Zhouhao Yang

Email: zyang145@jh.edu Mobile: +1 443-378-6940 Personal Website: zhouhao-yang.github.io

PROFILE

- Current Ph.D. Student in Department of Applied Mathematics and Statistics, Johns Hopkins University.
- Research Interest: AI for Science, LLMs, diffusion models, Reinforcement Learning, stochastic control.
- Diverse research experiences: Ph.D. Candidate of Applied Mathematics and Statistics at JHU, One-Year Ph.D. Student at NUS Computing, Caltech Visiting Undergraduate Research Program.

EDUCATION

Johns Hopkins University Aug 2024 – Present
Ph.D. Candidate, Department of Applied Mathematics and Statistics
Advisor: Haoyang Cao
Research Direction: Memory-Efficient LLMs, Reinforcement Learning, Impulse Control Problems.

National University of Singapore Aug 2023 – Jun 2024
Ph.D. Candidate, Department of Computer Science
Advisor: Kenji Kawaguchi and Johnathon Scarlett
Research Direction: AI for Science, Physics-Informed Neural Networks, Memory-Efficient LLMs.

Shanghai Jiao Tong University Sep 2019 – Jun 2023
Bachelor of Mathematics and Applied Mathematics (Zhiyuan Honors Program)
Selected Coursework: Probability, Statistics, Stochastic Process, Foundations of data science, Numerical Analysis and Scientific Computing, Partial Differential Equations, Fourier Analysis and Real Analysis, etc.

PUBLICATIONS

1. Zhouhao Yang, Xingyu Xu and Yuantao Gu, "A General Framework for Accurate and Private Mean Estimation," in IEEE Signal Processing Letters, vol. 29, pp. 2293-2297, 2022, doi: 10.1109/LSP.2022.3219356.
2. Zhouhao Yang*, Yihong Guo*, Pan Xu, Anqi Liu, and Anima Anandkumar, "Distributionally robust policy gradient for offline contextual bandits," in International Conference on Artificial Intelligence and Statistics, pages 6443–6462. PMLR, 2023.
3. Zheyuan Hu*, Zhouhao Yang*, Yezhen Wang*, George Em Karniadakis, Kenji Kawaguchi, "Bias-variance trade-off in physics-informed neural networks with randomized smoothing for high-dimensional PDEs." *SIAM Journal on Scientific Computing* 47.4 (2025): C846-C872.
4. Qianli Shen, Yezhen Wang, Zhouhao Yang, Kenji Kawaguchi et al., "Memory-Efficient Gradient Unrolling for Large-Scale Bi-level Optimization," The Thirty-eighth Annual Conference on Neural Information Processing Systems. 2024.
5. Yan Hao Ling, Zhouhao Yang, Jonathan Scarlett, "Statistical Mean Estimation with Coded Relayed Observations," Under review in IEEE Transactions on Information Theory.
6. Yezhen Wang*, Zhouhao Yang*, Kenji Kawaguchi et al, "Memory-Efficient LLM Training by Various-Grained Low-Rank Projection of Gradients." Under review in ICLR 2026.
7. Zhouhao Yang*, Nan Chen*, Soufiane Hayou, "Fast Intent Classification for LLM Routing via Statistical Analysis of Representations." Under review in ICLR 2026.
8. Haoyang Cao*, Yuchao Dong*, Zhouhao Yang*, "A Two-fold Randomization Framework for Impulse Control Problems." Under review in Siam Journal on Control and Optimization.

RESEARCH EXPERIENCE

JHU: Fast Intent Classification for LLM Routing via Statistical Analysis of Representations

Aug 2025-Sep 2025

- Professor of guidance: Soufiane Hayou.
- We introduce a training-free statistical method that perform intent classification directly within the LLM prefill phase, which only requires considerably lower computational and memory cost compared to direct LLM call.
- We also introduce a MLP-head method where a small classifier is trained to help LLM routing within the prefill phase.
- The statistical method is based on the statistical analysis of the neural features of the weight matrices.
- We provide theoretical analysis showing when each method excels: classifying with different statistics of neural features has tradeoff between memory, FLOPs, and calibration error.
- We validate our methods across seven LLMs (1B-32B parameters) on both coarse-grained and fine-grained intent classification tasks, demonstrating that statistical methods provide superior uncertainty quantification for mixed-intent prompts compared to overconfident training-based approaches, while achieving competitive accuracy with minimal computational overhead.

JHU: A Two-fold Randomization Framework for Impulse Control Problems Dec 2024-Aug 2025

- Professor of guidance: Haoyang Cao.
- Studied the reinforcement learning framework for impulse control problem, a class of stochastic control problems where the controller cannot continuously adjust the system but instead intervenes at discrete times.
- Designed two randomization methods for impulse control problem (Poisson compound measure/ Relation to optimal stopping) which produces the same HJB equation.
- Proved convergence theorem for randomization approach.
- Designed an actor-critic continuous-time RL algorithm and proved the policy improvement theorem.

