Runtian Zhai

GHC 5105, 5000 Forbes Ave, Pittsburgh, PA, USA 15213 ⋄ (1-)412-933-9633
www.runtianzhai.com ⋄ rzhai@cmu.edu ⋄ Github: RuntianZ

EDUCATION

Carnegie Mellon University, Pittsburgh, PA, USA
Ph.D. Student at Computer Science Department *Aug 2020 - May 2025 (expected)*

· Co-advised by Zico Kolter and Pradeep Ravikumar.
· Research focus: Machine learning theory, representation learning, OOD generalization, domain adaptation, continual learning, algorithmic fairness.

Peking University, Beijing, China
Bachelor of Science (honor, double degree) *Sep 2016 - Jul 2020*

· Major in Computer Science. Major GPA: 3.75/4.0
· Double major in Applied Mathematics. Major GPA: 3.57/4.0

SKILLS

· Over 200k lines of code in Python and PyTorch. Familiar with Huggingface.
· Programming Languages: Python, C/C++, MATLAB, HTML/CSS/JS, PHP.

LIST OF PUBLICATIONS

R. Zhai, C. Dan, J.Z. Kolter, P. Ravikumar. *Understanding Why Generalized Reweighting Does Not Improve Over ERM*. ICLR 2023.

R. Zhai, S. Schroel, A. Galstyan, A. Kumar, G. Ver Steeg, P. Natarajan. *Online Continual Learning for Progressive Distribution Shift (OCL-PDS): A Practitioner's Perspective*. ICLR DG workshop 2023.

Y. Lu, Z. Wang, **R. Zhai**, S. Kolouri, J. Campbell, K.P. Sycara. *Predicting Out-of-Distribution Error with Confidence Optimal Transport*. ICLR Trustworthy ML workshop 2023.

R. Zhai, C. Dan, A.S. Suggala, J.Z. Kolter, P. Ravikumar. *Boosted CVaR Classification*. NeurIPS 2021.

R. Zhai*, C. Dan*, J.Z. Kolter, P. Ravikumar. *DORO: Distributional and Outlier Robust Optimization*. ICML 2021.

R. Zhai*, C. Dan*, D. He*, H. Zhang, L. Wang, P. Ravikumar, B. Gong, C.J. Hsieh. *MACER: Attack-free and Scalable Robust Training via Maximizing Certified Radius*. ICLR 2020.

R. Zhai*, T. Cai*, D. He*, C. Dan, K. He, J.E. Hopcroft, L. Wang. *Adversarially Robust Generalization Just Requires More Unlabeled Data*. Preprint, arXiv: 1906.00555.

EMPLOYMENT

Amazon Alexa AI
Applied Scientist Intern. Managers: Stefan Schroedl, Pradeep Natarajan. *May 2022 - Aug 2022*

· Developed benchmarks and methods for OCL-PDS: Online continual learning for progressive distribution shift.
· Studied the distribution shift and model drift problems in real industrial applications.

Microsoft Research Asia (MSRA)
Research Intern in Machine Learning (ML) Group. Mentor: Di He. *Sep 2019 - Jun 2020*

· Conducted research on adversarial robustness, semi-supervised learning and transfer learning.

RESEARCH EXPERIENCE

Peking University
Research Assistant. Advisor: Professor Liwei Wang. *Jul 2018 - Jun 2020*

· Studied adversarial training in the semi-supervised setting. Proposed a robust training algorithm that learns with unlabeled data.

University of California, Los Angeles
Research Assistant. Advisor: Professor Cho-Jui Hsieh. *Feb 2019 - Sep 2019*

· Studied adversarial robustness certification of neural networks. Proposed MACER, an attack-free and scalable provable adversarial defense that trains provable l_2-robust models.