

John C. Duchi

Stanford University
126 Sequoia Hall
390 Jane Stanford Way
Stanford, CA 94305

Phone: (650) 498-5224
Email: jduchi@stanford.edu
Homepage: <http://web.stanford.edu/~jduchi>

Academic and Research Employment

Associate Professor of Statistics and Electrical Engineering, Stanford University, September 1, 2021–present.

Assistant Professor of Statistics and Electrical Engineering, Stanford University, September 1, 2014–August 31, 2021.

Research Assistant, Michael I. Jordan and Martin J. Wainwright, University of California, Berkeley, 2008–2014.

Research Intern/Engineer for Yoram Singer, Google, summers 2009–2014.

Software Engineer, Google, 2007–2008.

Research Assistant, Daphne Koller, Stanford University. 2006–2007.

Education

Ph.D. Computer Science, University of California, Berkeley. Advisors: Michael I. Jordan and Martin J. Wainwright. Thesis: *Multiple Optimality Guarantees in Statistical Learning*.

M.A. Statistics, University of California, Berkeley, 2012.

M.S. Computer Science, Stanford University, 2007.

B.S. (with distinction) Computer Science, Stanford University, 2007.

Awards

Best student paper award for “A Pretty Fast Algorithm for Adaptive Private Mean Estimation” with R. Kuditipudi and S. Haque, *International Conference on Learning Theory (COLT)*, 2023.

Thomas and Polly Bredt Faculty Scholar, Stanford University, 2021–2025.

Society for Industrial and Applied Mathematics (SIAM) Early Career Prize in Optimization, 2020.

Society for Industrial and Applied Mathematics (SIAM) SIGEST award for “an outstanding paper of general interest” for “First-Order Methods for Nonconvex Quadratic Minimization,” 2020.

Office of Naval Research Young Investigator, 2019.

Best student paper award (as advisor) for “Statistics of Robust Optimization: a Generalized Empirical Likelihood Approach,” *INFORMS Applied Probability Society*, 2018.

Sloan Foundation Fellow in Mathematics, 2017.

Best paper award for “Variance Based Regularization with Convex Objectives” with H. Namkoong, *Neural Information Processing Systems 30 (NIPS)*, 2017.

NSF CAREER Award, 2016.

Terman Fellow, Stanford University, 2015.

Association for Computing Machinery Doctoral Dissertation Award, Honorable Mention, 2014.

C.V. Ramamoorthy Distinguished Research Award, University of California, Berkeley, 2014.

Facebook PhD Fellowship, 2012–2013.

Best student paper award for “On the Consistency of Ranking Algorithms” with L. W. Mackey and M. I. Jordan, *International Conference on Machine Learning (ICML)* 2010.

National Defense Science and Engineering Graduate Fellowship, 2009–2012.

Terman Scholar for Academic Excellence, Stanford University School of Engineering, 2006.

Phi Beta Kappa, inducted as a junior, 2005.

Grants and Funding

Office of Naval Research. 2021 – 2024.

NSF Robust Intelligence 2006777: Robustness and Confidence in Machine-Learned Systems. 2020 – 2023.

Office of Naval Research Young Investigator. 2018 – 2021.

NSF HDR Award 1934578: The Stanford Data Science Collaboratory (sub-PI). 2019 – 2021.

DAWN Consortium, Stanford, 2018 – 2022.

Sloan Foundation Fellowship in Mathematics. 2017 – 2018.

Toyota Research Grant: Uncertainty on Uncertainty and Robustness (PI, joint with Peter Glynn). 2015 – 2018.

NSF CAREER Award 1553086: The Optimal Use of Data (PI). February 2016 – January 2021.

DARPA: Binder-Finder through Machine Learning (sub-PI). 2016 – 2018.

NSF EAGER: Design of Generative Product Behavior using Morphing Algorithms (sub-PI). 2015 – 2018.

Terman Fellowship, Stanford University (PI). 2015 – 2018.

Okawa Foundation Research Award. 2015.

Google Faculty Award: Randomized contextual kernels for data-dependent feature engineering. (Co-PI with Benjamin Recht.) 2015.

