Ziheng Cheng

ziheng_cheng@berkeley.edu Homepage github.com/Alexczh1

EDUCATION

Ph.D. in Industrial Engineering and Operations Research Sep, 2024-

UC Berkeley, CA, USA

B.S. in Mathematics Sep, 2020-Jun, 2024

Peking University, Beijing, China

RESEARCH INTEREST

Generative AI, Reinforcement Learning, Machine Learning Theory, Game Theory, Statistics

PUBLICATIONS & MANUSCRIPTS

(* stands for equal contribution)

- Multi-Objective Learning for Diffusion Models: A Statistical Theory under Semi-Supervised Learning
  Ziheng Cheng*, Yixiao Huang*, Hanlin Zhu, Haoran Geng, Somayeh Sojoudi, Jitendra Malik, Pieter Abbeel, Xin Guo*
  (preprint, under review)
- Bridging Discrete and Continuous RL: Stable Deterministic Policy Gradient with Martingale Characterization
  ***Ziheng Cheng**, Xin Guo, Yufei Zhang*
  (preprint, under review)
- OVERT: A Benchmark for Over-Refusal Evaluation on Text-to-Image Models
  Ziheng Cheng*, Yixiao Huang*, Hui Xu, Somayeh Sojoudi, Xuandong Zhao, Dawn Song, Song Mei*
  *(**NeurIPS 2025**)*
- Provable Sample-Efficient Transfer Learning Conditional Diffusion Models via Representation Learning
  ***Ziheng Cheng**, Tianyu Xie, Shiyue Zhang, Cheng Zhang*
  *(**NeurIPS 2025**)*
- Data-Efficient Training by Evolved Sampling
  ***Ziheng Cheng**, Zhong Li, Jiang Bian*
  (preprint, under review)
- Semi-Implicit Functional Gradient Flow for Efficient Sampling
  Shiyue Zhang, **Ziheng Cheng***, Cheng Zhang*
  (preprint, under review)
- Convergence of Distributed Adaptive Optimization with Local Updates
  ***Ziheng Cheng**, Margalit Glasgow*
  *(**ICLR 2025**)*
- Functional Gradient Flows for Constrained Sampling
  Shiyue Zhang,Longlin Yu*, **Ziheng Cheng***, Cheng Zhang*
  *(**NeurIPS 2024**)*
- Kernel Semi-Implicit Variational Inference
  Ziheng Cheng*, Longlin Yu*, Tianyu Xie, Shiyue Zhang, Cheng Zhang*
  *(**ICML 2024**)*
- Reflected Flow Matching
  Tianyu Xie, Yu Zhu*, Longlin Yu*, Tong Yang, **Ziheng Cheng**, Shiyue Zhang, Xiangyu Zhang, Cheng Zhang*
  *(**ICML 2024**)*
- The Limits and Potentials of Local SGD for Distributed Heterogeneous Learning with Intermittent Communication
  *Kumar Kshitij Patel, Margalit Glasgow, Ali Zindari, Lingxiao Wang, Sebastian U Stich, **Ziheng Cheng**, Nirmit Joshi, Nathan Srebro*
  *(**COLT 2024**)*
- Momentum Benefits Non-IID Federated Learning Simply and Provably
  Ziheng Cheng*, Xinmeng Huang*, Pengfei Wu, Kun Yuan*
  *(**ICLR 2024**)*
- Particle-based Variational Inference with Generalized Wasserstein Gradient Flow
  Ziheng Cheng*, Shiyue Zhang*, Longlin Yu, Cheng Zhang*
  *(**NeurIPS 2023**)*
- Joint Graph Learning and Model Fitting in Laplacian Regularized Stratified Models
  Ziheng Cheng*, Junzi Zhang*, Akshay Agrawal, Stephen Boyd*
  (preprint, under review)

RESEARCH EXPERIENCE

Reinforcement Learning & Stochastic Control May, 2025 -

Advisor: Prof. Xin Guo, UC Berkeley

- Studied reinforcement learning in continuous time with deterministic policy.

- Established a mathematical rigorous framework for SDE-driven reinforcement learning for the first time.
- Achieved the state-of-the-art performance in various benchmarks.

Safety Alignment in Generative AI Oct, 2024 - May, 2025

Advisor: Prof. Song Mei, UC Berkeley

- Studied the safety landscape in generative AI, especially on text-to-image models.
- Constructed an automatic pipeline based on LLM to generate large scale evaluation benchmarks.
- Analyzed over-refusal phenomenon and safety-utility tradeoff in T2I models for the first time.

Distributed Optimization & Federated Learning Theory Mar, 2023 - Oct, 2023

Advisor: Prof. Tengyu Ma, Stanford University; Prof. Kun Yuan, Peking University

- Studied the benefits of local iterations to reduce communication in distributed setting.
- Proposed a distributed adaptive optimization algorithm based on gradient-clipping and Adam.
- Achieved the convergence result of Adam in distributed setting for the first time.
- Proved that momentum can accelerate the convergence of classic FL algorithms without additional assumption.

Diffusion Model & Learning Theory & Sampling Algorithm May, 2022 - Oct, 2024

Advisor: Prof. Cheng Zhang, Peking University

- Studied sample complexity theories of transfer learning diffusion models.
- Studied particle-based variational inference methods to accelerate MCMC sampling.

INDUSTRY EXPERIENCE

Bytedance Seed, San Jose, CA, USA Jan, 2026 -

Research Intern

Manager: Lexing Ying

Project: Large Language Model Pre-training

Microsoft Research Asia, Beijing, China Oct, 2023 - May, 2024

Research Intern

Manager: Zhong Li, Jiang Bian

Project: Data Selection for Efficient Training LLM

AWARDS AND HONORS

Awards

- Honorable Mention in Alibaba Global Mathematics Competition 2022, 2023, 2024
- Bronze Medal in S.-T. Yau College Student Mathematics Contest 2022
- Meritorious Winner in Mathematical Contest in Modeling 2022

Honors

- May-Fourth Scholarship (top scholarship in Peking University, 0.5%) 2023
- National Scholarship (top 0.2% nation-wide) 2021
- Merit Student of Peking University 2021-2023