

Xinhua Zhang

Associate Professor
Department of Computer Science
University of Illinois at Chicago
SEO 1220, 851 S Morgan St, Chicago, IL 60607

Phone : 312-413-2416 (O)
Email: zhangx@uic.edu
Web: <https://www.cs.uic.edu/~zhangx>
US Permanent Resident

Education

- July 2010 **Ph.D. in Computer Science** (with specialization in Machine Learning)
Australian National University (ANU), Canberra, Australia
Thesis: Graphical Models: Modeling, Optimisation, and Hilbert Space Embedding
Advisors: Professor S.V.N. Vishwanathan and Professor Alex Smola
- April 2006 **M.Sc. in Computer Science**
National University of Singapore (NUS), Singapore
Thesis: Hyperparameter Learning for Graph Based Semi-supervised Learning
Advisor: Prof. Wee Sun Lee
- July 2003 **B.E. in Computer Science and Engineering**
Shanghai Jiao Tong University (SJTU), Shanghai, China

Professional Experience

- Aug 2021 – present **Associate Professor**
- Oct 2015 – Aug 2021 **Assistant Professor**
Department of Computer Science, University of Illinois at Chicago, Chicago, IL
- July 2014 – Oct 2015 **Senior Researcher**
- Oct 2012 – June 2014 **Researcher**
Machine Learning Research Group, National ICT Australia (NICTA)
Adjunct Research Fellow
Research School of Computer Science, The Australian National University
- Mar 2010 – Sept 2012 **Post-doctoral Research Fellow**
Alberta Innovates Center for Machine Learning (AICML), University of Alberta
Mentor: Prof. Dale Schuurmans
- Feb 2009 – Mar 2010 **Visiting Scholar**
Department of Statistics, Purdue University, West Lafayette, Indiana, USA
Supervisor: Prof. S. V. N. Vishwanathan
- Sept 2008 – Nov 2008 **Intern**
Microsoft Research Cambridge, UK
Topic: Bayesian online multilabel classification for text categorisation
Mentors: Thore Graepel and Ralf Herbrich
- June 2005 – Mar 2006 **Research Assistant**
Singapore-MIT Alliance
Topic: Semi-supervised machine learning
Mentor: Prof. Wee Sun Lee

Publications

• Scholarly Book Chapters

- [1] Xinhua Zhang. Seven articles: Support vector machines, kernel, regularization, empirical risk minimization, structural risk minimization, covariance matrix, Gaussian distribution. In Claude Sammut and Geoffrey Webb, editors, *Encyclopedia on Machine Learning*. Springer, 2010. [\[Link\]](#).

• Refereed Journal Articles

- [2] Elisa Tardini, Xinhua Zhang, Guadalupe Canahuate, Andrew Wentzel, Abdallah Mohamed, Lianne Van Dijk, Clifton Fuller, and G Elisabeta Marai. Optimal treatment selection in sequential systemic and locoregional therapy of oropharyngeal squamous carcinomas: Deep q-learning with a patient-physician digital twin dyad. *Journal of Medical Internet Research (JMIR)*, 24(4):e29455, 2022. [\[PDF\]](#).
- [3] Yaoliang Yu, Xinhua Zhang, and Dale Schuurmans. Generalized conditional gradient for sparse estimation. *Journal of Machine Learning Research (JMLR)*, 18(144):1–46, Dec, 2017. [\[PDF\]](#).
- [4] Xinhua Zhang, Ankan Saha, and S. V. N. Vishwanathan. Accelerated training of max-margin Markov networks with kernels. *Journal of Theoretical Computer Science (TCS)*, 519:88–102, Jan 2014. [\[PDF\]](#).
- [5] Xinhua Zhang, Ankan Saha, and S. V. N. Vishwanathan. Smoothing multivariate performance measures. *Journal of Machine Learning Research (JMLR)*, 13:3589–3646, Dec 2012. SmoothOPTZhaSahVis12.