Research Supervision

Former Students and Postdoctoral Researchers

Hilal Asi (Stanford University, Electrical Engineering PhD, 2018–2022). Research Scientist, Apple.

Maxime Cauchois (Stanford University, Statistics PhD, 2018–2022). Two Sigma.

Suyash Gupta (Stanford University, Statistics PhD, 2018–2022). Research Scientist, Amazon.

Almur Ali (Stanford University, post-doc, 2020–2022). Head of Machine Learning, Sisu Data.

Daniel Levy (Stanford University, Computer Science PhD, 2019–2022). OpenAI.

Yu Bai (Stanford University, Statistics PhD, 2015–2019). Research Scientist, Salesforce.

Yair Carmon (Stanford University, Electrical Engineering PhD, 2016–2020). Assistant Professor of Computer Science, Tel Aviv University.

Andrew Naber (Stanford University, Electrical Engineering PhD, 2016–2020). Engineer, Apple.

Hongseok Namkoong (Stanford University, Management Science and Engineering PhD, 2014–2019). Assistant Professor, Columbia University Business School.

Feng Ruan (Stanford University, Statistics PhD, 2015–2019). Assistant Professor, Northwestern University.

Tatsunori Hashimoto (Stanford University, post-doc, jointly supervised with Percy Liang, 2016–2019). Assistant Professor of Computer Science, Stanford University.

Aman Sinha (Stanford University, Electrical Engineering PhD, 2015–2020).

Fanny Yang (Stanford University, post-doc, jointly supervised with Percy Liang, 2019). Assistant Professor, ETH Zurich.

Current Students and Postdoctoral Researchers

Felipe Areces (Stanford University, Electrical Engineering PhD, 2022–current).

Karan Chadha (Stanford University, Electrical Engineering PhD, 2020–current).

Chen Cheng (Stanford University, Statistics PhD, 2020–current).

Gary Cheng (Stanford University, Electrical Engineering PhD, 2020–current).

Saminul Haque (Stanford University, Computer Science PhD, 2021–current).

Rohith Kuditipudi (Stanford University, Computer Science PhD, 2020–current).

Annie Marsden (Stanford University, Computer Science PhD, 2017–current).

Teaching Experience

Engineering 108, Introduction to Matrix Methods, Spring 2023. Stanford University.

Statistics 305a, Applied Statistics I (Linear Models), Autumn 2022. Stanford University.

Statistics 116, Introduction to Probability, Spring 2022. Stanford University.

Statistics 305a, Applied Statistics I (Linear Models), Autumn 2021. Stanford University.

Statistics 311/Electrical Engineering 377, Information Theory and Statistics, Autumn 2021. Stanford University.

Electrical Engineering 364a, Convex Optimization, Winter 2020. Stanford University.

Electrical Engineering 377/Statistics 311, Information Theory and Statistics, Winter 2019. Stanford University.

Statistics 300b, Theoretical Statistics, Winter 2019. Stanford University.

Statistics 101, Data Science, Autumn 2018. Stanford University.

Electrical 364b/Computational Mathematics and Engineering 364b, Convex Optimization II, Spring 2018. Stanford University.

Statistics 300b, Theoretical Statistics, Winter 2018. Stanford University.

Statistics 231/Computer Science 229T, Statistical Learning Theory, Spring 2017. Stanford University.

Statistics 231/Computer Science 229T, Statistical Learning Theory, Spring 2017. Stanford University.

Statistics 300b, Theoretical Statistics, Winter 2017. Stanford University.

Computer Science 229/Statistics 229, Machine Learning, Autumn 2016. Stanford University.

Computer Science 229/Statistics 229, Machine Learning, Spring 2016. Stanford University.

Statistics 311/Electrical Engineering 377, Information Theory and Statistics, Winter 2015. Stanford University.

Statistics 311/Electrical Engineering 377, Information Theory and Statistics, Fall 2014. Stanford University.

Electrical Engineering 364b/Computational Mathematics and Engineering 364b, Convex Optimization II, Spring 2015. Stanford University.

