

Han Zhao — Curriculum Vitae

3320 Siebel Center, 201 N Goodwin Ave, Urbana, IL – USA – 61801

🌐 <https://hanzhaoml.github.io/>

✉ hanzhao@illinois.edu

☎ (+1) 217-333-2064

Research Interests

Machine Learning: trustworthy machine learning, transfer and multitask learning, algorithmic fairness, domain adaptation/generalization

Artificial Intelligence: probabilistic circuits, graphical models

Professional Experience

University of Illinois at Urbana-Champaign

Department of Computer Science, Tenure-Track Assistant Professor

Aug. 2021 - Present

Department of Electrical and Computer Engineering, Assistant Professor (Affiliated)

Aug. 2021 - Present

Department of Computer Science, Adjunct Assistant Professor

Aug. 2020 - Aug. 2021

Amazon

Amazon Visiting Academic

May 2022 - Present

The D. E. Shaw Group

Machine Learning Researcher

Mar. 2020 - Aug. 2021

Education

Carnegie Mellon University

PhD, Machine Learning Department, School of Computer Science

Pittsburgh, PA, United States

Sep. 2015 - Aug. 2020

◦ Advisor: Prof. Geoff Gordon

◦ Thesis: Towards a Unified Framework for Learning and Reasoning, DOI: 10.13140/RG.2.2.10350.23363

◦ Thesis Committee: Geoffrey J. Gordon (chair), Ruslan Salakhutdinov, Barnabás Póczos and Tommi S. Jaakkola (Massachusetts Institute of Technology)

University of Waterloo

Master of Mathematics, Computer Science

ON, Canada

Sep. 2013 - May. 2015

◦ Advisor: Prof. Pascal Poupart

◦ Alumni Gold Medal Award, University of Waterloo

Tsinghua University

Bachelor of Engineering, Department of Computer Science

Beijing, China

Aug. 2009 - July. 2013

◦ Distinguished Graduate of Tsinghua University

University of Waterloo

Exchange Student, Computer Science

ON, Canada

Sep. 2012 - May. 2013

◦ Advisor: Prof. Pascal Poupart

Honors and Awards

◦ Google Research Scholar Award

2024

◦ AAAI New Faculty Highlights

2024

◦ Kavli Fellow, National Academy of Sciences

2023

◦ Teacher Ranked as Excellent, University of Illinois

2021-2023

◦ Facebook Research Award (Statistics for Improving Insights, Models, and Decisions)

2021

◦ ICLR Outstanding Reviewer Award

2021

◦ ICML Top Reviewer

2020

◦ NeurIPS Highest-Scoring Reviewer

2018-2019

◦ Alumni Gold Medal Award, University of Waterloo

2015

◦ David R. Cheriton Graduate Scholarship, University of Waterloo

2013-2015

- International Masters Student Award, University of Waterloo 2013-2015
- Graduate Research Scholarship, University of Waterloo 2013-2015
- Mathematics Graduate Experience Award, University of Waterloo 2013
- Distinguished Graduate of Tsinghua University 2013
- Tsinghua University Scholarship for Academic Excellence 2010-2012
- Google Excellence Scholarship 2012
- Liu Jimin Scholarship, Tsinghua University 2012
- Third Place in the poster competition of the 2012 REU Program, University of Notre Dame 2012
- Second Prize in Contributions for Laboratory Construction, Tsinghua University 2011
- ICML Travel Grants 2015-2016, 2019
- NIPS/NeurIPS Travel Grants 2016-2019
- ICLR Travel Grants 2018

Research and Industry Experience

Simons Institute for the Theory of Computing, Berkeley Aug. 2024 - Dec. 2024
Long-Term Visitor

- Program: Modern Paradigms in Generalization

Theoretical Sciences Visiting Program, Okinawa Institute of Science and Technology May 2024 - July 2024
Visiting Scholar

- Host: Prof. Makoto Yamada

Theoretical Sciences Visiting Program, Okinawa Institute of Science and Technology May 2023 - July 2023
Visiting Scholar

- Host: Prof. Makoto Yamada

SELECT Lab, Carnegie Mellon University Sep. 2015 - Apr. 2020
Research Assistant

Petuum, Inc. Jan. 2019 - May 2019
Parttime Research Intern

The D. E. Shaw Group May 2018 - Aug. 2018
Quantitative Research Intern
○ Mentor: Dr. Stephen Curran

Microsoft Research AI&R, Redmond May. 2017 - Aug. 2017
Research Intern
○ Mentors: Dr. Ivan Tashev and Dr. Shuayb Zarar

Baidu USA. Silicon Valley AI Lab May. 2016 - Aug. 2016
Research Intern
○ Director: Prof. Andrew Ng. and Dr. Adam Coates

Google Inc. May. 2015 - Aug. 2015
Software Engineer Intern
○ Host: Jerred Costanzo

Noah's Ark Lab & The Chinese University of Hong Kong. Sep. 2014 - Feb. 2015
Research Intern
○ Advisor: Dr. Zhengdong Lu and Dr. Hang Li

Artificial Intelligence Group, University of Waterloo Jan. 2013 - May. 2015
Research Assistant
○ Advisor: Prof. Pascal Poupart

Complex Network Lab, University of Notre Dame June. 2012 - Aug. 2012
Research Assistant
○ Advisor: Prof. Tijana Milenković

Peer-Reviewed Conference Publications

(* denotes equal contribution)

