JIAXI LI

University of Georgia, School of Computing
+1-(470)-263-8931 | e: jx.li@uga.edu
Homepage | Google Scholar | LinkedIn | GitHub

EDUCATION

University of Georgia — Georgia, USA
Ph.D. Student in Computer Science — Aug. 2024 – May 2028 (Expected)
Internship Availability: Spring / Summer 2027

Shandong University — Shandong, China
Bachelor of Science in Computer Science and Technology — Sep. 2020 – Jun. 2024

RESEARCH PROJECTS

Research Focus: Recursive Self-Improvement (RSI), Self-Evolving agents; Post-Training and Alignment; LLM Reasoning and Evaluation; Memory-Efficient Optimization for LLM Training.

Self-Evolving LLM Agents *Skill Evolution, Failure Attribution, Agentic Memory*

- **Learning from Success (SkillAny, NeurIPS-W'26)**: Formulated the **Online Skill Learning** setting for **agentic continual learning**, and proposed our method **SkillAny**, which induces executable skills from judged-successful trajectories and reuse skills at any intermediate steps **for future tasks**, enabling self-improving over time. Performances improved over the strongest baseline **by 10.6% / 10.0%** for **web agent tasks**.
- **Learning from Failure (Forth-and-Back)**: Developed **backward failure attribution** to identify **decisive error steps / root causes** in agent trajectories, enabling targeted skill repair and reusable failure experience acquisition. Improved WebArena-verified **success rates (+5.6 pp)** against the strongest baseline.

Foundation Model Post-Training and Alignment *RSI, Post-Training (PPO/GRPO/DPO, SFT), Synthetic Data*

- **Recursive Self-Improvement (RSI) with Synthetic Data (RST)**: Contributed to **recursive synthesis of 37,484 verified terminal tasks** across 15 rounds at approximately cost $0.05/task, supporting rejection-sampled **SFT** and agentic **PPO**. **SFT** improved **Qwen3.5-27B and Qwen3.5-122B-A10B** by 10% across terminal benchmarks; **PPO** achieved **49.44% on Terminal-Bench 2** with Qwen3.5-27B, a **20.0%** relative improvement over the base model.
- **Multimodal Preference Alignment (SymMPO, NeurIPS'25)**: Developed a theory-consistent symmetric DPO algorithm with preference margin consistency regularization to mitigate multimodal hallucinations. **DPO Training LLaVA-1.5-7B/13B** on **21.4K synthetic** image-text pairs; improved accuracy **from 71.3% to 82.6% (+11.3 pp)** and **reduced** hallucination rate **from 49.0% to 42.7%** versus standard DPO.

LLM Reasoning and Evaluation *Inference-time Scaling, LLM-as-Judge, Process Reward*

- **Efficient Inference-Time Scaling (MITS, PAKDD'26 Best Paper Finalist)**: Developed an **inference-time scaling** framework with **mutual information** scoring, entropy-adaptive sampling, and weighted voting. Improved StrategyQA accuracy from 65.32% to 68.45% while achieving 10× faster inference than rStar.
- Reasoning Evaluation and Benchmarking
  - **Multimodal Reasoning (GraphVerse, EMNLP'26)**: Co-developed a benchmark evaluating perception and visual reasoning across single-image and paired-image graph tasks. Introduced four graph-centric image-editing strategies and VGR-Score, a **process-sensitive** metric assessing **reasoning quality beyond final-answer accuracy**.
  - **Medical Reasoning and Knowledge Evaluation**: Developed **MedThink-Bench (npj Digital Medicine)** and a rationale-grounded LLM-as-a-Judge, achieving 0.87 correlation with experts at 1.4% of evaluation time; built the **MKJ** dataset to assess medical factual knowledge and calibration.

PUBLICATIONS AND PREPRINTS

(*Equal Contribution, †Corresponding Author.)

[1] SkillAny: Online Skill Learning for Language Agents at Any Intermediate State
Jiaxi Li, Ke Deng, Yun Wang, Jingyuan Huang, Yucheng Shi, Yuanfu Sun, Qiaoyu Tan, Jin Lu, Ninghao Liu
Accepted in ***NeurIPS 2026 TTCL Workshop*** (Towards Test-Time Continual Learning Agents); Under Review.

