Zixuan Wang

wang-zx23@mails.tsinghua.edu.cn https://wannabeyourfriend.github.io/

EDUCATION

Tsinghua University, Beijing, China **2023.9 - Present**

B.S. (in progress), Department of Physics (Advisor: Prof. Jifeng Dai)

- **GPA:** 3.9/4.0 **(rank: 9/69)**
- **Major:** Mathematics and Physics**, Minor:** Artificial Intelligence
- **Selected Course:** Deep Learning (Graduate Course) **A+,** 3D Visual Computing (Graduate Course) **A+,** Deep Reinforcement Learning (Graduate Course) **A,** Advanced Computer **Vision (Graduate C**ourse) **B+,** Computer Graphics **A,** Programming **A-,** Object Oriented Programming **A-,** Fundamental Physics **(I) A-, (II) A+, (III) A, Linear Algebr**a **B+,** Complex Analysis **A,** Statistical Inference **B+,** Advanced Calculus (I) **B+,** (II) **A-,** Probability Theory **A-,** Discrete Mathematics **A-**

University of California, San Diego, CA, USA **2025.9 - 2025.12**

Exchange Student at Marshall College, program funded by **UCEAP**

- **GPA:** 3.9/4.0
- **Major:** Computer Science
- **Selected Course:** MATH170A Linear Numerical Analysis **A+,** CSE151A Machine Learning Algorithms **A,** CSE110 Advanced Data Structures **A-**

PUBLICATIONS

[1] S^3-Sim: Simulating Humans for Personalized Language Modeling

In Submission [Paper] [Project]

Jinzhou Tang, Yufan Zhou*, **Zixuan Wang***, Zhaoxiang Feng, Xinle Yu, Steven Ngo, Zhengding Hu, Luoshang Pan, Lianhui Qin, Yufei Ding, Tianmin Shu, Jingbo Shang, Zhiting Hu, Zhen Wang†* **(*Indicates Equal Contribution)**

- **Proposed a human simulation–based data generation framework** for personalized language modeling, which synthesizes long-horizon conversational trajectories by explicitly maintaining and evolving latent user states (e.g., beliefs, psychological dynamics, and social relationships) as structural constraints during generation.
- **Demonstrated controllable scaling of high-fidelity personalized data**. Models fine-tuned on the generated data improved intent inference and theory-of-mind reasoning compared to baselines.

[2] Lightweight Neural Refinement for Drift Calibration in Eye Tracking Systems

In Submission [Paper] [Project]

Liu Jiaqi, **Zixuan Wang***, Yuhong Zhang, Dingkang Liang, Jane Hanqi Li, Tzyy-Ping Jung, Gert Cauwenberghsi†* **(*Indicates Equal Contribution)**

- **Proposed a lightweight neural refinement method for eye-tracking drift calibration**, modeling spatially heterogeneous errors via affine-conditioned residual prediction beyond single global mappings.
- **Achieved state-of-the-art monocular gaze accuracy** on both a newly collected dataset and JuDo1000, with validated improvements in a real-world eye-tracking application.

[3] RC3: Rollout Chunking with Context Compression for Accelerating Long-Horizon Reinforcement Learning

Work in Progress [Project]

- **Proposed Rollout Chunking with Context Compression**, a training acceleration method for long-horizon on-policy RL in deep research agents, which reduces rollout decoding cost by exploiting redundancy and local dependency in long contexts via periodic context compression and trajectory chunking.
- **Achieved ~2× RL training speedup with no performance degradation**; a 30B-parameter agent attains 75% success on GAIA-Text-103 and ~45% on BrowseComp, establishing state-of-the-art results.

