

Shuibai Zhang

shuibai@cs.wisc.edu | [Personal Website](#) | [Blog](#) | [Google Scholar](#)

Madison, Wisconsin, USA

RESEARCH INTERESTS

I am broadly interested in **language modeling**, and I explore directions including **post-training** (e.g., reinforcement learning and on-policy distillation), **diffusion language models (DLMs)**, and **test-time training**.

As an impact-driven researcher, I aim to build scalable methods that are effective in practice and deliver real-world impact.

EDUCATION

- **University of Wisconsin–Madison** 2024 – Present
Ph.D. Student in Computer Science Madison, USA
 - Research focus: RL for LLMs, diffusion language models (DLMs), LLM reasoning, Mixtrure of Experts (MOE).
- **University of Electronic Science and Technology of China (UESTC)** 2019 – 2023
B.Eng. in Electronic Information Engineering Chengdu, China
 - Dual-degree program with University of Glasgow
- **University of Glasgow (UoG)** 2019 – 2023
B.Eng. in Electrical & Electronic Engineering Glasgow, United Kingdom
 - Dual-degree program with UESTC

RESEARCH HIGHLIGHTS

- **12 accepted papers** at top ML conferences (ICML, ICLR, NeurIPS, COLM, CVPR, ACL), including **4 oral presentations**; 4 additional manuscripts under review.
- **First or co-first author on 9 papers**, spanning diffusion LLMs, RL post-training, LLM reasoning, and large-scale evaluation.
- Experience training 0.5B–7B models; strong background in **Reinforcement Learning** for LLMs.
- Hands-on experience with large-scale LLM training using **Megatron**, **VeOmni**, and RL post-training with the **VERL** framework.
- Hands-on experience training and optimizing **Mixture-of-Experts (MoE)** models, including routing, load balancing, and scaling efficiency.
- Extensive experience developing and analyzing **diffusion LLMs**, covering pretraining, post-training, and inference acceleration.
- Reviewer for ICLR, ICML, NeurIPS, COLM, ACL, and CVPR.

OPEN-SOURCE CONTRIBUTIONS

- **Open-dCoder: The First Fully Open Diffusion LLM for Code** Jul 2025 – Sep 2025
Open Source Project [\[GitHub \(★ 649\)\]](#) [\[Project Page\]](#)
 - **Co-Lead**: Co-directed the development and release of the first fully open-source diffusion LLM for coding.
 - Built and released a 0.5B diffusion language model for code with full pipeline and evaluation suite.
 - Project has received strong community adoption with **649 GitHub stars** by Aug 2026.

PUBLICATIONS & MANUSCRIPTS

* = equal contribution

- **Shuibai Zhang***, Xinchu Liu*, Fred Zhangzhi Peng, Zhihan Yang, Shutong Wu, Yingzi Ma, Jiawei Zhang,
“Off-Manifold Collapse in Guided Protein Language Models”
Under Review — ICLR 2027
- Zunhai Su, Bohan Sun, Xialie Zhuang, **Shuibai Zhang**, He Xiao, Jing Xiong, Hengyuan Zhang, Zhongzhu Zhou,
Tiantian Zhang, Ngai Wong, Chuan-Wei Kuo, “Massive Activations in Hybrid Linear Attention Large Language
Models: Pre-Attention Spikes and Inter-Spike Plateaus”
Under Review — ICLR 2027

