RUI SHI

+852 1711-2889 | jerrysiri.shi@mail.utoronto.ca

Academic Homepage | GitHub | LinkedIn Profile

Kowloon, Hong Kong S.A.R.

EDUCATION

- **Nanjing Univeristy (NJU)** *Nanjing, China*

  Bachelor of Engineering in Artificial Intelligence, ***Outstanding Graduate*** Sept, 2021 - June, 2025

  - **Grade:** 89.5/100 , Rank: 9/101
  - **Core Courses:** Data Structures and Algorithms Analysis (94), Operating Systems (100), Distributed and Parallel Processing (96), Advanced Machine Learning (98), Deep Learning Platform (93.4), Knowledge Representation (100)
  - **Thesis**: *Hierarchy RL enhanced Diffusion LLMs for Safe Structured Guidance*

- **University of Hong Kong (HKU)** *Hong Kong, Hong Kong S.A.R.*

  Exchange Student in Computer Engineering, ***GPA equivalent to First Class Honours*** Jan, 2024 - June, 2024

  - **Core Courses:** Pattern Recognition and Machine Intelligence (A+), Natural Language Processing (A-)

WORK EXPERIENCE

- **Hong Kong AI Institute of AI for Science (CityUHK) [CityU] & Li Auto[Li]** *Kowloon, Hong Kong S.A.R.*

  Miao Lab, Full-time Research Assistant | ***Collaborated LLM Algorithm Intern*** Oct, 2025 - Now

  Supervisor: Prof. Ning Miao | *Dr. Hao Xu*

  - Agentic LLM Self-evolving
  - Decoding Behaviour Control (unmask and remask) of Diffusion LLMs

- **XM Capital Management [XM]** *Shenzhen, China*

  Quantitative Researcher Oct, 2025 - Now

  - Established a distributed RL alpha-discovery engine that addresses GP's exploding search space and overfitting by using an LSTM backbone with IQN and Quantiled Conditional Moment Networks to learn unbiased variance and stabilize the MDP, generating diversified formulaic alphas that achieved SOTA IC on HS300/CSI500.
  - Built an LLM-based trading agent that crawls real-time news using LangChain and Pocketflow, performs event/stance analysis, and converts the interpreted market impact into executable formulaic alphs.

- **Kuaishou Technology [Kuaishou]** *Beijing, China*

  Klear Team, LLM Algorithm Research Intern Apr, 2025 - Sep, 2025

  Leanabell-Prover-V2: Verifier-integrated Reasoning for Formal Theorem Proving

  - Established iterative RL training framwork via VeRL, SandFusion and verifier-integrated DAPO & GRPO, designed feedback masking and reward mechanisms to improve verifier usage and zero entropy collapse.
  - Implemented vLLM-based evaluation pipeline with Lean proof assistant feedback, outperformed SOTA 7B Provers on almost all the Formal Maths benchmarks by 2%-5.3%, boosted MiniF2F by 3.2% (Kimina) and 2.0% (DeepSeek).
  - Explored curriculum-based decomposed subgoal and partial rollout strategies to address token inefficiencies.

RESEACH INTERESTS

MORE DETAILS AVAILABLE ON MY ACADEMIC HOMEPAGE

- Open-endness AI, Self-evolving Agents, Reinforcement Learning with Verifies Rewards (RLVR)
- Diffusion LLMs (dLLMs), Neuro-Symbolic Learning and Reasoning

PUBLICATIONS & MANUSCRIPTS

C=CONFERENCE, J=JOURNAL, U=UNDER REVIEW, P=IN PREPARATION

[J.1] Weipeng Zhou, **Rui Shi**, Gui Yang, Anran Li, Hua Xu, Timothy A. Miller. **Impact of Context on Large Language Models for Clinical Named Entity Recognition [Link]**. In *AMIA ANNUAL SYSPOSIUM (2025)*.

[J.2] Zhiyuan Cao, Vipina K. Keloth, Qianqian, Xie, Lingfei Qian, Yuntian Liu, Yan Wang, **Rui Shi**, Weipeng Zhou, Gui Yang, Jeffrey Zhang, Xueqing Peng, Ethan Zhen, Ruey-Ling Weng, Qingyu Chen, Hua Xu. **The Development Landscape of Large Language Models for Biomedical Applications [Link]**. In *ANNUAL REVIEW OF BIOMEDICAL DATA SCIENCE*, Vol. 8 (2025).

[U.J.3] Xueqing Peng, Huan He, **Rui Shi**, Vipina K. Keloth, Lingfei Qian, Yan Hu, Jimin Huang, Qianqian Xie, Fares Alahdab, Erin Leahey, Brian Ondov, Qiaozhu Mei, Na Hong, and Hua Xu. **How Researchers Claim Novelty in Biomedical Science: A Taxonomy for Understanding Innovation [Link]**. Manuscript submitted for publication in *SCIENCE ADVANCES (2025)*. Under Review.

[P.C.1] Xingguang Ji, Yahui Liu, Qi Wang, Jingyuan Zhang, **Rui Shi**, Chenxi Sun, Fuzheng Zhang, Guorui Zhou, Kun Gai. **Leanabell-Prover-V2: Verifier-integrated Reasoning for Formal Theorem Proving via Reinforcement Learning** [ArXiv]. Manuscript Preparing for *Association for Computational Linguistics (ACL), 2026.*

RESEARCH EXPERIENCE

Veri-code Team, Shanghai AI Lab & LAMDA Group, NJU *Shanghai, China*

- **Hierarchy RL enhanced Diffusion LLMs for Safe Structured Guidance** Feb, 2025 - May, 2025

  Research Intern | *Supervisor: Prof. Jie Fu, Prof. Cunjing Ge*

  - Designed hierarchical Masked dLLMs framework by integrating block-level with RLVR; modeled state transitions via timesteps using Dirac-based multi-step MDP formulation, enabling controllable and structured reasoning.