Instructor, Distributed Optimization, Spring 2011 (co-taught with Mikael Johansson and Laurent El Ghaoui). University of California, Berkeley.

Publications

Students or postdoctoral researchers supervised denoted by an asterisk *.

Refereed Journal Publications

1. J. C. Duchi and F. Ruan*. The right complexity measure in locally private estimation: It is not the Fisher information. *Annals of Statistics*, To appear, 2023.
2. Y. Arjevani, Y. Carmon*, J. C. Duchi, D. J. Foster, N. Srebro, and B. Woodworth. Lower bounds for non-convex stochastic optimization. *Mathematical Programming, Series A*, 199:165–214, 2023.
3. J. C. Duchi, T. Hashimoto*, and H. Namkoong*. Distributionally robust losses for latent covariate mixtures. *Operations Research*, 71(2):649–664, 2022.
4. S. Yadlowsky*, H. Namkoong*, S. Basu, J. Duchi, and L. Tian. Bounds on the conditional and average treatment effect in the presence of unobserved confounders. *Annals of Statistics*, 50(5):2587–2615, 2022.
5. A. Kipnis* and J. C. Duchi. Mean estimation from one-bit measurements. *IEEE Transactions on Information Theory*, 68(9):6276–6296, 2022.
6. J. C. Duchi and F. Ruan*. Asymptotic optimality in stochastic optimization. *Annals of Statistics*, 49(1):21–48, 2021.
7. M. Cauchois*, S. Gupta*, and J. Duchi. Knowing what you know: valid and validated confidence sets in multiclass and multilabel prediction. *Journal of Machine Learning Research*, 22(81):1–42, 2021.
8. J. C. Duchi and H. Namkoong*. Learning models with uniform performance via distributionally robust optimization. *Annals of Statistics*, 49(3):1378–1406, 2021.
9. J. C. Duchi, P. W. Glynn, and H. Namkoong*. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 46(3):946–969, 2021.
10. Y. Carmon* and J. C. Duchi. First-order methods for nonconvex quadratic minimization. *SIAM Review*, 62(2):395–436, 2020. SIGEST Award Paper.
11. Y. Carmon*, J. C. Duchi, O. Hinder*, and A. Sidford. Lower bounds for finding stationary points II: First order methods. *Mathematical Programming, Series A*, to appear, 2019.
12. H. Asi* and J. C. Duchi. The importance of better models in stochastic optimization. *Proceedings of the National Academy of Sciences*, 116(46):22924–22930, 2019. URL <https://arXiv.org/abs/1903.08619>.
13. H. Asi* and J. C. Duchi. Stochastic (approximate) proximal point methods: Convergence, optimality, and adaptivity. *SIAM Journal on Optimization*, 29(3):2257–2290, 2019. URL <https://doi.org/10.1137/18M1230323>.
14. Y. Carmon*, J. C. Duchi, O. Hinder*, and A. Sidford. Lower bounds for finding stationary points I. *Mathematical Programming, Series A*, 184:171–120, 2020.
15. J. C. Duchi and H. Namkoong. Variance-based regularization with convex objectives. *Journal of Machine Learning Research*, 20(68):1–55, 2019.