• Refereed Articles in Conference Proceedings

- [6] Zishun Yu and Xinhua Zhang. Actor-critic alignment for offline-to-online reinforcement learning. In *Proceedings of International Conference on Machine Learning (ICML)*, 2023. [\[PDF\]](#).
- [7] Siteng Kang, Zhan Shi, and Xinhua Zhang. Poisoning generative replay in continual learning to promote forgetting. In *Proceedings of International Conference on Machine Learning (ICML)*, 2023. [\[PDF\]](#).
- [8] Hongwei Jin, Zishun Yu, and Xinhua Zhang. Certifying robust graph classification under orthogonal gromov-wasserstein threats. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2022. [\[PDF\]](#).
- [9] Yesu Li, Danyal Saeed, Xinhua Zhang, Brian D Ziebart, and Kevin Gimpel. Moment distributionally robust tree structured prediction. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2022. [\[PDF\]](#).

- [10] Hongwei Jin, Zishun Yu, and Xinhua Zhang. Orthogonal gromov wasserstein distance with efficient lower bound. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI)*, 2022. [\[PDF\]](#).
- [11] Yingyi Ma and Xinhua Zhang. Warping layer: Representation learning for label structures in weakly supervised learning. In *Proceedings of International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022. [\[PDF\]](#).
- [12] Yeshe Li, Zhan Shi, Xinhua Zhang, and Brian Ziebart. Distributionally robust structure learning for discrete pairwise markov networks. In *Proceedings of International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022. [\[PDF\]](#).
- [13] Mao Li, Kaiqi Jiang, and Xinhua Zhang. Implicit task-driven probability discrepancy measure for unsupervised domain adaptation. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2021. [\[Code\]](#) [\[PDF\]](#).
- [14] Mohammad Ali Bashiri, Brian Ziebart, and Xinhua Zhang. Distributionally robust imitation learning. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2021. [\[PDF\]](#).
- [15] Zac Cranko*, Zhan Shi*, Xinhua Zhang, Richard Nock, and Simon Kornblith. Generalised Lipschitz regularisation equals distributional robustness. In *Proceedings of International Conference on Machine Learning (ICML)*, 2021. [\[Code\]](#) [\[PDF\]](#).
- [16] Hongwei Jin*, Zhan Shi*, Ashish Peruri, and Xinhua Zhang. Certified robustness of graph convolution networks for graph classification under topological attacks. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2020. [\[Code\]](#) [\[PDF\]](#).
- [17] Mao Li, Yingyi Ma, and Xinhua Zhang. Proximal mapping for deep regularization. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2020. [\[Code\]](#) [\[PDF\]](#).
- [18] Yingyi Ma, Vignesh Ganapathiraman, Yaoliang Yu, and Xinhua Zhang. Convex representation learning for generalized invariance in semi-inner-product space. In *Proceedings of International Conference on Machine Learning (ICML)*, 2020. [\[PDF\]](#).
- [19] Hongwei Jin and Xinhua Zhang. Robust training of graph convolutional networks via latent perturbation. In *Proceedings of European Conference on Machine Learning (ECML)*, 2020. [\[PDF\]](#).
- [20] Parameswaran Raman, Sriram Srinivasan, Shin Matsushima, Xinhua Zhang, Hyokun Yun, and S.V.N Vishwanathan. Scaling multinomial logistic regression via hybrid parallelism. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2019. [\[PDF\]](#).
- [21] Yingyi Ma, Vignesh Ganapathiraman, and Xinhua Zhang. Learning invariant representations with kernel warping. In *Proceedings of International Conference on Artificial Intelligence and Statistics (AIS-*

