Yingtao Luo

Tel: +1 4129547651 • Pittsburgh, PA • Email Address: yingtaol@andrew.cmu.edu • 954 Citations (until Dec 31, 2025)

SUMMARY

- Ph.D. candidate in Machine Learning focused on LLM and Agents, specializing in long-context agent planning and tool-calling, LLM reasoning and post-training, reinforcement learning (RL), and retrieval (RAG).

EDUCATION

Carnegie Mellon University, PA, USA
Ph.D. in Machine Learning & in Public Policy 08/2022 – 05/2027 (exp.)
M.S. in Machine Learning Research (halfway Ph.D. in Machine Learning) 08/2022 – 05/2025
University of Washington, WA, USA
M.S. in Computer Science and Systems 10/2020 – 06/2022
Huazhong University of Science and Technology, China
B.S. in Applied Physics with Minor in Computer Science 09/2015 – 06/2019

EXPERIENCES

PhD Researcher (Multimodal LLM & Agentic AI), CMU (Pittsburgh, PA) 09/2023 – now

➢ ***Topic 1: Agentic Deep Research Co-Pilot*** *(Presented in NeurIPS 2025)*

- Led the design of a modular agentic orchestration framework with long-context planning, reasoning, retrieval, tool calling, memory, self-verification, achieving sub-30s latency and ~4× faster execution than comparable baselines.
- Built a dynamic MCP discovery–validation pipeline, enabling automatic integration of new domain tools.
- Outperformed all other open-source agent baselines on benchmarks including HealthBench and MedAgentsBench.

➢ ***Topic 2: Retrieval & Representation Learning*** *(Submitted to ICML 2026)*

- Proposed a Generative Embedding paradigm that derives embedding from post-reasoning hidden states, enabling embedding to encode semantic abstraction and reasoning outcomes rather than surface similarity.
- Demonstrated consistent improvements over embedding models on multimodal image-text retrieval benchmarks.

➢ ***Topic 3: Reasoning & LLM Post-Training*** *(Submitted to ICML 2026)*

- Designed an on-policy LLM post-training paradigm that stabilizes reasoning behavior by replacing sparse reward signals with dense trajectory-level latent consistency and contrastive state learning.
- Demonstrated superior reasoning stability and efficiency compared to token-level RLHF-tyle baselines (e.g., PPO, DPO, GRPO) and SFT baselines across multiple reasoning benchmarks.

➢ ***Topic 4: Reinforcement Learning & Application*** *(Ph.D. Dissertation and Fellowship Project)*

- Deployed an AI decision-support system in hospitals, collaborating cross-functionally with experts and stakeholders.
- Developed counterfactual offline RL frameworks (both model-free and model-based) to learn causal accept/reject policies from biased and censored data trajectories to support long-horizon organ offer acceptance decision-making.

Research Intern, Alibaba (Seattle, WA) 2023.05 – 2023.08

- Built a Bayesian online learning for streaming data, outperformed strong baselines under tight latency constraints.
- Deployed for physics-informed weather forecasting application, enabling continuous-resolution real-time inference.
- Demonstrated robust long-horizon generalization on irregularly sampled time series without additional compute cost.

Machine Learning Intern, Microsoft Research Asia (Beijing, China) 2021.09 – 2022.03

- Developed a neural differential operator achieving state-of-the-art performance in time series forecasting.
- Enabled model interpretability by representing temporal dynamics via learned functional bases and PDE structures.
- Improved model reliability under data shift by incorporating uncertainty to suppress overconfident predictions.

PROFICIENT SKILLS

- **Programming:** Python (Pytorch, TensorFlow, DeepSpeed, LangChain, LlamaIndex, Transformers), SQL, Spark
- **Machine Learning:** Large Language Models, Deep Learning, Reinforcement Learning, Multimodal Model
- **Data Mining**: Recommender Systems, A/B Testing, Information Retrieval, Graph Analysis, Time Series Analysis
- **Tools:** Git, Linux, Docker, Kubernetes, DistributedDataParallel, Weights & Biases, Azure, GCP Cloud

AWARDS

- CMLH Fellowship in Generative AI (2025)
- Finalist for Best Paper Award in AMIA (Top 1%, 2025)
- Spotlight paper in ICML (Top 2%, 2024)
- CMLH Translational Fellowship in Digital Health (2023)
- CMU Presidential Fellowship (2022)

PUBLICATIONS

First / Co-first / Corresponding author unless noted (Full publication list available on Google Scholar)

- LLM-Empowered Medical Patient Communication: A Data-Centric Survey From a Clinical Perspective.
  AACL 2025 (ACL Rolling Review 2025 July).
- Demo: Orchestrating Large Language Model Agents and Resources for Medical Investigation.
  NeurIPS 2025 GenAI4Health.
- Your Diffusion Model is Secretly a Noise Classifier and Benefits from Contrastive Training.
  NeurIPS 2024. (Second Author)
- A Billion-Parameter Foundation Model for Capturing Long-Range Gene Context in Single-Cell Transcriptomics.
  Recently accepted by Nature Communications. (Collaborative work; contributed to model training and evaluation)
- Fairness without Demographics through Learning Graph of Gradients.
  KDD 2025.
- Physics-Guided Learning of Meteorological Dynamics for Weather Downscaling and Forecasting.
  KDD 2025.
- Improving Sequential Recommendations via Bidirectional Temporal Data Augmentation with Pre-Training.
  IEEE TKDE 2025.
- Prediction of Mortality After Adult Heart Transplantation Using the Updated UNOS Registry.
  Under minor review at JAMA Cardiology.
- Benchmarking Waitlist Mortality Prediction Through Time-to-Event Modeling using New UNOS Dataset.
  AMIA 2025. ***Finalist for the Best Student Paper (Top 1%).***
- Bayesian Online Multivariate Time series Imputation with functional decomposition.
  ICML 2024. ***Spotlight Paper (Top 2%).*** *(Third Author)*
- Contextualized Policy Recovery: Modeling and Interpreting Medical Decisions with Adaptive Imitation Learning.
  ICML 2024.
- Fairness without Demographics on Electronic Health Records.
  AAAI 2024 Spring Symposium on Clinical Foundation Models. ***Contributed Talk (Top 20%).***
- GSLB: The Graph Structure Learning Benchmark.
  NeurIPS 2023. (Collaborative work; contributed to model evaluation)
- Physics-Guided Discovery of Highly Nonlinear Parametric Partial Differential Equations.
  KDD 2023.
- AdaMCT: Adaptive Mixture of CNN-Transformer for Sequential Recommendation.
  CIKM 2023. (Third Author)
- Reject-aware Multi-Task Network for Financial Credit Scoring.
  IEEE Transactions on Knowledge and Data Engineering (TKDE 2022). (Second Author)
- Deep Stable Representation Learning on Electronic Health Records.
  ICDM 2022.
- Learning Differential Operators for Interpretable Time Series Modeling.
  KDD 2022.
- Symbolic genetic algorithm (SGA) for discovering open-form partial differential equations.
  Physics Review Research. (Second Author)
- STAN: Spatio-Temporal Attention Network for the Next Location Recommendation.
  Proceedings of the Web Conference (WWW 2021).