16. Y. Carmon* and J. C. Duchi. Gradient descent efficiently finds the cubic-regularized non-convex Newton step. *SIAM Journal on Optimization*, 29(3):2146—2178, 2019. URL <https://epubs.siam.org/doi/abs/10.1137/17M1113898>.
17. J. C. Duchi and F. Ruan*. Stochastic methods for composite and weakly convex optimization problems. *SIAM Journal on Optimization*, 28(4):3229–3259, 2018.
18. J. C. Duchi and F. Ruan*. Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval. *Information and Inference*, iay015, 2018.
19. Y. Carmon*, J. C. Duchi, O. Hinder*, and A. Sidford. Accelerated methods for non-convex optimization. *SIAM Journal on Optimization*, 28(2):1751–1772, 2018.
20. J. C. Duchi, K. Khosravi*, and F. Ruan*. Multiclass classification, information, divergence, and surrogate risk. *Annals of Statistics*, 46(6b):3246–3275, 2018.
21. J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Minimax optimal procedures for locally private estimation (with discussion). *Journal of the American Statistical Association*, 113(521):182–215, 2018.
22. Z. Lin, C. Yang, Y. Zhu, J. Duchi, Y. Fu, Y. Wang, B. Jiang, M. Zamanighomi, X. Xu, M. Li, N. Sestan, H. Zhao, and W. H. Wong. Simultaneous dimension reduction and adjustment for confounding variation. *Proceedings of the National Academy of Sciences*, 113(51):14662–14667, 2016.
23. J. C. Duchi, M. I. Jordan, M. J. Wainwright, and A. Wibisono. Optimal rates for zero-order optimization: the power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5):2788–2806, 2015.
24. Y. Zhang, J. C. Duchi, and M. J. Wainwright. Divide and conquer kernel ridge regression: A distributed algorithm with minimax optimal rates. *Journal of Machine Learning Research*, 16:3299–3340, 2015.
25. J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Privacy aware learning. *Journal of the Association for Computing Machinery*, 61(6):38:1–38:57, 2014.
26. J. C. Duchi, L. Mackey, and M. I. Jordan. The asymptotics of ranking algorithms. *Annals of Statistics*, 41(5):2292–2323, 2013.
27. Y. Zhang, J. C. Duchi, and M. J. Wainwright. Communication-efficient algorithms for statistical optimization. *Journal of Machine Learning Research*, 14:3321–3363, 2013.
28. A. Agarwal and J. C. Duchi. The generalization ability of online algorithms for dependent data. *IEEE Transactions on Information Theory*, 59(1):573–587, 2013.
29. J. C. Duchi, A. Agarwal, M. Johansson, and M. I. Jordan. Ergodic mirror descent. *SIAM Journal on Optimization*, 22(4):1549–1578, 2012.
30. J. C. Duchi, P. L. Bartlett, and M. J. Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.
31. J. C. Duchi, A. Agarwal, and M. J. Wainwright. Dual averaging for distributed optimization: convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, 57(3):592–606, 2012.
32. J. C. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.
33. J. C. Duchi and Y. Singer. Efficient online and batch learning using forward-backward splitting. *Journal of Machine Learning Research*, 10:2873–2898, 2009.

Refereed Conference Publications

1. J. C. Duchi, S. Haque*, and R. Kuditipudi*. A pretty fast algorithm for adaptive private mean estimation. In *Proceedings of the Thirty Sixth Annual Conference on Computational Learning Theory*, 2023. Best student paper award.
2. G. Cheng*, K. Chadha*, and J. Duchi. Federated asymptotics: a model to compare federated learning algorithms. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, 2023.
3. H. Asi*, K. Chadha*, G. Cheng*, and J. C. Duchi. Private optimization in the interpolation regime: faster rates and hardness results. In *Proceedings of the 38th International Conference on Machine Learning*, 2022.
4. K. Chadha*, G. Cheng*, and J. Duchi. Accelerated, optimal, and parallel: Some results on model-based stochastic optimization. In *Proceedings of the 38th International Conference on Machine Learning*, 2022.
5. C. Cheng*, J. Duchi, and R. Kuditipudi*. Memorize to generalize: on the necessity of interpolation in high dimensional linear regression. In *Proceedings of the Thirty Fifth Annual Conference on Computational Learning Theory*, 2022.
6. H. Asi*, D. Levy*, and J. Duchi. Adapting to function difficulty and growth conditions in private optimization. In *Advances in Neural Information Processing Systems 35*, 2021.
7. H. Asi*, J. Duchi, A. Fallah, O. Javidbakht, and K. Talwar. Private adaptive gradient methods for convex optimization. In *Proceedings of the 37th International Conference on Machine Learning*, 2021.
8. J. C. Duchi and F. Ruan*. A constrained risk inequality for general losses. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, 2021. Awarded oral presentation.
9. J. Duchi, A. Marsden*, and G. Valiant. On misspecification in prediction problems and robustness via improper learning. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, 2021. Awarded oral presentation.
10. A. Sinha*, M. O’Kelly, R. Tedrake, and J. Duchi. Neural bridge sampling for evaluating safety-critical autonomous systems. In *Advances in Neural Information Processing Systems 34*, 2020.
11. D. Levy*, Y. Carmon*, J. C. Duchi, and A. Sidford. Large-scale methods for distributionally robust optimization. In *Advances in Neural Information Processing Systems 34*, 2020.
12. H. Asi* and J. C. Duchi. Instance-optimality in differential privacy via approximate inverse sensitivity mechanisms. In *Advances in Neural Information Processing Systems 34*, 2020.
13. H. Asi*, K. Chadha*, G. Cheng*, and J. C. Duchi. Minibatch stochastic approximate proximal point methods. In *Advances in Neural Information Processing Systems 34*, 2020.
14. J. C. Duchi, O. Hinder*, A. Naber*, and Y. Ye. Conic descent and its application to memory-efficient optimization over positive semidefinite matrices. In *Advances in Neural Information Processing Systems 34*, 2020.
15. A. Sinha*, M. O’Kelly, H. Zheng, R. Mangharam, J. Duchi, and R. Tedrake. FormulaZero: Distributionally robust online adaptation via offline population synthesis. In *Proceedings of the 36th International Conference on Machine Learning*, 2020.
16. A. Raghunathan*, S. M. Xie*, F. Yang, J. Duchi, and P. Liang. Understanding and mitigating the tradeoff between robustness and accuracy. In *Proceedings of the 36th International Conference on Machine Learning*, 2020.

