

Chaoyu Wang

Seeking RA position in LLM Alignment

Email : samuelwang997@gmail.com

Mobile : (+1)-858-247-8856

EDUCATION

- **Northwestern University** Evanston, IL
Master of Science: Engineering Science & Applied Math Sep 2024 - June 2025
- **University of California San Diego** San Diego, CA
Bachelor of Science: Applied Math Sep 2022 - June 2024

EXPERIENCE

- **GuruGame HK** Hongkong SAR
Machine Learning Engineer Intern Sep 2025- Mar 2026
 - **Agent Design:** Built Ad campaign Agent covering 4 workflows (creative sourcing, ad upload, performance monitoring, internal knowledge query) with chained tool calls, multi-turn slot-filling, and a Milvus-backed RAG pipeline for internal knowledge retrieval
 - **Data Construction:** Constructed a 3,000+ seed dataset via rule-based templates and GPT synthesis; applied slot perturbation, semantic paraphrasing, and multi-turn augmentation with 80/20 stratified split.
 - **Class Imbalance:** Resolved long-tail imbalance in multi-turn samples by progressively decomposing dialogues (1→N turns) and oversampling final-turn responses 3×, expanding training set from 3,030 to 7,600+ samples.
 - **SFT Fine-tuning:** Fine-tuned Qwen3-0.6B with LoRA SFT, jointly optimizing CE Loss on natural language responses and argument-level function call parameters.
 - **Evaluation:** Deployed after A/B validation across routing accuracy, function-call EM, end-to-end task success rate, and human quality score; established a data flywheel for continuous iteration.
 - **Results:** End-to-end task completion 86%+ (from <40%); routing accuracy +18pp; function-call EM +25pp; inference cost reduced 54% vs. pure API-chaining baseline.
- **Huatai Securities** Nanjing, China
Software Engineer Intern Apr 2025 - Aug 2025
 - **System Overview:** Built a RAG-based Q&A system over 2,000+ multimodal documents (PDF, PPT, scanned images) to deliver real-time client needs matching and advisory recommendations for sales and wealth management teams
 - **RAG Pipeline:** Implemented structure-aware chunking (section headers, table boundaries, clause numbering) with hierarchical tagging; fine-tuned BGE via contrastive learning on internal financial Q&A data to improve domain terminology retrieval
 - **Retrieval & Dialogue:** Deployed hybrid retrieval (BM25 + dense vector) with a fine-tuned Cross-Encoder reranker; applied Context Injection and Query Rewrite to resolve coreference in multi-turn dialogues; enforced source attribution in prompts for compliance traceability
 - **GRPO Alignment:** Curated 3,000+ prompts from real consultation logs and red-team corpora; designed a 3-dimensional Reward Function (grounded faithfulness via LLM-as-Judge, multi-turn context consistency, response naturalness) to drive GRPO policy optimization on DeepSeek-R1-7B with group-normalized reward signals
 - **Results:** Recall@5 52%→88%, MRR 0.38→0.72; hallucination rate 8.0%→1.2%; multi-turn context error rate 3.0%→0.5%; generalization task completion 81% on unseen scenarios

TECHNICAL SKILLS

- **Languages:** Python, JavaScript, SQL, MATLAB
- **ML & Alignment:** Synthetic Data Construction, Data Augmentation, SFT, PEFT / LoRA, RL Alignment (DPO/GRPO), RAG, Contrastive Learning, Dense Retrieval, Reranking
- **Web & Fullstack:** Next.js, React, Flask, REST API, Prisma ORM
- **Databases & Infra:** PostgreSQL, Milvus, Elasticsearch, Docker, Git, RunPod
- **Data & Engineering:** Pandas, NumPy, Third-party API Integration, Data Pipeline Design