

Yaoliang Yu

Curriculum Vitae

January 2025

📍: David R. Cheriton School of Computer Science,
University of Waterloo, Ontario, Canada.

✉: yaoliang.yu@uwaterloo.ca

🌐: <https://cs.uwaterloo.ca/~y328yu/>

Research

My research focuses on developing efficient, scalable, and robust algorithms for modern machine learning models and applications, with formal theoretical guarantees and analyses. I am also interested in applying machine learning techniques to vision and natural language applications.

Education

Nov 2013 PhD in Computing Science (with specialization in Statistical Machine Learning)
University of Alberta
Thesis: *Fast Gradient Algorithms for Structured Sparsity*

Positions & Awards

Mar 2024		Ontario Early Researcher Awards
Jul 2021	– Present	Associate Professor Cheriton School of Computer Science, University of Waterloo
Dec 2019	– Present	Canada CIFAR AI Chair at the Vector Institute
Sep 2019	– Present	Faculty Member, Vector Institute
Jun 2020	– May 2023	Cheriton Faculty Fellow
Sep 2016	– Jun 2021	Assistant Professor Cheriton School of Computer Science, University of Waterloo
Feb 2014	– Aug 2016	Post-doctoral Fellow Machine Learning Department, Carnegie Mellon University

Publications

All published papers can be viewed by clicking the title.













■ Optimization □ Generative Models ■ Robustness □ Reproducing Kernel
□ System ■ Application ■ Miscellaneous

Book Chapter










- _{system} [B1] S. Malekmohammadi*, K. Shaloudegi, Z. Hu*, and Y. Yu. “A Unifying Framework for Federated Learning”. In: *Federated and Transfer Learning*. Springer, 2023.
- _{opt} [B2] Y. Yu. “Online Learning and Optimization”. In: *Encyclopedia of Algorithms*. Ed. by M.-Y. Kao. Springer, 2015.






















Journal Articles

- _{app} [J1] Y. Lu*, G. Zhang*, S. Sun, H. Guo, and Y. Yu. “ f -MICL: Understanding and Generalizing InfoNCE-based Contrastive Learning”. *Transactions on Machine Learning Research* (2023). Short version also appeared in NeurIPS 2021 Workshop on Self-Supervised Learning.
- _{app} [J2] W. Li*, L. Kari, Y. Yu, and L. Hug. “MT-MAG: Accurate and interpretable machine learning for complete or partial taxonomic assignments of metagenome-assembled genomes”. *PLOS One* (8) (2023), e0283536.

-  [J3] G. Zhang, S. Malekmohammadi*, X. Chen, and Y. Yu. “Proportional Fairness in Federated Learning”. *Transactions on Machine Learning Research* (2023).
-  [J4] Y. Lu*, G. Kamath, and Y. Yu. “Indiscriminate Data Poisoning Attacks on Neural Networks”. *Transactions on Machine Learning Research* (2022). Short version also appeared in NeurIPS 2022 workshop on Trustworthy and Socially Responsible Machine Learning.
-  [J5] Z. Hu*, K. Shaloudegi, G. Zhang*, and Y. Yu. “FedMGDA+: Federated Learning meets Multi-objective Optimization”. *IEEE Transactions on Network Science and Engineering* 9(4) (2022), 2039–2051.
-  [J6] T. Fujiwara, J. Zhao, F. Chen, Y. Yu, and K.-L. Ma. “Network Comparison with Interpretable Contrastive Network Representation Learning”. *Journal of Data Science, Statistics, and Visualisation* 2(5) (2022), 1–35.
-  [J7] G. Zhang*, P. Poupart, and Y. Yu. “Optimality and Stability in Non-Convex Smooth Games”. *Journal of Machine Learning Research* 23(35) (2022), 1–71.
-  [J8] M. Marchetti-Bowick, Y. Yu, W. Wu, and E. Xing. “A Penalized Regression Model for the Joint Estimation of eQTL Associations and Gene Network Structure”. *Annals of Applied Statistics* 13(1) (2019), 248–270.
-  [J9] Y. Zhou, Y. Liang, Y. Yu, W. Dai, and E. Xing. “Distributed Proximal Gradient Algorithm for Partially Asynchronous Computer Clusters”. *Journal of Machine Learning Research* 19 (2018), 733–764. (Short version also appeared in AISTATS 2016).
-  [J10] Y. Yu, X. Zhang, and D. Schuurmans. “Generalized Conditional Gradient for Sparse Estimation”. *Journal of Machine Learning Research* 18 (2017), 1–46.
-  [J11] S. Xu, Y. Zhou, K. Yuan, Y. Yu, X. Ni, P. Xie, and E. Xing. “Inference of Multiple-wave Population Admixture by Modeling Decay of Linkage Disequilibrium With Polynomial Functions”. *Heredity* 118 (2017), 503–510.
-  [J12] X. Chang, Y. Yu, Y. Yang, and E. Xing. “Semantic Pooling for Complex Event Analysis in Untrimmed Videos”. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39(8) (2017), 1617–1632. (Short version also appeared in ICML 2015).
-  [J13] E. Xing, Q. Ho, W. Dai, J. Kim, J. Wei, S. Lee, X. Zheng, P. Xie, A. Kumar, and Y. Yu. “Petuum: A New Platform for Distributed Machine Learning on Big Data”. *IEEE Transactions on Big Data* 1(2) (2015), 49–67. (Short version also appeared in KDD 2015).
-  [J14] Y. Yu, J. Jiang, and L. Zhang. “Distance Metric Learning by Minimal Distance Maximization”. *Pattern Recognition* 44 (2011), 639–649.