- TATS), 2019. [\[PDF\]](#).
- [22] Rizal Fathony, Ashkan Rezaei, Mohammad Ali Bashiri, Xinhua Zhang, and Brian Ziebart. Distributionally robust graphical models. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2018. [\[PDF\]](#).
- [23] Vignesh Ganapathiraman, Zhan Shi, Xinhua Zhang, and Yaoliang Yu. Convex two-layer modeling with latent structure. In *Proceedings of International Conference on Machine Learning (ICML)*, 2018. [\[PDF\]](#).
- [24] Rizal Fathony*, Sima Behpour*, Xinhua Zhang, and Brian Ziebart. Efficient and consistent adversarial bipartite matching. In *Proceedings of International Conference on Machine Learning (ICML)*, 2018. [\[PDF\]](#).
- [25] Mohammad Bashiri and Xinhua Zhang. Decomposition-invariant conditional gradient for general polytopes with line search. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2017. [\[PDF\]](#).
- [26] Zhan Shi, Xinhua Zhang, and Yaoliang Yu. Bregman divergence for stochastic variance reduction: Saddle-point and adversarial prediction. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2017. [\[PDF\]](#).
- [27] Shin Matsushima, Hyokun Yun, Xinhua Zhang, and S.V.N. Vishwanathan. Distributed stochastic optimization of the regularized risk via saddle-point problem. In *Proceedings of European Conference on Machine Learning (ECML)*, 2017. [\[PDF\]](#).
- [28] Vignesh Ganapathiraman, Xinhua Zhang, Yaoliang Yu, and Junfeng Wen. Convex two-layer modeling with latent structure. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2016. [\[PDF\]](#).
- [29] Hao Cheng, Yaoliang Yu, Xinhua Zhang, Dale Schuurmans, and Eric Xing. Scalable and sound low-rank tensor learning. In *Proceedings of International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016. [\[PDF\]](#).
- [30] Parameswaran Kamalaruban, Robert C Williamson, and Xinhua Zhang. Exp-concavity of proper composite losses. In *Conference on Computational Learning Theory (COLT)*, 2015. [\[PDF\]](#).
- [31] Ozlem Aslan, Xinhua Zhang, and Dale Schuurmans. Convex deep learning via normalized kernels. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2014. [\[PDF\]](#).
- [32] Changyou Chen, Jun Zhu, and Xinhua Zhang. Robust Bayesian max-margin clustering. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2014. [\[PDF\]](#).
- [33] Xianghang Liu, Xinhua Zhang, and Tiberio Caetano. Bayesian models for structured sparse estimation via set cover prior. In *European Conference on Machine Learning (ECML)*, 2014. [\[PDF\]](#).

- [34] Hengshuai Yao, Csaba Szepesvari, Bernardo Avila Pires, and Xinhua Zhang. Pseudo-MDPs and factored linear action models. In *Symposium on Adaptive Dynamic Programming and Reinforcement Learning (IEEE ADPRL)*, 2014. [\[PDF\]](#).
- [35] Xinhua Zhang, Wee Sun Lee, and Yee Whye Teh. Learning with invariance via linear functionals on reproducing kernel Hilbert space. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2013. [\[PDF\]](#).
- [36] Xinhua Zhang, Yaoliang Yu, and Dale Schuurmans. Polar operators for structured sparse estimation. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2013. [\[Code\]](#) [\[PDF\]](#).
- [37] Ozlem Aslan, Hao Cheng, Xinhua Zhang, and Dale Schuurmans. Convex two-layer modeling. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2013. [\[PDF\]](#).
- [38] Hao Cheng, Xinhua Zhang, and Dale Schuurmans. Convex relaxations of Bregman divergence clustering. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI)*, 2013. [\[PDF\]](#).
- [39] Martha White, Yaoliang Yu, Xinhua Zhang, and Dale Schuurmans. Convex multi-view subspace learning. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2012. [\[Code\]](#) [\[PDF\]](#).
- [40] Xinhua Zhang, Yaoliang Yu, and Dale Schuurmans. Accelerated training for matrix-norm regularization: A boosting approach. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2012. [\[PDF\]](#).
- [41] Yaoliang Yu, James Neufeld, Ryan Kiros, Xinhua Zhang, and Dale Schuurmans. Regularizers versus losses for nonlinear dimensionality reduction. In *Proceedings of International Conference on Machine Learning (ICML)*, 2012. [\[PDF\]](#).
- [42] Yi Shi, Xinhua Zhang, Xiaoping Liao, Guohui Lin, and Dale Schuurmans. Protein-chemical interaction prediction via kernelized sparse learning SVM. In *Pacific Symposium on Biocomputing (PSB)*, 2013. [\[PDF\]](#).
- [43] Yi Shi, Xiaoping Liao, Xinhua Zhang, Guohui Lin, and Dale Schuurmans. Sparse learning based linear coherent bi-clustering. In *Workshop on Algorithms in Bioinformatics (WABI)*, 2012. [\[PDF\]](#).
- [44] Xinhua Zhang, Ankan Saha, and S. V. N. Vishwanathan. Smoothing multivariate performance measures. In *Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI)*, 2011. [\[PDF\]](#).
- [45] Xinhua Zhang, Ankan Saha, and S. V. N. Vishwanathan. Accelerated training of max-margin Markov networks with kernels. In *Proceedings of International Conference on Algorithmic Learning Theory (ALT)*, 2011. [\[PDF\]](#).

