# Han Zhang

Email: han.harry.zhang@gmail.com

## EDUCATION

**University of California, Los Angeles(UCLA)** *2024 - 2029(Expected)*
Computer Science PhD student, Advised by Prof. Wei Wang
GPA: 4.0/4.0

**Shanghai Jiao Tong University (SJTU)** *2020 - 2024*
Bachelor of Engineering in Computer Science
Member of **ACM Class**, a selective CS program for the top 5% students.
Major GPA: 91.87/10 (5/35)

## RESEARCH INTEREST

Agentic RL, Agentic AI system, reinforcement learning–based fine-tuning methods, and techniques for optimizing agent behavior and task performance.

## PUBLICATION

**ARLArena: A Unified Framework for Stable Agentic Reinforcement Learning**
*Xiaoxuan Wang*, **Han Zhang***, Haixin Wang*, ... Jason Cong, Yizhou Sun, Wei Wang ICML 2026 in submission, * denotes equal contribution*

**Think, Refine, and Diversify: Enabling LLM Agent to Retrieve Tools Dynamically**
***Han Zhang**, Kyle Zheng, Renliang Sun, Chenchen Ye, Wei Wang* *Manuscript*

**Protein Large Language Models: A Comprehensive Survey**
*Yijia Xiao, Wanjia Zhao, ..., **Han Zhang**,..., James Zou, Yizhou Sun, Wei Wang* *EMNLP-Findings*

**GFS: Graph-based Feature Synthesis for Prediction over Relational Databases**
***Han Zhang**, Gan Quan, David Wipf, Weinan Zhang* *VLDB 2024 TaDA workshop*

**4DBInfer: A 4D Benchmarking Toolbox for Graph-Centric Predictive Modeling on Relational DBs**
*Minjie Wang, Quan Gan, David Wipf, ..., Jiahang Li, **Han Zhang**, ..., Muhan Zhang, Weinan Zhang, Christos Faloutsos, Zheng Zhang* *Neurips 2024 Datasets and Benchmarks Track*

**Synthesize, Retrieve, and Propagate: A Unified Predictive Modeling Framework for Relational Databases**
*Ning Li, Kounianhua Du, **Han Zhang,** Quan Gan, Minjie Wang, David Wipf, Weinan Zhang Preprint*

## WORK EXPERIENCE

**Data and Research Intern, Optum AI**
*Mentored by Kimmo Kärkkäinen* *Jun 2025 - Sep 2025*

- Built a two-stage LLM finetuning pipeline for ICD-10-CM diagnosis prediction
- Developed specialized loss functions in LLM finetune stage—including hierarchy-aware contrastive learning and dynamic confidence loss with hard negative sampling—to better distinguish similar ICD-10-CM codes and control generation end-points
- Scaled dataset from 100K → 2.5M+ samples using PySpark for distributed preprocessing and augmentation, conducting controlled experiments on dataset size vs. user diversity, and analyzing precision/recall trade-offs

**Applied Scientist Intern, Amazon AI Lab Shanghai**
*Mentored by Dr. Minjie Wang and Quan Gan* *Mar 2024 - Aug 2024*

- Proposed *GFS* and integrated into Project 4DBInfer (Multi-table-benchmark) under an internal DGLAI repo, enhancing its capability in downstream prediction tasks.
- Explored and implemented novel strategies for restructuring relational databases (e.g., adding dummy tables, splitting large tables) before graph construction, improving performance in downstream prediction tasks.
- Developed new features and resolved multiple bugs in Project 4DBInfer, contributing to the robustness and efficiency of the framework.

