Ergan Shang

Email: eshang@andrew.cmu.edu — Phone: +1 (412) 579 3236 — Homepage

Education

Carnegie Mellon University *Pittsburgh, PA, USA*
Ph.D. in Statistics, Advisors and Collaborators: Kathryn Roeder, Yuting Wei, Linjun Zhang *08/2023 - (exp.) 05/2028*
University of Science and Technology of China *Hefei, Anhui, China*
B.S. in Mathematics, GPA: 4.07/4.30 (Top 2%) *08/2019 - 06/2023*

Work Experience

- **Research Scientist Intern, Ads Foundation Model and Ranking Engineering, Meta Platforms, Inc.** *05/2026 - 08/2026*
  Under supervision of Flavio Sales Truzzi in **Training Efficiency Team**. **Submitted to NeurIPS 2026** [Paper]
  - Proposed a **Forward-Forward Early Backward** training algorithm to decouple forward/backward architectures for parameter updates, validating correctness on **NanoChat transformer** models via GPU kernel/operator trace analysis.
  - Conducted **low-fidelity proxy evaluations** to screen training-efficiency proposals, validating that the proxy introduced only **0.0029% model-performance bias**.
  - Applied on Meta's **production-scale Recommendation System Dynamic Product Ads (DPA) CTR model** for high-traffic ads ranking, yielding a **7.38% QPS(query per second)** gain while keeping **model performance** neutral within **1.38%** measured by normalized entropy (NE).

Research Experience

- **Comparative-signal-based GRPO fine-tuning and post training for mathematical reasoning** *12/2025 - present*
  - Proposed an **EIF (efficient influence function)**-corrected **GRPO** objective for **LLM post-training**, incorporating pairwise **CoT** comparison signals to construct preference-based influence scores for **policy optimization**.
  - Developed a **multi-LLM evaluation framework** leveraging **GPT-4o-mini** as an auxiliary LLM judge to compare candidate mathematical solutions and provide preference signals for **reinforcement learning**.
  - Implemented an end-to-end **RL fine-tuning pipeline** for **Qwen-0.5B** on **GPQA and GSM8K**, achieving a **10% improvement in pass@k rate** over the naive GRPO baseline.
- **Evaluating LLM Agents in Autonomous Scientific Reasoning by ASI-bench** *Apex Intelligence | 07/2026 - 09/2026*
  [Hugging Face #4 Paper of the day, Aim for Nature] [Web]
  - Designed a scientific **reasoning benchmark** task for **LLM agents**, featuring high-dimensional feature selection, **automated reward-based grading**, and verifiable research outcomes to assess autonomous scientific problem-solving.
  - Developed **progressive difficulty settings** with decreasing **human methodological guidance** to evaluate agents' independent reasoning and method selection, revealing a drop in average performance from **50.91 to 26.62** as guidance was removed.
- **AI Agent Benchmarking for Multi-Step Scientific Reasoning** **Submitted to ICLR 2027[Paper]** *07/2026 - 09/2026*
  - Developed a **long-horizon** scientific reasoning task for **Long-Horizon Terminal-Bench (LHTB)**, leveraging parameterized synthetic datasets and **multi-step workflows** to evaluate **LLM agents'** planning and execution capabilities.
  - Implemented hidden verifiers and **dense reward-based grading** to evaluate intermediate progress and task completion, revealing significant limitations in frontier **LLM agents**, with the best model achieving only **15.2% pass@1** across 46 long-horizon tasks.
- **A Diffusion-Based Debiasing Framework for High-Dimensional Response Prediction** *11/2024 - 09/2025*
  Accepted by Proceedings of the National Academy of Sciences of the US(PNAS): [Paper]
  - Developed a latent **diffusion-based generative** framework integrating **VAE** and **representation learning** to predict high-dimensional response distributions under unseen perturbations, generating ∼**30,000 samples in minutes** through an efficient end-to-end sampling pipeline versus ∼**100 in conventional** biological experiments,.
  - Incorporated **statistical debiasing** and uncertainty quantification to improve generative reliability under distribution shifts, achieving **85% mean estimation coverage** versus ∼80% or lower for competing methods.
- **Transformer-Based Transfer Learning for High-Dimensional Response Prediction** *12/2025 - 05/2026*
  - Developed **SCARF**, a deep **generative transfer-learning** framework integrating **Transformer-based representation learning**, optimal transport, and **VAE** for high-dimensional response prediction, achieving **0.95 average PCC** across real-world benchmarks while reducing GPU memory usage by ∼50%.
- **Inference for Balance Theory in Time-Varying Signed Network** *11/2023 - 12/2025*
  Submitted to Journal of Machine Learning Research (JMLR): [Paper]
  - Implemented a **nonparametric statistical model** for dynamic signed network using **PyTorch-based kernel smooting** method, achieving **7x precision** improvement in simulations and identifying anomalous relationships in real-world international networks.
- **XConTest: Convergence Analysis of Perturbation Effects for Gene Function Similarity** *11/2023 - 12/2025*
  Submitted to Science Advances: [Paper], [Software]
  - Designed a **high-dimensional statistical inference framework** combining **two-stage LASSO, projection and cross-fitting** to identify gene functional similarities, achieving **over 90% accuracy** in simulations and validating results on real biological datasets.

Awards

- Mihaela Serban Award for the **Best Poster Competition** in ASA Pittsburgh Chapter; *2025*
- **Candidate of Guomoruo Scholarship (Highest Award in Undergraduates) (Final Round, Top 2%);** *2022*
- **Meritorious Award** in Mathematical Contest In Modeling and Interdisciplinary Contest In Modeling (MCM/ICM) (Top 8% in the participating teams **Worldwide**); *2021*

Technical Skills

- **5+ Years** of research experience in **data science, statistics, machine learning**, with expertise in **PyTorch**, generative and foundation models, **Transformers**, **U-net**, and LLM-based frameworks like **Reinforcement Learning, fine-tuning** Algorithm, PPO, GRPO.
- **5+ Years** of programming experience in **Python, R, and SQL** for data analytics, statistical modeling, and large-scale data processing.