- [46] Xinhua Zhang, Yaoliang Yu, Martha White, Ruitong Huang, and Dale Schuurmans. Convex sparse coding, subspace learning, and semi-supervised extensions. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 2011. [\[PDF\]](#).
- [47] Ankan Saha, S.V.N. Vishwanathan, and Xinhua Zhang. New approximation algorithms for minimum enclosing convex shapes. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2011. [\[PDF\]](#).
- [48] Xinhua Zhang, Ankan Saha, and S. V. N. Vishwanathan. Lower bounds on rate of convergence of cutting plane methods. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2010. [\[PDF\]](#).
- [49] Xinhua Zhang, Thore Graepel, and Ralf Herbrich. Bayesian online learning for multi-label and multi-variate performance measures. In *Proceedings of International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010. [\[Code\]](#) [\[PDF\]](#).
- [50] Xinhua Zhang, Le Song, Arthur Gretton, and Alex Smola. Kernel measures of independence for non-iid data. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2008. [\[PDF\]](#).
- [51] Le Song, Xinhua Zhang, Alex Smola, Arthur Gretton, and Bernhard Schölkopf. Tailoring density estimation via reproducing kernel moment matching. In *Proceedings of International Conference on Machine Learning (ICML)*, 2008. [\[PDF\]](#).
- [52] Li Cheng, S. V. N. Vishwanathan, and Xinhua Zhang. Consistent image analogies using semi-supervised learning. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008. [\[PDF\]](#).
- [53] Xinhua Zhang, Douglas Aberdeen, and S. V. N. Vishwanathan. Conditional random fields for multi-agent reinforcement learning. In *Proceedings of International Conference on Machine Learning (ICML)*, 2007. [\[Code\]](#) [\[PDF\]](#).
- [54] Xinhua Zhang and Wee Sun Lee. Hyperparameter learning for graph based semi-supervised learning algorithms. In *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2006. [\[Code\]](#) [\[PDF\]](#).
- [55] Xinhua Zhang and Peter K K Loh. A fault-tolerant routing strategy for Fibonacci-class cubes. In *Asia-Pacific Computer Systems Architecture Conference (ACSAC)*, 2005. [\[PDF\]](#).

• Theses and Refereed Workshop Papers

- [56] Xinhua Zhang. Graphical models: Modeling, optimization, and Hilbert space embedding. PhD Thesis, Department of Computer Science, Australian National University, 2010. [\[PDF\]](#).
- [57] Xinhua Zhang. Hyperparameter learning for graph based semi-supervised learning algorithms. Master's Thesis, Department of Computer Science, National University of Singapore, 2006. [\[PDF\]](#).