◦ Designed the *GRPO-diffu* algorithm to estimate approximate log-probabilities, followed by multi-round refinement updates, which significantly improved computational efficiency and reduced KL divergence instability.
◦ Performed distributed LoRA *GRPO-diffu* training on 7B models using DeepSpeed and benchmarked diffusion steps, showing MDMs achieve higher efficiency and accuracy than AR LLMs on math tasks.

Clinical NLP Lab, Yale University *New Haven, United States*

- **Lite-Me-LLaMA: The Resource-Efficient Large Language Models** Jul, 2024 - Dec, 2024

Research Intern | Supervisor: Prof. Hua Xu
◦ Led the construction of a public dataset with 602k sample for medical question answering, ensuring well-balanced and impurity-free data across multiple medical categories to optimize the model's performance in diverse scenarios.
◦ Developed a continual pre-training pipeline for the LLaMA3-8B model on a 72.47-billion-token biomedical corpus, leveraging DeepSpeed 5 for efficient training. Implemented fine-tuning scripts using the auto-train framework.
◦ Developed an vLLM-based inference and multi-tasks evaluation pipeline to extract performance metrics and deliver a fine-tuned Lite-Me-LLaMA.

- **Clinical Trail Matching for Patients Recruitment** Jul, 2024 – Nov, 2024

Research Intern | Supervisior: Prof. Hua Xu, Prof. Jiang Bian (Univeristy of Florida)
◦ Designed a Text-to-SQL pipeline: defined Pydantic I/O schemas, used LLaMA-3.1-70B to structure eligibility criteria, and generated PostgreSQL within a LangChain-OMOP stack to automate patient-eligibility extraction.
◦ Deployed a Lucene/Elasticsearch engine over OMOP with Boolean, kNN, and hybrid search; leveraged GPT-4o for synonym expansion and SQL slot-filling; achieved F1 = 0.85 on condition_occurrence.
◦ Built a clinical-trial RAG: fetched trials by NCTID, parsed section-aware chunks, used LLaMA-3.3-70B for splitting and hierarchical summarization, and integrated LangChain-Milvus for embedding retrieval with chunk filtering and re-ranking; produced section-specific summaries and reached AUC 0.82 on annotated trial documents.

Knwoledge Representation and Reasoning Group (KRistal), NJU *Nanjing, China*

- **BondSenti: BERT-Based Bond Default Sentiment Analysis** Aug, 2023 - Jun, 2024

Team Leader | Supervisor: Prof. Yizheng Zhao, Prof. Xuebin Chen
◦ Architected and deployed a real-time decision-support web application using Flask, Vue.js, Redis, and Logstash; integrated frontend and backend pipelines to visualize and stream complex financial data for executive dashboards.
◦ Designed a multi-encoder BERT–BiLSTM–CNN–CRF model in PyTorch for named entity recognition; incorporated semantic matching for entity disambiguation, boosting extraction F1 by 20%.
◦ Extended the BERT base with GPT-4 knowledge distillation and fine-tuned on a proprietary financial corpus to classify bond-default sentiment into pessimistic, neutral, and optimistic categories; coupled outputs with XGBoost to reduce RMSE in maturity and default predictions by 7%.

- **Automatic Construction of OWL Ontologies Based on Open Information Extraction** Aug, 2023 - Jan, 2024

Research Intern | Supervisor: Prof. Yizheng Zhao
◦ Refactored DetIE into a multi-task seq-to-seq model with task-specific heads and priority-weighted loss, boosting etag/rtag subclass extraction accuracy while 20% faster vs. multi-model baselines.
◦ Parsed syntactic features to design an intermediate n-ary-predicate language that automates OWL ontology construction; grouped etags/rtags into 12 categories and defined tag-to-OWL transformation rules.

Honors and Awards

- Outstanding Graduate of Class 2025, Top 5% NJU, 2025
- Li and Fung scholarships, Top 5% HKU, 2024
- Gang Zheng Overseas Study Scholarship, Top 5% NJU, 2024
- Second Prize (6th place nationalwide), 19th "Citi Cup" Financial Innovation Application Competition China, 2024
- Third Prize (National Wide), China Undergraduate Mathematical Contest in Modeling China, 2023
- Excellence Award (University Level), E Fund Asset Management Cup "AI+" Innovation Challenge NJU, 2023
- Second prize (University Level), The people's scholarship in China, Top 10% NJU, 2023

Skills

- **Programming Languages:** C, C++, Python, Java, MySQL, HTML5, CSS, JavaScript
- **Development Tools:** Linux, Unix, Git/Github/GitLab, LaTeX, Docker
- **Data Science & Machine Learning:** Scikit-Learn, Numpy, Pandas, Scipy, Matplotlib
- **Deep Learning & Reinforcement Learning:** PyTorch, LangChain, vLLM, DeepSpeed, VeRL, OpenRLHF
- **LLM-based Agent:** LangGraph, Autogen, CAMEL, PocketFlow