- [C1] R. Xian, Q. Li, G. Kamath, **H. Zhao**, “Differentially Private Post-Processing for Fair Regression”, In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)* (27.5% acceptance rate)
- [C2] S. Liu, D. Zou, **H. Zhao**, P. Li, “Pairwise Alignment Improves Graph Domain Adaptation”, In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)* (27.5% acceptance rate)
- [C3] Y. He, S. Zhou, G. Zhang, H. Yun, Y. Xu, B. Zeng, T. Chilimbi, **H. Zhao**, “Robust Multi-Task Learning with Excess Risks”, In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)* (27.5% acceptance rate)
- [C4] Z. Gong, B. Usman, **H. Zhao**, D. I. Inouye, “Towards Practical Non-Adversarial Distribution Alignment via Variational Bounds”, In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS 2024)* (27.6% acceptance rate)
- [C5] G. Houry, H. Bao, **H. Zhao**, M. Yamada, “Fast 1-Wasserstein distance approximations using greedy strategies”, In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS 2024)* (27.6% acceptance rate)
- [C6] X. Han, J. Chi, Y. Chen, Q. Wang, **H. Zhao**, N. Zou, X. Hu, “FFB: A Fair Fairness Benchmark for In-Processing Group Fairness Methods”, In *Proceedings of the 12th International Conference on Learning Representations (ICLR 2024)* (30.8% acceptance rate)
- [C7] R. Xian, H. Zhuang, Z. Qin, H. Zamani, J. Lu, J. Ma, K. Hui, **H. Zhao**, X. Wang, M. Bendersky, “Learning List-Level Domain-Invariant Representations for Ranking”. In *Proceedings of the 37th Advances in Neural Information Processing Systems (NeurIPS 2023)* (**Spotlight**) (3.6% acceptance rate)
- [C8] Y. Hu, R. Xian, Q. Wu, Q. Fan, L. Yin, **H. Zhao**, “Revisiting Scalarization in Multi-Task Learning: A Theoretical Perspective”. In *Proceedings of the 37th Advances in Neural Information Processing Systems (NeurIPS 2023)* (26.1% acceptance rate)
- [C9] S. Shin, I. Shomorony, **H. Zhao**, “Efficient Learning of Linear Graph Neural Networks via Node Subsampling”. In *Proceedings of the 37th Advances in Neural Information Processing Systems (NeurIPS 2023)* (26.1% acceptance rate)
- [C10] C. Mavromatis, V. N. Ioannidis, S. Wang, D. Zheng, S. Adeshina, J. Ma, **H. Zhao**, C. Faloutsos, G. Karypis, “Train Your Own GNN Teacher: Graph-Aware Distillation on Textual Graphs”. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECMLPKDD 2023)* (24.0% acceptance rate)
- [C11] R. Xian, L. Yin, **H. Zhao**, “Fair and Optimal Classification via Post-Processing Predictors”. In *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)* (27.9% acceptance rate)
- [C12] Y. Hu, F. Wu, H. Zhang, and **H. Zhao**, “Understanding the Impact of Adversarial Robustness on Accuracy Disparity”. In *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)* (27.9% acceptance rate)
- [C13] S. Liu, T. Li, Y. Feng, N. Tran, **H. Zhao**, Q. Qiu, and Pan Li, “Structural Re-weighting Improves Graph Domain Adaptation”. In *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)* (27.9% acceptance rate)
- [C14] Q. Jiang, C. Chen, **H. Zhao**, L. Chen, Q. Ping, S. Dinh Tran, Y. Xu, B. Zeng, T. Chilimbi, “Understanding and Constructing Latent Modality Structures in Multi-modal Representation Learning”. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2023)* (25.8% acceptance rate)
- [C15] S. Zeng, R. des Combes, **H. Zhao**, “Learning Structured Representations by Embedding Class Hierarchy”. In *Proceedings of the 11th International Conference on Learning Representations (ICLR 2023)* (32.0% acceptance rate)

