

Zixuan Wang

[Email](#) | [Website](#) | [GitHub](#) | [LinkedIn](#)

EDUCATION

Tsinghua University, Beijing, China

2023/09 – Present

B.S. (in progress) in Physics Department [\[Transcript\]](#)

- **GPA:** 3.9/4.0 (**Top 15%**) **Major:** Mathematics and Physics, **Minor:** Artificial Intelligence
- **Skills:** Claude Code, Codex, Python, C/C++, SGLang, VeRL, Unsloth, PyTorch, TRL, Slime
- **Languages:** IELTS 7.0 / TOEFL 5.0
- **Selected Course:** Deep Learning, Reinforcement Learning, Machine Learning, Computer Vision, Natural Language Processing, Data Structure, Probability, Statistic Inference, Numerical Analysis

University of California San Diego, La Jolla, CA, USA

2025/09 – 2026/01

Visiting Student at Marshall College, program funded by [UCEAP](#)

- **GPA:** 3.9/4.0 **Major:** Computer Science [\[Transcript\]](#)

RESEARCH INTERESTS

Agentic RL · Agent Infra and Evaluation · LLM Personalization · Machine Learning

RESEARCH EXPERIENCE & MANUSCRIPTS

[1] Mind2Dialogue: Training Human-Aware Language Models by Simulating User Mental States |

UCSD | Project Lead in Prof. Zhiting Hu's MixLab (Advisor: [Prof. Zhen Wang](#))

[In Submission to ICLR2027](#) [\[ArXiv\]](#) [\[Project\]](#)

2025/11 – 2026/04

[Zixuan Wang](#)^{}, Yufan Zhou^{*}, Jinzhou Tang^{*}, Zhaoxiang Feng, Xinle Yu, Steven Ngo, Zhengding Hu, Luoshang Pan, Lianhui Qin, Yufei Ding, Tianmin Shu, Jingbo Shang, Zhiting Hu, Zhen Wang⁺ (*Equal Contribution)*

- **Proposed a scalable data generation framework** for personalized language post-training, 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. Models fine-tuned on the generated data improved intent inference and theory-of-mind reasoning compared to baselines.

[2] Harness Learning Enables Generalizable Test-Time Adaptation | CMU | Undergraduate Researcher

(Advisor: [Prof. Andrea Zanette](#))

[In Submission to ICLR2027](#) [\[Arxiv\]](#) [\[Project\]](#)

[CoLM2026 LLA Workshop Spotlight Oral](#)

2026/04 – 2026/09

Alvin Zhang^{}, Xuecheng Liu^{*}, [Zixuan Wang](#)^{*}, Fahim Tajwar, Ruslan Salakhutdinov, Daniel Khashabi, Yuda Song⁺, Andrea Zanette⁺ (*Equal Contribution)*

- **We introduce harness learning:** a proposer model is trained to revise the harness from observed solver model rollouts.
- **We evaluate the method on reasoning and coding tasks and observe both generalization and test-time scaling.** What harness RL improves is thus not performance on a fixed task but the model's ability to improve itself, a capability we call meta-self-improvement.

[3] MEMPROBE: Probing Long-Term Agent Memory via Hidden User-State Recovery | UCSD | MixLab

(Advisor: [Prof. Zhen Wang](#)) [\[ArXiv\]](#) [\[Code\]](#)

2025/11 – 2026/04

Enze Ma, Yufan Zhou, Wei-Chieh Huang, Jie Yang, Huanhuan Ma, [Zixuan Wang](#), Chengze Li, Chunyu Miao, Philip S. Yu, Zhen Wang⁺

- **A benchmark for auditing long-term agent memory via hidden user-state recovery.** We design the benchmark to ask: after an assistant interacts with a simulated user across ordinary assistance tasks, what hidden user state can be reconstructed from the memory artifact the assistant leaves behind?

[4] RC3: Rollout Chunking with Context Compression for Accelerating Long-Horizon

Reinforcement Learning | [MiroMind AI](#) | Research Intern (Mentor: [Dr. Yuntao Chen](#))

[Technical Blog](#) [\[Project\]](#)

2025/06 – 2025/11

Chenyu Yang*, **Zixuan Wang***, Tianchen Zhao*, Jin Wang*, Yuntao Chen* (*Equal Contribution)

- **We 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.

[5] MiroThinker: Pushing the Performance Boundaries of Open-Source Research Agents via Model, Context, and Interactive Scaling | [MiroMind AI](#)

[Technical Report](#) [\[Arxiv\]](#) [\[Project\]](#)

2025/06 – 2025/11

MiroMind Team: Song Bai, Lidong Bing, Yuntao Chen, ..., **Zixuan Wang**, Chenyu Yang, Tianchen Zhao, Jifeng Dai et al.

(authors listed alphabetically by last name)

- **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.

[6] Drift Calibration in Geometric Eye Tracking Systems | UCSD | Swartz Center for Computational Neuroscience | Undergraduate Researcher (Advisor: [Prof. Tzyy-Ping Jung](#))

[In Submission to Workshop NeurIPS2026](#) [\[ArXiv\]](#) [\[Code\]](#)

Liu Jiaqi*, **Zixuan Wang***, Yuhong Zhang, Dingkan Liang, Jane Hanqi Li, Tzyy-Ping Jung, Gert Cauwenberghs* (*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 a collected dataset, with validated improvements in a real-world eye-tracking application.

SCHOLARSHIP

- Technological Innovation Scholarship, Tsinghua University, Awarded: RMB 3,000 (2026)
- Comprehensive Merit Scholarship, Tsinghua University, Awarded: RMB 15,000 (2025)
- Comprehensive Merit Scholarship, Tsinghua University, Awarded: RMB 10,000 (2024)
- Tsinghua Xuetang List, Awarded: RMB 5,000 (2026)
- Tsinghua Xuetang List, Awarded: RMB 5,000 (2025)
- Tsinghua Xuetang List, Awarded: RMB 3,000 (2024)
- Tsinghua Xuetang Pandeng Talent List, Awarded: RMB 3,000 (2023)

ACTIVITIES

Tsinghua University Social Practice Team, Team Member, Visit to Singaporean Research Institutions: A*STAR, Nanyang Technological University, and National University of Singapore.

[Tinker Robot Team](#), Future Robotics Club, Team Member, National Representative RoboCup 2025 @Home competition in Salvador, Brazil.

Tsinghua University [Student China-Europe Exchange Association](#), Team Member

Tsinghua University [Symphonic Band](#), Team Member