

## 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: Representation Learning, Generalization Theory, Out-Of-Distribution Generalization

### Peking University

Bachelor of Science (honor, double degree)

Sep 2016 - Jul 2020

Beijing, China

Majors: Computer Science, Applied Mathematics (Double major)

Member of PKU MOE Top-Notch Undergraduate Researcher Program

Research advised by Liwei Wang

## Professional Experiences

---

### Amazon Alexa AI

Applied Scientist Intern

May 2022 – Aug 2022

Sunnyvale, CA, USA

- Mentors: Aram Galstyan, Anoop Kumar, Stefan Schroedl
- Paper: *Online Continual Learning for Progressive Distribution Shift (OCL-PDS): A Practitioner's Perspective*

### Microsoft Research Asia (MSRA)

Research Intern, Machine Learning Group

Sep 2019 – Jun 2020

Beijing, China

- Mentor: Di He
- Paper: *Transferred Discrepancy: Quantifying the Difference Between Representations*

### UCLA

Research Assistant

Jun 2019 – Sep 2019

Los Angeles, CA, USA

- Advisor: Cho-Jui Hsieh
- Paper: *MACER: Attack-free and Scalable Robust Training via Maximizing Certified Radius*

## 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] 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)**
- [4] Yash Gupta, **Runtian Zhai**, Arun Suggala, Pradeep Ravikumar  
**Responsible AI (RAI) Games and Ensembles**  
Neural Information Processing Systems, **(NeurIPS 2023)**
- [5] **Runtian Zhai**, Chen Dan, Zico Kolter, Pradeep Ravikumar  
**Understanding Why Generalized Reweighting Does Not Improve Over ERM**  
International Conference on Learning Representations, **(ICLR 2023)**

- [6] **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)**
- [7] 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)**
- [8] **Runtian Zhai**, Chen Dan, Arun Sai Suggala, Zico Kolter, Pradeep Ravikumar  
**Boosted CVaR Classification**  
 Neural Information Processing Systems, **(NeurIPS 2021)**
- [9] **Runtian Zhai\***, Chen Dan\*, Zico Kolter, Pradeep Ravikumar  
**DORO: Distributional and Outlier Robust Optimization**  
 International Conference on Machine Learning, **(ICML 2021)**
- [10] **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

### Honors and Awards

---

<b>MSRA Award of Excellence</b>	2020
<b>MOE Top-Notch Undergraduate Researcher Program, Peking University (20 students each year)</b>	2019
<b>Changfei Scholarship for Outstanding Students, Peking University (top 5%)</b>	2017

### Teaching

---

<b>CMU 10-701: Introduction to Machine Learning</b>	Fall 2022
<ul style="list-style-type: none"> <li>• <b>Head TA:</b> Assignments, recitals and exams preparation, project mentoring, office hours</li> <li>• One mentored project converted into a NeurIPS publication [3]</li> </ul>	

### Professional Activities

#### Conference Reviewer

- International Conference on Learning Representations (ICLR) 2023-2024
- Neural Information Processing Systems (NeurIPS) 2022-2023
- 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

#### Journal Reviewer

- Journal of Machine Learning Research (JMLR)