

Tong Zheng

☎: 202-826-2908 | ✉: tzheng24@umd.edu | ☺: kidzheng.github.io

Research Focus

My research focuses on scalable AI systems that **close the discovery loop** through **reasoning, exploration, and recursive self-improvement**, spanning **long-horizon discovery, test-time computation, and parallel reasoning**.

Research Highlights: **10 first/equal-first-author papers** at ICLR, ICML, ACL, and EMNLP; **1.5K+ GitHub stars**. My work has helped establish new research directions in **recursive self-improvement, parallel reasoning, and automated test-time scaling**: Dream-RSI innovated **recursive self-improvement (RSI)** through **dreaming over simulators derived from discovery history**, Parallel-R1 pioneered RL for parallel reasoning, and AutoTTS introduced automatic policy discovery for test-time scaling. Dream-RSI generated **millions of impressions across 3,000+ posts** and received multiple famous media coverage.

Education

University of Maryland, College Park

Ph.D. in Computer Science; Advisor: Prof. Heng Huang

Aug. 2024 – Present

College Park, MD

Northeastern University

B.Eng. in Computer Science and Technology

Aug. 2017 – Jun. 2021

Shenyang, China

Research Experience

Google

Student Researcher

May 2026 – Aug. 2026

Mountain View, CA

- Led Dream-RSI, introducing **dreaming over history-derived replay world models** for RSI; advancing multiple scientific discovery domain including kernel/algorithm/math optimization.

Tencent AI Lab, Seattle

Research Scientist Intern; Advisors: Hongming Zhang, Wenhao Yu, Xiaoyang Wang

May 2025 – Aug. 2025

Bellevue, WA

- Introduced Parallel-R1, the **first RL framework for parallel reasoning** in LLMs.

NLP Laboratory, Northeastern University

Research Assistant; Advisor: Prof. Tong Xiao

Sep. 2021 – Jun. 2024

Shenyang, China

- Developed **efficient Transformer architectures** for sequence generation, leading to ICML and ACL publications.

Selected Publications (* *Equal contribution. Full publication list: Google Scholar.*)

1. **Dream-RSI: Recursive Self-Improvement through Evolving Worlds** arXiv 2026
Tong Zheng, Xidong Wu, Zheng Zhang, Zhankui He, Chaoyi Zhang, et al. Paper · Project · Code
2. **LLMs Improving LLMs: Agentic Discovery for Test-Time Scaling** Neurips 2026
Tong Zheng, Haolin Liu, Chengsong Huang, Huiwen Bao, Sheng Zhang, et al. Paper · Code
3. **Parallel-Probe: Towards Efficient Parallel Thinking via 2D Probing** ICML 2026
Tong Zheng, Chengsong Huang*, Runpeng Dai*, Yun He, Rui Liu, et al. Paper · Code
4. **Parallel-R1: Towards Parallel Thinking via Reinforcement Learning** ICLR 2026
Tong Zheng, Hongming Zhang, Wenhao Yu, Xiaoyang Wang, Xinyu Yang, et al. Paper · Code
5. **Learning to Reason via Mixture-of-Thought for Logical Reasoning** ICLR 2026
Tong Zheng, Lichang Chen, Simeng Han, R. Thomas McCoy, Heng Huang Paper · Code
6. **Asymmetric Conflict and Synergy in Post-training for LLM-based MMT** ACL Findings 2025
Tong Zheng, Yan Wen, Huiwen Bao, Junfeng Guo, Heng Huang Paper
7. **Towards Optimal Multi-draft Speculative Decoding** ICLR 2025
Zhengmian Hu*, Tong Zheng*, Vignesh Viswanathan, Ziyi Chen, Ryan A. Rossi, et al. Paper
8. **A Bayesian Approach to Harnessing the Power of LLMs in Authorship Attribution** EMNLP 2024
Zhengmian Hu*, Tong Zheng*, Heng Huang Paper
9. **EIT: Enhanced Interactive Transformer** ACL 2024
Tong Zheng, Bei Li, Huiwen Bao, Tong Xiao, Jingbo Zhu Paper · Code
10. **PartialFormer: Modeling Part Instead of Whole for Machine Translation** ACL Findings 2024
Tong Zheng, Bei Li, Huiwen Bao, Jiale Wang, Weiqiao Shan, Tong Xiao, Jingbo Zhu Paper
11. **Learning Multiscale Transformer Models for Sequence Generation** ICML 2022
Bei Li*, Tong Zheng*, Yi Jing*, Chengbo Jiao, Tong Xiao, Jingbo Zhu Paper