- [58] Mao Li, Yingyi Ma, and Xinhua Zhang. Meta-learning of structured representation by proximal mapping. In *Workshop on Meta-Learning at NeurIPS*, 2019. [\[PDF\]](#)
(4 pages, *refereed* workshop paper selected for 4-min spotlight talk, 10 out of 84, 11.9% acceptance rate).
- [59] Zac Cranko, Zhan Shi, Xinhua Zhang, Richard Nock, and Simon Kornblith. Certifying distributional robustness using Lipschitz regularisation. In *Workshop on Safety and Robustness in Decision Making at NeurIPS*, 2019. [\[PDF\]](#)
(9 pages, *refereed* workshop paper, *International*).
- [60] Yaoliang Yu, Hao Cheng, and Xinhua Zhang. Approximate low-rank tensor learning. In *Optimization Workshop (OPT) at Neural Information Processing Systems (NIPS)*, 2014. [\[PDF\]](#)
(4 pages, *refereed* workshop paper, *International*).
- [61] Jiazhong Nie, Manfred Warmuth, S.V.N. Vishwanathan, and Xinhua Zhang. Lower bounds for boosting with Hadamard matrices. In *Conference on Computational Learning Theory (COLT)*, 2013. [\[PDF\]](#)
(4 pages, *refereed* open problem at a leading conference in theoretical machine learning).
- [62] Xinhua Zhang, Douglas Aberdeen, and S. V. N. Vishwanathan. Conditional random fields for multi-agent reinforcement learning. In *Learning Workshop (Snowbird)*, 2007. [\[PDF\]](#)
(4 pages, *refereed* workshop oral presentation paper, *International*, 22 out of 176, 12.5% acceptance rate).

Teaching

• Lecturing at the University of Illinois at Chicago, USA

Fall 2023	Introduction to Machine Learning (CS 412): undergrad and graduate (58 students)
Spring 2023	Deep Representation Learning (CS 594): graduate (12 students)
Spring 2023	Introduction to Machine Learning (CS 412): undergrad and graduate (49 students)
Fall 2022	Introduction to Machine Learning (CS 412): undergrad and graduate (48 students)
Spring 2022	Advanced Machine Learning (CS 512): graduate (49 students)
Fall 2021	Introduction to Machine Learning (CS 412): undergrad and graduate (80 students)
Spring 2021	Introduction to Machine Learning (CS 412): undergrad and graduate (60 students)
Fall 2020	Introduction to Machine Learning (CS 412): undergrad and graduate (128 students)
Spring 2020	Advanced Machine Learning (CS 512): graduate (40 students)

Fall 2019	Introduction to Machine Learning (CS 412): undergrad and graduate (103 students)
Spring 2019	Advanced Machine Learning (CS 512): graduate (30 students)
Fall 2018	Introduction to Machine Learning (CS 412): undergrad and graduate (73 students)
Spring 2018	Advanced Machine Learning (CS 594): graduate (32 students)
Fall 2017	Introduction to Machine Learning (CS 412): undergrad and graduate (60 students)
Spring 2017	Artificial Intelligence I (CS 411): undergrad and graduate (50 students)
Fall 2016	Advanced Topics in Machine Learning (CS 594): graduate (26 students)
Spring 2016	Artificial Intelligence I (CS 411): undergrad and graduate (37 students)

- **Lecturing at the Australian National University (ANU), Australia**

July – Oct 2015	Information Theory (COMP2610/6261) Joint undergraduate and postgraduate Co-lectured evenly with Aditya Menon and Mark Reid
July – Nov 2014	Advanced Topics in Statistical Machine Learning (COMP4680/8650) Joint undergraduate and postgraduate. Topic: Graphical Models Co-lectured evenly with Stephen Gould and Justin Domke
July – Nov 2013	Advanced Topics in Statistical Machine Learning (COMP4680/8650) Joint undergraduate and postgraduate. Topic: Convex Optimization Co-lectured evenly with Stephen Gould and Justin Domke

- **Guest Lecturing at Purdue University, USA**

Nov 2009	Statistical Learning Theory (STAT 598Y) Two guest lectures , joint undergraduate and postgraduate Topic: Large scale optimisation
----------	------------------------------------------------------------------------------------------------------------------------------------------------

Supervision of Research Students

- **Ongoing**

Aug 2023 – present **Sarang Gawane**

Ph.D. Topic: Meta-Reinforcement Learning

Aug 2021 – present **Zishun Yu**
Ph.D. Topic: Offline Reinforcement Learning

Aug 2020 – present **Siteng Kang**
Ph.D. Topic: Robust Lifelong Learning against Adversarial Forgetting

Aug 2020 – present **Kaiqi Jiang**
Ph.D. Topic: Fair Domain Adaptation

Aug 2020 – present **Wenzhe Fan**
Ph.D. Topic: Kernel Warping for Adversarial Imitation Learning

