

Xinyu Tang

txy20010310@gmail.com

[[Google Scholar](#)] | [[Semantic Scholar](#)] | [[DBLP](#)] | [[LinkedIn](#)]



RESEARCH INTEREST

Large language models, particularly LLM post-training, intelligent agents, reinforcement learning, etc.

EDUCATION

Renmin University of China

Ph.D. Student in Computer Science

Beijing

Sept. 2023 – Present

Beijing Normal University

B.Sc. in Computer Science

Beijing

Sept. 2019 – June 2023

INTERNSHIP EXPERIENCE

Ant Group

LLM Post-training Research Intern

Beijing

Jan. 2025 – Present

PUBLICATIONS

Ring-Zero: Scaling Zero RL to a Trillion

Parameters for Emergent Reasoning

Xinyu Tang*, Gangqiang Cao*, Yurou Liu, Yuliang Zhan, Xiaochong Lan,
Yifan Li, Yuchen Yan, Han Peng, Zican Dong, Zhenduo Zhang,
Tianshu Wang, Xinyu Kong, Zujie Wen, Wayne Xin Zhao¹, Zhiqiang Zhang, Jun Zhou

Technical Report

Rethinking Sample Polarity in Reinforcement Learning with Verifiable Rewards

Xinyu Tang*, Yuliang Zhan*, Zhixun Li*, Wayne Xin Zhao,
Zhenduo Zhang, Zujie Wen, Zhiqiang Zhang, Jun Zhou

ACL 2026

Towards High Data Efficiency in Reinforcement Learning with Verifiable Reward

Xinyu Tang*, Zhenduo Zhang*, Yurou Liu*, Wayne Xin Zhao,
Zujie Wen, Zhiqiang Zhang, Jun Zhou

ICLR 2026

**Unlocking General Long Chain-of-Thought Reasoning Capabilities
of Large Language Models via Representation Engineering**

Xinyu Tang*, Xiaolei Wang*, Zhihao Lv, Yingqian Min,
Wayne Xin Zhao, Binbin Hu, Ziqi Liu, Zhiqiang Zhang

ACL 2025

**DAWN-ICL: Strategic Planning of Problem-solving Trajectories
for Zero-Shot In-Context Learning**

Xinyu Tang*, Xiaolei Wang*, Wayne Xin Zhao, Ji-Rong Wen

NAACL 2025

**Unleashing the Potential of Large Language Models as Prompt Optimizers:
An Analogical Analysis with Gradient-based Model Optimizers**

Xinyu Tang*, Xiaolei Wang*, Wayne Xin Zhao, Siyuan Lu, Yaliang Li, Ji-Rong Wen

AAAI 2025

- Agent Systems with Harness Engineering** arXiv 2026
Xinyu Tang*, Han Peng*, Guoxin Chen, Yuze Shi, Zitao Su,
 Peiyu Liu, Wayne Xin Zhao, Yawen Li, Zhe Xue
- Enhancing Cross-task Transfer of Large Language Models
 via Fourier Activation Steering** arXiv 2025
Xinyu Tang*, Zhihao Lv, Yu Liang Zhan, Xiaoxue Cheng, Junyi Li, Xin Zhao, Zhiqiang Zhang, Jun Zhou
- L2V-CoT: Cross-Modal Transfer of Chain-of-Thought
 Reasoning via Latent Intervention** AAAI 2026
 Yuliang Zhan*, **Xinyu Tang***, Han Wan, Jian Li, Ji-Rong Wen, Hao Sun
- Investigating the Pre-Training Dynamics of In-Context Learning:
 Task Recognition vs. Task Learning** ICLR 2025
 Xiaolei Wang*, **Xinyu Tang***, Wayne Xin Zhao, Ji-Rong Wen
- Rethinking the evaluation for conversational recommendation
 in the era of large language models** EMNLP 2023
 Xiaolei Wang*, **Xinyu Tang***, Wayne Xin Zhao, Jingyuan Wang, Ji-Rong Wen
- GraphPO: Graph-based Policy Optimization for Reasoning Models** arXiv 2026
 Yuliang Zhan*, **Xinyu Tang***, Jian Li, Dandan Zheng, Weilong Chai,
 Jingdong Chen, Jun Zhou, Ge Wu, Wenyue Tang, Hao Sun
- SearchSwarm: Towards Delegation Intelligence in Agentic LLMs for
 Long-Horizon Deep Research** arXiv 2026
 Xiaochong Lan, Quan Chen, Kun Tao, **Xinyu Tang**, Tianshu Wang,
 Qianggang Cao, Xinyu Kong, Zujie Wen, Zhiqiang Zhang, Jun Zhou
- DrugTrail: Interpretable Drug Discovery via Structured
 Reasoning and Druggability-Tailored Preference Optimization** ICLR 2026
 Yurou Liu, Mingyang Li, Xinyuan Zhu, Rui Jiao, Yiming Dong,
Xinyu Tang, Yang Liu, Jieping Ye, Bing Su, Zheng Wang
- Incentivizing Dual Process Thinking for Efficient Large Language Model Reasoning** NIPS 2025
 Xiaoxue Cheng, Junyi Li, Zhenduo Zhang, **Xinyu Tang**,
 Wayne Xin Zhao, Xinyu Kong, Zhiqiang Zhang
- Improving conversational recommendation systems via
 counterfactual data simulation** KDD 2023
 Xiaolei Wang, Kun Zhou, **Xinyu Tang**, Wayne Xin Zhao, Fan Pan, Zhao Cao, Ji-Rong Wen
- A survey of large language models** FCS 2026
 Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min,
 Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen,
 Jinhao Jiang, Ruiyang Ren, Yifan Li, **Xinyu Tang**, Zikang Liu, Peiyu Liu, Jian-Yun Nie, Ji-Rong Wen

