

QIAN LIU(刘乾)

✉ liuqian.sea@gmail.com · 💻 [SivilTaram](#) · 🔗 <https://siviltaram.github.io> · 🌐 [SivilTaram](#)

👤 EXPERIENCE

- xAI**, Singapore 2025.12 – Now
Member of Technical Staff
Working on coding models and agents, with a focus on data-centric training and evaluation.
- TikTok**, Singapore 2024.11 – 2025.12
Coding Team Lead
Built a 10+ person coding research team. Focus on code pre-training and agentic reinforcement learning.
- Sea AI Lab**, Singapore 2022.7 – 2024.11
Research Scientist
Led the development of the Sailor LLM family, with the full pre-training pipeline: data crawling and cleaning, deduplication, data mixture optimization (RegMix), and continual pre-training.

🎓 EDUCATION

- Beihang University – Microsoft Research Asia Joint Ph.D. Program** 2017.9 – 2022.6
Ph.D. in Computer Science and Engineering
Beihang Supervisor: Qingping Zhao ♣ MSRA Mentor: Jian-Guang Lou
- Beihang University** 2013.9 – 2017.6
B.S. in Computer Science and Technology, Ranking 7 / 233

💡 RESEARCH INTERESTS

- **Pre-training & Data:** *Data mixture optimization, data filtering and synthetic data for model training.*
- **Code Generation:** *Building reliable code models, from pre-training corpora to agentic evaluation.*

🔗 SELECTED PROJECTS

- SimpleTIR: End-to-End RL for Multi-Turn Tool-Integrated Reasoning** 2025.3 – 2025.9
We introduce SimpleTIR, a plug-and-play reinforcement learning algorithm that stabilizes multi-turn Tool-Integrated Reasoning (TIR) training. We identify that training instability in multi-turn TIR is primarily caused by distributional drift from external tool feedback, leading to catastrophic gradient norm explosions. SimpleTIR addresses this by filtering void turns – turns yielding neither a code block nor a final answer – from trajectories during policy updates. Starting from the Qwen2.5-7B base model without supervised fine-tuning, SimpleTIR elevates the AIME24 score from 22.1 to 50.5, achieving state-of-the-art performance. The model also discovers diverse reasoning patterns such as self-correction and cross-validation. As the **corresponding author**, I guided the research direction, experimental design, and contributed to the key findings on training dynamics.
- SimpleRL-Zoo: Investigating and Taming Zero RL for Open Base Models** 2024.12 – 2025.3
We systematically investigate zero reinforcement learning – applying RL directly on base models without prior supervised fine-tuning – inspired by the success of DeepSeek-R1. By testing across 10 diverse open base models (Llama3, Mistral, Qwen2.5, etc.), we uncover key design strategies including format rewards and query difficulty control. Notably, we observe the first emergence of verification behavior (“aha moment”) in non-Qwen small models. Our open-source framework SimpleRL-Zoo provides a comprehensive toolkit for zero RL research. As a **co-first author**, I co-led the research direction, experimental design, and contributed to the key findings on training dynamics.

RegMix formulates pre-training data mixture selection as a regression problem: we train hundreds of small proxy models on diverse random mixtures, fit a regression model from mixture weights to target metrics, and extrapolate to identify the optimal mixture before launching the large run. With only 10% of the compute of prior methods, RegMix discovers mixtures that outperform human expert selection and DoReMi, and the discovered mixtures transfer reliably from 1M-parameter proxies to billion-scale models. RegMix has been adopted in real production pre-training (e.g., the Sailor model family). As the **co-first author**, I proposed the original idea and led the project.

