



Zhengyang Tang

zhengyangtang@link.cuhk.edu.cn | <https://tangzhy.github.io>

 Google Scholar |  zhengyang_42





RESEARCH

My research enhances Large Language Model reasoning through **scalable data generation, critique model, and tool integration**. I design novel instruction tuning frameworks (MathScale, ICML'24; GLAN, TMLR'25; ALAN, ACL'25) to build strong foundational models. I further advance their capabilities with self-evolving critic models (SCRIT, COLM'25) and by enabling complex tool-integrated reasoning (Qwen3; ORLM, OR). My work is also grounded in efficient information access, where my contributions to dense retrieval (DPTDR, COLING'22) have achieved top rankings on competitive benchmarks like LIC 2022 (2nd/793) and MS MARCO (3rd/168).

EDUCATION

- **The Chinese University of Hong Kong, Shenzhen** Jan 2023 - Jul 2026 (Expected)
Ph.D. in Computer and Information Engineering Shenzhen, China
- **Stanford University** Jan 2018 - Jul 2019
Graduate Certificate (Non-Degree-Option) in Computer Science, SCPD Stanford, CA, USA
- **Tongji University** Sep 2012 - Jul 2016
Bachelor of Engineering in Urban Planning Shanghai, China

EMPLOYMENT EXPERIENCES

- **Alibaba Qwen Team**  Oct 2024 - Present
Research Intern Shanghai, China
 - Contributed to the development of tool-integrated reasoning for the **Qwen3** large language model.
 - Co-developed the **SCRIT** framework (COLM 2025), enabling a critique model to self-evolve its capabilities for scalable oversight without external supervision.
- **Microsoft Research Asia (MSRA)**  Jun 2023 - Dec 2023
Research Intern Beijing, China
 - Designed and implemented **MathScale** (ICML 2024), a scalable framework to generate over 2 million math instruction pairs, boosting open-source model performance to state-of-the-art on MwpBench.
 - Co-authored the **GLAN** framework (TMLR 2025) for generalized instruction tuning, synthetically generating diverse training data from a taxonomy of human knowledge.
- **Tencent**  Aug 2020 - Jan 2023
Senior Researcher (T10) Shanghai, China
 - Led dense retrieval R&D for QQ Browser's QA system, resulting in the **DPTDR** (COLING'22) method and top-3 rankings on MS MARCO and LIC'22 benchmarks.
- **Alibaba Quark Search**  May 2019 - Aug 2020
Algorithm Engineer II (P6) Hangzhou, China
 - Developed sparse, reranking and MRC models for the intelligent search engine's QA pipeline.
- **CreditX Technology** Aug 2016 - May 2019
Algorithm Engineer Shanghai, China
 - Engineered financial risk models for major banks, including a credit default model (KS 0.5+) and a GNN-based fraud detection system boosting precision by 2000%+.

SELECTED PUBLICATIONS

*EQUAL CONTRIBUTION; C=CONFERENCE, J=JOURNAL, P=PREPRINT

- [P.1] Qwen Team (My Contribution: Tool-integrated Reasoning) (2025). **Qwen3 Technical Report**. *Technical Report*.
- [P.2] Chengpeng Li*, **Zhengyang Tang***, Ziniu Li*, Mingfeng Xue, Keqin Bao, Tian Ding, Ruoyu Sun, Benyou Wang, Xiang Wang, and Junyang Lin (2025). **CoRT: Code-integrated Reasoning within Thinking**. *arXiv preprint arXiv:2506.09820*.
- [P.3] Tongxu Luo, Wenyu Du, Jiayi Bi, Stephen Chung, **Zhengyang Tang**, Hao Yang, Min Zhang, and Benyou Wang (2025). **Learning from Peers in Reasoning Models**. *arXiv preprint arXiv:2505.07787*.
- [J.1] Chenyu Huang*, **Zhengyang Tang***, Shixi Hu, Ruoqing Jiang, Xin Zheng, Dongdong Ge, Benyou Wang, and Zizhuo Wang (2025). **ORLM: A Customizable Framework in Training Large Models for Automated Optimization Modeling**. *Operations Research (OR)*.
 - **First LLM paper** accepted by *Operations Research* in the journal's 70+ year history.
 - Commercially deployed at Cardinal Operations, boosting modeling accuracy by **10-25%** in industrial sectors.