- [C16] S. Shin, **H. Zhao**, I. Shomorony, “Adaptive Power Method: Eigenvector Estimation from Sampled Data”. In *Proceedings of the 34th International Conference on Algorithmic Learning Theory (ALT 2023)* (36.1% acceptance rate)
- [C17] Y. Shen, J. Du, **H. Zhao**, Z. Ji, C. Ma, M. Gao, “FedMM: Saddle Point Optimization for Federated Adversarial Domain Adaptation”. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)* (23.3% acceptance rate)
- [C18] J. Chi, W. Shand, Y. Yu, Kai-Wei Chang, **H. Zhao**, and Y. Tian, “Conditional Supervised Contrastive Learning for Fair Text Classification”. In *Findings of the Association for Computational Linguistics: EMNLP 2022 (EMNLP 2022 Findings)* (14.0% acceptance rate)
- [C19] Z. Chen, R. Jiang, B. Duke, **H. Zhao**, and P. Aarabi, “Exploring Gradient-based Multi-directional Controls in GANs”. In *Proceedings of the European Conference on Computer Vision (ECCV 2022)* (**Oral**) (2.7% acceptance rate)
- [C20] H. Wang, B. Li, **H. Zhao**, “Understanding Gradual Domain Adaptation: Improved Analysis, Optimal Path and Beyond”. In *Proceedings of the 39th International Conference on Machine Learning (ICML 2022)* (21.9% acceptance rate)
- [C21] H. Wang, H. Si, B. Li, **H. Zhao**, “Provable Domain Generalization via Invariant-Feature Subspace Recovery”. In *Proceedings of the 39th International Conference on Machine Learning (ICML 2022)* (21.9% acceptance rate)
- [C22] R. Cheng, G. Balasubramaniam, Y. He, Y. H. Tsai, **H. Zhao**, “Greedy Modality Selection via Approximate Submodular Maximization”. In *Proceedings of the 38th conference on Uncertainty in Artificial Intelligence (UAI 2022)* (32.3% acceptance rate)
- [C23] H. Shao, Y. Yang, H. Lin, L. Lin, Y. Chen, Q. Yang, **H. Zhao**, “Rethinking Controllable Variational Autoencoders”. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2022)* (25.3% acceptance rate)
- [C24] R. Xian, H. Ji, **H. Zhao**, “Cross-Lingual Transfer with Class-weighted Language-Invariant Representations”. In *Proceedings of the 10th International Conference on Learning Representations (ICLR 2022)* (32.9% acceptance rate)
- [C25] Y. H. Tsai, T. Li, M. Q. Ma, **H. Zhao**, K. Zhang, L-P. Morency, R. Salakhutdinov, “Conditional Contrastive Learning with Kernel”. In *Proceedings of the 10th International Conference on Learning Representations (ICLR 2022)* (32.9% acceptance rate)
- [C26] J. Chi, J. Shen, X. Dai, W. Zhang, Y. Tian, **H. Zhao**, “Towards Return Parity in Markov Decision Processes”. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS 2022)* (29.2% acceptance rate)
- [C27] S. Zhou, **H. Zhao**, S. Zhang, L. Wang, H. Chang, Z. Wang, W. Zhu, “Online Continual Adaptation with Active Self-Training”. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS 2022)* (29.2% acceptance rate)
- [C28] B. Li, Y. Shen, Y. Wang, W. Zhu, C. Reed, J. Zhang, D. Li, K. Keutzer, **H. Zhao**, “Invariant Information Bottleneck for Domain Generalization”. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI 2022)* (15.0% acceptance rate)
- [C29] G. Zhang, **H. Zhao**, Y. Yu, and P. Poupart, “Quantifying and Improving Transferability in Domain Generalization”. In *Proceedings of the 35th Advances in Neural Information Processing Systems (NeurIPS 2021)* (25.7% acceptance rate)
- [C30] Z. Zhang, H. Wang, **H. Zhao**, H. Tong and H. Ji, “EventKE: Event-Enhanced Knowledge Graph Embedding”. In *Findings of the Association for Computational Linguistics: EMNLP 2021 (EMNLP 2021 Findings)* (11.8% acceptance rate)
- [C31] J. Chi, Y. Tian, G. Gordon and **H. Zhao**, “Understanding and Mitigating Accuracy Disparity in Regression”. In *Proceedings of the 38th International Conference on Machine Learning (ICML 2021)* (21.5% acceptance rate)

- [C32] H. Wang, **H. Zhao** and B. Li, “Bridging Multi-Task Learning and Meta-Learning: Towards Efficient Training and Effective Adaptation”. In *Proceedings of the 38th International Conference on Machine Learning* (ICML 2021) (21.5% acceptance rate)
- [C33] P. Liao*, **H. Zhao***, K. Xu*, T. S. Jaakkola, G. Gordon, S. Jegelka and R. Salakhutdinov, “Information Obfuscation of Graph Neural Networks”. In *Proceedings of the 38th International Conference on Machine Learning* (ICML 2021) (21.5% acceptance rate)
- [C34] B. Li*, Y. Wang*, S. Zhang*, D. Li, T. Darrell, K. Keutzer and **H. Zhao**, “Learning Invariant Representations and Risks for Semi-supervised Domain Adaptation”. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR 2021) (23.7% acceptance rate)
- [C35] Y. H. Tsai, M. Q. Ma, M. Yang, **H. Zhao**, L-P Morency, R. Salakhutdinov, “Self-supervised Representation Learning with Relative Predictive Coding”, In *Proceedings of the 9th International Conference on Learning Representations* (ICLR 2021) (28.7% acceptance rate)
- [C36] P. Li, Y. Wang, **H. Zhao**, P. Hong, H. Liu, “On Dyadic Fairness: Exploring and Mitigating Bias in Graph Connections”, In *Proceedings of the 9th International Conference on Learning Representations* (ICLR 2021) (28.7% acceptance rate)
- [C37] **H. Zhao***, J. Chi*, Y. Tian and G. Gordon, “Trade-offs and Guarantees on Adversarial Representation Learning for Information Obfuscation”, In *Proceedings of the 34th Advances in Neural Information Processing Systems* (NeurIPS 2020) (20.1% acceptance rate)
- [C38] **H. Zhao***, R. Combes*, Y.X. Wang and G. Gordon, “Domain Adaptation with Conditional Distribution Matching and Generalized Label Shift”. In *Proceedings of the 34th Advances in Neural Information Processing Systems* (NeurIPS 2020) (20.1% acceptance rate)
- [C39] J. Shen, **H. Zhao**, W. Zhang and Y. Yu. “Model-based Policy Optimization with Unsupervised Model Adaptation”. In *Proceedings of the 34th Advances in Neural Information Processing Systems* (NeurIPS 2020) (**Spotlight**) (3.0% acceptance rate)
- [C40] Y. H. Tsai, **H. Zhao**, M. Yamada, L-P. Morency, R. Salakhutdinov, “Neural Methods for Point-wise Dependency Estimation”. In *Proceedings of the 34th Advances in Neural Information Processing Systems* (NeurIPS 2020) (**Spotlight**) (3.0% acceptance rate)
- [C41] **H. Zhao**, J. Hu and A. Risteski, “On Learning Language-Invariant Representations for Universal Machine Translation”. In *Proceedings of the 37th International Conference on Machine Learning* (ICML 2020) (21.8% acceptance rate)
- [C42] W. Wang, **H. Zhao**, H. Zhuang, N. Shah and R. Padman, “DyCRS: Dynamic Interpretable Postoperative Complication Risk Scoring”. In *The World Wide Web Conference* (WWW 2020) (**Oral**) (19.2% acceptance rate)
- [C43] P. Li, **H. Zhao** and H. Liu, “Deep Fair Clustering for Visual Learning”. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR 2020) (22.1% acceptance rate)
- [C44] **H. Zhao**, A. Coston, T. Adel and G. Gordon, “Conditional Learning of Fair Representations”. In *Proceedings of the 8th International Conference on Learning Representations* (ICLR 2020) (**Spotlight**) (4.1% acceptance rate)
- [C45] T. Adel, **H. Zhao** and R. E. Turner, “Continual Learning with Adaptive Weights (CLAW)”. In *Proceedings of the 8th International Conference on Learning Representations* (ICLR 2020) (26.5% acceptance rate)
- [C46] **H. Zhao** and G. Gordon, “Inherent Tradeoffs in Learning Fair Representations”. In *Proceedings of the 33rd Advances in Neural Information Processing Systems* (NeurIPS 2019) (21.1% acceptance rate)
- [C47] **H. Zhao***, Y. H. Tsai*, R. Salakhutdinov and G. Gordon, “Learning Neural Networks with Adaptive Regularization”. In *Proceedings of the 33rd Advances in Neural Information Processing Systems* (NeurIPS 2019) (21.1% acceptance rate)