17. Y. Arjevani, Y. Carmon*, J. C. Duchi, D. J. Foster, A. Sekhari, and K. Sridharan. Second-order information in non-convex stochastic optimization: Power and limitations. In *Proceedings of the Thirty Third Annual Conference on Computational Learning Theory*, 2020.
18. D. Levy* and J. C. Duchi. Necessary and sufficient geometries for adaptive gradient methods. In *Advances in Neural Information Processing Systems 33*, 2019. Awarded oral presentation.
19. Y. Carmon*, A. Raghunathan*, L. Schmidt, P. Liang, and J. Duchi. Unlabeled data improves adversarial robustness. In *Advances in Neural Information Processing Systems 33*, 2019.
20. J. C. Duchi and R. Rogers. Lower bounds for locally private estimation via communication complexity. In *Proceedings of the Thirty Second Annual Conference on Computational Learning Theory*, 2019.
21. Y. Carmon*, J. Duchi, A. Sidford, and K. Tian*. A rank-1 sketch for matrix multiplicative weights. In *Proceedings of the Thirty Second Annual Conference on Computational Learning Theory*, 2019.
22. H. Asi* and J. C. Duchi. Modeling simple structures and geometry for better stochastic optimization algorithms. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, 2019.
23. Y. Carmon* and J. C. Duchi. Analysis of Krylov subspace solutions of regularized nonconvex quadratic problems. In *Advances in Neural Information Processing Systems 32*, 2018. Awarded oral presentation.
24. M. O’Kelly, A. Sinha*, H. Namkoong*, J. Duchi, and R. Tedrake. Scalable end-to-end autonomous vehicle testing via rare-event simulation. In *Advances in Neural Information Processing Systems 31*, 2018.
25. R. Volpi, H. Namkoong*, O. Sener*, J. Duchi, V. Murino*, and S. Savarese. Generalizing to unseen domains via adversarial data augmentation. In *Advances in Neural Information Processing Systems 32*, 2018.
26. J. C. Duchi, F. Ruan*, and C. Yun*. Minimax bounds on stochastic batched convex optimization. In *Proceedings of the Thirty First Annual Conference on Computational Learning Theory*, 2018.
27. A. Sinha*, H. Namkoong*, and J. Duchi. Certifying some distributional robustness with principled adversarial training. In *Proceedings of the Sixth International Conference on Learning Representations*, 2018. URL <https://arxiv.org/abs/1710.10571>. Awarded oral presentation.
28. T. Hashimoto*, S. Yadlowsky*, and J. Duchi. Reducing optimization to repeated classification. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, 2018.
29. H. Namkoong* and J. C. Duchi. Variance-based regularization with convex objectives. In *Advances in Neural Information Processing Systems 31*, 2017. Best paper award.
30. T. Hashimoto*, J. Duchi, and P. Liang. Unsupervised transformation learning via convex relaxations. In *Advances in Neural Information Processing Systems 31*, 2017.
31. Y. Carmon*, J. C. Duchi, O. Hinder*, and A. Sidford. Convex until proven guilty: dimension-free acceleration of gradient descent on non-convex functions. In *Proceedings of the 33rd International Conference on Machine Learning*, 2017.
32. A. Sinha* and J. C. Duchi. Learning kernels with random features. In *Advances in Neural Information Processing Systems 30*, 2016.
33. H. Namkoong* and J. C. Duchi. Stochastic gradient methods for distributionally robust optimization with f -divergences. In *Advances in Neural Information Processing Systems 30*, 2016.

