

Yaoliang Yu

Curriculum Vitae

May 2020

📍: David R. Cheriton School of Computer Science,
University of Waterloo, Ontario, Canada.
☎: +1 519-888-4567 x34469
✉: 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 financial applications.

Education

- 2013 Ph.D. in Computing Science (with specialization in Statistical Machine Learning)
University of Alberta
Thesis: Fast Gradient Algorithms for Structured Sparsity
PhD Dissertation Award from the Canadian Artificial Intelligence Association, 2015
PhD Outstanding Thesis Award Winner, University of Alberta, 2014
- 2008 M.Sc. Fudan University
2005 B.Sc. Fudan University

Positions

- Dec 2019 – Present Canada CIFAR AI Chair
Sep 2019 – Present Faculty Member, Vector Institute
Sep 2016 – Present Assistant Professor
David R. Cheriton School of Computer Science, University of Waterloo
- Feb 2014 – Aug 2016 Post-doctoral Fellow
Machine Learning Department, Carnegie Mellon University

Publications

Book Chapter

1. Y.-L. Yu (2015). “Online Learning and Optimization”. In: *Encyclopedia of Algorithms*. Ed. by M.-Y. Kao. Springer.

Journal Articles

1. Marchetti-Bowick, M., Y. Yu, W. Wu, and E. Xing (2019). A penalized regression model for the joint estimation of eQTL associations and gene network structure. *Annals of Applied Statistics* **13**(1), 248–270.
2. Zhou, Y., Y. Liang, Y. Yu, W. Dai, and E. Xing (2018). Distributed Proximal Gradient Algorithm for Partially Asynchronous Computer Clusters. *Journal of Machine Learning Research* **19**, 733–764.
3. Chang, X., Y.-L. Yu, Y. Yang, and E. Xing (2017). Semantic Pooling for Complex Event Analysis in Untrimmed Videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(8), 1617–1632. (Preliminary version appeared in ICML, 2015).
4. Xu, S., Y. Zhou, K. Yuan, Y. Yu, X. Ni, P. Xie, and E. Xing (2017). Inference of Multiple-wave Population Admixture by Modeling Decay of Linkage Disequilibrium With Polynomial Functions. *Heredity* **118**, 503–510.

5. Yu, Y., X. Zhang, and D. Schuurmans (2017). Generalized Conditional Gradient for Sparse Estimation. *Journal of Machine Learning Research* **18**, 1–46.
6. Xing, E., Q. Ho, W. Dai, J. Kim, J. Wei, S. Lee, X. Zheng, P. Xie, A. Kumar, and Y. Yu (2015). Petuum: A New Platform for Distributed Machine Learning on Big Data. *IEEE Transactions on Big Data* **1**(2), 49–67. (Preliminary version appeared in KDD, 2015).
7. Y.-L. Yu, J. Jiang, and L. Zhang (2011). Distance Metric Learning by Minimal Distance Maximization. *Pattern Recognition* **44**, 639–649.