Refereed Conference Proceedings

-  [C1] J. Dong, B. Wang, and Y. Yu. “Last-iterate Convergence in Regularized Graphon Mean Field Game”. In: *Association for the Advancement of Artificial Intelligence (AAAI)*. 2025.
-  [C2] W. Li* and Y. Yu. “One Sample Fits All: Approximating All Probabilistic Values Simultaneously and Efficiently”. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2024.
-  [C3] Y. Lu*, M. Yang, Z. Liu, G. Kamath, and Y. Yu. “Disguised Copyright Infringement of Latent Diffusion Models”. In: *International Conference on Machine Learning (ICML)*. 2024.
-  [C4] S. Malekmohammadi*, Y. Yu, and Y. Cao. “Noise-Aware Aggregation for Heterogeneous Differentially Private Federated Learning”. In: *International Conference on Machine Learning (ICML)*. 2024.
-  [C5] Y. Lu*, M. Yang, G. Kamath, and Y. Yu. “Indiscriminate Data Poisoning Attacks on Pre-trained Feature Extractors”. In: *2nd IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*. 2024.
-  [C6] J. Dong, B. Wang, and Y. Yu. “Convergence to Nash Equilibrium and No-regret Guarantee in (Markov) Potential Games”. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2024.
-  [C7] W. Li* and Y. Yu. “Faster Approximation of Probabilistic and Distributional Values via Least Squares”. In: *International Conference on Learning Representations (ICLR)*. 2024.
-  [C8] W. Li* and Y. Yu. “Robust Data Valuation with Weighted Banzhaf Values”. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2023.
-  [C9] D. Jiang*, S. Sun, and Y. Yu. “Functional Rényi Differential Privacy for Generative Modeling”. In: *Advances in Neural Information Processing Systems (NeurIPS)*. Short version also appeared in ICML 2023 workshop on Challenges of Deploying Generative AI. 2023.






