

Education

Carnegie Mellon University

Ph.D. in Computer Science

Aug 2020 – May 2025 (expected)

Pittsburgh, PA, USA

Co-advised by Zico Kolter and Pradeep Ravikumar

Research Focus: Machine Learning, Self-supervised Learning, Representation Learning, Distribution Shift

Peking University

Bachelor of Science (double degree, with summa cum laude honor)

Sep 2016 - Jul 2020

Beijing, China

Majors: Computer Science, Applied Mathematics (Double major)

Member of PKU MOE Top-Notch Undergraduate Researcher Program, advised by Liwei Wang

Professional Experiences

Trexquant Investment LP (hedge fund)

Alpha Researcher Intern

May 2024 – Present

Remote

- Working on machine learning algorithms for portfolio optimization

Amazon Alexa AI

Applied Scientist Intern

May 2022 – Aug 2022

Sunnyvale, CA, USA

- Worked on continuous distribution shift. Mentors: Aram Galstyan, Anoop Kumar, Stefan Schroedl

Microsoft Research Asia (MSRA)

Research Intern, Machine Learning Group

Sep 2019 – Jun 2020

Beijing, China

- Worked on representation learning. Awarded MSRA Award of Excellence. Mentor: Di He.

UCLA

Research Assistant

Jun 2019 – Sep 2019

Los Angeles, CA, USA

- Worked on adversarial robustness. Advisor: Cho-Jui Hsieh

Publications

Refereed Conference and Journal Publications

- [1] **Runtian Zhai**, Bingbin Liu, Andrej Risteski, Zico Kolter, Pradeep Ravikumar
Understanding Augmentation-based Self-Supervised Representation Learning via RKHS Approximation and Regression
International Conference on Learning Representations, **(ICLR 2024 Spotlight)**
- [2] **Runtian Zhai**, Rattana Pukdee, Roger Jin, Maria-Florina Balcan, Pradeep Ravikumar
Spectrally Transformed Kernel Regression
International Conference on Learning Representations, **(ICLR 2024 Spotlight)**
- [3] **Runtian Zhai**, Bingbin Liu, Andrej Risteski, Zico Kolter, Pradeep Ravikumar
Augmentation Alone Leads to Generalization
International Conference on Learning Representations Workshop on Reliable and Responsible Foundation Models, **(ICLR 2023 R2FM Workshop)**
- [4] Yuzhe Lu, Yilong Qin, **Runtian Zhai**, Andrew Shen, Ketong Chen, Zhenlin Wang, Soheil Kolouri, Simon Stepputtis, Joseph Campbell, Katia P. Sycara
Characterizing Out-of-Distribution Error via Optimal Transport
Neural Information Processing Systems, **(NeurIPS 2023)**

- [5] Yash Gupta, **Runtian Zhai**, Arun Suggala, Pradeep Ravikumar
Responsible AI (RAI) Games and Ensembles
Neural Information Processing Systems, (**NeurIPS 2023**)
- [6] **Runtian Zhai**, Chen Dan, Zico Kolter, Pradeep Ravikumar
Understanding Why Generalized Reweighting Does Not Improve Over ERM
International Conference on Learning Representations, (**ICLR 2023**)
- [7] **Runtian Zhai**, Stefan Schroel, Aram Galstyan, Anoop Kumar, Greg Ver Steeg, Pradeep Natarajan
Online Continual Learning for Progressive Distribution Shift (OCL-PDS): A Practitioner's Perspective
International Conference on Learning Representations Workshop on Domain Generalization, (**ICLR 2023 DG Workshop**)
- [8] Yuzhe Lu, Zhenlin Wang, **Runtian Zhai**, Soheil Kolouri, Joseph Campbell, Katia P. Sycara
Predicting Out-of-Distribution Error with Confidence Optimal Transport
International Conference on Learning Representations Workshop on Trustworthy Machine Learning, (**ICLR 2023 Trustworthy ML Workshop**)
- [9] **Runtian Zhai**, Chen Dan, Arun Sai Suggala, Zico Kolter, Pradeep Ravikumar
Boosted CVaR Classification
Neural Information Processing Systems, (**NeurIPS 2021**)
- [10] **Runtian Zhai***, Chen Dan*, Zico Kolter, Pradeep Ravikumar
DORO: Distributional and Outlier Robust Optimization
International Conference on Machine Learning, (**ICML 2021**)
- [11] **Runtian Zhai***, Chen Dan*, Di He*, Huan Zhang, Liwei Wang, Pradeep Ravikumar, Boqing Gong, Cho-Jui Hsieh
MACER: Attack-free and Scalable Robust Training via Maximizing Certified Radius
International Conference on Learning Representations, (**ICLR 2020**)

Preprints

- [1] Yunzhen Feng*, **Runtian Zhai***, Di He, Liwei Wang, Bin Dong
Transferred Discrepancy: Quantifying the Difference Between Representations
arXiv preprint, arXiv:2007.12446
- [2] **Runtian Zhai***, Tianle Cai*, Di He*, Chen Dan, Kun He, John E. Hopcroft, Liwei Wang
Adversarially Robust Generalization Just Requires More Unlabeled Data
arXiv preprint, arXiv:1906.00555

Teaching

CMU 10-701: Introduction to Machine Learning

Fall 2022

- **Head TA:** Assignments, recitals and exams preparation, project mentoring, office hours
- One mentored project converted into a NeurIPS publication [3]

Professional Activities

Conference Reviewer

- International Conference on Learning Representations (ICLR) 2023-2024
- Neural Information Processing Systems (NeurIPS) 2022-2024
- International Conference on Machine Learning (ICML) 2022-2024
- International Conference on Artificial Intelligence and Statistics (AISTATS) 2023-2024
- ACM SIGKDD International Conference on Knowledge Discovery and Data Mining 2023-2024
- SIAM International Conference on Data Mining (SDM) 2024
- International Conference on Computer Vision (ICCV) 2023
- European Conference on Computer Vision (ECCV) 2024

Journal Reviewer

- Journal of Machine Learning Research (JMLR)
- Nature Communications