

Zhouhao Yang

Location: Baltimore, MD

Email: zyang145@jh.edu

Mobile: +1 443-378-6940

Personal Links: [Personal Website](#) [LinkedIn](#) [Google Scholar](#) [Github](#) [Orcid](#)

PROFILE OVERVIEW

- 3rd year Ph.D. candidate in Applied Mathematics and Statistics at Johns Hopkins University, specializing in stochastic control, mean field games, reinforcement learning, and diffusion models.
- First-author research in generative modeling for heavy-tailed data, stochastic control and financial applications, and distributionally robust learning, combining mathematical guarantees with large-scale empirical evaluation.

EDUCATION

Johns Hopkins University

Aug 2024 – May 2028 (Expected)

Ph.D. Candidate, Department of Applied Mathematics and Statistics

Advisor: Haoyang Cao

Research Direction: Reinforcement Learning, Diffusion Models, Stochastic Control Theory, Mean Field Games.

Selected Coursework: Stochastic Control, Games, and Learning (A), Data Science Method for Large Scale Graphs (A), Nonlinear Optimization (A), Statistical Theory (A), Random Matrix Theory in Data Science and Statistics (A).

National University of Singapore

Aug 2023 – Jun 2024

Ph.D. Candidate, Department of Computer Science (Transferred to Johns Hopkins)

Advisor: Kenji Kawaguchi and Johnathan Scarlett

Research Direction: AI for Science, Physics-Informed Neural Networks, Memory-Efficient Finetuning for LLMs.

Shanghai Jiao Tong University

Sep 2019 – Jun 2023

Bachelor of Mathematics and Applied Mathematics (Zhiyuan Honors Program)

Advisor: Dan Hu

RESEARCH EXPERIENCE (SELECTED)

JHU: Heavy-tailed Flow Matching via Random Clocks

March 2026 - May 2026

- Advisor(s): Haoyang Cao and Vladimir Braverman.
- Studied the generation & prediction of heavy-tailed data via flow matching and diffusion models, especially useful in fields like finance and weather forecast.
- Developed a unified flow-matching framework using path-valued random clocks to generally model Gaussian, α -stable, and Student-t source distributions while preserving tractable conditional-Gaussian training objectives.
- Encoded random-clock paths with truncated logsignature features for future training.
- Evaluated the method on imbalanced α -stable mixtures, CIFAR10-LT, and 128×128 HRRR weather fields. Our method outperformed heavy-tailed diffusion and ODE baselines under low sampling budgets.

JHU: A Two-fold Randomization Framework for Impulse Control Problems

Dec 2024 - Sep 2025

- Advisor(s): Haoyang Cao.
- Impulse controls are typically used for modelling exchange rate, transaction cost, insurance, and cash management, where the controller cannot continuously adjust the system but instead intervenes at discrete optimal times.
- Introduced a two-fold randomization scheme over intervention time and jump size to enable exploration of potentially unknown environment.
- Designed an offline, model-free continuous-time RL algorithm and proved the policy improvement theorem.
- Derived compound-Poisson and randomized optimal-stopping formulations leading to the same HJB equation.
- Proved existence, uniqueness and regularity of the solution to randomized framework, and a convergence theorem between the randomized framework and its classical counterpart.

JHU: Training-free versus Training-based Intent Classification in LLMs: Accuracy, Robustness, and Failure Modes

Aug 2025 - Sep 2025

- Advisor(s): Soufiane Hayou.
- Studied intent classification for LLMs, which categorizes user prompts into predefined classes
- Proposed two lightweight training-free intent classification methods. Both methods operate on statistics of LLM prefill-phase feature vector. Developed theory explaining their respective strengths.
- Conducted a systematic empirical study comparing training-free and training-based intent classification across seven LLMs (1B--32B parameters). Both paradigms saturate on easy tasks, while training-based methods typically outperform training-free ones on harder, fine-grained distinctions. Meanwhile, training-free methods are typically more robust to mixed-intent prompts and adversarial rephrasings, making them potentially attractive for practical LLM routing settings where inputs are often noisy or ambiguous.

