

Jingyu LU

jingyu.lu@anu.edu.au [GitHub](#) [Google Scholar](#) [LinkedIn](#) [Website](#)

Education

Australian National University	Dec 2026 (expected)
Master of Computing (Advanced)	6.7/7.0
University College Dublin	Sep 2020 – Jul 2024
Beijing-Dublin International College	
Bachelor of Science in Software Engineering	First Class Honours

Research Interests

Computer Architecture and Systems, Parallel Computing, High-Performance Computing.

Publications

- Li, Tianhao, **Jingyu Lu**, Chuangxin Chu, Tianyu Zeng, Yujia Zheng, Mei Li, Haotian Huang, Bin Wu, Zuoxian Liu, Kai Ma, et al. SciSafeEval: a comprehensive benchmark for safety alignment of large language models in scientific tasks. *AAAI AICS Oral*. [arXiv:2410.03769](#). 2024.
- Tianhao Li, Tianyu Zeng, Yujia Zheng, Chulong Zhang, **Jingyu Lu**, Haotian Huang, Chuangxin Chu, Fang-Fang Yin, Zhenyu Yang. Synthetic Poisoning Attacks: The Impact of Poisoned MRI Image on U-Net Brain Tumor Segmentation. *AAPM Oral, ICLR Workshop Poster*. [arXiv:2502.03825](#). 2025.
- Yujia Zheng, Tianhao Li, Haotian Huang, Tianyu Zeng, **Jingyu Lu**, Chuangxin Chu, Yuekai Huang, Ziyong Jiang, Qian Xiong, Yuyao Ge, Mingyang Li. Are All Prompt Components Value-Neutral? Understanding the Heterogeneous Adversarial Robustness of Dissected Prompt in Large Language Models. *EACL Oral*. [arXiv:2508.01554](#). 2025.
- Jingyu Lu**, Yujia Zheng, Chuangxin Chu, Tianyu Zeng, Yuxuan Zhao, Tianhao Li. Towards Emotion-Aware Healthcare: A Comprehensive Review of Multimodal Emotion Recognition Technologies in Medical Practice. *ICDT Oral*. [DOI:10.1002/sdtp.18963](#). 2024.

Research Experience

- Research on Prefetcher Optimization and Memory Interference Modeling** Feb 2026 – Present
Supervisor: Prof. [Shoaib Akram](#) The Australian National University
Designed a cycle-accurate attribution framework to quantify and decompose prefetch-induced interference on demand requests in multi-core systems, distinguishing cross-core harm across the memory hierarchy and translating memory-side delays into ROB-visible performance impact, enabling fine-grained identification of harmful prefetch behaviors and providing insights for prefetcher optimization.
- Research on Multi-Core Extensions for MLP Stacks** Feb 2026 – Present
Supervisor: Prof. [Shoaib Akram](#) The Australian National University

Extended MLP stacks to multi-core scenarios and optimized the selection of the computational denominator. MLP stacks is an analytical method that decomposes memory-level parallelism (MLP) into core and prefetch, measures hit/miss parallelism and visualizes the CPI contribution of each memory hierarchy level.

Research on Deep Learning for Biomarker Identification in Flow Cytometry Jul 2025 – Dec 2025

Supervisor: Prof. [Dan Andrews](#)

The Australian National University

Developed a CNN-based deep learning framework to analyze flow cytometry data, identifying predictive biomarkers for disease diagnosis. Designed customized CNN models to address batch effects and data heterogeneity, enhancing feature extraction and clustering for improved interpretability and performance.

Research on the Security of LLMs in Scientific Tasks

Aug 2024 – Jul 2025

Supervisor: Prof. Qiang Zhang

Zhejiang University

Developed SciSAFEVAL, a benchmark to evaluate and enhance the safety of general and scientific large language models. Designed comprehensive metrics and test sets covering chemistry, biology, medicine, and physics. Implemented safety alignment strategies to ensure models adhered to predefined standards.

Projects

Memory Allocator & Garbage Collector

[Code](#)

Implemented and benchmarked a series of heap allocator optimizations—including metadata reduction, constant-time coalescing via boundary tags, and size-class segregated free lists—progressively improving allocation throughput and memory utilization. The allocator was further extended with a garbage collector.

xv6 Operating Systems

[Code](#)

Extended the xv6 kernel across four subsystems, adding custom system calls with tracing and process inspection, exploring process scheduling policies, extending virtual memory management, and introducing kernel-level thread support with thread lifecycle management.

Distributed, Shared-Memory and GPU Advection Solver

Designed and optimized a high-performance distributed advection solver using MPI, OpenMP, and CUDA to achieve efficient parallel computing, with performance modeling and validation conducted on [NCI Gadi](#).

Distributed Computing: Built an MPI-based solver with optimized halo exchanges, supporting 1D/2D process grids and non-blocking communication. [MPI Code](#)

Shared-Memory Optimization: Implemented OpenMP and CUDA solvers, enhancing performance via loop scheduling and GPU thread tuning. [OpenMP/CUDA Code](#)

Contribution to Nvidia [Garak](#)

Contributed to NVIDIA Garak, a generative AI red-teaming and assessment toolkit, by improving components for probing and evaluating LLM vulnerabilities.

Honors and Awards

SoCo Top Students, Semester 2, 2025

College Scholarship, 2023

College Scholarship, 2022

University Scholarship, 2022