-  [C10] Y. Lu*, Y. Yu, X. Li, and V. P. NIA. "Understanding Neural Network Binarization with Forward and Backward Proximal Quantizers". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2023.
-  [C11] A. Ghose, A. Gupta, Y. Yu, and P. Poupart. "Batchnorm Allows Unsupervised Radial Attacks". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2023.
-  [C12] J. Xin*, R. Tang, Z. Jiang, Y. Yu, and J. Lin. "Operator Selection and Ordering in a Pipeline Approach to Efficiency Optimizations for Transformers". In: *Findings of the Association for Computational Linguistics (ACL)*. 2023.
-  [C13] Y. Lu*, G. Kamath, and Y. Yu. "Exploring the Limits of Model-Targeted Indiscriminate Data Poisoning Attacks". In: *International Conference on Machine Learning (ICML)*. 2023.
-  [C14] H. Lu*, D. Herman*, and Y. Yu. "Multi-Objective Reinforcement Learning: Convexity, Stationarity and Pareto Optimality". In: *International Conference on Learning Representations (ICLR)*. 2023.
-  [C15] D. Jiang*, S. Sun, and Y. Yu. "Revisiting flow generative models for Out-of-distribution detection". In: *International Conference on Learning Representations (ICLR)*. 2022.
-  [C16] S. Qian, H. Pham*, T. Lutellier, Z. Hu*, J. Kim, T. Lin, Y. Yu, J. Chen, and S. Shah. "Are My Deep Learning Systems Fair? An Empirical Study of Fixed-Seed Training". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2021.
-  [C17] G. Zhang*, H. Zhao, Y. Yu, and P. Poupart. "Quantifying and Improving Transferability in Domain Generalization". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2021.
-  [C18] X. Li, B. Liu, Y. Yu, W. Liu, C. Xu, and V. NIA. "S³: Sign-Sparse-Shift Reparametrization for Effective Training of Low-bit Shift Networks". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2021.
-  [C19] T. Dockhorn*, Y. Yu, E. Sari, M. Zolnouri, and V. NIA. "Demystifying and Generalizing BinaryConnect". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2021.
-  [C20] J. Xin*, R. Tang*, Y. Yu, and J. Lin. "The Art of Abstention: Selective Prediction and Error Regularization for Natural Language Processing". In: *The Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*. 2021.
-  [C21] H. Cheng, X. Liu, L. Pereira, Y. Yu, and J. Gao. "Posterior Differential Regularization with f -divergence for Improving Model Robustness". In: *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*. 2021.
-  [C22] J. Xin*, R. Tang*, Y. Yu, and J. Lin. "BERxiT: Better-fine-tuned and Wider-applicable Early Exit for *BERT". In: *The 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. 2021.
-  [C23] H. Pham*, S. Qian, J. Wang, T. Lutellier, J. Rosenthal, L. Tan, Y. Yu, and N. Nagappan. "Problems and Opportunities in Training Deep-Learning Software Systems: An Analysis of Variance". In: *35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. ACM Distinguished Papers. 2020.
-  [C24] K. Wu*, H. Wang*, and Y. Yu. "Stronger and Faster Wasserstein Adversarial Attacks". In: *International Conference on Machine Learning (ICML)*. 2020.
-  [C25] P. Jaini*, I. Kobyzev, Y. Yu, and M. Brubaker. "Tails of Lipschitz Triangular Flows". In: *International Conference on Machine Learning (ICML)*. 2020.
-  [C26] Y. Ma, V. Ganapathiraman, Y. Yu, and X. Zhang. "Convex Representation Learning for Generalized Invariance in Semi-Inner-Product Space". In: *International Conference on Machine Learning (ICML)*. 2020.
-  [C27] X. Lian*, K. Jain, J. Truszkowski, P. Poupart, and Y. Yu. "Unsupervised Multilingual Alignment using Wasserstein Barycenters". In: *International Joint Conference on Artificial Intelligence (IJCAI)*. Also appeared at the third annual WeCNLP (West Coast NLP) Summit. 2020.
-  [C28] J. Xin*, R. Tang*, J. Lee, Y. Yu, and J. Lin. "DeeBERT: Dynamic Early Exiting for Accelerating BERT Inference". In: *Proceedings of the Association for Computational Linguistics (ACL)*. 2020.
-  [C29] R. Tang*, J. Lee, J. Xin*, X. Liu, Y. Yu, and J. Lin. "Showing Your Work Doesn't Always Work". In: *Proceedings of the Association for Computational Linguistics (ACL)*. 2020.
-  [C30] K. Wu*, W. Ding, R. Huang, and Y. Yu. "On Minimax Optimality of GANs for Robust Mean Estima-

