

XINRAN ZHU

<https://xinranzhu.com/> ◊ <https://www.linkedin.com/in/xinranzhu/>

(+1) 607-379-2787 ◊ xz584@cornell.edu

EDUCATION

Cornell University

2018 - Dec 2023 (expected)

Ph.D. student in Applied Mathematics
Minor in Computer Science
Advisor: Prof. [David Bindel](#)

Shanghai Jiao Tong University

2014 - 2018

B.S. in Mathematics
Zhiyuan Honors Degree

RESEARCH INTERESTS

Gaussian Processes, Bayesian Optimization, Statistical Machine Learning, Hyperparameter Optimization, Uncertainty Quantification, Numerical Methods for Data Science, Numerical Optimization.

INTERNSHIP EXPERIENCES

SigOpt, an Intel company

June 2021 - Dec 2021

Machine Learning Research Intern

Supervisor: [Harvey Cheng](#)

- Worked on a hyperparameter optimization research project. Developed **SigOpt Mulch**, a model-aware and cost-aware hyperparameter optimization system for Gradient Boosted Tree (GBT) models. Conducted a parameter cost study and fidelity correlation study of GBT models, and developed the multi-fidelity optimization algorithm in SigOpt Mulch.

Lawrence Berkeley National Laboratory (LBNL)

June 2020 - Aug 2020

Research Intern

Supervisor: [Xiaoye Sherry Li](#)

- Worked on an autotuning project with the scalable solvers group at LBNL. Involved in developing **GP-Tune**, a Bayesian-optimization-driven multi-task autotuning Python software for exascale applications. Analyzed the performance of GPTune and developed metrics for the autotuning performance evaluation. Developed a new multi-fidelity tuning algorithm and developed **GPTuneBand** as a multi-fidelity extension of GPTune.

PUBLICATIONS

Equal contribution is denoted by *.

Variational Gaussian Processes with Decoupled Conditionals

Xinran Zhu, Kaiwen Wu, Natalie Maus, Jacob Gardner, David Bindel

Advances in Neural Information Processing Systems 2023 (NeurIPS 2023).

SigOpt Mulch: An Intelligent System for AutoML of Gradient Boosted Trees

Aleksei Sorokin*, Xinran Zhu*, Eric Hans Lee, Bolong Cheng

Knowledge-Based Systems 273 (2023): 110604.

Bayesian Transformed Gaussian Processes

Xinran Zhu, Leo Huang, Eric Hans Lee, Cameron Ibrahim, David Bindel

Transactions on Machine Learning Research (2023).

GPTuneBand: Multi-task and Multi-fidelity Autotuning for Large-scale High Performance Computing Applications

Xinran Zhu, Yang Liu, Pieter Ghysels, David Bindel, Xiaoye S Li

Proceedings of the 2022 SIAM Conference on Parallel Processing for Scientific Computing (PP22).

Scaling Gaussian Processes with Derivative Information Using Variational Inference

Misha A Padidar, Xinran Zhu, Leo Huang, Jacob R. Gardner, David Bindel

Advances in Neural Information Processing Systems 2021 (NeurIPS 2021).

GPTune: Multitask Learning for Autotuning Exascale Applications

Yang Liu, Wissam M. Sid-Lakhdar, Osni Marques, Xinran Zhu, James W. Demmel, Xiaoye S. Li

Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 2021).

WORKSHOP PAPERS

Efficient Variational Gaussian Processes Initialization via Kernel-based Least Squares Fitting

Xinran Zhu, Jacob R. Gardner, and David Bindel

NeurIPS Workshop on Gaussian Processes, Spatiotemporal Modeling, Decision-making Systems (2022).

ML-based Performance Portability for Time-Dependent Density Functional Theory in HPC Environments

Adrian P. Dieguez, Min Choi, Xinran Zhu, Bryan M. Wong, Khaled Z. Ibrahim

2022 IEEE/ACM International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS).

PROFESSIONAL SERVICE

Reviewer for AISTATS 2022, 2023 and 2024.

Reviewer for SIAM Journal on Scientific Computing.

Organized a Gaussian processes reading group at Cornell.

TECHNICAL STRENGTHS

Machine Learning

kernel methods, statistical modeling, Bayesian optimization

Numerical Methods for Data Science

statistics, linear and nonlinear optimization

Programming Proficiencies

Python, PyTorch, GPyTorch, Julia, MATLAB, R, L^AT_EX

MENTORING EXPERIENCES

Undergraduate and Master projects:

- Bayesian optimization with derivatives. Mentee(s): Chad Yu, Tejal Nair
- Learning To Cluster. Mentee(s): Junyoung Lim (Google), Catherine Horng, Erik Louie
- Function Approximation with Bounded Native Space Norm. Mentee(s): Shengye Zang

TEACHING EXPERIENCE

CS 5220: Applications of Parallel Computers

TA, Cornell University

CS 1112: Introduction to Computing Using MATLAB

TA, Cornell University

CS 6241: Numerical Methods for Data Science

TA, Cornell University

CS 4220: Numerical Analysis: Linear and Nonlinear Problems

TA, Cornell University

MATH 2940: Linear Algebra for Engineers

TA, Cornell University

CS 259: Numerical Methods for Data Science

TA, Shanghai Jiao Tong University