

## Hongyang R. Zhang

Email address: ho.zhang@northeastern.edu

Academic homepage: <https://www.hongyangzhang.com/>

## Employment

- 08/2020-  
Present Assistant Professor, Khoury College of Computer Sciences, *Northeastern University*, Boston, MA
- 10/2019-  
07/2020 Postdoc, Statistics and Data Science, Wharton School, *University of Pennsylvania*, Philadelphia, PA
- 08/2012-  
07/2013 Research Assistant, School of Physical and Mathematical Sciences, *Nanyang Technological University*, Singapore

## Education

- 09/2013-  
09/2019 Ph.D., Computer Science, *Stanford University*, Stanford, CA
- 09/2008-  
06/2012 B.Eng., Computer Science and Technology, *Shanghai Jiao Tong University*, Shanghai, China

## Publications<sup>1</sup>

Total Citations: 1981; h-index: 18 (Source: [Google Scholar], as of May 2026)

### Journal Publications

- 2026 “One-Sided Matrix Completion from Ultra-Sparse Samples.” *H. R. Zhang, Z. Zhang, H. L. Nguyen, and G. Lan. Transactions on Machine Learning Research (TMLR) 2026, Featured Certification*
- 2025 “Precise High-Dimensional Asymptotics for Quantifying Heterogeneous Transfers.” *F. Yang\*, H. R. Zhang\*, S. Wu, C. Ré, and W. Su. Journal of Machine Learning Research (JMLR), 2025*
- 2024 “Noise Stability Optimization for Finding Flat Minima: A Hessian-based Regularization Approach.” *H. R. Zhang, D. Li, and H. Ju. Transactions on Machine Learning Research (TMLR), 2024*
- “Learning Tree-Structured Composition of Data Augmentation.” *D. Li, K. Chen, P. Radivojac, and H. R. Zhang. Transactions on Machine Learning Research (TMLR), 2024*
- 2023 “Improved Group Robustness via Classifier Retraining on Independent Splits.” *T. H. Nguyen, H. R. Zhang, and H. L. Nguyen. Transactions on Machine Learning Research (TMLR), 2023*

---

<sup>1</sup>\* Indicates alphabetical authorship or equal contribution

“Identification of Negative Transfers in Multitask Learning Using Surrogate Models.” D. Li, H. L. Nguyen, and H. R. Zhang. *Transactions on Machine Learning Research (TMLR)*, 2023, **Featured Certification**

2022 “Incentive Ratio: A Game Theoretical Analysis of Market Equilibria.” N. Chen<sup>\*</sup>, X. Deng<sup>\*</sup>, B. Tang<sup>\*</sup>, H. R. Zhang<sup>\*</sup>, and J. Zhang<sup>\*</sup>. *Information and Computation (Inf. Comput.)*, 2022

### Refereed Conference Proceedings

2026 “WinQ: Accelerating Quantization-Aware Training of Language Models around Saddle Points.” D. Li, Z. Liu, K. Yi, Z. Zhang, C. Zhao, R. Krishnamoorthi, H. Khaitan, H. R. Zhang, and S. Li. In *International Conference on Machine Learning (ICML)*, 2026

“Efficient Estimation of Kernel Surrogate Models for Task Attribution.” Z. Zhang, M. Duan, and H. R. Zhang. In *International Conference on Learning Representations (ICLR)*, 2026

“Efficiently Learning Branching Networks for Multitask Algorithmic Reasoning.” D. Li, Z. Zhang, M. Duan, E. Dobriban, and H. R. Zhang. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2026

“Learning Multimodal Embeddings for Traffic Accident Prediction and Causal Estimation.” Z. Zhang, M. Duan, H. N. Koutsopoulos, and H. R. Zhang. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2026, Datasets & Benchmarks Track

“Scalable Multi-Objective and Meta Reinforcement Learning via Gradient Estimation.” Z. Zhang, M. Duan, Y. Ye, and H. R. Zhang. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2026

2025 “Linear-Time Demonstration Selection for In-Context Learning via Gradient Estimation.” Z. Zhang<sup>\*</sup>, Z. Zhang<sup>\*</sup>, D. Li, L. Wang, J. Dy, and H. R. Zhang. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025

“Efficient Ensemble for Fine-tuning Language Models on Multiple Datasets.” D. Li, Z. Zhang, L. Wang, and H. R. Zhang. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2025

2024 “Scalable Fine-tuning From Multiple Data Sources: A First-order Approximation Approach.” D. Li, Z. Zhang, L. Wang, and H. R. Zhang. In *Findings of Empirical Methods in Natural Language Processing (EMNLP Findings)*, 2024

“Scalable Multitask Learning Using Gradient-based Estimation of Task Affinity.” D. Li, A. Sharma, and H. R. Zhang. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2024