## RESEARCH EXPERIENCE

**SCAI Lab, University of California, Los Angeles**
*Supervised by Prof. Wei Wang* *Oct 2024 - Now*

- **Agentic RL Benchmark (On Going)**
  - Designed and implemented an open Agentic RL benchmark to evaluate 10+ RL-based fine-tuning algorithms(GRPO, GIGPO, DAPO...) of LLM agents across horizon, memory, tool-use, and multi-agent axes in multi-turn agentic RL settings.
  - Covering diverse tasks—including search agents, multi-modal agent, webshop agents, embodied agents, math agent and coding agent.
  - Developed an agentic RL diagnostic framework to analyze multi-turn agent behavior across diverse tasks, enabling us to extract general training recipes and identify task-specific algorithmic strategies that most effectively improve agent performance.
  - Added the framwork with Asynchronous Agentic RL training support, and extended the diagnostic suite to analyze how different RL algorithms behave and converge under asynchronous training dynamics.
- Designed and implemented dynamic tool retrieval mechanisms that enable LLMs to iteratively generate, refine, and diversify tool descriptions, significantly improving retrieval completeness and performance in complex problem-solving tasks.

**Ma Lab, Carnegie Mellon University**
*Supervised by Prof. Jian Ma* *Aug 2023 - Mar 2024*

- Focused on applying advanced machine learning techniques to analyze complex biological data, particularly in computational genomics.
- Utilized Convolutional Neural Networks (CNNs) and transformer-based models to extract meaningful features from high-dimensional genomic sequences and multi-modal biological data.
- Develop novel attention layer and transformer variants to incorporate 2D Hi-C data, aiming to predict gene expression. Key achievements include interpreting model weights, benchmarking gene expression outputs, and identifying biological patterns.

**APEX Lab, Shanghai Jiao Tong University & AWS Shanghai AI Lab**
*Supervised by Prof. Weinan Zhang and Dr. David Wipf* *Jul 2022 - Jul 2023*

- (Jul 2022 - Dec 2022) Developed a feature space alignment method for tabular data using hard negative contrastive pretraining. This involved identifying hard negatives based on row similarities and employing contrastive learning with anchors, their corrupted views, and hard negatives to achieve a more uniform and well-aligned feature representation.
- (Dec 2022 - Jul 2023) Led the *GFS* project, innovating in deep learning over relational databases. Addressed the complex challenge of column predictions in target tables using surrounding relational data. Developed a novel framework that interprets databases as heterogeneous graphs, integrating single-table models for both embedding and prediction. This approach streamlined feature engineering, bypassing traditional labor-intensive processes. Demonstrated enhanced performance over conventional methods in experiments with four real-world datasets, showcasing GFS's effectiveness and overcoming previous methods' limitations.

## PROJECTS

**Java-and-C-like Language Compiler (∼10K lines in Java)** [github]

The compiler can convert codes of program to an AST, then after semantic check to LLVM-like IR, and eventually to RISC-V assembly. Few optimizations were done to the Back-end, including mem2reg, graph-coloring.

**Ray tracer(∼4K lines in Rust)** [github]

Construct a raytracer renderer in rust, you can set the position parameters of the objects and the camera parameters yourself to get the rendered real-world like image through the renderer.

**RISC-V CPU of Tomasulo Architecture (∼3K lines in Verilog)** [github]

I-cache, and a 2-bit saturating counter branch predictor implemented to improve the performance. The CPU could run successfully on an FPGA board.

## OTHER EXPERIENCE

**Participants of Chinese Physics Olympiad(CPhO), Xuejun High School**
*Jul 2017 - Oct 2019*

I Have won second prize twice in Zhejiang Province, and I have learned a lot of knowledge about physics. This experience greatly improved my learning ability and helped me get grades 100 and grade 95 in my Physics (1) and Physics (2) courses at SJTU

## HONORS AND AWARDS

**Ruiyuan-Hongshan Scholarship**, top 5% students in ACM class *2023*

**Jin Long Yu Scholarship** *2021*

**Zhiyuan Honorary Scholarship**, top 5% students in SJTU *2021, 2022, 2023*

## SKILLS

Programming Languages: Python, C/C++, Java , Rust

Hardware: Verilog

Math: Calculus, Linear Algebra, Mathematical Logic, Measure Theory, Graph and Combinatorics