

Education

Carnegie Mellon University

Ph.D. in Computer Science

Pittsburgh, PA, USA

Aug 2020 – May 2025

Dissertation: *Contextures: The Mechanism of Representation Learning*

Co-advised by Zico Kolter and Pradeep Ravikumar

Research Focus: Machine Learning, Representation Learning, Embedding Models, Learning Theory, Distribution Shift

Peking University

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

Beijing, China

Sep 2016 - Jul 2020

Majors: Computer Science, Applied Mathematics (Double major)

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

Professional Experiences

Headlands Tech

Research Developer

Mar 2026 – Now

New York, NY, USA

- Doing machine learning research

Jump Trading

Quantitative Researcher Intern

Sep 2025 – Nov 2025

New York, NY, USA

- Trained deep learning models for better monetization via reinforcement learning

Jane Street Capital

Machine Learning Researcher Intern

May 2025 – Aug 2025

New York, NY, USA

- Trained Transformers to understand order books and predict trades

Trexquant Investment LP

Alpha Researcher Intern

May 2024 – Aug 2024

Remote

- Worked on machine learning algorithms to combine alphas for better portfolio construction

Amazon Alexa AI

Applied Scientist Intern. Mentors: Aram Galstyan, Anoop Kumar, Stefan Schroedl

May 2022 – Aug 2022

Sunnyvale, CA, USA

- Applied continual learning to deal with continuous data distribution shift. Published paper [8]

Microsoft Research Asia (MSRA)

Research Intern, Machine Learning Group. Mentor: Di He

Sep 2019 – Jun 2020

Beijing, China

- Worked on quantitatively comparing neural network representations. Published preprint [1]

UCLA

Research Assistant. Advisor: Cho-Jui Hsieh

Jun 2019 – Sep 2019

Los Angeles, CA, USA

- Worked on certifiable robust training algorithms. Published paper [12]

Publications

Refereed Conference and Journal Publications

[1] **Runtian Zhai**, Kai Yang, Che-Ping Tsai, Burak Varici, Zico Kolter, Pradeep Ravikumar

Contextures: Representations from Contexts

International Conference on Machine Learning, **(ICML 2025)**

[2] **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)**

[3] **Runtian Zhai**, Rattana Pukdee, Roger Jin, Maria-Florina Balcan, Pradeep Ravikumar

Spectrally Transformed Kernel Regression

International Conference on Learning Representations, **(ICLR 2024 Spotlight)**

- [4] **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)**
- [5] 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)**
- [6] Yash Gupta, **Runtian Zhai**, Arun Suggala, Pradeep Ravikumar
Responsible AI (RAI) Games and Ensembles
 Neural Information Processing Systems, **(NeurIPS 2023)**
- [7] **Runtian Zhai**, Chen Dan, Zico Kolter, Pradeep Ravikumar
Understanding Why Generalized Reweighting Does Not Improve Over ERM
 International Conference on Learning Representations, **(ICLR 2023)**
- [8] **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)**
- [9] 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)**
- [10] **Runtian Zhai**, Chen Dan, Arun Sai Suggala, Zico Kolter, Pradeep Ravikumar
Boosted CVaR Classification
 Neural Information Processing Systems, **(NeurIPS 2021)**
- [11] **Runtian Zhai***, Chen Dan*, Zico Kolter, Pradeep Ravikumar
DORO: Distributional and Outlier Robust Optimization
 International Conference on Machine Learning, **(ICML 2021)**
- [12] **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 15-750: Algorithms in the Real World** *Fall 2024*
- Assignments and exams preparation, office hours
- 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 [5]

Professional Activities

Journal Reviewer

- Nature Communications
- Journal of Machine Learning Research (JMLR)
- Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
- Transactions on Machine Learning Research (TMLR)

Conference Reviewer

- International Conference on Learning Representations (ICLR) 2023-2025
- Neural Information Processing Systems (NeurIPS) 2022-2025
- International Conference on Machine Learning (ICML) 2022-2025
- International Conference on Artificial Intelligence and Statistics (AISTATS) 2023-2025
- ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) 2023-2025
- Conference on Computer Vision and Pattern Recognition (CVPR) 2025
- Association for the Advancement of Artificial Intelligence (AAAI) 2025
- SIAM International Conference on Data Mining (SDM) 2024
- International Conference on Computer Vision (ICCV) 2023-2025
- European Conference on Computer Vision (ECCV) 2024
- Asian Conference on Computer Vision (ACCV) 2024