- [J.2] Haoran Li*, Qingxiu Dong*, **Zhengyang Tang***, Chaojun Wang*, Xingxing Zhang, Haoyang Huang, Shaohan Huang, Xiaolong Huang, Zeqiang Huang, Dongdong Zhang, Yuxian Gu, Xin Cheng, Xun Wang, Si-Qing Chen, Li Dong, Wei Lu, Zhifang Sui, Benyou Wang, Wai Lam, and Furu Wei (2025). **Synthetic Data (Almost) from Scratch: Generalized Instruction Tuning for Language Models**. *Transactions on Machine Learning Research (TMLR)*.
- [C.1] **Zhengyang Tang***, Ziniu Li*, Zhenyang Xiao*, Tian Ding, Ruoyu Sun, Benyou Wang, Dayiheng Liu, Fei Huang, Tianyu Liu, Bowen Yu, and Junyang Lin (2025). **Enabling Scalable Oversight via Self-Evolving Critic (SCRIT)**. In *Proceedings of the 2nd Conference on Language Modeling (COLM)*.
- [C.2] Jianqing Zhu*, Huang Huang*, Zhihang Lin*, Juhao Liang*, **Zhengyang Tang***, Khalid Almubarak, Abdulmohsen Alharthik, Bang An, Juncai He, Xiangbo Wu, Fei Yu, Junying Chen, Zhuoheng Ma, Yuhao Du, He Zhang, Emad A. Alghamdi, Lian Zhang, Ruoyu Sun, Haizhou Li, Benyou Wang, and Jinchao Xu (2025). **Second Language (Arabic) Acquisition of LLMs via Progressive Vocabulary Expansion**. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*. (Oral).
 - Selected for a prestigious **Oral & Panel Presentation**, a distinction given to only 25 of 3000+ accepted papers (**top 0.8%**).
- [C.3] **Zhengyang Tang**, Xingxing Zhang, Benyou Wang, and Furu Wei (2024). **MathScale: Scaling Instruction Tuning for Mathematical Reasoning**. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*.
- [C.4] **Zhengyang Tang**, Benyou Wang, and Ting Yao (2022). **DPTDR: Deep Prompt Tuning for Dense Passage Retrieval**. In *Proceedings of the 29th International Conference on Computational Linguistics (COLING)*.
- [C.5] Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Chenghao Ma, Shanghaoran Quan, Liang Chen, Qingxiu Dong, Runxin Xu, Yubo Wang, Zhaoxu Wang, **Zhengyang Tang**, Zikai Liu, Furu Wei, Haoyang Zou, Jiayi Yang, Jian Yang, Jiatong Han, Jinze Bai, Junchi Yan, Leilei Sun, Lei Li, Lidong Zhou, Lingpeng Kong, Ruoyu Sun, Tianyu Liu, Yibing Zhan, Yi Wu, and Zhaowei He (2025). **Omni-MATH: A Universal Olympiad Level Mathematic Benchmark for Large Language Models**. In *The 13th International Conference on Learning Representations (ICLR)*.
- [C.6] Wei Min, **Zhengyang Tang**, Min Zhu, Yuxi Dai, Yan Wei, and Ruinan Zhang (2018). **Behavior language processing with graph based feature generation for fraud detection in online lending**. In *Proceedings of the KDD 2018 Workshop on Misinformation and Misbehavior Mining on the Web (MIS2)*.

HONORS AND AWARDS

- **Runner-up (2nd Place / 793 Teams)** Jul 2022
LIC 2022 Passage-Ranking Competition, organized by Baidu and CCF [🌐]
 - As a solo participant, secured 2nd place, competing directly with and outranking top-tier industry labs and academic groups.
 - Placed ahead of prominent teams from OPPO (3rd place) and Tsinghua University's IR Group (4th place), with iFlytek securing 1st place.
- **Top-3 Ranking on Official Leaderboard** Dec 2022
MS MARCO Passage Ranking Leaderboard, maintained by Microsoft Research [🌐]
 - Achieved 3rd place on the world's most recognized and influential academic benchmark for passage ranking.
 - Validated the state-of-the-art performance of the developed retrieval systems against top research labs and industry teams globally.

PATENTS

- **A kind of Default Probability Forecasting Methodology of the unstructured data based on deep learning** May 2018
China National Intellectual Property Administration, Patent No. CN107992982A [🌐]
 - Inventors: **Zhengyang Tang**, Chunying Zhou, Mingjie Zhu, Min Zhu, and Yan Wei.

INVITED TALKS & MEDIA

- **Media Coverage for ORLM** 2024
Featured by Cardinal-AI and CUHK-Shenzhen
 - Research on ORLM featured in official reports by **Cardinal-AI** and **CUHK-Shenzhen**, highlighting its real-world industrial impact.
- **Invited Talk on MathScale** Jul 2024
41st International Conference on Machine Learning (ICML) [🌐]
 - Presented *MathScale: Scaling Instruction Tuning for Mathematical Reasoning* to a global audience of top AI researchers.
- **Invited Talk on Dense Passage Retrieval** 2022
Baidu Search Department
 - Presented the DPTDR framework, demonstrating a novel deep prompt tuning approach for efficient and effective dense retrieval.

SERVICES

- **Conference Reviewer:**

- The Conference on Language Modeling (COLM), 2025
- ACL Rolling Review (ARR), 2024 - Present

- **Teaching Assistant:**

- DDA4210 - Advanced Machine Learning, The Chinese University of Hong Kong, Shenzhen, Spring 2023

APPENDIX

Note on the Stanford Graduate Certificate: The Stanford Center for Professional Development (SCPD) provides formal graduate-level programs allowing students to earn academic credits and certificates. To ensure academic rigor equivalent to on-campus studies, all coursework and exams are completed under the supervision of a university-authorized proctor. The credential referenced in this CV was earned through the Foundations in Computer Science Graduate Program.