34. S. Chatterjee, J. Duchi, J. Lafferty, and Y. Zhu. Local minimax complexity of stochastic convex optimization. In *Advances in Neural Information Processing Systems 30*, 2016.
35. A. Raganathan*, R. Frostig*, J. Duchi, and P. Liang. Estimation from indirect supervision with linear moments. In *Proceedings of the 32nd International Conference on Machine Learning*, 2016.
36. S. Chaturapruek*, J. C. Duchi, and C. Ré. Asynchronous stochastic convex optimization: the noise is in the noise and SGD don't care. In *Advances in Neural Information Processing Systems 29*, 2015.
37. J. Steinhardt* and J. C. Duchi. Minimax rates for memory-bounded sparse linear regression. In *Proceedings of the Twenty Eighth Annual Conference on Computational Learning Theory*, 2015.
38. J. C. Duchi, M. I. Jordan, and H. B. McMahan. Estimation, optimization, and parallelism when data is sparse. In *Advances in Neural Information Processing Systems 27*, 2013.
39. Y. Zhang, J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Information-theoretic lower bounds for distributed estimation with communication constraints. In *Advances in Neural Information Processing Systems 27*, 2013. Awarded oral presentation.
40. J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy and minimax bounds: Sharp rates for probability estimation. In *Advances in Neural Information Processing Systems 27*, 2013.
41. J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy and statistical minimax rates. In *54th Annual Symposium on Foundations of Computer Science*, 2013.
42. T. Kraska, A. Talwalkar, J. C. Duchi, R. Griffith, M. Franklin, and M. I. Jordan. MLbase: A distributed machine-learning system. In *Sixth Biennial Conference on Innovative Data Systems Research (CIDR)*, 2013.
43. J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Privacy aware learning. In *Advances in Neural Information Processing Systems 26*, 2012.
44. J. C. Duchi, M. I. Jordan, M. J. Wainwright, and A. Wibisono. Finite sample convergence rates of zero-order stochastic optimization methods. In *Advances in Neural Information Processing Systems 26*, 2012.
45. Y. Zhang, J. C. Duchi, and M. J. Wainwright. Communication-efficient algorithms for statistical optimization. In *Advances in Neural Information Processing Systems 25*, 2012.
46. J. C. Duchi, P. L. Bartlett, and M. J. Wainwright. Randomized smoothing for (parallel) stochastic optimization. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.
47. A. Agarwal and J. C. Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems 25*, 2011.
48. J. C. Duchi, A. Agarwal, M. Johansson, and M. I. Jordan. Ergodic mirror descent. In *The 49th Allerton Conference on Communication, Control, and Computing*, pages 701–706, 2011.
49. A. Agarwal, J. C. Duchi, P. Bartlett, and C. Levrard. Oracle inequalities for computationally budgeted model selection. In *Proceedings of the Twenty Fourth Annual Conference on Computational Learning Theory*, 2011.
50. J. C. Duchi, A. Agarwal, and M. J. Wainwright. Distributed dual averaging in networks. In *Advances in Neural Information Processing Systems 24*, 2010.
51. J. C. Duchi, L. Mackey, and M. I. Jordan. On the consistency of ranking algorithms. In *Proceedings of the 27th International Conference on Machine Learning*, 2010. Best paper award.