- Xinyu Jessica Wang, Haoyue Bai, Yiyu Sun, Haorui Wang, **Shuibai Zhang**, Wenjie Hu, Mya Schroder, Bilge Mutlu, Dawn Song, Robert D. Nowak, “The Long-Horizon Task Mirage? Diagnosing Where and Why Agentic Systems Break”
Under Review — EMNLP 2026
- **Shuibai Zhang**, Fred Zhangzhi Peng, Yiheng Zhang, Jin Pan, Grigorios G. Chrysos, “Corrective Diffusion Language Model”
Under Review — NeurIPS 2026
- Zhihan Yang, Wei Guo, **Shuibai Zhang**, Subham Sekhar Sahoo, Yongxin Chen, Arash Vahdat, Morteza Mardani, John Thickstun, “Continuous Diffusion Scales Competitively with Discrete Diffusion for Language”
COLM 2026 Workshop on Non-Autoregressive Language Models — **Oral Presentation**
- **Shuibai Zhang***, Caspian Zhuang*, Chihan Cui*, Zhihan Yang, Fred Zhangzhi Peng, Yanxin Zhang, Haoyue Bai, Zack Jia, Yang Zhou, Guanhua Chen, Ming Liu, “Expert-Choice Routing Enables Adaptive Computation in Diffusion Language Models”
COLM 2026
- Yuchen Zeng*, **Shuibai Zhang***, Wonjun Kang*, et al., Dimitris Papailiopoulos, Kangwook Lee, “ReJump: A Tree-Jump Representation for Analyzing and Improving LLM Reasoning”
ICML 2026
- **Shuibai Zhang***, Junhyuck Kim*, Gyeongman Kim, Jaewoong Cho, “Is the Importance Ratio Necessary for Stable Reinforcement Learning in LLMs?”
ICLR 2026 SPOT Workshop: Scaling Post-training for LLMs
- Fred Zhangzhi Peng, Zachary Bezemek, Jarrod Rector-Brooks, **Shuibai Zhang**, Anru R. Zhang, Michael Bronstein, Avishek Joey Bose, Alexander Tong, “Planner Aware Path Learning in Diffusion Language Models Training”
ICLR 2026 — **Oral Presentation**
- Wonjun Kang, Kevin Galim, Seunghyuk Oh, Minjae Lee, Yuchen Zeng, **Shuibai Zhang**, Coleman Hooper, Yuezhou Hu, Hyung Il Koo, Nam Ik Cho, Kangwook Lee, “ParallelBench: Understanding the Tradeoffs of Parallel Decoding in Diffusion LLMs”
ICLR 2026
- Yingzi Ma, Yulong Cao, Wenhao Ding, **Shuibai Zhang**, Yan Wang, Boris Ivanovic, Ming Jiang, Marco Pavone, Chaowei Xiao, “dVLM-AD: Enhance Diffusion Vision-Language-Model for Driving via Controllable Reasoning”
CVPR 2026 (Findings)
- Thomas Zeng, **Shuibai Zhang**, et al., Kannan Ramchandran, Dimitris Papailiopoulos, Kangwook Lee, “VersaPRM: Multi-Domain Process Reward Model via Synthetic Reasoning Data”
ICML 2025 — **Oral Presentation**
- Yue Yang*, **Shuibai Zhang***, Wenqi Shao*, Kaipeng Zhang, Yi Bin, Yu Wang, Ping Luo, “Dynamic Multimodal Evaluation with Flexible Complexity by Vision-Language Bootstrapping”
ICLR 2025 — **Oral Presentation**
- Weiyun Wang*, **Shuibai Zhang***, Yiming Ren*, Yuchen Duan*, Tiantong Li*, et al., Wenhao Wang, “Needle In A Multimodal Haystack”
NeurIPS 2024
- Linyi Yang*, **Shuibai Zhang***, Zhuohao Yu*, et al., Xing Xie, Weizhu Chen, Yue Zhang, “Supervised Knowledge Makes Large Language Models Better In-Context Learners”
ICLR 2024
- Linyi Yang*, **Shuibai Zhang***, et al., Jindong Wang, Xing Xie, Yue Zhang, “GLUE-X: Evaluating Natural Language Understanding Models from an Out-of-distribution Generalization Perspective”
ACL 2023 (Findings)

RESEARCH EXPERIENCE

- **Krafton AI** Jul 2025 – Sep 2025
Research Intern Madison, USA
- **Shanghai Artificial Intelligence Laboratory** May 2024 – Aug 2024
Research Assistant Shanghai, China
- **WestlakeNLP Lab, Westlake University** Jun 2023 – Apr 2024
Research Assistant Hangzhou, China

TEACHING EXPERIENCE

- **COMP SCI 220: Data Science Programming I** Spring & Fall 2025
Teaching Assistant University of Wisconsin–Madison
- **COMP SCI 300: Programming II** Fall 2024
Teaching Assistant University of Wisconsin–Madison

TECHNICAL SKILLS

- **Language Models:** Diffusion LLMs, Autoregressive LLMs, RLHF/RLVR, GRPO/PPO/SFT, MOE Design.
- **Machine Learning:** Deep Learning, Reinforcement Learning, Optimization, Evaluation Design.
- **Large-Scale Training:** Distributed training (DDP/FSDP); proficient with **Megatron** and the **VeOmni** pretraining framework, and the **VERL** RL post-training framework.
- **Tools & Infrastructure:** Weights & Biases, Docker, Git, Linux.