Refereed Conference Proceedings

1. Lian, X., K. Jain, J. Truszkowski, P. Poupart, and Y. Yu (2020). Unsupervised Multilingual Alignment using Wasserstein Barycenters. In: *IJCAI*.
2. Tang, R., J. Lee, J. Xin, X. Liu, Y. Yu, and J. Lin (2020). Showing Your Work Without Bias. In: *ACL short paper*.
3. Wu, K., W. Ding, R. Huang, and Y. Yu (2020). On Minimax Optimality of GANs for Robust Mean Estimation. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
4. Xin, J., R. Tang, J. Lee, Y. Yu, and J. Lin (2020). DyBERT: Dynamic Early Exiting for Accelerating BERT Inference. In: *ACL short paper*.
5. Zhang, G. and Y. Yu (2020). Convergence of Gradient Methods on Bilinear Zero-Sum Games. In: *International Conference on Learning Representations (ICLR)*.
6. Jaini, P., K. Selby, and Y. Yu (2019). Sum-of-squares Polynomial Flow. In: *International Conference on Machine Learning (ICML)*.
7. Sun, S. and Y. Yu (2019). Least-Squares Estimation of Weakly Convex Functions. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
8. Wang, J., S. Sun, and Y. Yu (2019a). Multivariate Triangular Quantile Maps for Novelty Detection. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
9. Xin, J., J. Lin, and Y. Yu (2019). What Part of the Neural Network Does This? Understanding LSTMs by Measuring and Dissecting Neurons. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
10. Ganapathiraman, V., Z. Shi, X. Zhang, and Y. Yu (2018). Inductive Two-Layer Modeling with Parametric Bregman Transfer. In: *International Conference on Machine Learning (ICML)*.
11. Jaini, P., P. Poupart, and Y. Yu (2018). Deep Homogeneous Mixture Models: Representation, Separation and Approximation. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
12. Xie, P., J. K. Kim, Q. Ho, Y. Yu, and E. Xing (2018). Orpheus: Efficient Distributed Machine Learning via System and Algorithm Co-design. In: *ACM Symposium on Cloud Computing (SoCC)*.
13. Chang, X., Y. Yu, and Y. Yang (2017). Robust Top- k Multiclass SVM for Visual Category Recognition. In: *ACM Conference on Knowledge Discovery and Data Mining (KDD)*.
14. Law, M., Y. Yu, R. Urtasun, R. Zemel, and E. Xing (2017). Efficient Multiple Instance Metric Learning using Weakly Supervised Data. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
15. Ma, X., Y. Gao, Z. Hu, Y. Yu, Y. Deng, and E. Hovy (2017). Dropout with Expectation-Linear Regularization. In: *International Conference on Learning Representations (ICLR)*.
16. Shi, Z., X. Zhang, and Y. Yu (2017). Bregman Divergence for Stochastic Variance Reduction Methods: Adversarial Prediction and Saddle-Point Problems. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
17. Xie, P., Y. Deng, Y. Zhou, A. Kumar, Y. Yu, J. Zou, and E. Xing (2017). Analyzable Diversity-Promoting Latent Space Models. In: *International Conference on Machine Learning (ICML)*.
18. Yin, J. and Y. Yu (2017). Convex-constrained Sparse Additive Modeling and Its Extensions. In: *Conference on Uncertainty in Artificial Intelligence (UAI)*.
19. Chang, X., Y.-L. Yu, Y. Yang, and E. Xing (2016). Not Equally Reliable: Semantic Event Search using Differentiated Concept Classifiers. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

20. Cheng, H., Y. Yu, X. Zhang, E. Xing, and D. Schuurmans (2016). Scalable and Sound Low-Rank Tensor Learning. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
21. Ganapathiraman, V., X. Zhang, Y. Yu, and J. Wen (2016). Convex Two-Layer Modeling with Latent Structure. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
22. Kandasamy, K. and Y. Yu (2016). Additive Approximations in High Dimensional Nonparametric Regression via the SALSA. In: *International Conference on Machine Learning (ICML)*.
23. Law, M., Y. Yu, M. Cord, and E. Xing (2016). Closed-Form Training of Mahalanobis Distance for Supervised Clustering. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
24. Xie, P., J. Kim, Y. Zhou, Q. Ho, A. Kumar, Y. Yu, and E. Xing (2016). Lighter-Communication Distributed Machine Learning via Sufficient Factor Broadcasting. In: *Conference on Uncertainty in Artificial Intelligence (UAI)*.
25. Y.-L. Yu and E. Xing (2016). Exact Algorithms for Isotonic Regression and Related. In: *Journal of Physics: Conference Series*.
26. Zhou, Y., Y. Yu, W. Dai, Y. Liang, and E. Xing (2016). On Convergence of Model Parallel Proximal Gradient Algorithm for Stale Synchronous Parallel System. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
27. Chang, X., Y.-L. Yu, Y. Yang, and A. Hauptmann (2015). Searching Persuasively: Joint Event Detection and Evidence Recounting with Limited Supervision. In: *ACM Conference on Multimedia (MM)*.
28. Chang, X., Y. Yang, A. Hauptmann, E. Xing, and Y.-L. Yu (2015). Semantic Concept Discovery for Large-Scale Zero-Shot Event Detection. In: *International Joint Conference on Artificial Intelligence (IJCAI)*.
29. Y.-L. Yu, X. Zheng, M. Marchetti-Bowick, and E. Xing (2015). Minimizing Nonconvex Non-Separable Functions. In: *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
30. Zheng, X., Y.-L. Yu, and E. Xing (2015). Linear Time Samplers for Supervised Topic Models using Compositional Proposals. In: *ACM Conference on Knowledge Discovery and Data Mining (KDD)*.
31. Yu, A., W. Ma, Y. Yu, J. Carbonell, and S. Sra (2014). Efficient Structured Matrix Rank Minimization. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
32. Y.-L. Yu (2013a). On Decomposing the Proximal Map. In: *Advances in Neural Information Processing Systems (NeurIPS)*. (Oral presentation, 20/1420).
33. Y.-L. Yu (2013b). Better Approximation and Faster Algorithm Using the Proximal Average. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
34. Y.-L. Yu, H. Cheng, D. Schuurmans, and C. Szepesvári (2013). Characterizing the Representer Theorem. In: *International Conference on Machine Learning (ICML)*.
35. Zhang, X., Y. Yu, and D. Schuurmans (2013). Polar Operators for Structured Sparse Estimation. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
36. White, M., Y. Yu, X. Zhang, and D. Schuurmans (2012). Convex Multi-view Subspace Learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
37. Y.-L. Yu and C. Szepesvári (2012). Analysis of Kernel Mean Matching under Covariate Shift. In: *International Conference on Machine Learning (ICML)*.
38. Yu, Y., Ö. Aslan, and D. Schuurmans (2012). A Polynomial-time Form of Robust Regression. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
39. Yu, Y., J. Neufeld, R. Kiros, X. Zhang, and D. Schuurmans (2012). Regularizers versus Losses for Nonlinear Dimensionality Reduction. In: *International Conference on Machine Learning (ICML)*.
40. Zhang, X., Y. Yu, and D. Schuurmans (2012). Accelerated Training for Matrix-Norm Regularization: A Boosting Approach. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
41. Y.-L. Yu and D. Schuurmans (2011). Rank/Norm Regularization with Closed-Form Solutions: Application to Subspace Clustering. In: *Conference on Uncertainty in Artificial Intelligence (UAI)*.
42. Zhang, X., Y. Yu, M. White, R. Huang, and D. Schuurmans (2011). Convex Sparse Coding, Subspace Learning, and Semi-Supervised Extensions. In: *Association for the Advancement of Artificial Intelligence (AAAI)*.

