Junzhe Shen

CONTACT INFORMATION

Phone & WeChat	(+86)188-1131-1728
Email	junzheshen@sjtu.edu.cn
Homepage	https://junzheshen.github.io
OpenReview	https://openreview.net/profile?id= Junzhe_Shen1
Github	https://github.com/JunzheShen

RESEARCH INTERESTS

Research Areas: Diffusion Language Models.

My research interest is to explore diffusion-based language modeling as a scalable and controllable alternative to autoregressive generation. Concretely, I am interested in the following directions:

- **Better and scalable pretraining techniques for diffusion LMs:** Developing training objectives, noise schedules, parameterizations, and architectures that improve sample quality while reducing the training/inference cost, bridging the gap to strong autoregressive baselines and making diffusion LMs practical at scale.

- **Self-correction for improved reasoning and precise generation:** Using the denoising trajectory as a mechanism for revision: enabling models to diagnose mistakes, refine intermediate drafts, and perform targeted edits under constraints, improving faithfulness and controllability in generation.

- **Any-order and parallel generation for downstream applications:** Exploiting the flexibility of diffusion processes: generating and editing tokens/spans in arbitrary orders, supporting constraint satisfaction, planning-like behaviors, and interactive editing loops that are beneficial for agents and reasoning systems.

EDUCATION

Shanghai Jiao Tong University 2025 – 2030 (expected)
School of Computer Science
Ph.D. Candidate, Advisor: Prof. Zhouhan Lin

Shanghai Jiao Tong University 2021 – 2025
School of Computer Science
Information Security (IEEE pilot class)

INTERNSHIP EXPERIENCE

Beijing Institute for General Artificial Intelligence Sep. 2024 – May. 2025
Student Researcher

PUBLICATIONS

- Tong Wu[*], **Junzhe Shen**[*], Zixia Jia, Yuxuan Wang, Zilong Zheng. TokenSwift: Lossless Acceleration of Ultra Long Sequence Generation. ICML, 2025

- Zixia Jia, Jiaqi Li, Yipeng Kang, Yuxuan Wang, Tong Wu, Quansen Wang, Xiaobo Wang, Shuyi Zhang, **Junzhe Shen**, Qing Li, Siyuan Qi, Yitao Liang, Di He, Zilong Zheng, Song-Chun Zhu. The AI Hippocampus: How Far are We From Human Memory? TMLR, 2026

- **Junzhe Shen**, Jieru Zhao, Ziwei He, Zhouhan Lin. CoDAR: Continuous Diffusion Language Models are More Powerful Than You Think. https://junzheshen.github.io/files/CODAR.pdf

HONORS AND AWARDS

- Academic Scholarship C (twice)

- Second Prize, National College Student Information Security Competition (CISCN)