- [C48] **H. Zhao**, O. Stretcu, A. Smola and G. Gordon, “Efficient Multi-task Feature and Relationship Learning”. In *Proceedings of the 35th conference on Uncertainty in Artificial Intelligence* (UAI 2019) (26.0% acceptance rate)
- [C49] **H. Zhao***, Y. Xu*, X. Shi and N. B. Shah, “On Strategyproof Conference Peer Review”. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence* (IJCAI 2019) (17.9% acceptance rate)
- [C50] **H. Zhao**, R. Combes, K. Zhang and G. Gordon, “On Learning Invariant Representation for Domain Adaptation”. In *Proceedings of the 36th International Conference on Machine Learning* (ICML 2019) (**Long Oral**) (2.3% acceptance rate)
- [C51] **H. Zhao**, J. Hu, Z. Zhu, A. Coates and G. Gordon, “Deep Generative and Discriminative Domain Adaptation”. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems* (AAMAS 2019) (24.3% acceptance rate)
- [C52] C. Liang, J. Ye, **H. Zhao**, B. Pursel and C. Lee Giles, “Active Learning of Strict Partial Orders: A Case Study on Concept Prerequisite Relations”. In *Proceedings of the 12th International Conference on Educational Data Mining* (EDM 2019) (34.6% acceptance rate)
- [C53] **H. Zhao***, S. Zhang*, G. Wu, J. Costeira, J. Moura and G. Gordon, “Adversarial Multiple Source Domain Adaptation”. In *Proceedings of the 32nd Advances in Neural Information Processing Systems* (NeurIPS 2018) (20.8% acceptance rate)
- [C54] **H. Zhao** and G. Gordon, “Frank-Wolfe Optimization for Symmetric-NMF under Simplicial Constraint”, In *Proceedings of the 34th conference on Uncertainty in Artificial Intelligence* (UAI 2018) (30.8% acceptance rate)
- [C55] **H. Zhao**, S. Zarar, I. Tashev and C.-H. Lee, “Convolutional-Recurrent Neural Networks for Speech Enhancement”. In *Proceedings of 2018 IEEE International Conference on Acoustics, Speech and Signal Processing* (ICASSP 2018) (**Oral**) (49.7% acceptance rate)
- [C56] **H. Zhao** and G. Gordon, “Linear Time Computation of Moments in Sum-Product Networks”. In *Proceedings of the 31st Advances in Neural Information Processing Systems* (NIPS 2017) (20.9% acceptance rate)
- [C57] T. Adel, **H. Zhao** and A. Wong, “Unsupervised Domain Adaptation with a Relaxed Covariate Shift Assumption”. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence* (AAAI 2017) (24.6% acceptance rate)
- [C58] **H. Zhao**, P. Poupart and G. Gordon, “A Unified Approach for Learning the Parameters of Sum-Product Networks”. In *Proceedings of the 30th Advances in Neural Information Processing Systems* (NIPS 2016) (23.6% acceptance rate)
- [C59] P. Jaini, A. Rashwan, **H. Zhao**, Y. Liu, E. Banijamali, Z. Chen and P. Poupart, “Online Algorithms for Sum-Product Networks with Continuous Variables”. In *Proceedings of the 8th International Conference on Probabilistic Graphical Models* (PGM 2016)
- [C60] **H. Zhao**, T. Adel, G. Gordon and B. Amos, “Collapsed Variational Inference for Sum-Product Networks”. In *Proceedings of the 33rd International Conference on Machine Learning* (ICML 2016) (24.0% acceptance rate)
- [C61] A. Rashwan, **H. Zhao** and P. Poupart, “Online and Distributed Bayesian Moment Matching for SPNs”. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics* (AISTATS 2016) (30.7% acceptance rate)
- [C62] **H. Zhao**, M. Melibari and P. Poupart, “On the Relationship between Sum-Product Networks and Bayesian Networks”. In *Proceedings of the 32nd International Conference on Machine Learning* (ICML 2015) (26.0% acceptance rate)
- [C63] **H. Zhao**, Z. Lu and P. Poupart, “Self-Adaptive Hierarchical Sentence Model”. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence* (IJCAI 2015) (28.6% acceptance rate)

