Chongyang Shi

Gainesville, Florida | c.shi@ufl.edu | 352 226 9039 | Homepage | Google Scholar | LinkedIn | GitHub

Summary

My research focus on **scalable machine learning systems** including **reinforcement learning**, **representation learning**, and **large language models** in **PyTorch**. Experienced in **large-scale GPU** experiments (A100/B200), offline/online evaluation, and end-to-end pipelines for learning from multimodal (camera + lidar) data. Published in leading AI venues including **IJCAI, AAMAS, CDC, ICCPS, and L-CSS**.

Education

University of Florida, Ph.D. in Electrical and Computer Engineering — Aug 2022 – Present

Southern University of Science and Technology, B.S. in Statistics — Sept 2018 – June 2022

Experience

University of Florida, Research Assistant, advisor: Dr. Jie Fu — Aug 2022 – Present
Focus: large-scale deep learning and reinforcement learning systems.

Upcoming Pinterest Labs Research Intern, Machine Learning Engineer, — May 2026 – Aug 2026
Large language model reinforcement learning fine tuning for ads generation.

Selected Projects

Automatic Prompt Optimization for Large Language Models, advisor: Dr. Jie Fu — 2025

- Implemented an **automatic prompt discovery algorithm** based on **Greedy Coordinate Gradient (GCG)** to optimize target-token likelihoods in a frozen **GPT-2** language model.
- Formulated prompt search as a **discrete optimization problem** minimizing negative log-likelihood of specified keywords under teacher-forced continuation.
- Designed a **GPU-accelerated** (NVIDIA T4) optimization pipeline in **PyTorch**, enabling large-scale candidate evaluation with top-k token substitution.

Multimodal Contextual Reinforcement Learning, advisor: Dr. Jie Fu — 2025 – Present

- Developed a **policy-based RL algorithm** for problems where the environment has an **unobservable context** and designed an optimal policy for different contexts based on high-dimensional observations.
- Designed **multimodal encoders** and **latent representations** for decision-making under uncertainty.
- Implemented large-scale training and evaluation pipelines in **PyTorch** using **Nvidia A100/B200 GPUs**.
- Improve the success rate of task by 68% compared to the RDPG-RNN, RDPG-LSTM and POMCP baselines.

Reasoning and Inference in Active Perception, advisor: Dr. Jie Fu — 2024 – 2025

- Designed **information-theoretic learning objectives** for active perception problems, explicitly modeling the trade-off between **control reward and inference accuracy**.
- Developed end-to-end inference pipelines and **policy-gradient-based training** frameworks in PyTorch.
- Achieved a **28.35% improvement in inference accuracy** over the standard policy gradient method.

Multi-Agent Learning, Selection and Coordination, advisor: Dr. Jie Fu — 2025

- Studied **multi-agent coordination** under decentralized and asymmetric information and designed a greedymax algorithm which utilizes the submodularity for **agent selection** and policy optimization.
- Built scalable multi-agent learning benchmarks and evaluated **MAPPO, MADDPG, and IPG** methods.

Sequence Modeling for Investment Analysis, advisor: Dr. Bingyi Jing — 2019 – 2022

- Designed and evaluated sequence models for time-series forecasting using **RNNs, LSTMs, and Transformer architectures** for the stock prediction, investment analysis tasks.
- Conducted systematic comparisons of recurrent and **attention-based** models, analyzing convergence behavior and generalization performance.

Relevant Skills

Deep Learning: Representation learning, multimodal encoders, variational autoencoders, generative models
Tools: PyTorch, GPU training, experiment tracking, data analysis
Programming: Python, Java, C/C++, CUDA (basic), MATLAB
Courses: Mathematical statistics, Bayesian statistics, information theory, time series analysis, stochastic processes, machine learning, deep learning

Publications

Journal Papers

Active Perception With Initial-State Uncertainty: A Policy Gradient Method
Chongyang Shi, Shuo Han, Michael R. Dorothy, Jie Fu
IEEE Control Systems Letters (L-CSS), PDF

Conference Papers

Quantitative Planning with Action Deception in Concurrent Stochastic Games
Chongyang Shi, Shuo Han, Jie Fu
International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2023), PDF

Synthesis of Opacity-Enforcing Winning Strategies Against Colluded Opponent
Chongyang Shi, Abhishek N. Kulkarni, Hazhar Rahmani, Jie Fu
IEEE Conference on Decision and Control (CDC 2023), PDF

Information-Theoretic Opacity-Enforcement in Markov Decision Processes
Chongyang Shi, Yuheng Bu, Jie Fu
International Joint Conference on Artificial Intelligence (IJCAI 2024), PDF

Covert Planning against Imperfect Observers
Haoxiang Ma, **Chongyang Shi**, Shuo Han, Michael R. Dorothy, Jie Fu
International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2024), PDF

Synthesis of Dynamic Masks for Information-Theoretic Opacity in Stochastic Systems
Sumukha Udupa, **Chongyang Shi**, Jie Fu
ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS 2025), PDF

Active Inference through Incentive Design in Markov Decision Processes
Xinyi Wei, **Chongyang Shi**, Shuo Han, Ahmed H. Hemida, Charles A. Kamhoua, Jie Fu,
International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2026), PDF

IMAS^2: Joint Agent Selection and Information-Theoretic Coordinated Perception In Dec-POMDPs
Chongyang Shi, Wesley A. Suttle, Michael Dorothy, Jie Fu,
International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2026), PDF

Preprints

Integrated Control and Active Perception in POMDPs for Temporal Logic Tasks and Information Acquisition
Chongyang Shi, Michael R. Dorothy, Jie Fu, arXiv

Policy Gradient Methods for Information-Theoretic Opacity in Markov Decision Processes
Chongyang Shi, Sumukha Udupa, Michael R. Dorothy, Shuo Han, Jie Fu, arXiv

C-IDS: Solving Contextual POMDP via Information-Directed Objective
Chongyang Shi, Michael R. Dorothy, Jie Fu, arXiv

Honors and Awards

Student Travel Award, AAMAS 2023

Freshmen Scholarship, SUSTech 2018