- tion". In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2020.
-  [C31] G. Zhang^{*} and Y. Yu. "Convergence of Gradient Methods on Bilinear Zero-Sum Games". In: *International Conference on Learning Representations (ICLR)*. 2020.
-  [C32] J. Wang^{*}, S. Sun, and Y. Yu. "Multivariate Triangular Quantile Maps for Novelty Detection". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2019.
-  [C33] J. Xin^{*}, J. Lin, and Y. Yu. "What Part of the Neural Network Does This? Understanding LSTMs by Measuring and Dissecting Neurons". In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2019.
-  [C34] P. Jaini^{*}, K. Selby, and Y. Yu. "Sum-of-squares Polynomial Flow". In: *International Conference on Machine Learning (ICML)*. 2019.
-  [C35] S. Sun and Y. Yu. "Least-Squares Estimation of Weakly Convex Functions". In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2019.
-  [C36] P. Jaini^{*}, P. Poupart, and Y. Yu. "Deep Homogeneous Mixture Models: Representation, Separation and Approximation". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2018.
-  [C37] V. Ganapathiraman, Z. Shi, X. Zhang, and Y. Yu. "Inductive Two-Layer Modeling with Parametric Bregman Transfer". In: *International Conference on Machine Learning (ICML)*. 2018.
-  [C38] P. Xie, J. Kim, Q. Ho, Y. Yu, and E. Xing. "Orpheus: Efficient Distributed Machine Learning via System and Algorithm Co-design". In: *ACM Symposium on Cloud Computing (SoCC)*. 2018.
-  [C39] Z. Shi, X. Zhang, and Y. Yu. "Bregman Divergence for Stochastic Variance Reduction Methods: Adversarial Prediction and Saddle-Point Problems". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2017.
-  [C40] J. Yin and Y. Yu. "Convex-constrained Sparse Additive Modeling and Its Extensions". In: *Conference on Uncertainty in Artificial Intelligence (UAI)*. 2017.
-  [C41] P. Xie, Y. Deng, Y. Zhou, A. Kumar, Y. Yu, J. Zou, and E. Xing. "Analyzable Diversity-Promoting Latent Space Models". In: *International Conference on Machine Learning (ICML)*. 2017.
-  [C42] X. Chang, Y. Yu, and Y. Yang. "Robust Top-*k* Multiclass SVM for Visual Category Recognition". In: *ACM Conference on Knowledge Discovery and Data Mining (KDD)*. 2017.
-  [C43] M. Law, Y. Yu, R. Urtasun, R. Zemel, and E. Xing. "Efficient Multiple Instance Metric Learning using Weakly Supervised Data". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
-  [C44] X. Ma, Y. Gao, Z. Hu, Y. Yu, Y. Deng, and E. Hovy. "Dropout with Expectation-Linear Regularization". In: *International Conference on Learning Representations (ICLR)*. 2017.
-  [C45] V. Ganapathiraman, X. Zhang, Y. Yu, and J. Wen. "Convex Two-Layer Modeling with Latent Structure". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2016.
-  [C46] P. Xie, J. Kim, Y. Zhou, Q. Ho, A. Kumar, Y. Yu, and E. Xing. "Lighter-Communication Distributed Machine Learning via Sufficient Factor Broadcasting". In: *Conference on Uncertainty in Artificial Intelligence (UAI)*. 2016.
-  [C47] K. Kandasamy and Y. Yu. "Additive Approximations in High Dimensional Nonparametric Regression via the SALSAs". In: *International Conference on Machine Learning (ICML)*. 2016.
-  [C48] X. Chang, Y. Yu, Y. Yang, and E. Xing. "They Are Not Equally Reliable: Semantic Event Search using Differentiated Concept Classifiers". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
-  [C49] M. Law, Y. Yu, M. Cord, and E. Xing. "Closed-Form Training of Mahalanobis Distance for Supervised Clustering". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
-  [C50] H. Cheng, Y. Yu, X. Zhang, E. Xing, and D. Schuurmans. "Scalable and Sound Low-Rank Tensor Learning". In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2016.
-  [C51] Y. Yu and E. Xing. "Exact Algorithms for Isotonic Regression and Related". In: *Journal of Physics: Conference Series*. 2016.
-  [C52] X. Chang, Y. Yu, Y. Yang, and A. Hauptmann. "Searching Persuasively: Joint Event Detection and Evidence Recounting with Limited Supervision". In: *ACM Conference on Multimedia (MM)*. 2015.
-  [C53] X. Zheng, Y. Yu, and E. Xing. "Linear Time Samplers for Supervised Topic Models using Compositional

Proposals". In: *ACM Conference on Knowledge Discovery and Data Mining (KDD)*. 2015.

