Chenggong Zhang

+1 (310) 486-7361 | chenggong61@g.ucla.edu |

Education

University of California, Los Angeles (UCLA) — Los Angeles, CA
M.S. Electrical and Computer Engineering (GPA: 3.967/4.00) — *Sep 2024 – Dec 2025*

University of California, Los Angeles (UCLA) — Los Angeles, CA
B.S. Computer Engineering (GPA: 3.616/4.00) — *Sep 2021 – Jun 2024*

Research Interests

Agent, Reinforcement Learning, Generative Model (Diffusion, LLMs).

Research Experience

Graduate Researcher, AGI Lab — Apr. 2024 – Present
Advisor: Prof. Quanquan Gu, UCLA — *Los Angeles, CA*

- Developing Linear Maillard Sampling for linear bandits: implemented a linear ridge-regression parameter estimator and a Boltzmann exploration mechanism.
- Conducting theoretical regret analysis for algorithm using self-normalized concentration arguments and elliptical potential lemma; drafting theorem proof notes and preparing a manuscript.
- Reproduced diffusion-based energy-guided sampling baselines for offline RL with large-scale experiments on computing clusters, and evaluated generation quality with FID on ImageNet.
- Lead development of the EurekaClaw project (GitHub / Website): designed and implemented the frontend interface, coordinated open-source community contributions, and maintained codebase for research and public use.

Undergraduate Researcher — Sep. 2022 – Dec. 2023
Advised by Prof. Lin Yang, UCLA — *Los Angeles, CA*

- **Model-Based RL wih Generative Model Policy Learnin):** Studied perturbed model-based planning for discounted MDPs; re-derived core lemmas and compiled theorem-level proof notes.
- **Zero-Sum Markov Games with Boltzmann exploration:** Investigated extending POLITEX mixture policy iteration to simultaneous-move two-player zero-sum Markov games in an online learning setting; formulated regret against best-response opponents and drafted a Boltzmann policy update with approximate value evaluation.
- **Local Sampling Access in Discounted MDPs:** Formalized a practical sampling method that permits queries at previously visited states with resets; proposed an exploration-and-planning algorithm and outlined an analysis roadmap toward near-oracle sample efficiency on chain-structured MDPs.

Selected Coursework

Stochastic Processes; Information Theory; High-Dimensional Statistics; Statistical Machine Learning; Theory of Reinforcement and Online Learning; Convex Optimization; Large-Scale Optimization Theory; Applied Deep Learning.

Research Projects

Web World Models (WWM): | *Princeton AI Lab, Prof.Mengdi Wang* — Dec. 2025

- Prototyped a **Web World Model** framework that generates unlimited context via LLMs in text/code environments without data storage.
- Built multiple WWM-style environments spanning real-world and fictional settings (e.g., infinite world travel, infinite galaxy exploration), demonstrating controllable and open-ended interaction for users and agents.

TENG++: Time-Evolving Natural Gradient for Boundary-Enhanced PINNs | *Advisor: Prof.Di Luo (UCLA)*

- Extended the Time-Evolving Natural Gradient (TENG) framework with Dirichlet boundary constraints in PINNs; implemented a Jacobian-based natural-gradient least-squares stepper and integrated it with Euler/Heun time-stepping.
- Validated on heat-equation benchmarks with analytical solutions; observed accuracy with Heun (errors down to $\sim 10^{-4}$) and stronger efficiency with Euler in simpler regimes (GitHub:code link).

Hallucination Detection and Evaluation of Large Language Model | *Advisor: Prof.Yuan Tian (UCLA)*

- Built a KnowHalu-style hallucination evaluation pipeline: Replaced KnowHalu's judge with HHEM to accelerate evaluation (10-minute judging) while maintaining QA detection performance (76.9% accuracy); explored a non-fabrication-checking variant reaching 82.2% accuracy.
- Improved summarization hallucination detection via segment-level verification; Write "expert" prompt templates to strengthen agentic reasoning.

Predicting Text from Intracranial Neural Signals | *Advisor: Prof.Jonathan Kao (UCLA)*

- Implemented RoPE positional embeddings in a CNN-Transformer decoder for intracortical speech-to-text; added Gaussian-smoothing and structured data augmentation to improve performance. Improved training stability by implementing AdamW with warmup/decay learning-rate scheduling and gradient clipping; implemented the CR-CTC loss, reducing CER from 0.234 to 0.1636 (GitHub: code link).

Deep RL Algorithms in Cryptocurrency Trading | *Advisor: Prof.Jonathan Kao (UCLA)*

- Built an ETH trading environment (continuous state/action) with trading features and trading actions through hourly Bitfinex ETH dataset.
- Re-implemented and debugged off-policy actor–critic agents (DDPG, TD3, SAC) with replay/target networks; benchmarked against random walk trading policy and found SAC achieved the best returns while TD3 trained more stably than DDPG. (GitHub: code link)

Teaching Experience

Teaching Assistant:
ECE 132B: Data Communications and Telecommunication Networks (Prof. Izhak Rubin, Jan 2025 – Apr 2025)
ECE 241: Stochastic Process (Prof. Lin Yang, Jan 2025 – Apr 2025)
Learning Assistant: ENGR 96A: Machine Learning With Python(Prof. Jacob Schmidt, Sep 2023 – Jun 2024)
Reader: CS 51A: Logic Design of Digital Systems; ECE 131A: Probability Theory; CS M146: Machine Learning.

Awards & Honors & Social

Awards: Full Tuition Grant (Fall 2024, Winter 2025); AWS Computing Fund for BigANN Competition (NeurIPS 2023).
Honors: Dean's List(Spring 2023); Eta Kappa Nu (HKN) Honor Society (Jun 2022); Upsilon Pi Epsilon(UPE) Honor Society (Jun 2022).
Social: RSS 2025; Neurips 2025

Skills

Languages: Chinese (native), English (proficient), Spanish (principiante)
Tools: C/C++, Python, MATLAB, LaTeX, GCP, Docker, Conda, Git, Bash Scripting, Tmux
ML: JAX, PyTorch, Weights & Biases, Huggingface/Transformers, NumPy, Matplotlib, Scikit-learn, Pandas, Gymnasium, VeRL