


YIPING WANG

✉ ypwang61@cs.washington.edu  <https://ypwang61.github.io/>

(Latest update: June 16, 2026)

📖 RESEARCH INTEREST

My long-term research goal is to develop safe and scalable AI systems with super-human capabilities that can drive significant scientific progress. My recent research interests focus on reinforcement learning, long-horizon coding agents, self-evolving AI, and AI for science. Before that, I have also explored data selection, multimodal learning, video generation, and machine learning theory.

👤 PROFESSIONAL EXPERIENCES

- Member of Technical Staff @ xAI, Reasoning Org** Dec 2025 - Mar 2026
Project: Coding agent of Grok 4.2. Terminal Use. Synthetic Data.
- Research Intern @ Microsoft, Weizhu Chen's Group** June 2024 - Nov 2025
Mentor: [Yelong Shen](#) and [Shuohang Wang](#)
Project: AlphaEvolve-style Evolution; RL for LLM; Video Generation

🎓 EDUCATION

University of Washington, Seattle Sept. 2023 - Present

Ph.D. student in Paul G. Allen School of Computer Science & Engineering

Adviser: [Prof. Simon Shaolei Du](#)

Zhejiang University Sept. 2019 - June 2023

B.Eng. in Computer Science & Technology at College of Computer Science and Technology

- Minor in Mathematics and Applied Mathematics, earned credits: 45.5.
- Rank top 1% in Chu Kochen Honors College.

📦 RECENT PROJECTS

ScaleAutoResearch is a self-built system that runs many autonomous frontier-LLM coding agents in parallel, scaling both width and depth. Each agent iterates indefinitely against an immutable verifier, and stronger agents (ranked by a measurable frontier quality such as conflict count) share and seed their experimental state to the rest through GitHub branches. The same pipeline has driven results in both mathematics and LLM training:

- ★ *Open Math Problems: New Ramsey-Number Lower Bounds, Beating a 32-Year Record & AlphaEvolve* [[Github](#)]
The agents discovered and open-sourced machine-verified graph witnesses for two new Ramsey-number lower bounds. $R(3, 17) \geq 93$ is the first improvement to the 1994 Wang-Wang-Yan bound in 32 years, a result that Google DeepMind's AlphaEvolve matched but did not beat; $R(4, 15) \geq 160$ surpasses AlphaEvolve's 159. All of this ran on nothing more than Claude Code and Codex agents and a single CPU server.
- ★ *LLM Optimizer Discovery: nanoGPT Speedrun (Track 3)* [[X](#)]
Transferred the same pipeline (Claude Code on Opus 4.8 and Codex on GPT-5.5) to Track 3, optimizer design, of Keller Jordan's modded-nanoGPT Speedrun. On one or two A40 nodes (~300 experiments, ~5k A40-GPU-hours), the agents cut the non-interpolation step-count SOTA from 2875 to 2755 for the 3.28 FineWeb target ($n = 8$), below the published Track-3 record of 2925. The recipe stacks a few simple, well-reasoned tweaks: a longer β_2 memory for the 1-D "aux" parameters, SOAP on every hidden matrix refreshed each step, a shorter LR cooldown, and pruning many now-redundant Muon and SOAP variants.

📖 KEY HONORS AND AWARDS

<i>Amazon AI PhD Fellowship</i> [link]	2025
<i>Chu Kochen Scholarship</i> (Highest honor scholarship in Zhejiang University)	2022
<i>National Scholarship</i> in Chu Kochen Honor College	2020
<i>1st Prize</i> for Academic Excellence in Chu Kochen Honor College	2020&2021&2022
<i>1st Prize</i> in Zhejiang Division of National Mathematics Competition for College Students	2020

📄 PAPERS

(* denotes equal contribution or alphabetical ordering, ★ denotes recent featured works.)

[[Google Scholar](#)]

