Zheng Luo

Los Angeles, CA | luozheng@usc.edu | (213) 762-9926
Portfolio: https://zhengluo-ai.com/

Research Profile

AI researcher focused on agentic large language models, tool-use robustness, reinforcement learning, and fine-grained credit assignment. Experienced in building end-to-end research systems, including benchmark construction, large-scale model evaluation, LoRA/RL fine-tuning, error analysis, and simulation environments.

Education

University of Southern California, Viterbi School of Engineering | Los Angeles, CA
M.S. in Computer Science, Artificial Intelligence | Expected May 2027

- GPA: 3.92/4.0
- Relevant coursework: Machine Learning, Computer Vision, Analysis of Algorithms
- Grader: Trustworthy Foundation Models; Deep Learning and Its Applications

University of Michigan, College of Engineering | Ann Arbor, MI
Bachelor of Science in Engineering in Computer Science | Graduated May 2025

- UM-SJTU dual-degree program; separate degrees awarded by each university.
- GPA: 4.0/4.0
- Relevant coursework: Operating Systems, Embedded Systems, Compiler Construction
- Grader: Introduction to Embedded Systems

Shanghai Jiao Tong University, UM-SJTU Joint Institute | Shanghai, China
Bachelor's degree in Electrical and Computer Engineering | Graduated July 2025

- UM-SJTU dual-degree program; completed the final summer semester at SJTU in 2025.
- GPA: 3.6/4.0
- Relevant coursework: Circuit Design, Logic Design, Signals and Systems

Publications and Manuscripts

Peer-Reviewed Publications

- **Zheng Luo**, T. Pranav Kutralingam, Ogochukwu N. Okoani, Wanpeng Xu, Hua Wei, and Xiyang Hu. "Lost in Execution: On the Multilingual Robustness of Tool Calling in Large Language Models." *ACL 2026*.
- Xiaolin Zhou, Aojie Yuan, **Zheng Luo**, Zipeng Ling, Xixiao Pan, Yicheng Gao, Haiyue Zhang, Jiate Li, Shuli Jiang, Prince Zizhuang Wang, Zixuan Zhu, Jinbo Liu, Ryan A. Rossi, Hua Wei, and Xiyang Hu. "When Simulation Lies: A Sim-to-Real Benchmark and Domain-Randomized RL Recipe for Tool-Use Agents." *NeurIPS 2026 Evaluations and Datasets Track*, poster; also presented at the *FAGEN Workshop at ICML 2026*.

Manuscripts and Submissions

- **Zheng Luo**, Xiaolin Zhou, Yi Nian, Yue Zhao, and Xiyang Hu. "TreeMPO: Fine-Grained Credit Assignment for LLM Training via Token-Level Trajectory Branching." Manuscript in preparation.
- **Zheng Luo**, Kexin Wang, Xiyang Hu, Yue Zhao, and Ruolin Li. "Street-Level Competitive Ride-Hailing with Fixed-Rival Defensive PSRO." Submitted to *AAAI 2027*.
- Xiaolin Zhou, **Zheng Luo**, Jiashu Qu, Anglin Cai, Yicheng Gao, Ruishan Liu, and Xiyang Hu. "SkiLT: What Agent Reinforcement Learning Actually Learns from a Latent Skill Library." Unpublished manuscript.
- Xiaolin Zhou, **Zheng Luo**, Yicheng Gao, Qixuan Chen, Xiyang Hu, Yue Zhao, and Ruishan Liu. "Fairness or Fluency? An Investigation into Language Bias of Pairwise LLM-as-a-Judge." *NeurIPS 2026 Workshop JUDGe* submission.
- Gan Wang∗, **Zheng Luo**∗, Huiying Wu, Jiaojiao Wang, Wen Ma, and Li Luo. "Construction of a Machine Learning-Based Model for Predicting Specialist Nurse Transfer Tendencies." *BMC Health Services Research*, accepted for publication. (∗Equal contribution.)

Research Experience

TreeMPO: Fine-Grained Credit Assignment for LLM Training | First Author
Advisors: Prof. Xiyang Hu, Prof. Yue Zhao | Apr. 2026 - Present

- Designed TreeMPO, a reference-policy-free training method that constructs reasoning trees through guided token-level branching and infers segment-level credit from final-answer correctness.
- Formulated a joint maximum-a-posteriori estimator of signed segment contributions and uncertainty, converting inferred credit into token-level advantages while normalizing shared-prefix contributions.
- Built a split research pipeline for rollout collection, trajectory-tree construction, advantage-weighted data generation, LoRA fine-tuning, and checkpoint validation under constrained GPU resources.
- Evaluated mathematical reasoning across five model families against a GRPO-style baseline and branching/credit-assignment ablations; analyzed modest validation gains, mixed held-out transfer, and optimization sensitivity.