2023 “Graph Neural Networks for Road Safety Modeling: Datasets and Evaluations for Accident Analysis.” A. Nippani, D. Li, H. Ju, H. N. Koutsopoulos, and H. R. Zhang. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023, Datasets & Benchmarks Track

“Boosting Multitask Learning on Graphs through Higher-Order Task Affinities.” D. Li, H. Ju, A. Sharma, and H. R. Zhang, in *ACM SIGKDD Conf. on Knowledge Discovery and Data Mining (KDD)*, 2023

- “Generalization in Graph Neural Networks: Improved PAC-Bayesian Bounds on Graph Diffusion.” H. Ju, D. Li, A. Sharma, and H. R. Zhang. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2023
- “Optimal Intervention on Weighted Networks via Edge Centrality.” D. Li, T. Eliassi-Rad, and H. R. Zhang. In *SIAM International Conference on Data Mining (SDM)*, 2023
- 2022 “Robust Fine-Tuning of Deep Neural Networks with Hessian-based Generalization Guarantees.” H. Ju, D. Li, and H. R. Zhang. In *International Conference on Machine Learning (ICML)*, 2022
- “Correct-N-Contrast: A Contrastive Approach for Improving Robustness to Spurious Correlations.” M. Zhang, N. Sohoni, H. R. Zhang, C. Finn, and C. Ré, In *International Conference on Machine Learning (ICML)*, 2022, **Long Presentation** (2.1%)
- 2021 “Improving Regularization and Robustness for Fine-tuning in Neural Networks.” D. Li and H. R. Zhang. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021
- “Observational Supervision for Medical Image Classification using Gaze Data.” K. Sabb, S. Hooper, N. Sohoni, J. Parmar, B. Pogatchnik, S. Wu, J. Dunnmon, H. R. Zhang, D. Rubin, and C. Ré. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2021
- 2020 “Learning Over-Parametrized Two-Layer ReLU Neural Networks beyond NTK.” Y. Li<sup>\*</sup>, T. Ma<sup>\*</sup>, and H. R. Zhang<sup>\*</sup>. In *Conference on Learning Theory (COLT)*, 2020
- “On the Generalization Effects of Linear Transformations in Data Augmentation.” S. Wu<sup>\*</sup>, H. R. Zhang<sup>\*</sup>, G. Valiant, and C. Ré. In *International Conference on Machine Learning (ICML)*, 2020
- “Understanding and Improving Information Transfer in Multi-Task Learning.” S. Wu<sup>\*</sup>, H. R. Zhang<sup>\*</sup>, and C. Ré. In *International Conference on Learning Representations (ICLR)*, 2020
- 2019 “Pruning based Distance Sketches with Provable Guarantees on Random Graphs.” H. R. Zhang, H. Yu, and A. Goel. In *The Web Conference (WWW)*, 2019, Oral Presentation
- “Recovery Guarantees for Quadratic Tensors with Sparse Observations.” H. R. Zhang, V. Sharan, M. Charikar, and Y. Liang. In *Intl. Conf. on Artificial Intelligence and Statistics (AISTATS)*, 2019
- 2018 “Algorithmic Regularization in Over-parameterized Matrix Sensing and Neural Networks with Quadratic Activations.” Y. Li<sup>\*</sup>, T. Ma<sup>\*</sup>, and H. R. Zhang<sup>\*</sup>. In *Conference on Learning Theory (COLT)*, 2018, **Best Paper Award**
- 2016 “Approximate Personalized PageRank on Dynamic Graphs.” H. R. Zhang, P. Lofgren, and A. Goel. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2016, Long Talk
- “Incentives for Strategic Behavior in Fisher Market Games.” N. Chen<sup>\*</sup>, X. Deng<sup>\*</sup>, B. Tang<sup>\*</sup>, and H. R. Zhang<sup>\*</sup>. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2016
- 2015 “A Note on Modeling Retweet Cascades on Twitter.” A. Goel<sup>\*</sup>, K. Munagala<sup>\*</sup>, A. Sharma<sup>\*</sup>, and H. R. Zhang<sup>\*</sup>. In *Workshop on Algorithms and Models for the Web Graph (WAW)*, 2015

- “Connectivity in Random Forests and Credit Networks.” A. Goel<sup>\*</sup>, S. Khanna<sup>\*</sup>, S. Raghvendra<sup>\*</sup>, and H. R. Zhang<sup>\*</sup>. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015
- 2012 “Computing the Nucleolus of Matching, Cover and Clique Games.” N. Chen<sup>\*</sup>, P. Lu<sup>\*</sup>, and H. R. Zhang<sup>\*</sup>. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2012, *Oral Presentation*
- “Incentive Ratios of Fisher Markets.” N. Chen<sup>\*</sup>, X. Deng<sup>\*</sup>, H. R. Zhang<sup>\*</sup>, and J. Zhang<sup>\*</sup>. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, 2012
- 2011 “On Strategy-proof Allocation without Payments or Priors.” L. Han<sup>\*</sup>, C. Su<sup>\*</sup>, L. Tang<sup>\*</sup>, and H. R. Zhang<sup>\*</sup>. In *International Conference on Web and Internet Economics (WINE)*, 2011