- [C64] **H. Zhao**, P. Poupart, Y. Zhang and M. Lysy, “SoF: Soft-Cluster Matrix Factorization for Probabilistic Clustering”. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence* (AAAI 2015) (26.7% acceptance rate)
- [C65] T. Milenković, **H. Zhao** and F. Faisal, “Global Network Alignment in the Context of Aging”. In *Proceedings of the 4th ACM International Conference on Bioinformatics, Computational Biology and Biomedicine* (ACM-BCB 2013) (29.0% acceptance rate)

Journal Publications

(* denotes equal contribution)

- [J1] H. Wang, H. Si, H. Shao, **H. Zhao**, “Enhancing Compositional Generalization via Compositional Feature Alignment”. In *Transactions on Machine Learning Research* (TMLR 2024).
- [J2] Y. Yang; M. Lin, **H. Zhao**, Y. Peng, Z. Lu, “A Survey of Recent Methods for Addressing AI Fairness and Bias in Biomedicine”, In *Journal of Biomedical Informatics* (JBI 2024).
- [J3] Y. He, R. Cheng, G. Balasubramaniam, Y. H. Tsai, **H. Zhao**, “Efficient Modality Selection in Multimodal Learning”, In *Journal of Machine Learning Research* (JMLR 2024).
- [J4] X. Wang, **H. Zhao**, K. Nahrstedt, S. Koyejo, “Personalized Federated Learning with Spurious Features: An Adversarial Approach”, In *Transactions on Machine Learning Research* (TMLR 2024).
- [J5] J. Shen, H. Lai, M. Liu, **H. Zhao**, Y. Yu, and W. Zhang, “Adaptation Augmented Model-based Policy Optimization”. In *Journal of Machine Learning Research* (JMLR 2023).
- [J6] **H. Zhao**, “Costs and Benefits of Fair Regression”. In *Transactions on Machine Learning Research* (TMLR 2023).
- [J7] J. Dong, S. Zhou, B. Wang, **H. Zhao**, “Algorithms and Theory for Supervised Gradual Domain Adaptation”. In *Transactions on Machine Learning Research* (TMLR 2022).
- [J8] **H. Zhao***, C. Dan*, B. Aragam, T. Jaakkola, G. Gordon, and P. Ravikumar, “Fundamental Limits and Tradeoffs in Invariant Representation Learning”. In *Journal of Machine Learning Research* (JMLR 2022).
- [J9] **H. Zhao** and G. Gordon, “Inherent Tradeoffs in Learning Fair Representations”. In *Journal of Machine Learning Research* (JMLR 2022).
- [J10] S. Zhao, X. Yue, S. Zhang, B. Li, **H. Zhao**, B. Wu, R. Krishna, J. E. Gonzalez, A. L. Sangiovanni-Vincentelli, S. A. Seshia and K. Keutzer, “A Review of Single-Source Deep Unsupervised Visual Domain Adaptation”. In *IEEE Transactions on Neural Networks and Learning Systems* (IEEE TNNLS 2020).
- [J11] F. Faisal, **H. Zhao** and T. Milenković, “Global Network Alignment in the Context Of Aging”, In *IEEE/ACM Transactions on Computational Biology and Bioinformatics* (IEEE/ACM TCBB 2014).

Workshop Papers

(* denotes equal contribution)

- [W1] W. Chu, C. Xie, B. Wang, L. Li, L. Yin, **H. Zhao**, B. Li, “FOCUS: Fairness via Agent-Awareness for Federated Learning on Heterogeneous Data” In *International Workshop on Federated Learning in the Age of Foundation Models* (NeurIPS 2023) (**Oral**)
- [W2] Z. Chen, **H. Zhao**, P. Aarabi, R. Jiang, “SC2GAN: Rethinking Entanglement by Self-correcting the Correlated GAN Space”. In *The 2nd Workshop and Challenges for Out-of-Distribution Generalization in Computer Vision* (ICCV 2023)
- [W3] G. Balasubramaniam, H. Wang, and **H. Zhao**, “Invariant Feature Subspace Recovery for Multi-Class Classification”. In *Workshop on Distribution Shifts: Connecting Methods and Applications* (NeurIPS 2022)

