Yufeng Yang, Ph.D. Student

Seeking Engineer and Research Internship Opportunities in the Field of Machine Learning, Data Science, and Quantitative Finance.

yufeng.yang@tamu.edu Linkedin Personal Website Google Scholar

Education

2023.08 – now **Ph.D. , Computer Science**, Texas A&M University, *College Station, TX*.
2023.08-2024.08 **Ph.D. ,Electrical Engineering** (Transfer out with advisor), University of Utah, *Salt Lake City, UT*.
Advisor: Dr. Yi Zhou
Selected course work: Information Retrieval; Distributed System and Cloud Computing; Deep Reinforcement Learning. Machine Learning; Stochastic Calculus; Game Theory; Multi-agent Reinforcement Learning.

2021.08 – 2023.05 **M.Sc. ,Computational Science and Engineering,** Rice University, *Houston, TX*.
Selected course work: Deep Learning for Vision and Language; Convex Optimization; Optimization under Uncertainty; Algorithms and Data Structure; Numerical Algebra; Scientific Computing; Statistical Signal Processing.

2017.08 – 2021.07 **B.Sc. , Applied Mathematics,** The Chinese University of Hong Kong, *Shenzhen, China*
2018 Summer Exchange Program, University of California, Irvine.
Selected course work: Linear and Integer Optimization; Stochastic Process; Partial Differential Equations; Regression Analysis; Mathematical Statistics; Measure Theoretical Probability; Multivariate statistics; Web Data Analytics; Mathematical Analysis; Information Theory.

Employment History

2024.08 – now **Research Assistant,** Texas A&M University, *College Station, TX*.
Research centers on the design of scalable algorithms for stochastic optimization and reinforcement learning. The primary theme is to develop resilient, robust, and human-value-aligned AI systems capable of overcoming practical challenges such as learning under uncertainty, resolving conflicting objectives, and training with noisy data.
Advanced research in Distributionally Robust Optimization (DRO) by providing rigorous convergence analysis on proposed algorithms under different training paradigm (i.e., Adversarial training, Multi-task training and LLM-post training), and conducting large-scale experiments that demonstrated strong robustness and convergence performance on different tasks. Results are summarized in multiple first-author papers under review top conference and journal including *JMLR* and *ICML*.

2024.06-2024.07 **Summer Visiting PhD Student**, Griffiss Institute, *Rome, NY*
Established "three phase" convergence rate for Adaptive Normalization Gradient Descent (AN-GD) on problems with irregular geometries characterized by generalized smoothness and the PŁ condition, providing theoretical insights into the practical benefits of adaptive normalization on training convergence.
Proposed Independent Sampling and Adaptive Normalized SGD (IAN-SGD); Established convergence results under heavy-tailed gradient noise assumptions.
Achieved state-of-the-art performance on diverse applications in Computer Vision, with results summarized in a first-author paper published in *TMLR*.

2023.08 – 2024.06 **Research Assistant,** University of Utah, *Salt Lake City, UT*.
Formulated the generalized Sinkhorn distance-regularized DRO framework into a nested contextual stochastic programming.
Proposed to solve aforementioned formulation via Nested Stochastic Gradient Descent (SGD) with provable guarantee. These findings were summarized in a first-author paper published in *NeurIPS OPT workshop*.

Employment History (continued)

2020.04 – 2021.01

Research Assistant. Shenzhen Institute of Artificial Intelligence and Robotics for Society, *Shenzhen, China*.
Designed a novel object detection framework for an industrial partner (China Resources Sanjiu) by creating custom loss functions and fusing Generative Adversarial Networks (GANs) with a Convolutional Block Attention Module (CBAM).
Enhanced model robustness against out-of-distribution samples by implementing a data processing pipeline that incorporated advanced data augmentation and unsupervised learning techniques.

Research Publications

Papers with code

1. Y. Yang, F. Zhuo, Z. Chen, H. Huang, and Y. Zhou, "Provably efficient stochastic distributionally-robust multi-objective optimization," *Under Review*, 2026.
2. Y. Yang, E. Tripp, Y. Sun, S. Zou, and Y. Zhou, "Adaptive gradient normalization and independent sampling for (stochastic) generalized-smooth optimization," *Transactions on Machine Learning Research*, 2025. URL: `https://github.com/ynyang94/Gensmooth-IAN-SGD.git`.
3. Y. Yang, Y. Zhou, and Z. Lu, "Nested stochastic algorithm for generalized sinkhorn distance-regularized distributionally robust optimization," *Submitted to Journal of Machine Learning Research*, 2025. URL: `https://github.com/ynyang94/GeneralSinkhorn-Regularized-DRO.git`.
4. Y. Yang, Y. Zhou, and Z. Lu, "A stochastic algorithm for sinkhorn distance-regularized distributionally robust optimization," *NeurIPS 2024 OPT workshop*, 2024.

Other Project Experience

2026

Process Reward-Guided Tree-of-Thoughts-MCTS for Mathematical Reasoning
Integrate Tree-of-Thought (ToT) reasoning with Monte Carlo Tree Search (MCTS) to enhance the multi-step mathematical problem-solving capabilities of a 1.5B parameter Large Language Model.
Developed and evaluated self-critic verification and log-likelihood-based process rewards—to provide granular, step-wise feedback during the reasoning process, leading to a 5% increase in accuracy over terminal-reward baselines on GSM8K.

2025

LUMINA, Multi-modal AIGC image-retrieval system
Constructed a scalable and diverse AIGC image-retrieval dataset associated with text captions.
Participated in designing LUMINA, which **combines multi-modal deep neural networks with a pre-trained VLM verifier** to simulate a human-like retrieval system, capable of retrieving images or videos based on various query forms.

2023

Machine Learning and Optimizer Library
Developed from scratch a **Python library** containing attributes for classical machine learning models, including decision trees, ensemble learning, and parametric methods.
Developed from scratch a **first-order optimizer library**, including frequently-used first-order optimization algorithms such as Nesterov acceleration, mirror descent, proximal gradient descent, and ADMM.

Fair Awareness Image Classification
Addressed imbalanced image classification challenges by evaluating techniques including AUC maximization, transfer learning, contrastive learning and zero-shot learning through a vison-language model. Results demonstrate that leveraging multi-modal information from pre-trained vision-language models is a highly effective strategy for long-tailed datasets.

2022

Distributed Cubic Regularization Optimization
Proposed an **inexact cubic regularization algorithm for distributed logistic regression** on static, undirected graphs, achieving a faster convergence speed than existing distributed first-order methods.

Skills

Programming Language	Python, LaTeX, MATLAB, Linux command, Git, C/C++, R, SQL.
Package/Software	PyTorch, Numpy, Pandas, Matplotlib, JAX, CVX, Gurobi.
Modelling	Machine Learning, Continuous and Discrete Optimization, Differential Equations, Stochastic Modeling and Simulation.