NUS: Statistical Mean Estimation with Coded Relayed Observations

May 2024-Aug 2024

- Professor of guidance: Jonathan Scarlett.
- Studied the problem of statistical mean estimation, in which the samples are not observed directly, but are instead observed by a relay ("teacher") that transmits information through a memoryless channel to the decoder ("student"), who then produces the final estimate.
- Focused on Bernoulli sources and binary symmetric channels, and then sub-Gaussian and heavy-tailed settings along with arbitrary discrete memoryless channels.
- Designed a block-structured strategy, proved its superiority on error rate compared with baselines.

NUS: Bias-Variance Trade-off in Physics-Informed Neural Networks with Randomized Smoothing for High-Dimensional PDEs Aug 2023-Nov 2023

- Professor of guidance: George Em Karniadakis from Brown University, Kenji Kawaguchi from NUS.
- Conducted in-depth research for physics-informed neural network with randomized smoothing, a backpropagation-free method for high-dim PINN, which addresses the memory and time costs brought by curse of dimension.
- Analyzed the bias-variance trade-off in Randomized Smoothing-PINN training.
- Corrected the bias brought by randomized smoothing in non-linear MSE loss and PDE terms.
- Proposed a hybrid loss function of biased and unbiased version to accelerate convergence of training and improve final accuracy. Biased training loss provides faster convergence, while unbiased training loss ensures better accuracy.
- Conducted extensive experiments on various PDEs including Fokker-Planck, HJB, Allen-Cahn, Sine-Gorden, etc.

NUS: Gradient-Enhanced Physics-Informed Neural Networks with Variance Reduced Randomized Smoothing for High-Order and High-dimensional PDEs Aug 2023-Jan 2024

- Professor of guidance: George Em Karniadakis from Brown University, Kenji Kawaguchi from NUS.
- Generalized RS-PINN to high-order derivatives and various differential operators.
- Proposed variance reduction methods for randomly smoothed derivative estimators, including antithetic variate method, importance sampling, and Gaussian quadrature.
- Proved that gradient-enhanced techniques are basically free lunch under randomized smoothing, which greatly improves the efficiency of training.
- Conducted thorough experiments on low-dim and high-dim PDEs with RS-GPINN.

Caltech VURP (Visiting Undergraduate Research Program): Distributionally Robust Policy Gradient for Offline Contextual Bandits Mar 2022-Oct 2022

- Professor of guidance: Anima Anandkumar from Caltech, Pan Xu from Duke University.
- Proposed a distributionally robust policy optimization method (DROPO) for offline contextual bandits, addressing the distributional shift between the static logging policy and the learning policy in policy gradient.
- Modelled the distributionally robust policy optimization problem as a minimax optimization problem, which has an analytical framework of solution.
- Generalized the DROPO method to contextual bandits with limited online exploration.
- Provided a generalization bound and a convergence analysis for DROPO.
- Designed and implemented experiments on *UCI Optdigits* and *MNIST*.

A General Framework for Accurate and Private Mean Estimation July 2021-Jun 2022

- Professor of guidance: Yuantao Gu from Tsinghua University.
- Proposed a differentially private mean estimation algorithm for general types of distributions with given cumulative distribution function.
- Proved an upper bound for the algorithm's sample complexity, which is tailored to the cumulative distribution function of underlying population.
- Illustrated that our complexity bound is better in order for heavy-tailed distribution and at least enjoys the same order for light-tailed distribution.
- Designed and implemented experiments on Gaussian distribution and Levy-stable distribution.
- Independently studied the book *"High-dimensional Probability"* written by Roman Vershynin.

Bachelor's Thesis: Machine Learning Methods-Based Inverse Lithography Technology
Aug 2021-May 2023

- Professor of guidance: Dan Hu from Shanghai Jiao Tong University.
- Conducted a thorough literature review for the development of Inverse Lithography Technology over past decades.
- Proposed a complete machine learning methods-based ILT framework is proposed, which consists of a sequential framework of machine learning-based ILT and simulated annealing method-based ILT.
- Designed a convolutional neural network (UNet) to learn the mapping from on-wafer patterns to masks.
- Introduced a simulated annealing method based-ILT framework to refine the quasi-optimized masks and achieve better mask quality.
- Implemented the ML-based ILT framework and conducted experiments on ICCAD 2013 dataset.

HONORS AND AWARDS

- Outstanding students Scholarship of Shanghai Jiao Tong University 2019-2023 (four times)
- Honors scholarship for students of Zhiyuan College 2019-2023 (four times)
- Travel Award for Poster Presentation at ECMF9 Sep 2025

TEACHING EXPERIENCE

- Teaching Assistant for Introduction to Computational Mathematics (EN.553.385) Aug 2024-May 2025
- Instructor of Probability Theory at MSE Orientation 2025 Aug 2025

SKILLS

- Python (PyTorch, pandas, numpy, scipy, etc.), MATLAB, Linux