

Lin Zheng

✉ linzheng@connect.hku.hk • 🌐 lzhengisme.github.io

Education

The University of Hong Kong <i>Ph.D. in Computer Science</i> Advisor: Lingpeng Kong	2021.09 – 2026.08
Sun Yat-sen University <i>B.E. in Software Engineering</i> Advisor: Qinliang Su	2017.09 – 2021.06

Research Experience

Google DeepMind <i>Student Researcher, with Joshua Maynez</i>	2025.09 – 2025.12
ByteDance TikTok AIIC <i>Research Intern, with Qian Liu</i>	2025.03 – 2025.08
SambaNova Systems <i>Research Intern, with Urmish Thakker, Mingran Wang, and Yun Du</i>	2024.02 – 2024.12
ByteDance AML Research <i>Research Intern, with Jianbo Yuan and Zhi Zhang</i>	2023.04 – 2024.01
ByteDance AML Research <i>Research Intern, with Chong Wang and Jianbo Yuan</i>	2021.06 – 2022.08
The University of Hong Kong <i>Graduate Research Assistant, with Lingpeng Kong</i>	2020.11 – present
Sun Yat-sen University <i>Undergraduate Research Assistant, with Qinliang Su</i>	2019.03 – 2021.05

Selected Publications

- [1] Scratchpad Patching: Decoupling Compute from Patch Size in Byte-Level Language Models
Lin Zheng, Vasilisa Bashlovkina, Timothy Dozat, Dan Garrette, Laura Rimell, and Joshua Maynez
arXiv preprint, 2026
- [2] Proxy Compression for Language Modeling
Lin Zheng*, Xinyu Li*, Qian Liu, Xiachong Feng, and Lingpeng Kong
ICML 2026
- [3] DreamOn: Diffusion Language Models For Code Infilling Beyond Fixed-size Canvas
Zirui Wu*, **Lin Zheng***, Zhihui Xie, Jiacheng Ye, Jiahui Gao, Yansong Feng, Zhenguo Li, Victoria W., Guorui Zhou, and Lingpeng Kong
ICLR 2026
- [4] Dream-Coder 7B: An Open Diffusion Language Model for Code
Zhihui Xie*, Jiacheng Ye*, **Lin Zheng***, Jiahui Gao*, Jingwei Dong, Zirui Wu, Xueliang Zhao, Shansan Gong, Xin Jiang, Zhenguo Li, and Lingpeng Kong
arXiv preprint, 2025
- [5] Dream 7B: Diffusion Large Language Models
Jiacheng Ye*, Zhihui Xie*, **Lin Zheng***, Jiahui Gao*, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng Kong
arXiv preprint, 2025
- [6] EvaByte: Efficient Byte-level Language Models at Scale

Lin Zheng, Xueliang Zhao, Guangtao Wang, Chen Wu, David Dong, Angela Wang, Mingran Wang, Yun Du, Haige Bo, Amol Sharma, Bo Li, Kejie Zhang, Changran Hu, Urmish Thakker, and Lingpeng Kong
Blog post, 2025

- [7] Beyond Autoregression: Discrete Diffusion for Complex Reasoning and Planning
Jiacheng Ye, Jiahui Gao, Shansan Gong, **Lin Zheng**, Xin Jiang, Zhenguo Li, and Lingpeng Kong
ICLR 2025
- [8] Scaling Diffusion Language Models via Adaptation from Autoregressive Models
Shansan Gong*, Shivam Agarwal*, Yizhe Zhang, Jiacheng Ye, **Lin Zheng**, Mukai Li, Chenxin An, Peilin Zhao, Wei Bi, Jiawei Han, Hao Peng, and Lingpeng Kong
ICLR 2025
- [9] Self-Infilling Code Generation
Lin Zheng, Jianbo Yuan, Zhi Zhang, Hongxia Yang, and Lingpeng Kong
ICML 2024
- [10] Diffusion of Thoughts: Chain-of-Thought Reasoning in Diffusion Language Models
Jiacheng Ye*, Shansan Gong*, Liheng Chen*, **Lin Zheng**, Jiahui Gao, Han Shi, Chuan Wu, Zhenguo Li, Wei Bi, and Lingpeng Kong
NeurIPS 2024
- [11] A Reparameterized Discrete Diffusion Model for Text Generation
Lin Zheng, Jianbo Yuan, Lei Yu, and Lingpeng Kong
COLM 2024
- [12] Efficient Attention via Control Variates
Lin Zheng, Jianbo Yuan, Chong Wang, and Lingpeng Kong
ICLR 2023 (oral)
- [13] Linear Complexity Randomized Self-attention Mechanism
Lin Zheng, Chong Wang, and Lingpeng Kong
ICML 2022
- [14] Ripple Attention for Visual Perception with Sub-quadratic Complexity
Lin Zheng, Huijie Pan, and Lingpeng Kong
ICML 2022
- [15] Cascaded Head-colliding Attention
Lin Zheng, Zhiyong Wu, and Lingpeng Kong
ACL 2021

Professional Service

Reviewer: ACL 2021, NAACL 2021, ICML 2022-2024, ICLR 2025, and NeurIPS 2022-2023.

Teaching

- o TA for COMP3314B Machine Learning, 2022 Spring and 2023 Spring, HKU

Selected Awards

- 2022:** M. Braun Postgraduate Prize at the University of Hong Kong.
- 2021:** The University Postgraduate Fellowship at the University of Hong Kong.
- 2019:** The first-grade scholarship at Sun Yat-sen University.
- 2018:** The second-grade scholarship at Sun Yat-sen University.

Skills

Programming Language & Tools: Python, PyTorch, Triton, JAX