Tianyang Zhou

(+1)412-906-3071 tzhou3@andrew.cmu.edu

Education

Carnegie Mellon University **Pittsburgh, PA**

Ph.D. Student in Information Systems and Management *Aug. 2025 – June 2030 (Expected)*

GPA: 3.95/4.00; **Coursework:** Generative AI, Optimization for ML, Intermediate Statistics, Econometrics, Microeconomics

Research Interest: Trustworthy AI, LLM interpretability, Prompt Optimization, Agentic LLM, LLM-RL, Bandit Learning

Tsinghua University **Beijing, China**

Bachelor of Economics and Finance (with Distinction) *Aug. 2021 – June 2025*

- **GPA:** 3.93/4.00; **Ranking:** 9/193; **Honors:** Comprehensive Excellence Scholarship (1%), Technological Innovation Scholarship
- **Minor in Statistics:** 4.00/4.00; **CS Coursework:** Deep Learning, Artificial Intelligence, Data Structure, OOP, C++ Programming

University of Pennsylvania **Philadelphia, PA**

International Guest Student *Aug. 2023 – Dec. 2023*

- **GPA:** 4.00/4.00; **Coursework:** Real Analysis (PhD, A+), Machine Learning (Graduate, A), Analysis of Algorithms, Game Theory

Research Experience

Structured Prompt Optimization Meets RL for Global and Local Interpretability over Complex Text **Pittsburgh, PA**

Advisor: Dr. Leman Akoglu, Heinz College, Carnegie Mellon University *Aug. 2025 – Present*

- Designed a 3-stage LLM pipeline (rulebook learning + distillation + RL) producing both global and instance level explanations for complex text classification across legal, academic and clinical domains; manuscript submitted to **EMNLP 2026 (1st Author)**
- Introduced two novel algorithms: **SPO4SOP**, a structured prompt optimization method that learns interpretable and composable rule sets in linear time, and **BD-GRPO**, a classification aware RL variant robust to severe class imbalance
- Outperformed Prompt Optimization baselines at Stage 1 and surpassed end-to-end SFT with full pipeline on both classification performance and reasoning quality across 3 benchmarks; BD-GRPO reaches peak performance **3-5x faster than vanilla GRPO**

Confidence as Control: A Survey of Confidence Utilization in Large Language Models **Pittsburgh, PA**

- Co-authored a comprehensive survey synthesizing **130+ papers**; proposed a novel taxonomy that reframes LLM confidence scores as an actionable decision signal for training, inference, model selection, RAG, and risk control
- Led the writing and synthesis of inference and risk control sections; manuscript under review at **TMLR (2nd Author)**

Learning-aided Non-Stationary Bandit Algorithm **Atlanta, GA**

Advisor: Dr. Yao Xie, School of Industrial and Systems Engineering, Georgia Institute of Technology *May 2024 – Dec. 2024*

- Developed a hybrid approach integrating a time-series foundation model with non-stationary bandit algorithms to exploit reward variation information in a distribution-free setting; applied adaptive weighting to balance the optimal trade-off
- Proved the method's ability to achieve tight regret upper bound in relation to the black-box model's error budget
- Validated **order-wise regret improvements** over SOTA baselines through extensive simulations

Achieving Statistically Reliable LFQA in LLMs through Calibration and Iterative Refinement **Philadelphia, PA**

Advisor: Dr. Weijie Su, Department of Statistics and Data Science, University of Pennsylvania *Oct. 2023 – May. 2024*

- Built an iterative, distribution-free sampling algorithm to generate reliable, high-quality LLM outputs for Long-form QA
- Implemented a multi-role LLM system (generator, splitter, scorer, feedback provider) to iteratively refine outputs
- Proved theoretical guarantees for the correctness probability of selected answer components
- **Improved answer length by 1.5× and achieved ~20% ROUGE-L gains** over single-step baselines on FActScore and LiveQA

Work Experience

Renaissance-Era Investment **Beijing, China**

Quantitative Research Intern *June 2023 – Aug. 2023*

- Developed alpha strategies for China's A-share market with 10-minute level data; contributed 4 factors for the company
- Trained XGB models and achieved **6%** annual excess returns in out-of-sample, long-only backtests under high transaction costs

Leadership & Activities

Kaggle Machine Learning Competition: Multimodal Single-Cell Integration **Beijing, China**

Leader *Aug. 2022 – Nov. 2022*

- Built ML models to predict stem cell differentiation; handled >30GB datasets by sparse matrices and TSVD; leveraged GBDT's feature importance for informed dimensional reduction; ensembled MLP, XGB, and KRR models. **Earned a silver medal**

Skills

Language: Chinese (Native); English (Fluent, TOEFL: 118/120, GRE: V165+Q170+AW4.0)

Technical Skills: Python (PyTorch, Transformers, Scikit-learn, vLLM, VeRL, liteLLM, NumPy, Pandas), C++ (STL), R, LaTeX, Git