- [W4] M. Q. Ma, Y.-H H. Tsai, P. P. Liang, **H. Zhao**, K. Zhang, R. Salakhutdinov, L.-P. Morency, “Conditional Contrastive Learning for Improving Fairness in Self-Supervised Learning”. In *Workshop on Self-Supervised Learning – Theory and Practice* (NeurIPS 2022)
- [W5] Y. He, H. Wang, and **H. Zhao**, “Generative Gradual Domain Adaptation with Optimal Transport”. In *Principles of Distribution Shift (PODS)* (ICML 2022)
- [W6] **H. Zhao**, A. Coston, T. Adel and G. Gordon, “Conditional Learning of Fair Representations”. In *Workshop on Machine Learning with Guarantees* (NeurIPS 2019).
- [W7] **H. Zhao***, J. Chi*, Y. Tian and G. Gordon, “Adversarial Privacy Preservation under Attribute Inference Attack”. In *Workshop on Machine Learning with Guarantees* (NeurIPS 2019).
- [W8] **H. Zhao***, Y. H. Tsai*, R. Salakhutdinov and G. Gordon, “Approximate Empirical Bayes for Deep Neural Networks”. In *Uncertainty in Deep Learning workshop* (UAI 2018).
- [W9] **H. Zhao***, S. Zhang*, G. Wu, J. Costeira, J. Moura and G. Gordon, “Multiple Source Domain Adaptation with Adversarial Learning”. In *Proceedings of the 6th International Conference on Learning Representations* (ICLR 2018, workshop track).
- [W10] Y. H. Tsai, **H. Zhao**, R. Salakhutdinov and N. Jojic, “Discovering Order in Unordered Datasets: Generative Markov Networks”. In *Time Series workshop* (NIPS 2017, arXiv:1711.03167).
- [W11] **H. Zhao**, O. Stretcu, R. Negrinho, A. Smola and G. Gordon, “Efficient Multi-task Feature and Relationship Learning”. In *Learning with Limited Labeled Data: Weak Supervision and Beyond workshop* (NIPS 2017, arXiv:1702.04423).
- [W12] **H. Zhao** and P. Poupart, “A Sober Look at Spectral Learning”. In *Method of Moments and Spectral Learning workshop* (ICML 2014, arXiv:1406.4631).

Patents

- [P1] I. Tashev, S. Zarar, C-H. Lee, Y-H. Tu and **H. Zhao**, “Systems, Methods, and Computer-Readable Media for Improved Real-Time Audio Processing”. US Patent. US 15/952,353.

Conference and Invited Talks

Revisiting Scalarization in Multi-Task Learning

- Deep Learning: Theory, Applications, and Implications (DL2024), Tokyo, Japan Mar. 2024

Towards Foundation Models for Geospatial Data

- Department of Architecture and Urban Design, Kyushu University Mar. 2024

Directional Preference Alignment with Multi-Objective Rewards for LLMs

- AI & Search Science Talk Series, Amazon Mar. 2024

Algorithmic Fairness and Robust Generalization from a Causal Perspective

- Technical Seminar, Radix Trading, LLC. Feb. 2024
- Spatial Omics Initiative, IGB Center for Artificial Intelligence and Modeling Feb. 2024

Trustworthy Machine Learning: Theory, Algorithms and Applications

- Machine Learning Summer School, Okinawa, Japan Mar. 2024

Robust Learning under Distribution Shifts

- Advanced Controls Research Laboratory, University of Illinois Urbana-Champaign Apr. 2024
- Trustworthy AI workshop, Osaka University, Japan Mar. 2024
- Machine Learning Summer School, Okinawa, Japan Mar. 2024
- Kavli Frontiers of Science: Japanese-American-German Frontiers of Science Symposium (JAGFOS), Dresden, Germany Oct. 2023

Fair and Optimal Prediction via Post-Processing

- New Faculty Highlights, AAAI Conference on Artificial Intelligence (AAAI), Vancouver, Canada Feb. 2024

- Dagstuhl Seminars: Emerging Issues in Bioimaging AI Publications, Germany Jan. 2024
- Waterloo Artificial Intelligence Institute (Waterloo.ai), University of Waterloo Nov. 2023
- Methods and Theory Seminar Series, Department of Statistical Sciences, University of Toronto Nov. 2023
- Lunch & Learn Seminar Series, Modiface Inc., Toronto, Canada Nov. 2023
- Computational Biology Branch, National Center for Biotechnology Information Nov. 2023
- The Institute for Data, Econometrics, Algorithms, and Learning (IDEAL), Northwestern University Oct. 2023
- Computer Engineering Seminar, Purdue University Oct. 2023
- The Information-Based Induction Sciences and Machine Learning (IBISML) Workshop, OIST, Japan July 2023

- Robust Multi-Task Learning with Excess Risks**
- AI & Search Science Talk Series, Amazon May 2023
- AWS Tech-Talk Series, Amazon June 2023

- Learning Structured Representations by Embedding Class Hierarchy**
- M5 Tech-Talk Series, AI & Search Science, Amazon Oct. 2023

- Understanding and Constructing Latent Modality Structures in Multi-modal Representation Learning**
- AI & Search Science Talk Series, Amazon Aug. 2022

- Provable Domain Generalization via Invariant-Feature Subspace Recovery**
- TrustML Young Scientist Seminars, RIKEN, Japan Mar. 2023
- Statistical Artificial Intelligence and Learning Group, Tsinghua University, China July 2022

- Understanding Gradual Domain Adaptation: Improved Analysis, Optimal Path and Beyond**
- Dagstuhl Seminars: Recent Advancements in Tractable Probabilistic Inference, Germany Apr. 2022
- Microsoft Research, Distinguished Talk Series, USA May 2022

- Bridging Multi-Task Learning and Meta-Learning: Towards Efficient Training and Effective Adaptation**
- AI & Search Science Talk Series, Amazon Mar. 2022