52. J. C. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. In *Proceedings of the Twenty Third Annual Conference on Computational Learning Theory*, 2010.
53. J. C. Duchi, S. Shalev-Shwartz, Y. Singer, and A. Tewari. Composite objective mirror descent. In *Proceedings of the Twenty Third Annual Conference on Computational Learning Theory*, 2010.
54. J. C. Duchi and Y. Singer. Efficient learning using forward-backward splitting. In *Advances in Neural Information Processing Systems 23*, 2009. Awarded oral presentation.
55. J. C. Duchi and Y. Singer. Boosting with structural sparsity. In *Proceedings of the 26th International Conference on Machine Learning*, 2009.
56. J. C. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *Proceedings of the 25th International Conference on Machine Learning*, 2008.
57. V. Ganapathi, D. Vickrey, J. Duchi, and D. Koller. Constrained approximate maximum entropy learning of Markov random fields. In *Proceedings of the 24th Conference Conference on Uncertainty in Artificial Intelligence (UAI)*, 2008.
58. J. C. Duchi, S. Gould, and D. Koller. Projected subgradient methods for learning sparse gaussians. In *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2008.
59. J. Duchi, D. Tarlow, G. Elidan, and D. Koller. Using combinatorial optimization within max-product belief propagation. In *Advances in Neural Information Processing Systems 20*, 2006.

Books, book chapters, and lecture notes

1. J. C. Duchi. Introductory lectures on stochastic convex optimization. In *Park City Mathematics Institute Graduate Summer School: Collected Lectures*. American Mathematical Society (forthcoming), 2017.

Preprints, papers under review, and technical notes

1. A. Ali*, M. Cauchois*, and J. Duchi. The lifecycle of a statistical model: Model failure detection, identification, and refitting. *arXiv:2202.04166 [stat.ME]*, 2022.
2. M. Cauchois*, S. Gupta*, and A. A. J. Duchi. Predictive inference with weak supervision. *arXiv:2201.08315 [stat.ML]*, 2022.
3. H. Asi* and J. C. Duchi. Near instance-optimality in differential privacy. *arXiv:2005.10630 [cs.CR]*, 2020.
4. H. Asi*, J. Duchi, and O. Javidi. Element level differential privacy: The right granularity of privacy. *arXiv:1912.04042 [cs.LG]*, 2019.
5. A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers. Protection against reconstruction and its applications in private federated learning. *arXiv:1812.00984 [stat.ML]*, 2018.
6. S. Chatterjee, J. Duchi, J. Lafferty, and Y. Zhu. Local minimax complexity of stochastic convex optimization. *arXiv:1605.07596 [stat.ML]*, 2016.
7. R. F. Barber and J. C. Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv:1412.4451 [math.ST]*, 2014.
8. J. C. Duchi, M. I. Jordan, M. J. Wainwright, and Y. Zhang. Information-theoretic lower bounds for distributed statistical estimation with communication constraints. *arXiv:1405.0782 [cs.IT]*, 2014.

9. J. C. Duchi and M. J. Wainwright. Distance-based and continuum Fano inequalities with applications to statistical estimation. *arXiv:1311.2669 [cs.IT]*, 2013.
10. A. Agarwal, P. L. Bartlett, and J. Duchi. Oracle inequalities for computationally adaptive model selection. *arXiv:1208.0129 [stat.ML]*, 2012.

