YIJIA XIAO

yijia-xiao.com | linkedin.com/in/yijia-xiao | Google Scholar | yijia.xiao@cs.ucla.edu

OBJECTIVE

Seeking to apply my expertise in **reasoning LLMs** and **verifiable agents** along two complementary directions:

I. Engineering: Agent harnesses that steer long-horizon verifiable tasks: specification-driven implementation, deployment, testing, and experimental verification, from deploying research codebases to backtesting quantitative strategies.

II. Research: Self-evolving agents automating the research loop: deep research, structured knowledge-base construction, ideation, and hypothesis calibration, across financial (fundamental and quantitative) and scientific research.

EDUCATION

- **University of California, Los Angeles** — *Sep. 2022 – Present*
  Ph.D. Candidate in Computer Science — Los Angeles, CA, USA
  ◦ Advisor: Professor Wei Wang

- **Tsinghua University** — *Sep. 2018 – June 2022*
  Bachelor of Computer Science and Technology — Beijing, China
  ◦ Advisor: Professor Jie Tang

EXPERIENCE

- **Jump Trading LLC** — *June 2026 – Sep. 2026*
  Quantitative Research Intern — Chicago, IL, USA
  ◦ Built automated research systems for alpha mining, generating, validating, and ensembling candidate signals.
  ◦ Researched transformer architectures and reinforcement learning methods for equity trading.

- **Google Cloud AI Research** — *Jan. 2026 – June 2026*
  Student Researcher — Sunnyvale, CA, USA
  ◦ Developed FinanceHarness, the layered agent harness automating financial deep research end to end; accepted at EMNLP Findings 2026 (Publication [C.1]).
  ◦ Built FinanceGym, a point-in-time benchmark of 400 expert-validated questions whose rubrics separate pre-cutoff retrieval from post-cutoff reasoning.

- **Point72** — *June 2025 – Dec. 2025*
  Quantitative Research Intern — New York, NY, USA
  ◦ Applied LLMs to ingest and process large-scale internal notes (e.g., company earnings), building analytics and signal-generation tools for analysts and portfolio managers.
  ◦ Developed an earnings forecast signal outperforming the production signal in hit rate, correlation, and PnL.

- **Amazon Web Services** — *June 2024 – Sep. 2024*
  Applied Scientist Intern — Seattle, WA, USA
  ◦ Developed CSR-Agents, a multi-agent framework toward automated research that deploys CS repositories, calibrating scripts via environment feedback, shell drafting, log analysis, tool use, and issue-database RAG.
  ◦ Built CSR-Bench, the companion benchmark scoring deployment accuracy, efficiency, and script quality across NLP, CV, and ML repositories; Oral at NAACL 2025 Main (Publication [C.2]).

- **Sequoia Capital China** — *July 2022 – Sep. 2022*
  Investor Intern — Beijing, China
  ◦ Focused on investments in renewable energy (photovoltaic, hydrogen), AI infrastructure, and AI4Science.
  ◦ Contributed to 20+ deals from Seed to Series B+, spanning market analysis, technical roadmap reviews, due diligence, roadshows, and term sheets.

- **Tsinghua University** — *Dec. 2020 – Aug. 2021*
  Project Leader — Beijing, China
  ◦ Led pre-training of the world's **2nd-largest** (2021–2023) Protein LLM, **Wen Su** (*$1.5M* funding; 504× NVIDIA V100 32G cluster on a Megatron-LM framework), improving contact-prediction accuracy from 36% to 75%.
  ◦ Adopted by researchers at 30+ institutions; featured on the Wu Dao Wikipedia page (Publication [C.8]).

