors

Oct 2024 – Dec 2024

Team Programmer

- Evaluated CNN (ResNet-50, GoogleNet, EfficientNet-b0) and transformer (ViT-b/16) robustness against GAN- and diffusion-generated images—finding up to 35% accuracy drop under texture-based distortions (edge features 40% more vulnerable than color cues)—and used Grad-CAM to expose architectural differences in background vs. texture attention.

Work Experience

Xpeng Robotics Limited

Jun 2023 – Sep 2023

Product Department Intern, Shenzhen

- Conducted market research on 15+ robotic products (two competitive analysis reports) and worked with a cross-functional team to translate two product requirements into technically feasible solutions and prototype projects.