[2] Forth-and-Back: Self-Evolving Skill Learning through Back-propagated Failure Attribution
Jiaxi Li, Huaqin Zhao, Zhisheng Qi, Haowen Xu, Bing He, Zhongzhi Li, Yu Wang, Ninghao Liu, Jin Lu
Under Review.

[3] MITS: Enhanced Tree Search Reasoning for LLMs via Pointwise Mutual Information
Jiaxi Li, Yucheng Shi, Xiao Huang, Jin Lu, Ninghao Liu
Published in ***PAKDD 2026 (Best Paper Finalist Award; Oral Presentation)*** (The 30th Pacific-Asia Conference on Knowledge Discovery and Data Mining).

[4] HELENE: Hessian Layer-wise Clipping and Gradient Annealing for Accelerating Fine-tuning LLM with Zeroth-order Optimization
Huaqin Zhao*, ***Jiaxi Li****, Yi Pan, Shizhe Liang, Xiaofeng Yang, Fei Dou, Tianming Liu, Jin Lu
Published in ***EMNLP 2025 Main Conference*** (The 2025 Conference on Empirical Methods in Natural Language Processing Main Conference).

[5] Fact or Guesswork? Evaluating Large Language Models' Medical Knowledge with Structured One-Hop Judgments
Jiaxi Li, Yiwei Wang, Kai Zhang, Yujun Cai, Bryan Hooi, Nanyun Peng, Kai-Wei Chang, Jin Lu
Under Review.

[6] Recursive Synthesis for Long-Horizon Terminal Tasks
Zhongzhi Li*, Yucheng Shi*, Zongxia Li*, Ruhan Wang, ***Jiaxi Li***, Anhao Li, Junyao Yang, Lei Ke, Zixun Huang, Yu Wang, Jingyuan Huang, Ninghao Liu, Haitao Mi, Leo Liang
Under Review.

[7] Mitigating Hallucination Through Theory-Consistent Symmetric Multimodal Preference Optimization
Wenqi Liu, ***Jiaxi Li***, Yinwei Wei, Na Zheng, Jianhua Yin, Liqiang Nie, Xuemeng Song†
Published in ***NeurIPS 2025*** (The Thirty-Ninth Annual Conference on Neural Information Processing Systems).

[8] Automating Expert-Level Medical Reasoning Evaluation of Large Language Models
Shuang Zhou*, Wenya Xie*, ***Jiaxi Li****, Zaifu Zhan, …, Yucheng Shi, Ninghao Liu, Zirui Liu, Rui Zhang
Published in ***npj Digital Medicine, 2025***.

[9] GraphVerse: A Comprehensive Visual Graph Reasoning Benchmark for Multimodal Large Language Models
Yuanfu Sun, Yuanhang Ren, Kang Li, Chuanhao Ji, ***Jiaxi Li***, Jiajin Liu, Ninghao Liu, Qiaoyu Tan
Published in ***EMNLP 2026 Main Conference*** (The 2026 Conference on Empirical Methods in Natural Language Processing Main Conference).

[10] Co-adaptive Rule Routing and Constraining for Structured Knowledge Retrieval
Zhisheng Qi, Haoyu Han, ***Jiaxi Li***, Utkarsh Sahu, Ryan A. Rossi, Franck Dernoncourt, Huiyuan Chen, Bing He, Mahantesh M Halappanavar, Jin Lu, Yu Zhang, Tianming Liu, Yu Wang
Under Review.

[11] PRISM: From Feasible Geometry to Routing Capacity in Learning to Optimize
Ke Deng, ***Jiaxi Li***, Geng Yuan, Fei Dou, Jin Lu
Under Review.

[12] Proximal Federated Learning for Body Mass Index Monitoring using Commodity WiFi
Jiaxi Li, Kiran Davuluri, Khairul Mottakin, Zheng Song, Fei Dou, Jin Lu
Published in ***ICASSO@MobiCom 2024*** (ICASSO Workshop of the 30th Annual International Conference on Mobile Computing and Networking).

TECHNICAL SKILLS

Programming: Python, PyTorch, C/C++, Shell Scripting, MySQL.
Infrastructure for Large Language Models (LLMs): verl, vLLM, SGLang, rllm, TRL, Ollama.
Infrastructure for Machine Learning (ML): Linux, Git, Docker, Slurm, Distributed Training Framework (DeepSpeed, Ray, Accelerate).