

Zihao Wang

404-952-0911 | ziw140@ucsd.edu | gdfwj.github.io | github.com/gdfwj

EDUCATION

University of California, San Diego

Master of Science in Computer Science

La Jolla, CA

Sep. 2024 – Now

- GPA: 4.0/4.0
- Key courses: Systems for LLMs and AI Agents, Differentiable Programming (A+), ML: Learning Algorithms (A+)

Beihang University

Bachelor of Engineering in Computer Science and Technology

Beijing, China

Sep. 2020 – Jun. 2024

- GPA: 3.92/4.0 (WES)
- Key courses: Computer Organization (96), Operating Systems (95), Object Oriented Design Construction (98), Compiler Technology (95), Signal and System (98), Computer Network Experiments (97)

RESEARCH INTERESTS

Edge/IoT systems; human-IoT interaction; reliability and efficiency; privacy-aware deployment; LLM-enabled interactive systems.

Currently: edge-client collaborative health agents in System Energy Efficiency Lab (UCSD), advised by Prof. **Tajana Šimunić Rosing** and **Ye Tian**.

SCHOLARSHIPS & AWARDS

- **2021, 2022, 2023** Academic Merit Scholarship (Awarded to undergraduate students for outstanding GPA throughout the school year)
- **2022, 2023** Academic Competition Scholarship (Awarded to undergraduate students who achieved outstanding rankings in national or university-level discipline competitions)

PUBLICATIONS

1. Ye Tian*, **Zihao Wang***, Onat Gungor, Xiaoran Fan, Tajana Rosing. “LifeAgentBench: A Multi-dimensional Benchmark and Agent for Personal Health Assistants in Digital Health.” *International Conference on Learning Representations (IJCAI)*, 2026 (under review)
2. Ye Tian, Xiaoyuan Ren, **Zihao Wang**, Onat Gungor, Xiaofan Yu, Tajana Rosing. “DailyLLM: Context-Aware Activity Log Generation Using Multi-Modal Sensors and LLMs.” *22nd IEEE International Conference on Mobile Ad-Hoc and Smart Systems (MASS)*, 2025

RESEARCH EXPERIENCE

Efficient Edge-Client Collaborative Agent Systems

Oct. 2025–Present

Research Assistant, Advisors: Ye Tian and Tajana Rosing

System Energy Efficiency Lab, University of California, San Diego

- Designed an edge-client collaborative agent architecture that separates raw sensor data on edge devices from LLM inference on a local client, enabling parallel data access and reasoning.
- Developed a compact tool interface for communication and a tool-prompting mechanism that enables small models to reliably retrieve and reason over long-term, multi-domain health data under limited context length and latency constraints.
- Built end-to-end evaluation on a health QA benchmark to measure correctness across tool execution, data retrieval, and final answers, linking model performance with system-level constraints.
- Exploring alternative deployment strategies for severely constrained client resources, including a cloud-based planner with a lightweight local model for response synthesis.
- Privacy-aware design: in the default setup both edge and client are local, and in the alternative setup the cloud only plans from the query while tool outputs are visible only to the local client.

Multidimensional QA Benchmark for Health Agent (LifeAgentBench)

Jul. 2025–Sep. 2025

Research Assistant, Advisors: Ye Tian and Tajana Rosing

System Energy Efficiency Lab, University of California, San Diego

- Built **LifeAgentBench**, a large-scale benchmark (22,573 questions) spanning single-domain retrieval to long-horizon, cross-dimensional reasoning over lifestyle health records in a structured database.
- Built an automated evaluation pipeline for raw-table reasoning and schema-to-SQL execution on 11 LLMs, highlighting failures that impact downstream health analysis.
- Implemented **LifeAgent**, a training-free multi-step agent for health data retrieval, outperforming baselines on challenging long-horizon, cross-dimensional questions;
- Conducted a case study showing effective question decomposition into sub-queries and that retrieval accuracy is important for high-quality assessments and recommendations.
- Released code on Github and paper on arXiv, also submitted to IJCAI (under review).

DailyLLM: Multimodal Activity Log Generation

Jan. 2025–Apr. 2025

Research Assistant, Advisors: Ye Tian and Tajana Rosing

System Energy Efficiency Lab, University of California, San Diego

- Developed **DailyLLM**, a lightweight LLM-based system for context-aware activity log generation and summarization from multimodal smartphone and smartwatch sensors.
- Organized IMU-based datasets for daily activity understanding, including data preprocessing, augmentation and data construction for model training.
- Built multi-task LoRA fine-tuning and evaluation pipelines to balance model capability and system efficiency.
- Deployed the system on Raspberry Pi for real-time inference, and resolved an on-device quantized-kernel issue by switching to a supported quantization bit-width.
- Released code, data, and a project website (GitHub); paper accepted to **IEEE MASS 2025**.

Reconstruction of perceived face images from fMRI

Apr. 2023–Dec. 2023

Research Assistant, Advisors: Hui Zhang

Medical-Engineering Innovation Research Institute

- Developed a GAN-based pipeline to reconstruct face stimulus images from fMRI signals.
- Proposed a pretraining strategy on CelebA with VGGFace feature conditioning to mitigate limited paired fMRI-image data.
- Designed a double-flow discriminator to improve reconstruction fidelity.
- Manuscript released on arXiv.

SELECT COURSE PROJECTS

Compiler Design and Implementation

Sep. 2022 – Dec. 2022

- Implemented a SysY-to-MIPS compiler covering lexical and grammatical analysis, IR generation, object code generation, and classic optimizations; realized full functionality.
- Gained experience structuring a multi-file codebase with complex logic.

Computer Organization Project: Five-Stage Pipelined MIPS CPU

Sep. 2021 – Jan. 2022

- Implemented a five-stage pipelined MIPS CPU in Verilog with interrupt/exception handling and simple I/O; built a single-cycle baseline then upgraded to a pipelined design.
- Strengthened understanding of CPU microarchitecture and hardware coding.

Operating Systems Project: MIPS-based OS

Mar. 2022 – Jun. 2022

- Built a simple OS from scratch on MIPS: boot, printf, memory management, process management, user/kernel modes, system calls, file system, and shell.
- Implemented basic threading management and synchronization system calls.