-  [C54] X. Chang, Y. Yang, A. Hauptmann, E. Xing, and Y. Yu. "Semantic Concept Discovery for Large-Scale Zero-Shot Event Detection". In: *International Joint Conference on Artificial Intelligence (IJCAI)*. 2015.
-  [C55] Y. Yu, X. Zheng, M. Marchetti-Bowick, and E. Xing. "Minimizing Nonconvex Non-Separable Functions". In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2015.
-  [C56] A. Yu, W. Ma, Y. Yu, J. Carbonell, and S. Sra. "Efficient Structured Matrix Rank Minimization". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2014.
-  [C57] Y. Yu. "Better Approximation and Faster Algorithm Using the Proximal Average". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2013.
-  [C58] Y. Yu. "On Decomposing the Proximal Map". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2013. (Oral presentation, 20/1420).
-  [C59] X. Zhang, Y. Yu, and D. Schuurmans. "Polar Operators for Structured Sparse Estimation". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2013.
-  [C60] Y. Yu, H. Cheng, D. Schuurmans, and C. Szepesvári. "Characterizing the Representer Theorem". In: *International Conference on Machine Learning (ICML)*. 2013.
-  [C61] X. Zhang, Y. Yu, and D. Schuurmans. "Accelerated Training for Matrix-Norm Regularization: A Boosting Approach". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2012.
-  [C62] M. White, Y. Yu, X. Zhang, and D. Schuurmans. "Convex Multi-view Subspace Learning". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2012.
-  [C63] Y. Yu, Ö. Aslan, and D. Schuurmans. "A Polynomial-time Form of Robust Regression". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2012.
-  [C64] Y. Yu, J. Neufeld, R. Kiros, X. Zhang, and D. Schuurmans. "Regularizers versus Losses for Nonlinear Dimensionality Reduction". In: *International Conference on Machine Learning (ICML)*. 2012.
-  [C65] Y. Yu and C. Szepesvári. "Analysis of Kernel Mean Matching under Covariate Shift". In: *International Conference on Machine Learning (ICML)*. 2012.
-  [C66] Y. Yu and D. Schuurmans. "Rank/Norm Regularization with Closed-Form Solutions: Application to Subspace Clustering". In: *Conference on Uncertainty in Artificial Intelligence (UAI)*. 2011.
-  [C67] X. Zhang, Y. Yu, M. White, R. Huang, and D. Schuurmans. "Convex Sparse Coding, Subspace Learning, and Semi-Supervised Extensions". In: *Association for the Advancement of Artificial Intelligence (AAAI)*. 2011.
-  [C68] Y. Yu, M. Yang, L. Xu, M. White, and D. Schuurmans. "Relaxed Clipping: A Global Training Method for Robust Regression and Classification". In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2010.
-  [C69] Y. Yu, Y. Li, D. Schuurmans, and C. Szepesvári. "A General Projection Property for Distribution Families". In: *Advances in Neural Information Processing Systems (NeurIPS)*. Short version also appeared in *Multidisciplinary Symposium on Reinforcement Learning 2009*. 2009.
-  [C70] P. Guan, Y. Yu, and L. Zhang. "A Novel Facial Feature Point Localization Method on 3D Faces". In: *IEEE Conference on Image Processing (ICIP)*. 2007.

Workshop Papers

-  [W1] H. Lu^{*}, Y. Lu^{*}, D. Jiang^{*}, S. Szabados^{*}, S. Sun, and Y. Yu. "CM-GAN: Stabilizing GAN Training with Consistency Models". In: *ICML Workshop on Structured Probabilistic Inference & Generative Modeling*. 2023.
-  [W2] T. Dockhorn^{*}, R. Rombach, A. Blatmann, and Y. Yu. "Distilling the Knowledge in Diffusion Models". In: *CVPR workshop on Generative Models for Computer Vision*. 2023.
-  [W3] J. Sun^{*}, D. Jiang^{*}, and Y. Yu. "Conditional Generative Quantile Networks via Optimal Transport". In: *ICLR Workshop on Deep Generative Models for Highly Structured Data*. 2022.
-  [W4] Z. Shen, W. Li, J. Zhao, Y. Yu, and M. Dell. "OLALA: Object-Level Active Learning based Layout Annotation". In: *5th workshop on Natural Language Processing and Computational Social Science at EMNLP*. 2022.
-  [W5] G. Zhang^{*}, K. Wu^{*}, P. Poupard, and Y. Yu. "Newton-type Methods for Minimax Optimization". In: *ICML Workshop on Beyond First-Order Methods in ML Systems*. 2021.

- [W6] H. Pham^{*}, M. Kim, T. Lin, Y. Yu, and N. Nagappan. “DEVIATE: A Deep Learning Variance Testing Framework”. In: *ASE Tool Demonstrations*. 2021.
- [W7] J. Xin^{*}, R. Nogueira, Y. Yu, and J. Lin. “Early Exiting BERT for Efficient Document Ranking”. In: *Proceedings of the First Workshop on Simple and Efficient Natural Language Processing (SustaiNLP 2020)*. 2020.
- [W8] T. Dockhorn^{*}, J. Ritchie, Y. Yu, and I. Murray. “Density Deconvolution with Normalizing Flows”. In: *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*. 2020.
- [W9] K. Wu^{*} and Y. Yu. “Understanding Adversarial Robustness: The Trade-off between Minimum and Average Margin”. In: *NeurIPS Workshop on Matching Learning with Guarantees*. 2019.
- [W10] Y. Yu, Y. Zhang, and C. Szepesvári. “Online TD(1) Meets Offline Monte Carlo”. In: *Multidisciplinary Symposium on Reinforcement Learning*. 2009.

