

Yixiao Wang

 yixiao.site |  yixiaowang@sjtu.edu.cn |  +8613390606996

ACADEMIC BACKGROUND

School of Computer Science, Shanghai Jiao Tong University — Software Engineering

Core GPA: 3.93 / 4.3 (Ranking: 7 / 89)

Performance in Fundamental Science Courses: Circuit Theory (Honors, 99) Linear Algebra (Honors, 97) University Physics (Honors I, 96; Honors II, 97) Mathematical Analysis (Honors, 95)

Performance in Major Courses: Software Engineering Practice (99) Data Structures (97) Internet Application Development Technologies (96) Fundamentals of Computer Systems (I, 93; II, 94)

RESEARCH EXPERIENCE

ELORA: Efficient LoRA and KV Cache Management for Multi-LoRA LLM Serving

- ELORA is a high-performance cache management system designed for serving multi-LoRA large language models. It significantly optimizes inference efficiency, particularly by reducing TTFT (Time-To-First-Token) and TPOT (Tokens-Per-Output-Token). By unifying the management of LoRA adapters and their corresponding KV caches, ELORA addresses common inefficiencies in existing systems, such as ineffective cache usage and unreasonable memory fragmentation.
- The system introduces a Dependency-aware Cache Manager and a Performance-driven Cache Swapper, which dynamically manage LoRA and KV caches based on usage dependencies. This enables smarter allocation and utilization of HBM resources, thereby reducing latency and improving throughput.

PROJECT EXPERIENCE

LSM-KV: KVStore using Log-structured Merge Tree

[Link to Repo](#)

- Built a high-performance key-value store, with skip lists & Bloom filters to optimize access and queries.
- Applied key-value separation and multi-level indexing to reduce write amplification and disk I/O, improving performance under high concurrency.

Tiger Compiler: A Complete Compiler Implementation

[Link to Repo](#)

- Developed a full compiler pipeline (lexical, syntax, semantic analysis, and code generation).
- Implemented frontend with **Flexc++**/**Bisonc++** and backend with **LLVM**, producing optimized **x86** assembly.

HONORS & AWARDS

Zhiyuan Honor Scholarship, Shanghai Jiao Tong University

Dec 2022, Dec 2023, Dec 2024

Class C Scholarship, Shanghai Jiao Tong University

Dec 2023, Dec 2024

OTHERS

- **Teaching Assistant**: Programming Design and Methodology (Honor, C++), Fall 2024–2025
- **Skills**: Python, C++, LaTeX, vLLM
- **Certificate**: CET-6 (Score: 583)