[4] MiroThinker: Pushing the Performance Boundaries of Open-Source Research Agents via Model, Context, and Interactive Scaling

Preprint [Paper] [Project]

MiroMind Team, Song Bai, Lidong Bing, Carson Chen, Guanzheng Chen, Yuntao Chen, Zhe Chen, Ziyi Chen, Jifeng Dai, Xuan Dong, Wenhan Dou, Yue Deng, Yunjie Fu, Junqi Ge, Chenxia Han, Tammy Huang, Zhenhang Huang, Jerry Jiao, Shilei Jiang, Tianyu Jiao, Xiaoqi Jian, Lei Lei, Ruilin Li, Ryan Luo, Tiantong Li, Xiang Lin, Ziyuan Liu, Zhiqi Li, Jie Ni, Qiang Ren, Pax Sun, Shiqian Su, Chenxin Tao, Bin Wang, Hellen Wang, Haonan Wang, James Wang, Jin Wang, Jojo Wang, Letian Wang, Shizun Wang, Weizhi Wang, ***Zixuan Wang****, Jinfan Xu, Sen Xing, Chenyu Yang, Hai Ye, Jiaheng Yu, Yue Yu, Muyan Zhong, Tianchen Zhao, Xizhou Zhu, Yanpeng Zhou, Yifan Zhang, Zhi Zhu* **(The listing of authors is in alphabetical order based on their last names)**

- **Contributed to MiroThinker v1.0**, focusing on the reinforcement learning stage for interaction scaling, where the model is trained to sustain deep, multi-turn agent–environment interactions and long-horizon tool-augmented reasoning.
- **Led the design of the agent search tool set**, enabling efficient external information retrieval during multi-step reasoning which supports up to 600 tool calls per task under a 256K context window and plays a key role in the agent's strong performance across GAIA, HLE, and BrowseComp benchmarks.

EXPERIENCE

MiroMind AI, Shanda Group, Beijing Office, China **2025.6 – 2025.11**

Intern (Mentor: Dr. Yuntao Chen), focused on agentic LLM reinforcement learning. Designed agent context compression frameworks to improve inference efficiency in long-horizon rollouts.

Swartz Center for Computational Neuroscience, UCSD, CA, USA **2025.11 – 2025.12**

Part-time Undergraduate Researcher, focused on eye-tracking–based human-computer interaction research; data-driven machine learning methods for improving gaze estimation.

MixLab, Halıcıoğlu Data Science Institute, UCSD, CA, USA **2025.11 – Present**

Undergraduate Researcher (Advisor by Dr. Zhen Wang and Prof. Zhiting Hu), focusing on personalized LLM development.

RESEARCH INTERESTS

My research interests lie in advancing **LLM algorithms** that enable more reliable, adaptive, and autonomous interaction with humans. I am particularly interested in **LLM alignment**, with a focus on leveraging **synthetic data**, structured supervision, and preference-based training signals to improve models' understanding of human intent while ensuring user privacy and safety. I am also interested in agentic LLM systems with long-horizon capabilities, including **long-context reasoning**, planning, **tool use**, and self-refinement in **real-world grounded** environments. More broadly, my goal is to develop learning and inference frameworks that allow LLMs to **collaborate effectively with humans** and perform **sustained autonomous work**.

AWARDS

- **Comprehensive Merit Scholarship, Tsinghua University, Awarded:** ***RMB 15,000*** **(2025)**
- **Comprehensive Merit Scholarship, Tsinghua University, Awarded:** ***RMB 10,000*** **(2024)**
- **Second Prize, Interdisciplinary Challenge of Basic Science. Zhili College, Tsinghua University (2024)**
- **Third Prize, China College Physics Competition (2024)**

ACTIVITIES

Tsinghua University Social Practice Team, Team Member – Conducted research visits to leading Singaporean institutions: A*STAR, Nanyang Technological University, and National University of Singapore.

Tinker Robot Team, Future Robotics Club, Team Member / National Representative - Achieved 8th place globally in RoboCup 2025 @Home competition in Salvador, Brazil.

Team Member, Tsinghua University Symphonic Band – Participated as an active member, contributing to performances and campus events.

SKILLS

Programming & Framework: Python, C/C++, LaTeX, VeRL, PyTorch

Languages: English (TOEFL iBT 96, IELTS 7.0), Chinese (Native Speaker)