

Qingcheng Zeng

Ph.D. Candidate · LLM Agents, Reinforcement Learning & Information Retrieval

Evanston, IL · qcz@u.northwestern.edu · qcznlp.github.io · [Google Scholar](#)

Profile

Ph.D. candidate in Computational Linguistics with industry experience in agentic reinforcement learning and production retrieval at Alibaba International, Snowflake, and Tencent AI Lab. Research spans instruction-following retrieval, embeddings, trustworthy LLMs, multilingual NLP, and benchmark design, with publications at ACL, EMNLP, NeurIPS, EACL, CIKM, and IJCAI.

Technical Expertise

Applied ML: Python, PyTorch, Hugging Face Transformers; LLM post-training and experimental analysis.

Agentic RL: GRPO, PPO, process reward modeling; [slime](#) and [Miles](#) frameworks.

Search & Evaluation: information retrieval, dense embeddings, instruction following, multilingual retrieval, calibration, robustness, and benchmark design.

Industry Experience

Applied Scientist Intern — Alibaba International, Accio Agent Team

June 2026 – Sep 2026

Agentic Reinforcement Learning

- Trained long-horizon agents with GRPO/PPO and helped deliver [Occamy-1.0](#), raising Claw-Eval Pass³ from 50.0 to 71.3 over its base model.

Research Scientist Intern — Snowflake, Cortex Search AI Team

March 2026 – June 2026

Remote

- Built a five-dataset SQL schema-retrieval benchmark and corpus-adaptive embeddings, raising average Recall@10 from 60.4 to 75.6 ([paper](#)).

Research Scientist Intern — Snowflake, Cortex Search AI Team

Sep 2025 – Dec 2025

Menlo Park, CA

- Shipped instruction-following retrieval for Cortex Search and led [Dual-View Training](#) (*EMNLP 2026 Main*), improving FollowIR p -MRR by 57.5% on both tested backbones.

Research Scientist Intern — Tencent AI Lab America

June 2025 – Aug 2025

Bellevue, WA

- Developed verified Critical Step Optimization for LLM agents, improving GAIA-Text-103 accuracy by 37% over SFT while supervising only 16% of trajectory steps ([Findings of ACL 2026](#)).

Selected Research

- Retrieval alignment & embedding quality:** built instruction-following retrievers with dual-view training and diagnosed spread-out regularization failures in dense embeddings (*EMNLP 2026 Main*).
- Reliable agents & RAG:** analyzed and mitigated miscalibration in tool-use agents (*ACL 2026 Oral*) and studied noise-aware confidence calibration for RAG (*Findings of EMNLP 2026*).
- Multilingual & code-switching search:** created cross-lingual benchmarks to expose failure modes in retrievers and deep-research agents (*Findings of ACL 2026*; *arXiv 2026*).
- Model behavior & risk analysis:** studied linguistic bias across encoder- and decoder-only models (*EMNLP 2026 Main*) and confidence reliability in reasoning models (*EMNLP 2025*).

Education & Recognition

Northwestern University — Ph.D. in Computational Linguistics, expected 2026

Advisors: Rob Voigt, Kaize Ding

University of Manchester — B.A. in Linguistics, 2021 · **Zhejiang University** — B.A. in Linguistics, 2021

Recognition: Gold Reviewer, ICML 2026; Excellent Reviewer, KDD 2025 Applied Data Science Track. **Languages:** Mandarin Chinese (native); English (professional proficiency).