

# Zimu Wang

Berkeley, CA | zm.wang@berkeley.edu | 510-603-9827 | www.linkedin.com/in/zimu-wang-32b62a289

## Education

University of California, Berkeley, BA in Computer Science

Graduation: May 2027

- GPA: 3.94/4.0

## Skills

Languages: Python, Java, SQL

ML/LLM: PyTorch, HuggingFace (Transformers, Datasets), VeRL, vLLM, RLHF

Data/Tools: NumPy, Pandas, scikit-learn, Matplotlib, Git, Linux, LaTeX

Technical Foundations: Machine Learning, Reinforcement Learning, Neural Networks, Computer Vision, Data Structures, Algorithms, Probability, Optimization Models, Linear Algebra, Object-Oriented Programming

## Selected Publications

- *EffiSkill: Automated Code Efficiency Optimization via Skill Mining* [arXiv]  
**Zimu Wang**, Yuling Shi, Mengfan Li, Zijun Liu, Jie M. Zhang, Chengcheng Wan, Xiaodong Gu.  
Under review at ASE 2026.
- *CodeEffiJudge: Can LLMs Reason about Code Efficiency?*  
Mengfan Li, Yuling Shi, **Zimu Wang**, Zijun Liu, Chengcheng Wan, Yuting Chen, Xiaodong Gu.  
Under review at ICML 2026.

## Experiences and Projects

Researcher Shanghai Jiao Tong University

Spring 2026

- Developed **EffiSkill**, a two-stage LLM framework for automated code-efficiency optimization that mines reusable optimization skills from slow/fast program pairs for execution-free rewriting.
- Built scalable inference and evaluation pipelines on **EffiBench-X** (623 tasks), supporting Top- $k$  candidate generation, public/private test ranking, and offline runtime evaluation across Python and C++.
- Improved optimization success rate over the strongest baseline by **3.69–12.52 points** across model and language settings; **arXiv**: 2603.27850

Research Assistant, University of California, Berkeley

Fall 2025

- Built an end-to-end data processing pipeline on a **1.2M-row** gig-economy dataset; produced **2 clean, analysis-ready datasets** with validation and feature engineering
- Conducted **30+ linear regressions** and generated **40+ analysis plots** (statsmodels / scikit-learn) with robustness checks to support collaborator-facing findings.

Researcher, Shanghai Jiao Tong University

Summer 2025

- Implemented an **SFT + GRPO training pipeline** for code LLMs using KodCode (48k examples), including prompt standardization, reward parsing, and automated evaluation harnesses
- Trained for **100k+ steps on 8×A100 GPUs**; ran controlled ablations; established a reproducible RL fine-tuning workflow and evaluation stack

Teaching Assistant, University Of California, Berkeley

Fall 2024

- Mentored students in foundational Data Science and Python programming, leading discussion sections and providing technical guidance on homework and projects.

Computer Vision with Machine Learning Portfolio

[GitHub]

- Implemented diffusion/flow-matching models and NeRF pipelines, covering conditional UNets, CFG sampling, inpainting, camera pose estimation, and volume rendering.