

YANCHEN LIU

University of Southern California, Los Angeles, CA, USA
+1-213-921-6117 | ✉ liuyanch@usc.edu | 🏠 <https://jamesnulliu.github.io>
Google Scholar: https://scholar.google.com/citations?user=_nKcsagAAAAJ

EDUCATION HISTORY

M.Sc. in Computer Science <i>University of Southern California</i> , LA, CA, US	Sep. 2025 – Jun. 2027
B.Eng. in Computer Science <i>Shanghai University</i> , Shanghai, CN	Sep. 2021 – Jun. 2025

RESEARCH EXPERIENCE

Graduate Research Assistant	Jun. 2025 - Present
------------------------------------	----------------------------

University of Southern California - INK Lab (Los Angeles, CA, US)

Topic: Signals and Environments for Agentic RL; LLM Reasoning

1. Led the research of LLM-based game-playing agents, using replay-anchored judge as the long-horizon reward signal in reinforcement learning; Developed a text-based game-playing harness; Patched slime for training.
2. Led the research of user simulation analysis, designed a benchmark system for individual fidelity of user simulators; Proved that current methods are mixing individualization with population and recency.
3. Participated in the research of LLM-based recommendation agent; Proposed a new token-level self-distillation method for agentic RL on different recommendation tasks.
4. Proposed Soft-Thinking Policy Optimization, which explores the paradigm of LLM's latent-reasoning, including patching vLLM and SGLang, building an optimized Tree-Search generation and Training engine.
5. Enhanced Chain-of-Thought data with segment and token importance evaluated by different methods; Utilized the data for Supervised Fine-Tuning of various reasoning models with unsloth.

Undergraduate Research Assistant	Mar. 2023 - Apr. 2025
---	------------------------------

Shanghai University Cyber Security Lab (Shanghai, CN)

Topic: Vehicle Modeling, Simulation, and Intrusion Detection

1. Deep learning research on cyber security and functional safety in in-vehicle communication system.
2. Developed and optimized a novel gradient descent solver for Multi-Dimensional Hawkes Process; Introduced MDHP-Net for improved timing feature extraction in in-vehicle communication data.

Undergraduate Research Intern	Jun.2023 – Jul. 2024
--------------------------------------	-----------------------------

Shanghai University & East China Air Traffic Control Bureau (Shanghai, CN)

Topic: Super Resolution for Meteorological Data

1. Deployed DL models including PanGu and FourCastNet to improve meteorological monitoring, forecasting, and super-resolution for the East China Air Traffic Control Bureau.
2. Utilized key-point and semantic constraints, combined with feature map fusion, to improve the up-sampling process in super-resolution models.

INDUSTRY EXPERIENCE

Software Development Engineer Intern, AI/ML	Jun. 2026 – Aug. 2026
--	------------------------------

Amazon Lab 126 - Kindle (San Jose, CA, US)

Topic: Benchmark System for Agentic Coding

1. Designed and built a benchmark system for Kindle coding agents. It auto-crawls the history of internal packages into a reproducible task dataset, benchmarks each agent configuration inside an execution harness, and reports performance metrics and analysis.
2. Designed the workflow as a three-stage pipeline: 1) sample generation (provide each sample with a leak-free spec plus a held-out hidden test), 2) agent run (develop on a clean-base tree behind a read/edit/build/self-test loop), and 3) evaluation (an oracle on f2p/p2p, an advisory LLM judge, and a behavioral-equivalence rescue).
3. Accelerated the system with a workspace pool and scheduler for isolation and parallel builds; Developed an anchor algorithm to reuse caches across different dependency snapshots, which turns hundreds of cold multi-minute builds into second-level warm rebuilds.

4. Built a knowledge-update evaluation gate for a GraphRAG knowledge base. Candidate memory and experience are evaluated by multi judges and coding harness, with human-in-the-loop for the final decision.

