Zedong Peng

zedongpeng1@gmail.com zedongpeng.com

Education

M.Eng., Electronic and Information Engineering, *Shanghai Jiao Tong University* [transcript]
School of Computer Science. Advised by Prof. Jieru Zhao.
Shanghai, China
2024 – 2027 (expected)

B.Eng., Information Engineering, *Shanghai Jiao Tong University* [transcript]
School of Electronic Information and Electrical Engineering. Outstanding Graduate.
Shanghai, China
2020 – 2024

Experience

Research Intern, *National Center of Technology Innovation for EDA*
Advised by Prof. Qiang Xu and Prof. Jieru Zhao.
Nanjing, China
July 2024 – 2025

- The HLS part of LLM4EDA: LLM post-training and data construction for a domain-specific language.

Machine Learning Intern, *Intel*
Advised by Yuan Zhou.
Shanghai, China
Nov 2023 – July 2024

- Spark SQL native engine Gluten: optimizing the CI pipeline.
- LLM framework optimization: Intel device support for vLLM.
- LLM applications: text to SQL.

Winter School, *Data Science Institute, Imperial College London*
London, UK
Jan 2024 – Feb 2024

- Applied vector representations to analyze language features in COVID-19 text.

Research Intern, *Shanghai Jiao Tong University*
Advised by Prof. Liyao Xiang.
Shanghai, China
July 2022 – July 2024

- Security-focused machine learning research.

Publications

MemoryTax: The Bitter Lesson of Agent Memory [paper] [arxiv] [code] 2026
Zedong Peng
arXiv preprint arXiv:2610.29823

DiffHLS: Differential Learning for High-Level Synthesis QoR Prediction with GNNs and LLM Code Embeddings [arxiv] [pdf] [code] [poster] 2026
Zedong Peng, Zeju Li, Qiang Xu, Jieru Zhao
IEEE/ACM International Symposium of EDA (ISEDA)

HIPPO: A Hierarchy-Preserving and Noise-Tolerant Pre-HLS Power Modeling Framework for FPGA [paper] [pdf] 2025
Zefan Lin, **Zedong Peng**, Mingzhe Gao, Jieru Zhao, Zhe Lin
Proceedings of the 44th IEEE/ACM International Conference on Computer-Aided Design (ICCAD)

ForgeHLS: A Large-Scale, Open-Source Dataset for High-Level Synthesis [arxiv] [pdf] [code] [data] 2025
Zedong Peng, Zeju Li, Mingzhe Gao, Qiang Xu, Chen Zhang, Jieru Zhao
arXiv preprint arXiv:2507.03255

DeepCircuitX: A Comprehensive Repository-Level Dataset for RTL Code Understanding, Generation, and PPA Analysis [paper] [arxiv] [pdf] [data] 2025
Zeju Li, Changran Xu, Zhengyuan Shi, **Zedong Peng**, Yi Liu, Yunhao Zhou, Lingfeng Zhou, Chengyu Ma, Jianyuan Zhong, Xi Wang, Jieru Zhao, Zhufei Chu, Xiaoyan Yang, Qiang Xu
IEEE International Conference on LLM-Aided Design (ICLAD)

A verifiable and privacy-preserving federated learning training framework [paper] [pdf] 2024
Haohua Duan, **Zedong Peng**, Liyao Xiang, Yuncong Hu, Bo Li
IEEE Transactions on Dependable and Secure Computing

Projects

llama.cpp FPGA backend [blog] In progress
A llama.cpp backend that runs LLM inference on an AMD Alveo U280 FPGA card; twelve models from four families run on it. The code is private for now.

Robotic Dog Development Shanghai, China
Advised by Prof. Weiming Zhao and Xiaomi Technology. Nov 2023 – Nov 2024

- Implemented algorithms that improved locomotion control and user interaction.

Honors and Awards

First-class Postgraduate Academic Scholarship, Shanghai Jiao Tong University 2025
Outstanding Graduate, Shanghai Jiao Tong University 2024
2nd Prize, 13th National University Mathematics Competition 2021
1st Prize, SJTU “Yunhan Cup” RobotMaster Intramural Competition 2021
Liu Yongling Foundation Scholarship 2021
Advanced Individual in Social Practice, Shanghai Jiao Tong University 2021

Skills

Programming: C++, Python, SQL, Verilog, Shell, Docker, MATLAB
Languages: Chinese (native), English (CET-6, 481), German (A1)

Updated October 2026