Kehan Guo

kguo2@nd.edu kehanguo2.github.io github.com/KehanGuo2 Google Scholar

RESEARCH SUMMARY

Ph.D. candidate working on RL post-training, reward and verifier design, and evaluation for reliable LLM agents. Research experience at Apple, Amazon AWS AI, and ByteDance, with work spanning post-training, tool-using agents, and scientific reasoning. First or co-first author of two NeurIPS papers, including a Spotlight.

RESEARCH EXPERIENCE

Apple Jun 2026 – present

Research Intern · Cupertino, CA

- Research on post-training for large language models.

ByteDance Apr 2026 – Jun 2026

Research Scientist Intern · San Jose, CA

- Built a tool-using LLM agent for complex information-seeking and risk-investigation tasks, orchestrating domain-specific tools over heterogeneous data sources.
- Developed a feedback-driven skill-evolution pipeline using offline evaluations and real user-query patterns to improve task-completion rates.

Amazon AWS AI May 2025 – Aug 2025

Applied Scientist Intern · New York, NY

- Developed rank-based reward shaping for group-relative policy optimization (GRPO), converting sparse rollout outcomes into denser within-group feedback for long-horizon reasoning.
- Combined rule-based verification with calibrated model critiques to improve training stability on ambiguous, multi-step tasks.
- Evaluated the approach across mathematical, coding, and logical-reasoning benchmarks.

SELECTED RESEARCH

University of Notre Dame (MINE Lab) Aug 2022 – Present

Ph.D. Research · Advisor: Prof. Xiangliang Zhang

- Developed **MolPuzzle**, a multimodal benchmark for molecular structure elucidation; evaluated language and multimodal models against human experts, revealing a large exact-match gap (1.4% best model vs. 66.7% human). NeurIPS 2024 Spotlight.
- Developed **Reward Transport**, using property-aligned noise-data coupling to steer flow-matching models without an inference-time oracle, reward model, gradient, or additional forward pass.
- Contributed to **TRACE**, which compiles natural-language user corrections into runtime checks for coding agents, cutting out-of-distribution preference violations on ClawArena from 100% to 2%.
- Led chemistry question collection and evaluation for SDE, a cross-domain benchmark of scenario-grounded scientific-discovery tasks.
- Co-first-authored a comprehensive benchmark of LLM capabilities across eight chemistry tasks. NeurIPS 2023.

EDUCATION

University of Notre Dame Aug 2022 – May 2027 (expected)

Ph.D. in Computer Science

Boston University 2020 – 2022

M.S. in Computer Science

SELECTED PUBLICATIONS

* equal contribution. Full list at Google Scholar.

[1] Reward Transport: Property Control in Flow Matching via Noise-Space Alignment.
K. Guo, Y. Shen, Y. Zhou, Y. Huang, C. Gao, S. Du, X. Zhang. *Under review 2026* [Code]

[2] Can LLMs Solve Molecule Puzzles? A Multimodal Benchmark for Molecular Structure Elucidation.
K. Guo, B. Nan, Y. Zhou, et al. *NeurIPS 2024* **[Spotlight]**

[3] What Can Large Language Models Do in Chemistry? A Comprehensive Benchmark on Eight Tasks.
T. Guo*, **K. Guo***, B. Nan, et al. *NeurIPS 2023*

[4] Evaluating Large Language Models in Scientific Discovery.
Z. Song*, J. Lu*, Y. Du*, B. Yu*, T. M. Pruyn*, Y. Huang*, **K. Guo***, X. Luo*, Y. Qu*, et al.
arXiv 2025 [Paper]

[5] Getting Better at Working With You: Compiling User Corrections into Runtime Enforcement for Coding Agents (TRACE).
Y. Zhou, **K. Guo**, H. Zhuang, et al. *arXiv 2026* [arXiv]

[6] AdaReasoner: Adaptive Reasoning Enables More Flexible Thinking.
X. Wang, Y. Huang, Y. Wang, X. Luo, **K. Guo**, Y. Zhou, X. Zhang. *NeurIPS 2025* **[Spotlight]**

[7] Artificial Intelligence in Spectroscopy: Advancing Chemistry from Prediction to Generation and Beyond.
K. Guo, Y. Shen, et al. *IJCAI 2025* **[Survey Track]**

[8] ChemOrch: Empowering LLMs with Chemical Intelligence via Synthetic Instructions.
Y. Huang, Z. Jiang, X. Luo, **K. Guo**, et al. *NeurIPS 2025*

TECHNICAL SKILLS

Training & Inference: Python, PyTorch, JAX, Megatron-LM, DeepSpeed ZeRO, VeRL, Ray, vLLM, Transformers

Methods: RL post-training, GRPO, PPO, DPO, reward modeling, verifier design, LLM evaluation

Agent Systems: Tool use, retrieval, runtime verification, evidence-grounded agents

HONORS, TALKS & SERVICE

Honors: NeurIPS 2024 Spotlight (MolPuzzle, first author); NeurIPS 2025 Spotlight (AdaReasoner, co-author); OpenAI Researcher Access Program (2024)

Reviewer: NeurIPS (2024–25), ICLR (2025–26), ICML, AAAI, IJCAI, KDD, WWW, ACL Rolling Review (ACL/EMNLP)

Tutorial: From VAEs to Diffusion and LLMs: Modern Generative Models for Molecular Discovery, KDD 2026

Organizer: Workshop on Reliable Scientific Foundation Models (RelSciFM), KDD 2026