SELECTED PUBLICATIONS

T=TECHNICAL REPORT / PREPRINT, C=CONFERENCE

I. AGENTIC LLM SYSTEMS

[C.1] **Yijia Xiao**, Rujun Han, Yanfei Chen, Zifeng Wang, Ke Jiang, Zhongying CuiZhu, Vishy Tirumalashetty, Wei Wang, Burak Gokturk, Tomas Pfister, Chen-Yu Lee. (2026). FinanceHarness: Autonomous Financial Deep Research Framework. In *EMNLP Findings 2026*.

- FinanceHarness automates financial deep research end to end, paired with FinanceGym, a point-in-time benchmark separating pre- from post-cutoff reasoning; lifts an open-weight backbone from 25.3% to 32.4%.

[T.1] **Yijia Xiao**, Edward Sun, Di Luo, Wei Wang. (2025). TradingAgents: Multi-Agents LLM Financial Trading Framework.

- TradingAgents synthesizes insights across multiple data sources with LLM-based reasoning to generate systematic, interpretable investment decisions aligned with institutional research standards.

- **100K+ Stars** on GitHub; customized and adopted by securities and equity research firms.

[T.2] **Yijia Xiao**, Edward Sun, Tong Chen, Fang Wu, Di Luo, Wei Wang. (2025). Trading-R1: Financial Trading with LLM Reasoning via Reinforcement Learning.

- Trading-R1 combines supervised and reinforcement fine-tuning to make models financially aware, producing structured, evidence-based decisions.

- Generates structured investment theses that support disciplined and interpretable trading decisions.

[C.2] **Yijia Xiao**, Runhui Wang, Luyang Kong, Davor Golac, Wei Wang. (2024). CSR-Bench: Benchmarking LLM Agents in Deployment of Computer Science Research Repositories. In *NAACL 2025* (**Oral**).

- CSR-Agents is a multi-agent LLM framework for automated research codebase deployment.

[C.3] Yiqiao Jin, Qinlin Zhao, Yiyang Wang, Hao Chen, Kaijie Zhu, **Yijia Xiao**, Jindong Wang. (2024). AgentReview: Exploring Peer Review Dynamics with LLM Agents. In *EMNLP 2024* (**Oral**).

II. MULTIMODAL FOUNDATION MODELS FOR SCIENCE

[C.4] **Yijia Xiao**, Edward Sun, Yiqiao Jin, Qifan Wang, Wei Wang. (2024). ProteinGPT: Multimodal LLM for Protein Property Prediction and Structure Understanding. In *Machine Learning for Genomics, ICLR 2025* (**Spotlight**). *In collaboration with Meta AI.*

- ProteinGPT integrates protein sequence and structure encoders with an LLM for precise property analysis and structure-aware question answering.

[C.5] **Yijia Xiao**, Wanjia Zhao, Junkai Zhang, et al., Pan Lu, Xiao Luo, Yu Zhang*, James Zou*, Yizhou Sun*, Wei Wang*. (2025). Protein Large Language Models: A Comprehensive Survey. In *EMNLP 2025*.

[C.6] Tianyu Liu, **Yijia Xiao**, Xiao Luo, Hua Xu, W. Jim Zheng, Hongyu Zhao. (2023). Geneverse: A Collection of Open-source Multimodal Large Language Models for Genomic and Proteomic Research. In *EMNLP 2024*.

[C.7] **Yijia Xiao**, Edward Sun, Yiqiao Jin, Wei Wang. (2024). RNA-GPT: Multimodal Generative System for RNA Sequence Understanding. In *Machine Learning for Structural Biology Workshop, NeurIPS 2024*.

[C.8] **Yijia Xiao**, Jiezhong Qiu, Ziang Li, Chang-Yu Hsieh, Jie Tang. (2021). Modeling Protein Using Large-scale Pretrain Language Model. In *Pretrain Workshop, KDD 2021*. World's **2nd-largest** Protein LLM (2021–2023), *Wen Su*; part of the Wu Dao project, see Wikipedia page.

