

Personal Information

Name: Yuxuan Liu

Phone: (+86) 139-0923-7776

Email: l-yx23@mails.tsinghua.edu.cn

GitHub: github.com/Rosist-Sallina

Education

Tsinghua University, B.Eng. & B.Econ. Sep. 2023 – Jul. 2027 (expected)

- B.Eng. in Computer Science and Technology & B.Econ. in Economics and Finance (dual degree) **GPA: 3.3 / 4.0**
- **Selected coursework:** Software Engineering (A), Fundamentals of Computer Graphics (A), Principles of Computer Networks (A⁻), Network Security Engineering and Practice (A⁻), Machine Learning in Finance (A⁻), Student Research Training (SRT) (A)



Research Interests

Efficient inference for LLMs; agent memory and self-evolution

Publications & Research Experience

Boosting Fine-Grained Visual Perception via Multi Granularity Data Synthesis and Local-to-Global Curriculum Learning

- Co-first author | **EMNLP 2026 (CCF-B)**, under revision Feb. 2026 – Present
- Proposed multi-granularity data synthesis and a local-to-global curriculum learning strategy to improve fine-grained visual perception in multimodal LLMs
- Led experiment design, data synthesis and evaluation; co-wrote the paper

From Cache to Belief: Online Probabilistic Memory for Long-Context Language Models

- First author (student) | **NeurIPS 2026 (CCF-A)**, under review Feb. 2026 – Jul. 2026
- Probabilistically modeled cached memory as an online-updated belief, improving the reliability of long-range memory in long-context settings
- Designed the evaluation protocol, conducted system-level experiments and result analysis

A Scalable Benchmark Test Suite for Dynamic Multi-Objective Optimization with a Changing Number of Objectives

- Third author | **PPSN 2026 (CCF-B)**, accepted Aug. 2025 – Dec. 2025
- Identified gaps in scalability and problem-feature coverage in existing dynamic multi-objective benchmarks and proposed a solution
- Built test suites; designed and evaluated comparative experiments

CSI-Diff: A Contextual Generative Foundation Model for Wireless Channel Prediction

- Third author | **IEEE TMC (CCF-A)**, under review Oct. 2025 – Mar. 2026
- Use diffusion models for wireless channel prediction and a MoE architecture to improve model
- Contributed to experiment design, data processing and evaluation; co-wrote the paper

A Device-Cloud Collaborative Multimodal Inference Method and System

- **First inventor** | Chinese invention patent (under substantive examination) App. No. 2026105681825
- Dynamically allocated perception and reasoning across device and cloud, cutting on-device inference cost while strengthening privacy protection

Projects

Graph-Generating LLMs for Drug Discovery, SRT Project, Dept. of CST Jun. 2024 – Jun. 2025

- Built the training-data and prompt synthesis pipeline; ran RL-based model optimization
- Post-trained DeepSeek-R1 and other models for cross-model comparison; proposed a new benchmark for LLMs on rule-constrained graph generation

Internships

- ZTE Corporation**, AI Research Intern Aug. 2025 – Mar. 2026
- Researched inference acceleration for multimodal LLMs and finetuned visual perception
 - Independently built a data synthesis pipeline and evaluation framework, supporting multiple rounds of ablation studies
- Meituan**, Intern, Autonomous Vehicle Business / World Model Jun. 2026 – Jul. 2026
- Contributed to autonomous-vehicle world-model development, including multimodal data processing and model training
 - Evaluated and optimized the perception module, improving system robustness in complex environments

Service & Leadership

- Student Association for Science and Technology, Dept. of CST**, Vice President (AI Agent Division) Jun. 2025 – Present
- Led (as team leader) the development of the official game player/visualizer for the 29th AI Agent Contest, used by all participants university-wide
 - Organized university-wide AI agent competitions, coordinating problem design, scheduling and technical support
- Participated in research fieldwork in Japan and volunteer teaching in Qinghai, serving as the lead for communications and teaching activities.

★ Honors & Awards

- Outstanding Social Work Scholarship, Tsinghua University 2025
- Third Prize, Youth Digital Innovation Competition (Beijing Xicheng) 2025
- Winning Prize, 7th “Bambu Lab Cup” Software Design Competition, Dept. of EE 2024

Skills

- Proficient in Python; familiar with C/C++/C# and Rust.
- Skilled in PyTorch and Transformers, with end-to-end engineering experience in LLM training, fine-tuning and evaluation.
- Able to independently deliver full-stack projects; familiar with multi-GPU distributed training and experiment management under Linux, with experience in Git collaboration and code review.
- Chinese (native), English (proficient academic reading and writing).

Summary

Third-year undergraduate at Tsinghua University pursuing a dual degree in Computer Science and Technology and Economics & Finance. Research interests include LLM reasoning, agentic capabilities, agent memory and self-evolution. Contributed to four research projects over two years, producing three international conference papers (one accepted at PPSN, one under revision for EMNLP, and one under review at NeurIPS) and a journal paper (IEEE TMC under review) and a national invention patent as first inventor. Experienced in the full research cycle from problem formulation and experimental design to paper writing, with solid engineering skills in LLM training and evaluation. Seeking graduate research opportunities in efficient inference and memory mechanisms for large language models.