1. ★ *ThetaEvolve: Test-time Learning on Open Problems* [[Arxiv](#)]
Yiping Wang, Shao-Rong Su, Zhiyuan Zeng, Eva Xu, Liliang Ren, Xinyu Yang, Zeyi Huang, Xuehai He, Luyao Ma, Baolin Peng, Hao Cheng, Pengcheng He, Weizhu Chen, Shuohang Wang, Simon Shaolei Du[†], Yelong Shen[†]
ICML 2026.
2. ★ *RLVE: Scaling Up Reinforcement Learning for Language Models with Adaptive Verifiable Environments* [[Arxiv](#)]
Zhiyuan Zeng*, Hamish Ivison*, **Yiping Wang***, Lifan Yuan*, Shuyue Stella Li, Zhuorui Ye, Siting Li, Jacqueline He, Runlong Zhou, Tong Chen, Chenyang Zhao, Yulia Tsvetkov, Simon Shaolei Du, Natasha Jaques, Hao Peng, Pang Wei Koh, Hannaneh Hajishirzi
ICML 2026.
3. ★ *Spurious Rewards: Rethinking Training Signals in RLVR* [[Arxiv](#)]
Rulin Shao*, Shuyue Stella Li*, Rui Xin*, Scott Geng*, **Yiping Wang**, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, Yulia Tsvetkov, Hannaneh Hajishirzi, Pang Wei Koh, Luke Zettlemoyer
ICML 2026.
4. ★ *Reinforcement Learning for Reasoning in Large Language Models with One Training Example* [[Arxiv](#)]
Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Liyuan Liu, Baolin Peng, Hao Cheng, Xuehai He, Kuan Wang, Jianfeng Gao, Weizhu Chen, Shuohang Wang, Simon Shaolei Du, Yelong Shen
NeurIPS 2025.
#1 Paper of the day on Huggingface Daily Papers.
5. *Latent Recurrent Transformer: Architecture Exploration, Training Strategies, and Scaling Behavior* [[Arxiv](#)]
Zeyi Huang, Xuehai He, Liliang Ren, **Yiping Wang**, Baolin Peng, Hao Cheng, Shuohang Wang, Pengcheng He, Jianfeng Gao, Yong Jae Lee, Yelong Shen
Preprint.
6. *Is Your World Simulator a Good Story Presenter? A Consecutive Events-Based Benchmark for Future Long Video Generation* [[Arxiv](#)]
Yiping Wang, Xuehai He, Kuan Wang, Luyao Ma, Jianwei Yang, Shuohang Wang, Simon Shaolei Du, Yelong Shen
CVPR 2025.
7. *FloE: On-the-Fly MoE Inference* [[Arxiv](#)]
Yuxin Zhou, Zheng Li, Jun Zhang, Jue Wang, **Yiping Wang**, Zhongle Xie, Ke Chen, Lidan Shou
ICML 2025.
8. *Mojito: Motion Trajectory and Intensity Control for Video Generation* [[Arxiv](#)]
Xuehai He, Shuohang Wang, Jianwei Yang, Xiaoxia Wu, **Yiping Wang**, Kuan Wang, Zheng Zhan, Olatunji Ruwase, Yelong Shen, Xin Eric Wang

Preprint.

9. *SHARP: Accelerating Language Model Inference by SHaring Adjacent layers with Recovery Parameters* [[Arxiv](#)]
Yiping Wang, Hanxian Huang, Yifang Chen, Jishen Zhao, Simon S. Du, Yuandong Tian
Preprint.
10. *Infer Human's Intentions Before Following Natural Language Instructions* [[Arxiv](#)]
Yanming Wan, Yue Wu, **Yiping Wang**, Jiayuan Mao, Natasha Jaque
AAAI 2025.
11. *CLIPLoss and Norm-Based Data Selection Methods for Multimodal Contrastive Learning* [[Arxiv](#)]
Yiping Wang*, Yifang Chen*, Wendan Yan, Alex Fang, Wenjing Zhou, Kevin Jamieson, Simon S. Du
NeurIPS 2024 (Spotlight).
12. *JoMA: Demystifying Multilayer Transformers via JOint Dynamics of MLP and Attention* [[Arxiv](#)]
Yuandong Tian, **Yiping Wang**, Zhenyu Zhang, Beidi Chen, Simon S. Du
ICLR 2024.
13. *Scan and Snap: Understanding Training Dynamics and Token Composition in 1-layer Transformer* [[Arxiv](#)]
Yuandong Tian, **Yiping Wang**, Beidi Chen, Simon S. Du
NeurIPS 2023.
Oral presentation at High-dimensional learning dynamics workshop at ICML 2023
14. *Improved Active Multi-Task Representation Learning via Lasso* [[Arxiv](#)]
Yiping Wang, Yifang Chen, Kevin Jamieson, Simon S. Du
ICML 2023.
15. *C-Mixup: Improving Generalization in Regression* [[Arxiv](#)]
Huaxiu Yao*, **Yiping Wang***, Linjun Zhang, James Zou, Chelsea Finn
NeurIPS 2022.

⚙️ PROFESSIONAL ACTIVITIES

- Paper Reviewer: NeurIPS(23,24,25), ICLR(24,25,26), ICML(23,24,25,26), CVPR(25), (TF2M,DMLR)@ICML24.
- UW CSE Ph.D. Admission Reviewer: 2024, 2025.
- TA: CSE 446 Machine Learning (25sp), CSE 543 Deep Learning (24Au).