SELECTED PUBLICATIONS

- Lin Zheng, Xinyu Li, **Qian Liu**, Xiachong Feng, Lingpeng Kong, [Proxy Compression for Language Modeling](#). (ICML-2026)
- Zhenghai Xue*, Longtao Zheng*, **Qian Liu**, Yingru Li, Xiaosen Zheng, Zejun Ma, Bo An, [SimpleTIR: End-to-End Reinforcement Learning for Multi-Turn Tool-Integrated Reasoning](#). (ICLR-2026)
- Qian Liu***, Xiaosen Zheng*, Niklas Muennighoff, Guangtao Zeng, Longxu Dou, Tianyu Pang, Jing Jiang, Min Lin, [RegMix: Data Mixture as Regression for Language Model Pre-training](#). (ICLR-2025 **Spotlight**, * = equal contribution)
- Longxu Dou*, **Qian Liu***, Fan Zhou*, Changyu Chen*, et al., [Sailor2: Sailing in South-East Asia with Inclusive Multilingual LLMs](#). (Preprint-2025, * = equal contribution, **project lead**)
- Tongyao Zhu, **Qian Liu**, Haonan Wang, Shiqi Chen, Xiangming Gu, Tianyu Pang, Min-Yen Kan, [Sky-Ladder: Better and Faster Pretraining via Context Window Scheduling](#). (ICML-2025)
- Jinjie Ni, **Qian Liu**, Longxu Dou, Chao Du, Zili Wang, Hang Yan, Tianyu Pang, Michael Qizhe Shieh, [Diffusion Language Models are Super Data Learners](#). (Preprint-2025)
- Weihao Zeng*, Yuzhen Huang*, **Qian Liu***, Wei Liu, Keqing He, Zejun Ma, Junxian He, [SimpleRL-Zoo: Investigating and Taming Zero Reinforcement Learning for Open Base Models in the Wild](#). (COLM-2025, * = equal contribution)
- Xinyi He, **Qian Liu**, Mingzhe Du, Lin Yan, Zhijie Fan, Yiming Huang, Zejian Yuan, Zejun Ma, [SWE-Perf: Can Language Models Optimize Code Performance on Real-World Repositories?](#) (Preprint-2025)
- Fangyu Lei*, Jixuan Chen*, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin Su, Zhaoqing Suo, Hongcheng Gao, Wenjing Hu, Pengcheng Yin, Victor Zhong, Caiming Xiong, Ruoxi Sun, **Qian Liu**, Sida Wang, Tao Yu, [Spider 2.0: Evaluating Language Models on Real-World Enterprise Text-to-SQL Workflows](#). (ICLR-2025 **Oral**)
- Siming Huang*, Tianhao Cheng*, J.K. Liu, Jiaran Hao, Liuyihan Song, Yang Xu, J. Yang, J.H. Liu, Chenchen Zhang, Linzheng Chai, Ruifeng Yuan, Zhaoxiang Zhang, Jie Fu, **Qian Liu**, Ge Zhang, Zili Wang, Yuan Qi, Yinghui Xu, Wei Chu, [OpenCoder: The Open Cookbook for Top-Tier Code Large Language Models](#). (ACL-2025)
- Chaofan Tao, **Qian Liu**, Longxu Dou, Niklas Muennighoff, Zhongwei Wan, Ping Luo, Min Lin, Ngai Wong, [Scaling Laws with Vocabulary: Larger Models Deserve Larger Vocabularies](#). (NeurIPS-2024)
- Chengsong Huang*, **Qian Liu***, Bill Yuchen Lin*, Tianyu Pang, Chao Du, Min Lin, [LoraHub: Efficient Cross-Task Generalization via Dynamic LoRA Composition](#). (COLM-2024, * = equal contribution)
- Niklas Muennighoff, **Qian Liu**, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, Shayne Longpre, [OctoPack: Instruction Tuning Code Large Language Models](#). (ICLR-2024 **Spotlight**)
- Yiheng Xu*, Hongjin Su*, Chen Xing*, Boyu Mi, **Qian Liu**, Weijia Shi, Binyuan Hui, Fan Zhou, Yitao Liu, Tianbao Xie, Zhoujun Cheng, Siheng Zhao, Lingpeng Kong, Bailin Wang, Caiming Xiong, Tao Yu, [Lemur: Harmonizing Natural Language and Code for Language Agents](#). (ICLR-2024 **Spotlight**, * = equal contribution)
-  BigCode, [StarCoder: May the Source be with You!](#) (TMLR-2023, **2.3K** Like in HuggingFace Hub)
- Xinyu Pi*, **Qian Liu***, Bei Chen, Morteza Ziyadi, Zeqi Lin, Yan Gao, Qiang Fu, Jian-Guang Lou, Weizhu Chen, [Reasoning Like Program Executors](#). (EMNLP-2022, MLADS-2022 **Distinguished Contribution Award**, * = equal contribution)
- Qian Liu**, Bei Chen, Jiaqi Guo, Morteza Ziyadi, Zeqi Lin, Weizhu Chen, Jian-Guang Lou, [TAPEX: Table](#)

Pre-training via Learning a Neural SQL Executor. (ICLR-2022)

♡ SERVICE & AWARDS

Co-Organizer of Deep Learning for The Third Table Representation Learning Workshop	2024
KAUST Rising Stars in AI 2024 (30 Worldwide)	2024
Beijing Outstanding Doctoral Thesis Nomination Awards	2023
Baidu Scholarship Nomination 2020 (20 Worldwide)	2020
Graduate Student National Scholarship	2019, 2021