43. Yu, Y., M. Yang, L. Xu, M. White, and D. Schuurmans (2010). Relaxed Clipping: A Global Training Method for Robust Regression and Classification. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
44. Y.-L. Yu, Y. Li, D. Schuurmans, and C. Szepesvári (2009). A General Projection Property for Distribution Families. In: *Advances in Neural Information Processing Systems (NeurIPS)*.
45. Guan, P., Y. Yu, and L. Zhang (2007). A Novel Facial Feature Point Localization Method on 3D Faces. In: *IEEE Conference on Image Processing (ICIP)*.

Workshop Posters

1. Wu, K. and Y. Yu (2019). Understanding Adversarial Robustness: The Trade-off between Minimum and Average Margin. In: *NeurIPS Workshop on Machine Learning with Guarantees*.
2. Yu, Y., H. Cheng, and X. Zhang (2014). “Approximate Low-Rank Tensor Learning”. In: *NeurIPS workshop on Optimization*.
3. Y.-L. Yu, Y. Zhang, and C. Szepesvári (2009). “Online TD(1) Meets Offline Monte Carlo”. In: *Multidisciplinary Symposium on Reinforcement Learning*.

Manuscripts

1. Fujiwara, T., J. Zhao, F. Chen, Y. Yu, and K.-L. Ma (2020). “Removed for double-blind review”. submitted.
2. Hu, Z., Y. Zhu, I. Ilyas, and Y. Yu (2020). “Removed for double-blind review”. submitted.
3. Ma, Y., V. Ganapathiraman, Y. Yu, and X. Zhang (2020). “Convex Representation Learning for Generalized Invariance in Semi-Inner-Product Space”. submitted.
4. Pham, H. V., S. Qian, J. Wang, T. Lutellier, J. Rosenthal, L. Tan, Y. Yu, and N. Nagappan (2020). “Removed for double-blind review”. submitted.
5. Wu, K., H. Wang, and Y. Yu (2020). “Removed for double-blind review”. submitted.
6. Zhang, G., P. Poupart, and Y. Yu (2020). “Optimality and Stability in Non-Convex-Non-Concave Min-Max Optimization”. submitted.
7. Jaini, P., I. Kobyzev, M. Brubaker, and Y. Yu (2019). “Tails of Triangular Flows”. submitted.
8. Ram, A., J. Xin, M. Nagappan, Y. Yu, R. Lozoya, A. Sabetta, and J. Lin (2019). “Exploiting Token and Path-based Representations of Code for Identifying Security-Relevant Commits”. submitted.

