

YUANPENG LI

☎ +1) 949-618-8992 ✉ yuanpl2@uci.edu 🌐 yuanpeng-li.github.io 📄 github.com/Yuanpeng-Li 🌐 linkedin.com/in/yuanpengli

EDUCATION

University of California, Irvine

Doctor of Philosophy (PhD) in Statistics

Irvine, CA

Sep. 2024 – Present

Jilin University – School of Mathematics

Bachelor of Science Degree, Major in Mathematics and Applied Mathematics

Changchun, China

Aug. 2020 – Jul. 2024

RESEARCH EXPERIENCE

Sample-Efficient Distillation of Large LLMs into Mixture-of-Experts

Graduate Researcher, UC Irvine

Irvine, CA

May 2026 – Present

- Developing a sample-efficient pipeline to distill an ultra-scale dense LLM into a smaller Mixture-of-Experts (e.g. Qwen3 35B-A3B-style low-activation models) by distilling experts and training the router.
- Framing distillation as imitation learning (Dagger, DART, goal-conditioned IL) with RL-from-rubric rewards (RLRR), built and trained on the VERL framework.

TOPPO: Rethinking PPO for Multi-Task RL with Critic Balancing

First author; under review (arXiv:2605.11473); Advisors: Prof. Annie Qu, Prof. Rui Miao

Irvine, CA

Jul. 2023 – Apr. 2026

- Designed *Critic Balancing* for multi-task PPO — per-task PopArt value normalization, critic LayerNorm, and per-side gradient combiners (PCGrad / CAGrad / FairGrad) — to fix critic-side gradient ill-conditioning where tail tasks stall.
- On Meta-World+ MT50, beat published SAC- and ARS-family baselines on mean and worst- k tail-task success with up to $22.7\times$ fewer parameters and far fewer environment steps.

PROJECTS

gradescope-mcp | Python, Model Context Protocol, uv

- Built an MCP server exposing Gradescope to AI agents — 34 tools across course management, grading, rubric CRUD, and statistics, with all write operations gated behind explicit confirmation.
- Added a reusable assisted-grading skill for human-approved batch grading, with a custom roster parser and 30 automated tests.

StudyGround | JavaScript, Node, Claude Code Plugin, Pyodide

- Built a Claude Code plugin for AI-assisted self-study: a web reader with inline Ask buttons, in-browser Python via Pyodide, and exercises that deep-link into VSCode.
- Added NotebookLM-style RAG over course materials: PDF/OCR extraction and BM25 plus embedding search with page-anchored citations.

TEACHING

- **Graduate Teaching Assistant**, UC Irvine Statistics (Sep. 2024 – Present): STATS 7, 67, 110, 120C, and 205P (Bayesian Data Analysis, MDS).

INTERNSHIP EXPERIENCE

Synkrotron

Intern, AI Research Department

Xi'an, China

Dec. 2022 – Jan. 2023

- Controlled remote Linux devices through FRP to configure network information in the automated road patrol project.
- Researched the use of computer-generated simulation datasets to enhance real-world performance in autonomous driving algorithm testing projects.

AWARDS

- Scholarship in the academic year 2022-2023 Oct. 2023
- The Chinese Mathematics Competitions: **National 3rd Prize** Dec. 2021
- Jilin Province Mathematics Competitions: **Provincial 2nd Prize** Dec. 2021

SKILLS

- **Programming Languages:** Python, C, MATLAB, Wolfram Mathematica, \LaTeX
- **ML Frameworks:** PyTorch, Tianshou, StableBaselines3, Gymnasium, MetaWorld, WandB, Scikit-learn, Pandas
- **Dev Tools & Systems:** Claude Code, OpenAI Codex CLI, GitHub Copilot, Cursor, Git, Docker, Linux Server Administration, VMware ESXi, OpenWrt, Network Administration