

Jingxi Lu

✉ jingxil@usc.edu | in LinkedIn | GitHub | Google Scholar

EDUCATION

University of Southern California

Master of Electrical and Computer Engineering **GPA: 3.94/4.0**

- Courses: Linear Algebra, Probability, Deep Learning, Reinforcement Learning and LLM, Digital Signal Processing

HongKong Baptist University

Bachelor of Computer Science and Technology **GPA: 3.51/4.0 (Top 10%)**

- Courses: Data Structures, Data Mining, Database Systems, Machine Learning, Computer Vision, NLP, Data Visualization

California, United States

Sept. 2024 – Expected May 2026

HongKong, China

Sept. 2020 – Jul. 2024

RESEARCH INTERESTS

Reinforcement Learning, Learning-based Control, Sequential Decision Making, Embodied AI, Optimization, Robotics

WORKING EXPERIENCE

Lunar Lab, Georgia Institute of Technology

Research Assistant, Prof. Lu Gan

May. 2025 – Present

Atlanta, United States

- Proposed a **latent-state RL framework** for fault-tolerant quadruped locomotion, removing reliance on explicit fault detection; deployed on **Unitree Go2** to enable end-to-end **resilient walking** under actuator failures.
- Enabled generalization to **multiple simultaneous motor faults** via a **curriculum learning strategy** that gradually increases fault complexity beyond the single-joint scenarios handled in prior work.
- Integrated **SwAV-style contrastive learning** to improve encoder fault representations, enhancing policy robustness in **unseen fault conditions**.

LIDAR, Georgia Institute of Technology

Research Assistant, Prof. Zhao Ye

May. 2025 – Aug. 2025

Atlanta, United States

- Contributed to the design and real-world deployment of a **stability-aware bipedal navigation system** on the **Digit robot**, bridging simulation and hardware for dynamic locomotion.
- Developed custom **ROS1/ROS2 pipelines** for LiDAR and pose data integration, and generated real-time, robot-centric **elevation maps** for terrain-adaptive planning.
- Implemented and optimized a **Transformer-based traversability network** on hardware, achieving **low-latency inference** for **risk-sensitive navigation planning** in unstructured environments.

Dynamic Robotics and Control Laboratory, USC

Research Assistant, Prof. Quan Nguyen

Oct. 2024 – May. 2025

California, United States

- Developed a **PPO-based locomotion policy** for the **bipedal robot Hector** in **Isaac Gym**, leveraging **asymmetric actor-critic** architectures and a β -VAE to reconstruct velocity and environment observations for locomotion.
- Achieved **stable walking behaviors** in simulation through extensive training, hyperparameter tuning, and reward shaping under a high-fidelity physics engine.
- Worked on **whole-body imitation learning** experiments on the **humanoid robot G1**, designing pipelines for retargeting and reproducing mocap-based movements using behavior cloning.

Shenzhen Research Institution of Big Data

Research Assistant

Mar. 2024 – Jul. 2024

Shenzhen, China

- Proposed a hybrid framework combining **reinforcement learning** and **column generation** to address the **aircraft recovery problem**, where the policy learns to generate high-quality initial columns; work **accepted at LESCDT 2024**.
- Built on the **PPO algorithm**, integrating **Graph Attention Networks (GAT)** to encode flight topology graphs and **Pointer Networks** to sequentially decode optimized recovery plans.
- Demonstrated strong **generalization capability** from small-scale training to large real-world scenarios, achieving up to **50% faster convergence** compared to rule-based baselines.

Efficient Large Vision Language Model Inference

Student Researcher, Prof. Peter A. Beerel

Jun. 2024 – Present

Los Angeles, CA

- Developed a **training-free token pruning** framework for **large vision-language models (LVLMs)**, eliminating the need for fine-tuning while significantly improving inference speed.
- Introduced a **multi-scale token selection strategy** that adaptively retains fine-grained visual tokens for efficient and accurate downstream reasoning.
- Improved inference throughput on **LLaVA-1.5 7B** from **18.66** to **27.99 tokens/sec**, outperforming existing token pruning baselines in both speed and accuracy.

Hybrid Learning for Microservice Scheduling in Edge Computing

Research Assistant, Prof. Jianxiong Guo

May. 2023 – Jan. 2024

Zhuhai, China

- Proposed a hybrid learning framework for **container scheduling** in **dynamic edge environments**, optimizing **latency** and **energy consumption**; work published at **MSN 2024**, with an extended version under review at **IEEE TMC**.
- Designed a **GRU-based soft actor-critic (SAC)** model with a hybrid policy to capture inter-task dependencies and adapt to real-time resource variations, achieving up to **50% higher cumulative reward** over SAC baselines.
- Integrated a rule-based expert into an **offline imitation learning** pipeline, pretraining the policy via **behavior cloning**, which led to **70% faster convergence** and improved early-stage training stability.

PUBLICATIONS & MANUSCRIPTS

- **Jingxi Lu***, Wenhao Li, Jianxiong Guo, Xingjian Ding, Zhiqing Tang, Tian Wang, Weijia Jia.
Hybrid Learning for Cold-Start-Aware Microservice Scheduling in Dynamic Edge Environments.
Under review at **IEEE Transactions on Mobile Computing (TMC)**.
- **Jingxi Lu***, Wenhao Li, Jianxiong Guo, Xingjian Ding, Zhiqing Tang, Tian Wang, Weijia Jia.
Container Scheduling with Dynamic Computing Resource for Microservice Deployment in Edge Computing.
Accepted at MSN 2024.
- **Jingxi Lu***, Xiongwen Qian.
A Reinforcement Learning Approach for Initialization of Column Generation with Application to Aircraft Recovery Problem.
Accepted at LESCDT 2024.

SKILLS

Programming: Python (expert), C++, C, Java, MATLAB, R, SQL

ML Tools: PyTorch, HuggingFace, Scikit-learn

Robotics Stack: ROS1/2, Mujoco, Unitree SDK, NVIDIA Isaac Gym

DevOps & Tools: Docker, Git, Linux, LaTeX