NUS: Memory-Efficient Gradient Unrolling for Large-Scale Bilevel Optimization

March 2024 - May 2024

- Advisor(s): Kenji Kawaguchi.
- Proposed (FG)²U, a Hessian-free forward-gradient unrolling method that produces an unbiased stochastic estimate of the bilevel meta-gradient without storing the full inner-optimization trajectory.
- Conducted analysis on variance and convergence of (FG)²U and to a parallelizable two-phase optimization strategy that combines efficient approximate updates with accurate late-stage meta-gradient refinement.
- Evaluated the method on image data condensation, online adaptation of language models, and black-box PDE discovery; demonstrated scalability to GPT-2 XL with 1.5B parameters and substantially longer unrolling depths.

Caltech: Distributionally Robust Policy Gradient for Offline Contextual Bandits

Mar 2022 - Oct 2022

- Advisor(s): Anima Anandkumar from Caltech, Pan Xu from Duke University.
- Developed a minimax distributionally robust policy-gradient method to address logging-to-target distribution shift in offline contextual bandits.
- Established generalization and convergence guarantees and extended the framework to settings with limited online exploration.
- Implemented experiments on MNIST and UCI Optdigits to evaluate robustness under policy and distribution shift.

WORKING AND TEACHING EXPERIENCE

- Teaching Assistant at IBM's Thomas J. Watson Research Center July 2026
Brief Description: I served as a teaching assistant for the **Dynamical Systems for Machine Learning and AI** graduate summer workshop, where I prepared lecture notes, delivered several tutorial sessions, and mentored students on their projects.
- Teaching Assistant for Nonlinear Optimization Aug 2025 - Jan 2026
- Instructor of Probability Theory at AMS Master Student Orientation Aug 2025
- Teaching Assistant for Introduction to Computational Mathematics Aug 2024 - May 2025

HONORS AND AWARDS

- Rufus P. Isaacs Graduate Fellowship at Johns Hopkins University Academic Year 2026-2027
- Duncan Research Fund at Johns Hopkins University June 2026, Sep 2026
- Travel Award for Poster Presentation at ECMF9 Sep 2025
- Outstanding students Scholarship of Shanghai Jiao Tong University 2019-2023 (four times)

PUBLICATIONS

In Preparation

1. Haoyang Cao*, **Zhouhao Yang***, "Mean Field Games for Randomized Impulse Control." In preparation (2026).

Published / To Appear

2. **Zhouhao Yang**, Yezhen Wang, Kenji Kawaguchi, Vladimir Braverman, Haoyang Cao, "Heavy-Tailed Flow Matching via Random Clocks." Accepted and to appear in Neurips 2026.
3. **Zhouhao Yang***, Nan Chen*, Soufiane Hayou, "Training-free versus Training-based Intent Classification in LLMs: Accuracy, Robustness, and Failure Modes." Accepted and to appear in COLM 2026.
4. Haoyang Cao*, Yuchao Dong*, **Zhouhao Yang***, "A Two-fold Randomization Framework for Impulse Control Problems." To appear in Siam Journal on Control and Optimization, 2026.
5. Yan Hao Ling, **Zhouhao Yang**, Jonathan Scarlett, "Statistical Mean Estimation with Coded Relayed Observations," in IEEE Transactions on Information Theory, doi: 10.1109/TIT.2026.3716635.
6. 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.
7. Qianli Shen, Yezhen Wang, **Zhouhao Yang**, Xiang Li, Haonan Wang, Yang Zhang, Jonathan Scarlett, Zhanxing Zhu, Kenji Kawaguchi, "Memory-Efficient Gradient Unrolling for Large-Scale Bi-level Optimization," The Thirty-eighth Annual Conference on Neural Information Processing Systems. 2024.
8. **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.
9. **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.

SKILLS

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