Lost in Execution: Multilingual Robustness of Tool Calling in LLMs | First Author
Advisors: Prof. Xiyang Hu, Prof. Hua Wei | Sept. 2025 - Jan. 2026

- Built MLCL, a diagnostic benchmark for evaluating tool-calling robustness across English, Chinese, Hindi, and Igbo user queries.
- Designed a fine-grained error taxonomy separating intent understanding, tool selection, parameter extraction, and execution-level language mismatches.
- Combined rule-based and LLM-assisted analysis to identify parameter-value language mismatch as a major failure mode in multilingual tool calling.
- Evaluated inference-time mitigation strategies and showed that they reduce language-induced execution errors but do not fully recover English-level performance.

Street-Level Competitive Ride-Hailing with Fixed-Rival Defensive PSRO | First Author
Advisors: Prof. Ruolin Li, Prof. Xiyang Hu, Prof. Yue Zhao | Feb. 2026 - Aug. 2026

- Developed a SUMO-backed street-level competitive ride-hailing simulator with per-vehicle control, multiple autonomous fleets, mixed human traffic, and demand calibrated from NYC TLC trip records.
- Designed a feature-engineered controller for relocation, pricing, and service decisions, and evaluated U-Net-style residual extensions for local policy refinement.
- Proposed Fixed-Rival Defensive PSRO to reweight fixed opponent mixtures toward difficult rivals; improved mean revenue by 4.0% over fixed-rival standard PSRO across four validation scenarios and four seeds using validation-selected checkpoints.
- Measured a 23.9% head-to-head revenue advantage over the strongest tested zone-level controller and 5.8% higher role-balanced revenue than a fixed-price greedy baseline in Manhattan simulations.

When Simulation Lies: Sim-to-Real Robustness for Tool-Use Agents
Advisors: Prof. Xiyang Hu, Prof. Yue Zhao | Jan. 2026 - July 2026

- Implemented user-query perturbations on ACEBench for RobustBench-TC and evaluated tool-use robustness under the resulting input noise.

Fairness or Fluency? Language Bias of Pairwise LLM-as-a-Judge
Advisors: Prof. Ruishan Liu, Prof. Yue Zhao, Prof. Xiyang Hu | Sept. 2025 - Jan. 2026

- Designed and implemented experiments analyzing the relationship between pairwise judge preference and answer perplexity across languages.
- Compared nested regression models using incremental explained variance and F-tests to assess language-associated judge preferences beyond answer perplexity.
- Found that language identity explained additional variance in judge preference after controlling for log-perplexity differences, particularly for low-resource languages.
- Implemented style-aligned answer construction and masked perplexity scoring over core answer tokens to reduce confounding from response framing.

Construction of a Machine Learning-Based Model for Predicting Specialist Nurse Transfer Tendencies | Co-First Author
Advisor: Prof. Li Luo | July 2024 - Aug. 2024

- Developed machine-learning models to predict specialist nurses' transfer intention from structured healthcare survey data.
- Built data preprocessing and validation workflows to improve reliability of model evaluation and interpretation.
- Collaborated with healthcare researchers to connect model outputs with domain-specific factors relevant to nurse workforce management.

Technical Skills

- **Programming:** Python, Rust, C++, C#, TypeScript
- **Machine learning:** PyTorch, Scikit-learn, LoRA fine-tuning, RL fine-tuning, model evaluation, ablation studies
- **LLM systems:** vLLM, SGLang, tool calling, benchmark pipelines, multilingual evaluation, agent simulation
- **Software and infrastructure:** Git, Docker, Slurm/HPC workflows, Apptainer, Linux
- **Simulation and graphics:** SUMO, Unity, Unreal Engine 5, Fyrox, Blender

Industry and Applied Experience

Data Science Intern, School of Public Health, Fudan University | Shanghai, China
June 2024 - Aug. 2024 | 40 hours/week

- Designed and implemented a rule-based extraction algorithm that converted textual and spatial information in medical reports into structured data.
- Improved robustness to distorted report images and achieved approximately 80% field-level extraction accuracy on 50 reports, each spanning 2-3 pages with approximately 6-8 fields.

Selected Projects

Automatic PCB Routing Using Discrete Bayesian Inference | Team Leader | May 2025 - July 2025

- Proposed and implemented a discrete Bayesian routing scheduler in Rust, using sampled trace interactions to guide routing order and backtracking.
- Led integration of custom hybrid grid/shape-based A$*$ search, collision detection, and PCB file conversion; implemented routing visualizations and compared the system with FreeRouting.

Crafty Piggies 3D, Physics-Based Game Development | Team Leader | Sept. 2024 - Dec. 2024

- Led a four-person team building a 3D physics game in Unity, coordinating gameplay design, level design, and implementation milestones.
- Built a vehicle dynamics system and supported logistics for interactive level design and gameplay testing.

Table Tennis Feeding Machine | Team Leader | Feb. 2024 - Apr. 2024

- Engineered and fabricated mechanical components for a training machine capable of launching table-tennis balls with variable spin.

Awards and Honors

- Dean's Honor List, University of Michigan, 2024
- Successful Participant, Mathematical Contest in Modeling, COMAP, 2022