• Graduated

Oct 2016 – Apr 2022 **Hongwei Jin**
Ph.D. Topic: Robust Learning on Graphs
Current: Postdoc researcher in Mathematics and Computer Science (MCS) division at Argonne National Laboratory

Aug 2016 – Nov 2021 **Yingyi Ma (female)**
Ph.D. Dissertation: Learning Invariant Representation through Warping Layers
Current: Machine Learning Scientist at Facebook

Aug 2016 – Nov 2021 **Mao Li**
Ph.D. Dissertation: Learning Structural Priors with Optimization-based Modeling
Current: Applied Scientist II at Amazon

Aug 2015 – July 2021 **Zhan Shi**
Ph.D. Dissertation: Distributionally Adversarial Learning
Current: Senior Software Engineer in Autonomous Driving at Baidu USA

Aug 2015 – Oct 2020 **Vignesh Ganapathiraman**
Ph.D. Dissertation: Convex Latent Representation Learning with Generalized Invariance
Current: Research Scientist at Amazon

Jan 2020 – Mar 2021 **Elisa Tardini (female)**
MS Thesis: Recursive Meta-Reinforcement Learning for Personalized Sequential Dynamic Treatment Policies
Current: Applied Scientist I at Amazon Search

July 2018 – Mar 2019 **Amlaan Bhoi**
MS Thesis: Invariant Kernels for Few-shot Learning
Current: Applied Scientist at Amazon (WW Returns, ReCommerce & Sustainability)

- July 2018 – May 2019 **Ganesh Jagadeesan**
MS Thesis: Faster Inductive Training of Convex Two-Layer Models
Current: Software Engineer II in AI at Dataminr
- July 2018 – May 2019 **Nimeesha Zhe Wey Chan (female) and Bhavana Iaxmi Radharapu (female)**
Undergrad Capstone Thesis: Representation Learning on Heterogeneous Graphs
Current: Chan: Ph.D. student at Johns Hopkins University
Radharapu: Software Development Engineer at MasterCard
- Jan 2013 – May 2014 **Xianghang Liu** (PhD at Univ. of New South Wales, de facto co-supervision)
Topic: Bayesian sparse learning (one chapter of PhD thesis)
Current: Machine Learning Engineer at Facebook
- Mar 2011 – Sept 2012 **Hao Cheng** (Master at University of Alberta, de facto co-supervision)
Topic: Convex Bregman divergence clustering (whole Master thesis)
Current: PhD student at University of Washington

• **Supervised in Australian National University until Oct 2015 (leaving for US)**

- May 2014 – present **Kamal Parameswaran** (PhD at Australian National University, co-supervision)
Topic: Sublinearity for machine learning
- Oct 2012 – present **Suvash Sedhain** (PhD at Australian National University, co-supervision)
Topic: Social recommendation systems

• **Summer Scholar**

- May – Aug 2014 **Thanard Kurutach** (Summer scholar from MIT, USA)
Topic: Structured Modeling for Solar Prediction

Invited Talks

- June 2021 Annual Meeting of Canadian Applied and Industrial Mathematics Society
Certified Robustness of Graph Convolution Networks for Graph Classification under Topological Attacks
- Oct 2019 Northrop Grumman University Research Symposium (Anaheim, CA)
AI-assisted Dynamic Adaptive Planning for Human-in-the-Loop Multi-Agent Systems (with S. Mou and J. Biswas)
- July 2018 International Symposium on Mathematical Programming (Bordeaux, France)
Talk 1: Generalized Conditional Gradient and Convex Neural Networks
Talk 2: Bregman Divergence for Stochastic Variance Reduction: Saddle-Point and Adversarial Prediction