- Costs and Benefits of Invariant Representation Learning**
- University of California, Berkeley, ML & CV Seminar July 2020
- Google Research July 2020
- Tsinghua University, XuetangX Seminar June 2020
- (Virtual) Brown University, Department of Computer Science Apr. 2020
- University of California, Santa Barbara, Department of Computer Science Mar. 2020
- University of North Carolina, Chapel Hill, Department of Computer Science Mar. 2020
- (Virtual) Hong Kong University of Science and Technology, Department of Computer Science Mar. 2020
- University of Illinois at Urbana-Champaign, Department of Computer Science Feb. 2020
- University of Virginia, Department of Computer Science Feb. 2020
- Dartmouth College, Department of Computer Science Feb. 2020
- University of Waterloo, David R. Cheriton School of Computer Science Feb. 2020
- Borealis AI Lab, Toronto Feb. 2020
- Pennsylvania State University, Department of Computer Science and Engineering Feb. 2020
- Brandeis University Feb. 2020

- Inherent Tradeoffs in Learning Invariant Representations**
- Massachusetts Institute of Technology Oct. 2019
- Brandeis University Nov. 2019

- Learning Neural Networks with Adaptive Regularization**
- ShanghaiTech University Aug. 2019
- New York University, Shanghai Aug. 2019
- Tencent AI Lab Aug. 2019
- Noah's Ark Lab, Huawei Aug. 2019
- Toutiao AI Lab, ByteDance Aug. 2019
- Amazon AWS AI Lab, Shanghai Sep. 2019
- Borealis AI Lab, Waterloo Sep. 2019
- University of Waterloo, Waterloo Sep. 2019

On Strategyproof Conference Peer Review

- International Joint Conference on Artificial Intelligence Aug. 2019

On Learning Invariant Representations for Domain Adaptation

- Microsoft Research Seminar series, Microsoft Research Montreal Lab, Canada May. 2019
- International Conference on Machine Learning June 2019

Multiple Source Domain Adaptation with Adversarial Learning

- AI seminar, Carnegie Mellon University Apr. 2018
- Microsoft Research Seminar series, Microsoft Research Montreal Lab, Canada Sep. 2018
- Technical seminar, Petuum Inc. Oct. 2018
- Technical seminar, Pony. AI Nov. 2018

High-Accuracy Neural-Network Models for Speech Enhancement

- Microsoft Research Seminar series, Microsoft Research Redmond Lab Aug. 2017
- IEEE International Conference on Acoustics, Speech and Signal Processing Apr. 2018

Sum-Product Networks: A New Probabilistic Inference Machine

- AI seminar, Carnegie Mellon University Mar. 2017

Collapsed Variational Inference for Sum-Product Networks

- International Conference on Machine Learning June 2016

On the Relationship between Sum-Product Networks and Bayesian Networks

- AI seminar, University of Waterloo Mar. 2015
- International Conference on Machine Learning July 2015

Self-Adaptive Hierarchical Sentence Model

- AI seminar, University of Waterloo Mar. 2015
- Technical seminar, Google Apr. 2015
- International Joint Conference on Artificial Intelligence July 2015
- AI lunch, Jump Trading Nov. 2017

SoF: Soft-Cluster Matrix Factorization for Probabilistic Clustering

- Association for the Advancement of Artificial Intelligence Jan. 2015

Professional Activities**National Science Foundation (NSF) review panelist**

- CISE: Core Programs: IIS 2022
- OAC: OAC Core 2023
- GEO/RISE: CAIG 2024

Mentoring Program

- Mentorship Program (ICLR) 2024

Conference Area Chair/Senior Program Committee

- Area Chair, International Conference on Artificial Intelligence and Statistics (AISTATS) 2024
- Senior Program Committee, AAAI Conference on Artificial Intelligence (AAAI) 2021-2024
- Area Chair, International Conference on Machine Learning (ICML) 2020-2024
- Area Chair, Neural Information Processing Systems (NeurIPS) 2021-2023
- Area Chair, International Conference on Learning Representations (ICLR) 2024

Conference Reviewer/Program Committee Member

- IEEE International Symposium on Information Theory (ISIT) 2021
- Advances in Neural Information Processing Systems (NIPS/NeurIPS) 2016, 2018-2021
- International Conference on Machine Learning (ICML) 2017-2021
- International Conference on Artificial Intelligence and Statistics (AISTATS) 2017-2021
- International Conference on Learning Representations (ICLR) 2019-2021
- International Conference on Computer Vision (ICCV) 2019
- Conference on Computer Vision and Pattern Recognition (CVPR) 2020

- European Conference on Computer Vision (ECCV) 2020
- Time Series Workshop (ICML) 2021
- IEEE Workshop on Statistical Signal Processing (SSP) 2016
- Workshop on Adaptive & Multitask Learning: Algorithms & Systems (ICML) 2019
- Workshop on Statistical Deep Learning in Computer Vision (ICCV) 2019
- AAAI Conference on Artificial Intelligence (AAAI) 2019-2020
- International Joint Conference on Artificial Intelligence (IJCAI) 2020-2021
- Asian Conference on Machine Learning (ACML) 2019
- Deep Generative Models Workshop, ICML 2018
- Workshop on Tractable Probabilistic Models, ICML 2018
- Workshop on Principled Approaches to Deep Learning, ICML 2017
- Workshop on Learning with Rich Experience: Integration of Learning Paradigms, NeurIPS 2019

Journal Reviewer

- SIAM Journal on Mathematics of Data Science (SIMODS)
- IEEE Computational Intelligence Magazine
- Neural Computing
- Transactions on Audio, Speech and Language Processing (TASLP)
- Pattern Analysis and Applications
- PLoS One
- IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
- Machine Learning Journal (MLJ)
- IEEE Access
- Journal of Artificial Intelligence Research (JAIR)
- Journal of Machine Learning Research (JMLR)
- Transactions on Machine Learning Research (TMLR)