Funding

My research is supported by

- Ontario Early Researcher Awards
- Natural Sciences and Engineering Research Council of Canada (NSERC)
- Canada CIFAR AI Chairs Program

Talks (Selected)

Nov-2024	Vector Institute Workshop on Machine Learning Theory <i>Is Explainability A Game?</i>	Waterloo
Jul-2024	International Symposia on Mathematical Programming (ISMP) <i>Understanding Neural Network Binarization with Forward and Backward Proximal Quantizers</i>	Montreal
Nov-2023	Vector Institute Workshop on Machine Learning Theory <i>Data, Model and Values</i>	Toronto
Aug-2023	Joint Statistical Meetings <i>Indiscriminate Data Poisoning Attacks on Neural Networks</i>	Toronto
Jul-2023	Vector Institute Workshop on Machine Learning Security and Privacy <i>Data Poisoning: Hype or Real?</i>	Toronto
Oct-2022	24th Midwest Optimization Meeting <i>Demystifying and generalizing binaryConnect</i>	Waterloo
Jun-2022	Statistical Society of Canada Annual Meeting <i>Demystifying and generalizing binaryConnect</i>	online
Apr-2022	SIAM Conference on Uncertainty Quantification <i>A triangular approach to generative modeling</i>	online
Jun-2021	Canada Mathematical Society Summer Meeting <i>Splitting algorithms for federated learning</i>	online
Mar-2020	2nd Edge Intelligence workshop <i>A Triangular Approach to Probabilistic Modeling</i>	Montreal
Dec-2019	Canada Mathematical Society winter meeting <i>Least-squares Estimation of Weekly-convex Functions</i>	Toronto
Sep-2018	Fields Institute <i>Generalized Conditional Gradient for Sparse Estimation</i>	Toronto

Dec-2017	Canada Mathematical Society winter meeting <i>Minimizing the Sum of Non-separable Functions</i>	Waterloo
Sep-2017	BIRS workshop <i>Splitting Algorithms, Modern Operator Theory, and Applications</i> <i>On Decomposing the Proximal Map</i>	Oxaca

Supervision of Research Students

Jan-2023	Spencer Szabados [✉]	MMath
Sep-2021	Haoye Lu [✉] GO-Bell Scholarship, Cheriton Scholarship, NSERC Post-Graduate Scholarship committee: Wenhui Chen, Pascal Poupart	PhD
Jan-2021	Saber Malekmohammadi [✉] committee: Gautam Kamath, Hong Zhang	PhD
Sep-2020	Yiwei Lu [✉] 2024 Cheriton Scholarship committee: Gautam Kamath, Hongyang Zhang	co-supervised with Sun Sun, PhD
Sep-2018	Zeou Hu [✉] committee: Kimon Fountoulakis, Justin Wang	PhD

Teaching

Number	Course Title	Times	Type	Material
CS341	Algorithms	1	undergrad	
CS480/680	Introduction to Machine Learning	7	mixed	web & notes
CS475	Computational Linear Algebra	1	undergrad	web & notes
CO673/CS794	Optimization for Data Science	4	grad	web & notes
CS886	Theory of Deep Learning	1	grad	web & notes
CS886	Causal Inference in Machine Learning	1	grad	
CS886	Diffusion models	1	grad	notes

Professional Service

I serve regularly as an area chair for International Conference on Machine Learning (ICML), International Conference on Learning Representation (ICLR), Neural Information Processing Systems (NeurIPS). Depending on my availability, I sometimes review for AACL, AISTATS, ALT, COLT, CVPR, IJCAI and UAI.

Reviewer for journals

I am an action editor for Transaction on Machine Learning Research and an associated editor for ACM Transactions on Probabilistic Machine Learning. I regularly review for Journal of Machine Learning Research. Other journals that I occasionally review for include IEEE Transactions (TPAMI, TAC, TKDE), Mathematical Programming, Computational Optimization and Application, Machine Learning Journal, Artificial Intelligence Journal, etc.