Machine Learning Engineer Intern

Jul. 2024 - Jun. 2025

Shanghai AI Laboratory (Shanghai, CN)

Topic: LLM Inference Engine, AI Compiler for Graph, Post-training

1. Extended vLLM for LLM inference on in-house TPUs, including researching and adapting Speculative Decoding, Paged-Attention and Continuous-Batching.
2. Developed high-performance kernels with MLIR to ensure seamless compatibility and optimal performance of LLMs on in-house TPUs.
3. Conducted research on LLM knowledge injection and fine-tuning for kernel fusion and translation across different hardware platforms.

HONORS AND AWARDS

1. [2024] *First Prize and Group Competition Award* in 2024 ASC Student Supercomputer Challenge Global Final.
2. [2022] *First-Class Academic Scholarship* for outstanding academic performance, Shanghai University.

PUBLICATIONS

† Equal contribution. * Corresponding author.

2026 -----

1. [ICLR Poster] S. Wang, **Y. Liu**, X. Ren. "Segment-Level Attribution for Selective Learning of Long Reasoning Traces".
2. **Y. Liu**†, X. Che†, B. Joshi†, R. Fukuda, J. Liu, R. Li, Y. Ge, H. Yan, X. Ren. "Replay-Grounded Signals for Long-Horizon Civilization Agents".
3. **Y. Liu**†, J. Liu†, B. Joshi†, X. Che, R. Li, X. Ren. "Attribution Trials: Auditing Whether User Simulators Represent the Intended Individual".
4. B. Joshi†, **Y. Liu**†, S. Swayamdipta, X. Ren, T. Paek, R. Koncel-Kedziorski. "Learning to Seek User Information for Personalization via Multi-turn Conversations".
5. J. Liu†, **Y. Liu**†, L. He, Z. Zhang, X. Che, L. Ge, F. Qian. "When Repairs Stop Being Recipes: Re-Grounding Repair Memory for Coding Agents".
6. [Arxiv] X. Che†, H. Yan†, **Y. Liu**†, H. Liu, R. Li, A. Shi, H. Wang, J. Liu. "Does Learning to Predict the World Help Agents Act? Auditing World-Model Post-Training".
7. [Arxiv] X. Che†, Y. Ge†, S. Li†, **Y. Liu**†, H. Yan†, X. Lei†, Y. Wang, Z. Dong, Y. Yao, Q. Xie, L. Zhu, J. Liu*. "GameLogicBench: Evaluating Coding Agents on Runtime Game Logic with Tick-Level State Assertions".

2025 -----

8. [Arxiv] J. Lv, X. He, **Y. Liu**, X. Dai, A. Shen, Y. Li, J. Hao, J. Ding, Y. Hu, S. Yin. "HPCTransCompile: An AI Compiler Generated Dataset for High-Performance CUDA Transpilation and LLM Preliminary Exploration".
9. [Arxiv] Q. Liu†, **Y. Liu**†, R. Li, C. Cao, Y. Li*, X. Li*, P. Wang, R. Feng, S. Bu. "MDHP-Net: Detecting an Emerging Time-Exciting Threat in IVN".

2024 -----

10. [IEEE Internet of Things Journal] Q. Liu, X. Li, K. Sun, Y. Li*, **Y. Liu***. "SISSA: Real-Time Monitoring of Hardware Functional Safety and Cybersecurity With In-Vehicle SOME/IP Ethernet Traffic".
11. [MDPI Future Internet] X. Li, R. Li, **Y. Liu**. "HP-LSTM: Hawkes Process-LSTM-Based Detection of DDoS Attack for In-Vehicle Network".

TECH STACK

Programming Languages

C++ (cmake, bazel, vcpkg, gcc, clang, msvc, msbuild, simd); CUDA (simt, nsight compute); Python (pip, conda, uv)

AI/ML

1. **Post-Training & (Agentic) RL** (lora, sft, opd, rlhf, ppo, grpo, daplo, opsd, sdpo, agentic rl; trl, verl, slime)
2. **Inference Acceleration** (TP, prefilling, decoding, prefix caching, continuous-batching, Chunked Prefill, Speculative Decoding, Quantization, Paged Attention, Flash Attention, MQA, GQA, MLA)
3. **Training Acceleration** (Mixed Precision, Gradient Checkpointing, ZeRO, DP, DDP, FSDP, MP, PP)