## Selected Talks

- 2023-2026 *Second-order Perspectives on Fine-tuning, Task Attribution, and Reinforcement Learning*, Northeastern University (Theory Seminar; Expeditions in AI Seminar), Yale (guest lecture), University of Iowa, Indiana University Bloomington
- 2020-2023 *Towards a Foundation for Multitask Learning*, AAI (New Faculty Highlight), Worcester Polytechnic Institute, Shanghai Jiao Tong University, INFORMS, MSOM, One World Seminar Series on the Mathematics of Machine Learning
- 2020-2021 *Principled Methods for Data Augmentation*. University of Pennsylvania (Workshop on Equivariance and Data Augmentation), Brown University (ML+X Seminar)
- 2018-2019 *Algorithms and Generalization for Large-Scale Data*, Hong Kong University of Science and Technology, Chinese University of Hong Kong, Google Research, Northeastern University, Peking University (Maths; Center for Frontier Computing), Rensselaer Polytechnic Institute, Salesforce Research, Tsinghua University, Stanford University (Theory lunch seminar and ML lunch seminar)
- 2018 *Algorithmic Regularization in Overparameterized Matrix Sensing*, COLT (Stockholm)

## Teaching Experience

Advanced Machine Learning	Spring 2021, 2022, 2023, 2025, 2026
Supervised Machine Learning	Fall 2021, 2022, 2024
Machine Learning	Fall 2023
Machine Learning and Data Mining I	Spring 2023
Principles of Artificial Intelligence	Fall 2025
Special Topics in AI: Algorithmic and Statistical Aspects of Deep Learning	Fall 2020

## PhD Students

- 2025- Minxuan Duan
- 2024- Ziniu (Michael) Zhang
- 2024- Zhenshuo Zhang
- 2021- Dongyue Li. Dissertation title: *Scalable Algorithms for Multitask Learning on Foundation Models: A Gradient-based Framework*. JP Morgan AI PhD Fellowship Awardee. Expected graduation: 2026

## Undergraduate and MS Students

- 2025 Youran Ye, MS student during Spring and Summer 2025. Co-author of an AAAI'26 paper. Will be a PhD student at Northeastern in Fall 2026
- 2024 Kailai Chen, undergraduate student visitor during Fall 2024. Co-author of a TMLR'24 paper. PhD student at Northwestern
- 2023 Abhinav Nippani, MS student during Spring and Summer 2023. First author of a NeurIPS'23 paper. Co-founder of WellDhan
- 2022 Minghao Liu, MS student 2020 - 2021. Software Engineer at Palantir

## Awards

- 2026 Featured Certification, *Transactions on Machine Learning Research* (TMLR)
- 2023 Featured Certification, *Transactions on Machine Learning Research* (TMLR)
- 2023 New Faculty Highlights, *Association for the Advancement of Artificial Intelligence* (AAAI)
- 2018 Best Paper Award, *Conference on Learning Theory* (COLT)
- 2013 School of Engineering Fellowship, *Stanford University*

## Professional Services

### Editorial Boards and Program Committees

**Action editor** for the *Journal of Data-Centric Machine Learning Research* (2024-) and *Transaction on Machine Learning Research* (2025-)

**Area chair** for ML/DM conferences, including: ICML 2024-2026, NeurIPS 2026, KDD 2026, AAAI 2025-2026 (senior program committee), AISTATS 2023-2026, ALT 2024

Program committee / Reviewer for COLT 2024-2025, NeurIPS 2019-2025 (*top 10% reviewers* at NeurIPS'20; *best reviewer* NeurIPS'25), ICML 2019-2023, ICLR 2019-2024, KDD 2022-2024, EMNLP 2024, AAAI 2019-2023, AISTATS 2021-2022, WSDM 2023-2024, ACML 2024-2025, WWW 2022

Journal referee for the *Journal of Machine Learning Research*, *IEEE Transactions on Information Theory*

*Conference Organization*: INFORMS Session Organizer Chair 2023-2026

**College Service**

Research and Scholarship Committee 2025

MS in Data Science Curriculum Committee 2024

PhD Curriculum Committee 2023

PhD Admissions Committee 2020-2022

**Proposal Review and Panel Service**

NSF panelist (2022, 2026) and NSF reviewer (2024, 2026)

Reviewer for NeurIPS Workshop Proposals (2024, 2025)

Last updated: May 15, 2026