

JIAXIN GUO

(+86)1821-7418-112 jx-guo21@mails.tsinghua.edu.cn

Education

Tsinghua University (THU)

Sep.2021 – Jun.2026_(expected)

Beijing, China

- Bachelor of Science in Mathematics and Physics + Measurement Control Technology and Instruments
- **GPA:** 3.9/4.0
- **Core Courses:** Advanced Topics in Linear Algebra (A), Convex Optimization (A), Data Structure (A-)

Swiss Federal Institute of Technology in Lausanne (EPFL)

Feb.2024 – Jul.2024

Lausanne, Switzerland

- Exchange Major: Computer Science
- **GPA:** 5.82/6.0
- **Core Courses:** Advanced probability and applications (A), Optimization for machine learning (A), Learning theory (A-)

Research Interest

My research interests focus on reasoning-centric alignment of large language models and exploring the upper limits of model reasoning and generalization capabilities. I aim to develop principled, reasoning-aware learning strategies that enhance the depth, reliability and robustness of model reasoning that are practically deployable in industrial settings.

Publications

- **Jiabin Guo***, Zewen Chi*, Li Dong*, Qingxiu Dong, Xun Wu, Shaohan Huang, Furu Wei, "Reward Reasoning Model", *NeurIPS 2025*, **Spotlight (top 3.5%)**.
- **Jiabin Guo***, Yuezhou Hu*, Xinyu Feng, Tuo Zhao, "AdaSPEC: Selective Knowledge Distillation for Efficient Speculative Decoders", *NeurIPS 2025*, **Poster**.
- Yihao Li*, Hanle Zheng*, **Jiabin Guo**, Lei Deng, "Decoupling the Contributions of Spatio-Temporal Coding: From ANNs to SNNs", *Neurips Workshop 2024*.

Research Experience

Diffusion Language Model

Jun.2025 – Present

Research Intern at MSRA, Beijing

Advisor: Li Dong, Principal Researcher

- Designed Uniform-State Diffusion Language Model for fast text generation with self-correction.
- Pretrained diffusion language model initialized from an autoregressive model, accelerating training and ensuring stable, single-stage optimization without multi-stage pipelines.
- Developed block-based architecture to speed up inference while preserving generation quality.
- Achieved next-token prediction accuracy comparable to autoregressive models, demonstrating competitive performance with efficient sampling.

Reward Reasoning Model

Feb.2025 – May 2025

Research Intern at MSRA, Beijing

Advisor: Li Dong, Principal Researcher

- Built Reward Reasoning Models (RRMs), a novel reward model paradigm that performs deliberate chain-of-thought reasoning before predicting final rewards.
- Implemented reinforcement learning framework to enable RRMs' reasoning to self-evolve without relying on supervised reasoning traces.
- Leveraged adaptive test-time computation to selectively engage reasoning for complex queries, boosting reward accuracy while maintaining efficiency.

- Achieved state-of-the-art performance across diverse reward modeling benchmarks, demonstrating effectiveness of reasoning-enhanced reward inference.
- Applied RRM in both RLHF and test-time RL settings, achieving stable alignment gains across multiple domains while **avoiding reward hacking** common in traditional scalar reward models.
- Published as Poster at NeurIPS 2025 (co-first author).

Energy-Based Diffusion Reasoning Model

Oct.2024 – Feb.2025

Research Intern at MSRA, Beijing

Advisor: Li Dong, Principal Researcher

- Designed energy-based diffusion framework in limited-token setting that used score matching to model energy functions for reasoning tasks.
- Developed Recursive Denoising Langevin Sampling method to scale solution accuracy with number of Langevin steps and recursion depth, enhancing scalability to complex problems.
- Achieved SOTA on Sudoku that reached **95.6% accuracy**, displaying effectiveness of energy-based iterative reasoning in discrete-space tasks.

Selective Knowledge Distillation for Efficient Speculative Decoders

Jul.2024 – Feb.2025

Research Assistant at H. Milton Stewart School of Industrial and Systems Engineering, Gatech

Advisor: Tuo Zhao, Associate Professor

- Examined speculative decoding for LLM inference acceleration to craft **lightweight, plug-and-play** selective knowledge distillation methods that enhance draft-target model alignment.
- Applied **AdaSPEC**, a novel knowledge distillation (KD) method integrating selective token filtering to optimize draft model alignment with target model by focusing on simpler tokens.
- Designed and implemented reference model to filter out tokens difficult to fit, significantly improving KD process and refining draft model's alignment with target model.
- Conducted performance evaluation across diverse tasks and model configurations of up to 32B, achieving state-of-the-art acceptance rates and **up to 3.3x inference speed-up**.
- Attained up to **8.9% faster generation** and **7.45% higher token throughput** when integrated with advanced SD methods like Eagle, ensuring **practicality for real-world industrial deployment**.
- Published as **Spotlight** at NeurIPS 2025 (co-first author).

Brain-inspired Intelligence Model

Jun.2022 – Feb.2024

Research Assistant at Center for Brain Inspired Computing Research, THU

Advisor: Lei Deng, Associate Professor

- Investigated emergence of functional communities in Spiking Neural Networks (SNNs), extending findings from Artificial Neural Networks where neurons form specialized clusters for different tasks.
- Applied extended modularity metric for negative links to detect structural communities in SNNs trained on specifically designed memory tasks.
- Functionally clustered neurons based on firing variance across task stages.
- Identified distinct functional communities in SNNs that align with modularity-based clusters, demonstrating task-specific selectivity and functional specialization in brain-inspired networks.

Awards and Honors

- | | |
|--|-------------|
| • First prize in Chinese High School Mathematics League | 2018 - 2021 |
| • Academic Excellence Scholarship, 2022 Comprehensive Excellence Scholarship | 2022 |

Skills & Others

Languages: Chinese (native); English (fluent-TOEFL: 108; GRE: 328)

Computer: Fluent in C++, Python, MATLAB; basic understanding of Julia.

Math: Proficiency in convex analysis and optimization algorithms