(Co-) Organizers

- CVPR 2021 Responsible Computer Vision Workshop
- KDD 2021 Machine Learning for Consumers and Markets Workshop

Professional Membership

- Member of the Association for the Advancement of Artificial Intelligence 2023 - Present
- Member of Association for Computing Machinery 2021 - Present

Student Advising

Ph.D. Students

- Weixin Chen (UIUC CS PhD)
- Yifei He (UIUC CS PhD)
- Yuzheng Hu (UIUC CS PhD)
- Seiyun Shin (UIUC ECE PhD, co-advised with Ilan Shomorony) Mavis Future Faculty Fellows, 2023-2024
- Haozhe Si (UIUC ECE PhD)
- Haoxiang Wang (UIUC ECE PhD, co-advised with Bo Li) Mavis Future Faculty Fellows, 2022-2023
- Ruicheng Xian (UIUC CS PhD)
- Siqi (Cindy) Zeng (UIUC CS PhD)

Thesis-based Master Students

- Gargi Balasubramaniam (UIUC MSCS → Google DeepMind) Siebel Scholar, Class of '23
- Yifei He (UIUC MSCS → UIUC CS PhD)
- Haozhe Si (UIUC ECE MS → UIUC ECE PhD)
- Aditya Sinha (UIUC MSCS)
- Qilong Wu (UIUC MSCS)

Undergraduate Mentoring

- Samuel Schapiro (UIUC)

- Sixian Du (Peking University → Stanford MSEE)
- Siqi (Cindy) Zeng (CMU → UIUC CS PhD)
- Haozhe Si (UIUC → UIUC ECE Master) Thesis: *ISR: Invariant Subspace Recovery*
- Peiyuan Liao (CMU → CTO of Cyber Manufacture Co.)
- Bo Li (Harbin Institute of Technology → Ph.D. student, National University of Singapore)

Ph.D. Thesis Committee

- | | |
|---|---------------------------|
| ○ Qingyun Wang, University of Illinois Urbana-Champaign | Advisor: Heng Ji |
| ○ Liliang Ren, University of Illinois Urbana-Champaign | Advisor: Chengxiang Zhai |
| ○ Kung-Hsiang (Steeve) Huang, University of Illinois Urbana-Champaign | Advisor: Heng Ji |
| ○ Tianshi Wang, University of Illinois Urbana-Champaign | Advisor: Tarek Abdelzaher |
| ○ Qian Jiang, University of Illinois Urbana-Champaign | Advisor: Minh N. Do |
| ○ Jun Wu, University of Illinois Urbana-Champaign | Advisor: Jingrui He |
| ○ Jian Kang, University of Illinois Urbana-Champaign | Advisor: Hanghang Tong |
| ○ (Vicki) Qi Zeng, University of Illinois Urbana-Champaign | Advisor: Heng Ji |
| ○ Olawale Elijah Salaudeen, University of Illinois Urbana-Champaign | Advisor: Sanmi Koyejo |
| ○ Shiji Zhou, Tsinghua University | Advisor: Wenwu Zhu |
| ○ Jianfeng Chi, University of Virginia | Advisor: Yuan Tian |

Teaching Experience

Instructor, University of Illinois at Urbana-Champaign

- | | |
|---|-------------|
| ○ CS 446: Machine Learning | Spring 2024 |
| ○ CS 442: Trustworthy Machine Learning (Teacher Ranked as Excellent, University of Illinois) | Fall 2023 |
| ○ CS 598: Transfer Learning | Spring 2023 |
| ○ CS 498: Trustworthy Machine Learning | Fall 2022 |
| ○ CS 442: Trustworthy Machine Learning (Teacher Ranked as Excellent, University of Illinois) | Spring 2022 |
| ○ CS 598: Transfer Learning (Teacher Ranked as Excellent, University of Illinois) | Fall 2021 |

Instructor, TechX Academy

- | | |
|--|-------------|
| ○ Advanced Introduction to Deep Learning | Summer 2019 |
|--|-------------|

Teaching Assistant, Carnegie Mellon University

- | | |
|---|-------------|
| ○ Introduction to Machine Learning | Spring 2016 |
| ○ Convex Optimization | Fall 2017 |
| ○ Undergraduate Computational Complexity Theory | Spring 2018 |

Teaching Assistant, Tsinghua University

- | | |
|---|-----------|
| ○ Introduction to Information Retrieval | Fall 2012 |
|---|-----------|

Teaching Assistant, University of Waterloo

- | | |
|---------------------------------|-----------|
| ○ Designing Functional Programs | Fall 2013 |
|---------------------------------|-----------|

Department and University Service

University of Illinois at Urbana-Champaign

- | | |
|--|---------------|
| ○ Department of Computer Science, PhD Qualification Exam Committee | Fall 2021 - |
| ○ Department of Computer Science, PhD Admission Committee | Fall 2020 - |
| ○ Department of Computer Science, Academic Appeals Committee | Spring 2022 - |
| ○ Department of Computer Science, Faculty Search Sub-Committee | Spring 2023 - |

Carnegie Mellon University

- | | |
|--|-------------|
| ○ Machine Learning Department, PhD Admission Committee | 2017 |
| ○ Machine Learning Department, PhD Speaking Skills Committee | 2018 - 2020 |
| ○ School of Computer Science, coordinator of AI Seminar | 2018 - 2020 |