- June 2018 CSIRO Australian National University site (Canberra, Australia)
Learning Invariant Representations with Kernel Warping for Convolutional Kernel Networks
- Mar 2015 Microsoft Research Redmond (Redmond, WA)
Convex Deep Learning and Efficient Optimization
- Sept 2014 Chinese University of Science and Technology (Hefei, China)
Sublinearity for Machine Learning
- August 2013 Tsinghua University (Beijing, China)
Convex Subspace Learning and Efficient Optimization
- July 2013 Environmental Analytics Showcase (Brisbane, Australia)
Distributed Solar Prediction
- Sept 2012 University of Alberta (Edmonton, Canada)
Accelerated Training for Matrix-norm Regularization: A Boosting Approach
- April 2012 Microsoft Research Cambridge (Cambridge, UK)
Learning and Representing: a Jointly Optimal Approach
- March 2012 Virginia Tech (Blacksburg, USA)
Learning and Representing: a Jointly Optimal Approach
- Feb 2012 Aalto University and University of Helsinki (Helsinki, Finland)
Accelerated optimization for machine learning: A smoothing approach

Professional Services

Editorial board reviewer

2020 - present Journal of Machine Learning Research (JMLR)

Area Chair / Senior Program Committee Member

2017–2023, 2015 Neural Information Processing Systems (NeurIPS)
 2020–2024 International Conference on Machine Learning (ICML)
 2021–2024 International Conference on Representation Learning (ICLR)
 2020, 2019, 2016 AAAI Conference on Artificial Intelligence (AAAI)
 2018, 2015, 2011 International Joint Conference on Artificial Intelligence (IJCAI)

Program Committee Member or Reviewer

2019 International Conference on Representation Learning (ICLR)
 2011–2018 International Conference on Machine Learning (ICML)
 2016, 2011–2014 Neural Information Processing Systems (NeurIPS)
 2019 Uncertainty in Artificial Intelligence (UAI)
 2014–2019 Artificial Intelligence and Statistics (AISTATS)
 2017 International Symposium on Computational Geometry (SoCG)
 2016 International Joint Conference on Artificial Intelligence (IJCAI)
 2015, 2014 Conference on Learning Theory (COLT)

- 2015 AAAI Conference on Artificial Intelligence (AAAI)
- 2014 Knowledge Discovery in Databases (KDD)
- 2014, 2013 Computer Vision and Pattern Recognition (CVPR)
- 2013 International Conference on Computer Vision (ICCV)
- 2012 European Conference on Computer Vision (ECCV)

Reviewer for journals in various years

Journal of Machine Learning Research (JMLR), IEEE Transactions on Pattern Analysis and Machine Intelligence (IEEE-TPAMI), Machine Learning Journal (MLJ), Neurocomputing, Journal of Artificial Intelligence Research (JAIR), Neural Computation (NECO), Theoretical Computer Science (TCS), Artificial Intelligence Journal (AIJ), IEEE Transactions on Signal Processing (IEEE-SP), IEEE Transactions on Neural Networks and Learning Systems (IEEE-TNNLS), Pattern Recognition (PR), Transactions on Knowledge and Data Engineering (TKDE), Statistics and Computing (STCO).

Awards and Honors for Myself and My Advised Students

- 2023 UIC College of Engineering Teaching Award
- 2021 UIC College of Engineering (CoE) Faculty Research Awards
- 2021 Sole recipient of UIC Outstanding MS Thesis Award for my **student** Elisa Tardini
- 2021 UIC CoE Graduate Research Awards for my Ph.D. **student** Yingyi Ma
- 2019 Sole recipient of UIC Outstanding MS Thesis Award for my **student** Amlaan Bhoi
- 2019 UIC Provost's Graduate Research Award for my PhD **student** Yingyi Ma
- 2015 Outstanding Program Committee member for AAAI 2015 (amongst seven awardees)
- 2007 Best student paper award at International Conference on Machine Learning (ICML)
- 2007 Microsoft Research Asia Fellowship

Key Member of Applied Projects

Along with Justin Domke, I have been in charge of building machine learning models for the Distributed Solar Prediction project at National ICT Australia (NICTA), from January 2013 to June 2015. My major responsibility was to accurately predict the power output of rooftop PV panels in a 5-30 minute horizon, via fusing heterogeneous data collected from a distributed network of data loggers and low-resolution all-sky cameras.

Public report: <https://arena.gov.au/wp-content/uploads/2017/04/Final-Project-and-Public-Dissemination-Report-1-USO012.pdf>