Xiang (Shawn) Zhang

E-mail: xiangzhg@umich.edu
https://shawnzhg.github.io/
Address: Ann Arbor, MI, USA

EDUCATION

University of Michigan, Ann Arbor United States
Major: M. S. in Data Science *08/2025-06/2027*
Average GPA: **3.9**
Core Courses: Machine learning (A), Web System.

Capital University of Economics and Business Beijing, China
Major: B. S. in Data Science *09/2021-07/2025*
Average GPA: **4.0**
Core Courses: Natural language processing, Operational research, Advanced Algebra, Artificial Intelligence, Data Mining.

RESEARCH INTERESTS

Self-Evolving AI · NLP & LLM · Machine Learning · RL &post-training · Reasoning · Agents & Trustworthy AI

RESEARCH EXPERIENCE

LLM Post training and Latent Reasoning

Ann Arbor, USA

Research Assistant *at University of Michigan Liu lab (Jie Liu)* *01/2026-05/2026*

- Literature-writing agent (multi-turn reward design for **RL post training**): The writing task has no binary reward (unlike math/coding) and GRPO's single terminal reward barely propagates across a long generation, designed rubric-grounded, multi-turn, multi-dimensional rewards across the full trajectory.
- Biomedical **deep-research agent**: To curb hallucination and citation fabrication in literature review (where reliability is non-negotiable), co-designed a training-free, **multi-module agent** (compress-before-reason and verbatim-grounded full-text hybrid answers enforced by a mechanical faithfulness gate), raising fact-level completeness 0.442→0.645 and entity coverage 0.676→0.823 at 55% lower cost with zero fabricated citations.
  (https://github.com/shawnzhg/pubmed-deep-research-agent)
- **KV-cache compression and reuse** (latent-space reasoning): Since natural-language reasoning chains scale token/context cost with length, compressed reusable "Skills" into a small set of KV-cache states and trained them as RL-optimizable parameters, which improve the accuracy of SearchQA from 77.7→88.3.
  (These functions and agents will go live at https://glkb.org soon.)
- **Latent-space reasoning**: Designed MCIG, a training-free per-token redundancy signal, and a latent-feedback adapter (curriculum SFT + MCIG-guided RL, Qwen2.5-1.5B) to fold redundant reasoning into latent space; ablations isolated four mechanistic roadblocks (offline-discriminative-but-RL-useless signal, error accumulation, KV-cache poisoning/anchor collapse, RoPE-manifold mismatch) showing why pure latent skipping collapses at sub-2B scale (64.5%→55% at one latent step) and motivating selective latent↔explicit switching.
  (https://github.com/shawnzhg/adaptive-latent-reasoning/blob/main/REPORT.md)

Reasoning Language Models underthinking (Post training of LLM)

Beijing, China / Remote

Research Assistant *at Institute of Automation, Chinese Academy of Sciences(CASIA)* *06/2025- 04/2026*

- Depth-Quality Erosion (**RL training diagnosis**): Suspecting GRPO/RLVR's rising aggregate Pass@1 masked real regressions, analyzed post-training at per-reasoning-depth granularity and derived a closed-form variance-update identity — uncovering a Simpson's-paradox mode-seeking collapse where pass rate drops in every depth bin (up to −26%) even as Pass@1 rises (+2.0 pp).
- **STAGE-GRPO**: To enrich the sparse group signal left by GRPO's terminal-only reward — the root of the erosion above — proposed a stratified-advantage variant with online per-prompt difficulty estimation that provably decouples variance dynamics from geometric alignment, eliminating the per-bin degradation and lifting MATH-500 Pass@1 to 0.868 (+4.2 pp vs. base, +2.2 pp vs. GRPO).
- **Prefix-contrastive credit assignment:** Observed GRPO spreads one terminal advantage across all tokens though only a

few decide correctness (so the policy games length); designed a GRPO-identity-preserving per-token estimator that concentrates gradient on decision-critical tokens by contrasting, within a group, the branch positions where same-prefix correct/incorrect rollouts diverge; end-to-end ablations in progress.
(https://github.com/shawnzhg/grpo-underthinking/blob/main/REPORT.md)

Spatiotemporal Drought Forecasting via Transformer & GraphRAG Beijing, China

***Research Assistant**, Chinese Academy of Agricultural Sciences(CAAS)* *02/2025- 06/2025*

- Proposed **Cross-Spatiotemporal-Former (CSTF):** Designed a novel Transformer-based architecture with **Cross-Attention** by crossing the query of attention to fuse spatial (grid) and temporal (time-series) features, achieving 98.6% accuracy in drought classification.
- Outperformed Baselines: Validated on 12 years of ERA5-Land data (0.1 resolution); the model outperforming LSTM etc. baselines by capturing long-range dependencies in SPI/SPEI indices.
- Developed **GraphRAG System**: Constructed a domain-specific Knowledge Graph (Drought-Water-Wheat) using DistilBERT and FAISS; integrated with Baichuan2-13B to reduce LLM hallucinations and provide evidence-based agricultural solutions.

Information Retrieval (RAG) Based on Large Language Model Beijing, China

Student Researcher *04/2024-011/2024*

- Scraped bidding data using Python and built a project-location knowledge graph in Neo4j
- Fine-tuned Baichuan2 13B and LLAMA3 8B models with QLoRA for improved bidding data understanding
- Created and integrated a project-location knowledge graph using Python and Neo4j
- Optimized GraphRAG model by integrating external knowledge sources to enhance query accuracy
- Implemented natural language retrieval capabilities to process fuzzy location inputs and retrieve relevant bidding information efficiently.

PUBLICATION

- **Zhang, X.**, (co-first author) Zhaowei, H., Jie, L. (2026, July). GLKB literature writing agent. In preparation.
- Liu, E., **Zhang, X.**, Pan, P., Wang, Y., Liu, B., Zhang, L., & Mei, X. (2024, October). Agricultural technology gap affects climate change adaptation and mitigation based on patent data. *NPJ Climate and Atmospheric Science*, under review.
- Zhang, L., **Zhang, X.**, Gao, J., Lei, S., Yan, C., & Liu, E. (2024, September). An insight into the technological landscape and competitive situation of the biodegradable PBAT industry based on multidimensional patent analysis. *IEEE Transactions on Engineering Management*, under review.

SKILLS

- Distributed Training: **Deepspeed**, **DDP**, **FSDP**; RL infra: **TRL/veRL/OpenRLHF**; **vLLM**; **SGLang**
- Proficient in **Claude Code**, **Codex**, Python (PyTorch, Pandas, NumPy etc.), Java, R, SQL
- Skilled in using Linux servers with extensive experience, and proficient in shell scripting,
- Github: https://github.com/shawnzhg

INTERNSHIP

Baidu.com Times Technology (Beijing) Co., Ltd. Beijing, China

Algorithm Intern *07/2024-10/2024*

- Developed and optimized a speech recognition system using Transformer architecture to achieve high-precision audio-to-text transcription, particularly for dialect data
- Built Python-based web scraping tools to collect video data, extracted audio features, and used the features to improve speech-to-text accuracy through a customized encoder structure
- Applied dynamic time warping to align ASR and OCR outputs, improving text alignment accuracy for multi-language and dialect audio

HONORS AND AWARDS

- First-Class Merit-based Scholarship (Top 5%) 2023-2024
- First Prize, National College Student Market Research and Analysis Competition 2024
- Scientific Innovation Scholarship 2024
- Project Leader, Beijing University Student Innovation & Entrepreneurship Development Program 2024