Every Step Evolves: Scaling Reinforcement Learning for Trillion-Scale Thinking Model

Technical Report

Ling Team, Anqi Shen, Baihui Li, Bin Hu, Bin Jing, Cai Chen, Chao Huang,
Chao Zhang, Chaokun Yang, Cheng Lin, Chengyao Wen, Congqi Li, Deng Zhao,
Dingbo Yuan, Donghai You, Fagui Mao, Fanzhuang Meng, Feng Xu, Guojie Li,
Guowei Wang, Hao Dai, Haonan Zheng, Hong Liu, Jia Guo, Jiaming Liu, Jian Liu,
Jianhao Fu, Jiannan Shi, Jianwen Wang, Jianxin Lai, Jin Yang, Jun Mei,
Jun Zhou, Junbo Zhao, Junping Zhao, Kuan Xu, Le Su, Lei Chen, Li Tang,
Liang Jiang, Liangcheng Fu, Lianhao Xu, Linfeng Shi, Lisha Liao,
Longfei Zheng, Meng Li, Mingchun Chen, Qi Zuo, Qiang Cheng, Qianggang Cao,
Qitao Shi, Quanrui Guo, Senlin Zhu, Shaofei Wang, Shaomian Zheng,
Shuaicheng Li, Shuwei Gu, Siba Chen, Tao Wu, Tao Zhang, Tianyu Zhang,
Tianyu Zhou, Tiwei Bie, Tongkai Yang, Wang Hong, Wang Ren, Weihua Chen,
Wenbo Yu, Wengang Zheng, Xiangchun Wang, Xiaodong Yan, Xiaopei Wan,
Xin Zhao, Xinyu Kong, **Xinyu Tang**, Xudong Han, Xudong Wang,
Xuemin Yang, Xueyu Hu, Yalin Zhang, Yan Sun, Yicheng Shan, Yilong Wang,
Yingying Xu, Yongkang Liu, Yongzhen Guo, Yuanyuan Wang, Yuchen Yan,
Yuefan Wang, Yuhong Guo, Zehuan Li, Zhankai Xu, Zhe Li, Zhenduo Zhang,
Zhengke Gui, Zhenxuan Pan, Zhenyu Huang, Zhenzhong Lan, Zhiqiang Ding,
Zhiqiang Zhang, Zhixun Li, Zhizhen Liu, Zihao Wang, Zujie Wen

Yulan: An open-source large language model

arXiv 2024

Yutao Zhu, Kun Zhou, Kelong Mao, Wentong Chen, Yiding Sun, Zhipeng Chen, Qian Cao,
Yihan Wu, Yushuo Chen, Feng Wang, Lei Zhang, Junyi Li, Xiaolei Wang, Lei Wang,
Beichen Zhang, Zican Dong, Xiaoxue Cheng, Yuhan Chen, **Xinyu Tang**, Yupeng Hou,
Qiangqiang Ren, Xincheng Pang, Shufang Xie, Wayne Xin Zhao, Zhicheng Dou, Jiaxin Mao,
Yankai Lin, Ruihua Song, Jun Xu, Xu Chen, Rui Yan, Zhewei Wei, Di Hu, Wenbing Huang,
Ze-Feng Gao, Yueguo Chen, Weizheng Lu, Ji-Rong Wen