III. REASONING, POST-TRAINING & EVALUATION

[C.9] **Yijia Xiao**, Yiqiao Jin, Yushi Bai, Yue Wu, Xianjun Yang, Xiao Luo, Wenchao Yu, Xujiang Zhao, Yanchi Liu, Quanquan Gu, Haifeng Chen, Wei Wang, Wei Cheng. (2023). PrivacyMind: Large Language Models Can Be Good Privacy Protection Learners. In *EMNLP 2024*.

[C.10] Fang Wu, Xu Huang, Weihao Xuan, Zhiwei Zhang, **Yijia Xiao**, Guancheng Wan, Xiaomin Li, Bing Hu, Peng Xia, Jure Leskovec, Yejin Choi. (2025). Multiplayer Nash Preference Optimization. In *ICLR 2026* (**Oral**).

[T.3] **Yijia Xiao**, Edward Sun, Wei Wang. (2024). LogicVista: Multimodal LLM Logical Reasoning Benchmark in Visual Contexts.

[C.11] Wanjia Zhao, Qinwei Ma, Jingzhe Shi, Shirley Wu, Jiaqi Han, **Yijia Xiao**, Si-Yuan Chen, Xiao Luo, Ludwig Schmidt, James Zou. (2025). PRISM-Physics: Causal DAG-Based Process Evaluation for Physics Reasoning. In *ICLR 2026*.

[C.12] Yushi Bai, Jiahao Ying, Yixin Cao, et al., **Yijia Xiao**, Haozhe Lyu, Jiayin Zhang, Juanzi Li, Lei Hou. (2023). Benchmarking Foundation Models with Language-Model-as-an-Examiner. In *NeurIPS 2023*.

Honors and Awards

- **Amazon PhD Fellowship**, 2026.
- **Research Excellence Scholarship**, Tsinghua University, 2021.
- **Silver Medal**, International Collegiate Programming Contest (ICPC) Asia East Continent Final, 2020.
- **Gold Medal**, International Collegiate Programming Contest (ICPC) Asia Regional Contest, 2020.
- **First Prize**, Chinese Collegiate Physics Olympiad, 2019.
- **National Bronze Medal**, Chinese Physics Olympiad, 2017.

Selected Invited Talks

- **Verifiable and Agentic AI for Quantitative Research** *Sep. 2025*
  Jump Trading LLC Chicago, IL, USA
  - Presented Trading-R1, combining supervised fine-tuning and reinforcement learning for investment analysis.
  - Introduced agentic workflows that automate quantitative research and development, scaling researchers' coverage.

- **Agentic AI for Fundamental Research** *Oct. 2025*
  Trivariate Research, L.P. New York, NY, USA
  - Presented agentic systems that scale fundamental equity research coverage.
  - Introduced customizable pipelines that adapt to diverse data sources and encode analysts' priors.

- **TradingAgents: Multi-Agents LLM Financial Trading Framework** *Apr. 2025*
  LLMQuant London, UK
  - Introduced TradingAgents framework for automated trading systems.
  - Demonstrated LLM-based approaches to quantitative finance research.

Professional Services

- **Reviewer**, NeurIPS 2023/2024/2025, ICML 2023/2024, ICLR 2023, EMNLP, NAACL, AAAI, AISTATS, ACM Multimedia, PLOS ONE, etc.

Research Initiative

- **Tauric Research** *Mar. 2025 – Present*
  Founder Los Angeles, CA, USA
  - AI-driven trading intelligence research initiative focusing on financial data analytics and alpha discovery. Ongoing research projects include Agents, Post-Training, Deep Research, and RAG.
  - Established research collaborations with Stanford, UW, and MIT for quantitative research systems and data infrastructure. Leading a team across AI, finance, and systematic trading.
  - Supported by Google Cloud Startups Program and Amazon Web Services Startups Program.

Research Support

- My research has been supported by the Thinking Machines Lab Research Grant, Lambda Research Grant, Google Startup Funding, Amazon Startup Program, NVIDIA Inception, and Jina AI.