Funding

My research is supported by

- ▶ Natural Sciences and Engineering Research Council of Canada (NSERC)
- ▶ Borealis AI and Mitacs
- ▶ National Research Council (NRC)
- ▶ Huawei, through Waterloo-Huawei Joint Innovation Lab

Talks (Selected)

- August 2019 Waterloo ML + Security + Verification Workshop (Waterloo, ON)
Adversarial Machine Learning: Old and New
- November 2018 Vector Institute (Toronto, ON)
Deep Homogeneous Mixture Models: Representation, Separation and Approximation
- March 2018 Borealis AI (Toronto, ON)
Generalized Conditional Gradient for Structured Sparsity and Convex Neural Network
- December 2017 Canada Mathematical Society winter meeting (Waterloo, ON)
Minimizing the Sum of Non-separable Functions
- September 2017 BIRS workshop *Splitting Algorithms, Modern Operator Theory, and Applications*
(Oxaca, Mexico)
On Decomposing the Proximal Map

Supervision of Research Students

- Sep 2020 – Present Daniel Herman (MMath at UWaterloo)
- May 2020 – Present Wanxin Li (MMath at UWaterloo, co-supervised with Lila Kari)
- May 2020 – Present Tim Dockhorn (PhD at UWaterloo)
- May 2019 – Present Guojun Zhang (PhD at UWaterloo, co-supervised with Pascal Poupart)
- Jan 2019 – Present Ralph Tang (PhD at UWaterloo, co-supervised with Jimmy Lin)
- Nov 2018 – Present Hung Viet Pham (PhD at UWaterloo, co-supervised with Tan Lin)
- Sep 2018 – Present Zeou Hu (PhD at UWaterloo)
- Sep 2018 – Present Ji Xin (PhD at UWaterloo, co-supervised with Jimmy Lin)
- Sep 2018 – Present Kaiwen Wu (MMath at UWaterloo)
- Sep 2018 – Present Allen Wang (MMath at UWaterloo)
- Sep 2017 – Present Jingjing Wang (PhD at UWaterloo)

Alumni

- Dec 2017 – Dec 2019 Priyank Jaini (PhD at UWaterloo, co-supervised with Pascal Poupart)
Mathematics Faculty Doctoral Prize, 2nd Place
- Sep 2017 – Dec 2019 Xin Lian (MMath at UWaterloo)
- Jan 2018 – May 2019 Yingluo Xun (MMath at UWaterloo, co-supervised with Yuying Li)

Teaching

University of Waterloo

Spring 2020	<i>Introduction to Machine Learning</i> (CS480/680)
Fall 2019	<i>Introduction to Machine Learning</i> (CS480/680)
Winter 2019	<i>Theory of Deep Learning</i> (CS886)
Fall 2018	<i>Introduction to Machine Learning</i> (CS480/680)
Spring 2018	<i>Computational Linear Algebra</i> (CS475/675)
Fall 2017	<i>Introduction to Machine Learning</i> (CS489/698)
Winter 2017	<i>Algorithms</i> (CS341)
Fall 2016	<i>Optimization for Machine Learning</i> (CS798)

Professional Service

Program Committee Member or Reviewer for Leading Conferences

2018	International Conference on Learning Representation (ICLR)
2010-2019	International Conference on Machine Learning (ICML)
2009-2020	Neural Information Processing Systems (NeurIPS)
2013, 2016, 2020	International Joint Conference on Artificial Intelligence (IJCAI)
2015	Uncertainty on Artificial Intelligence (UAI)
2014-2015, 2018-2019	AAAI Conference on Artificial Intelligence (AAAI)
2011	Artificial Intelligence and Statistics (AISTATS)
2009, 2020	Conference on Learning Theory (COLT)

Reviewer for journals

Computational Optimization and Application (COA)
Journal of Machine Learning Research (JMLR)
Machine Learning Journal (MLJ)
Artificial Intelligence Journal (AIJ)
IEEE Transactions on Pattern Analysis and Machine Intelligence (IEEE-TPAMI)
IEEE Transactions on Automatic Control (IEEE-TAC)
IEEE Transactions on Knowledge and Data Engineering (IEEE-TKDE)
Theoretical Computer Science (TCS)
Neural Computation (NECO)

Reviewing Recognition

2018 NeurIPS Top 200 reviewers
2018 ICML Top 10 reviewers