Invited Conference Publications

1. A. Kipnis* and J. C. Duchi. Mean estimation from adaptive one-bit measurements. In *The 55th Allerton Conference on Communication, Control, and Computing*, 2017.
2. J. C. Duchi, K. Khosravi*, and F. Ruan*. Information measures, experiments, multi-category hypothesis tests, and surrogate losses. In *The 54th Allerton Conference on Communication, Control, and Computing*, 2016.
3. R. F. Barber and J. C. Duchi. Privacy: A few definitional aspects and consequences for minimax mean-squared error. In *53rd IEEE Conference on Decisions and Control*, 2014.
4. J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy, quantitative data processing inequalities, and statistical minimax rates. In *IEEE Information Theory Workshop*, 2013.
5. J. C. Duchi, A. Agarwal, and M. J. Wainwright. Dual averaging for distributed optimization. In *The 50th Allerton Conference on Communication, Control, and Computing*, 2012.
6. J. C. Duchi, P. L. Bartlett, and M. J. Wainwright. Randomized smoothing for stochastic optimization. In *51st IEEE Conference on Decisions and Control*, 2012.
7. J. C. Duchi. Commentary on “Toward a noncommutative arithmetic-geometric mean inequality: Conjectures, case-studies, and consequences”. In *Proceedings of the Twenty Fifth Annual Conference on Computational Learning Theory*, pages 11.25–11.27, 2012.

Invited and Award Talks

- 2023 Statistics Seminar, Harvard University
- 2023 Jordan Symposium (France)
- 2023 Workshop on Optimization and Statistical Learning (Les Houches, France)
- 2022 AMS-EMS-SMF International Meeting (Grenoble)
- 2022 Workshop on learning: optimization and stochastics (Ecole Polytechnique Federale de Lausanne)
- 2022 International Meeting on Statistics (London)
- 2021 SIAM Conference on Optimization (remote)
- 2021 Bernoulli-IMS 10th World Congress in Probability and Statistics (remote)
- 2021 Mathematical Foundations of Machine Learning (Oberwolfach, Germany, remote)
- 2020 Massachusetts Institute of Technology Machine Learning and Optimization Seminar
- 2020 McGill Operations Research Seminar
- 2020 Cornell Operations Research Symposium
- 2020 Conference on Uncertainty in Artificial Intelligence (UAI) Tutorial
- 2019 Optimization for Machine Learning: OPT 2019 (Vancouver, Canada)
- 2019 International Conference on Continuous Optimization (Berlin, Germany)
- 2019 Machine Learning Summer School (London, England)
- 2019 Recent Themes in Resource Tradeoffs (Institute for Mathematics and its Applications, Minnesota, USA)
- 2019 Statistical and Computational Aspects of Learning with Complex Structure (Oberwolfach, Germany)
- 2019 Workshop on Optimization and Statistical Learning (Les Houches, France)
- 2019 Rice University Symposium on Data Privacy (Rice University, USA)

2018 Massachusetts Institute of Technology IDSS Seminar
 2018 Workshop on Adaptive Data Analysis (Simons Institute, UC Berkeley, USA)
 2018 INFORMS Conference Tutorial
 2018 DIMACS/TRIPODS Conference on Optimization, Plenary Speaker (Lehigh University, USA)
 2018 Workshop on Adaptive Data Analysis (Simons Institute, UC Berkeley, USA)
 2018 International Symposium on Mathematical Programming, Learning Cluster (Bordeaux, France)
 2018 Big Data challenges: heterogeneity, model misspecification and changepoints (Isaac Newton Institute, Cambridge, England)
 2018 West Coast Optimization Meeting (University of Washington, USA)
 2017 University of Michigan Statistics Seminar
 2017 University of Washington CORE and Statistics Seminar
 2017 Rice University DSP and Machine Learning Seminar
 2017 Conference on Uncertainty in Artificial Intelligence (UAI) Tutorial
 2017 Neural Information Processing Systems (Long Beach, California)
 2017 Optimization, Statistics, and Uncertainty (Simons Institute, UC Berkeley, USA)
 2016 Park City Mathematics Institute Graduate Summer School Lecturer
 2016 Oberwolfach workshop on Computationally and Statistically Efficient Inference for Complex Large-scale Data (Oberwolfach, Germany)
 2016 Optimization without Borders (Les Houches, France)
 2015 Joint Statistical Meetings, Session on Tradeoffs in Statistics (Seattle, USA)
 2015 International Symposium on Mathematical Programming, Sparse Optimization Cluster (Pittsburgh, USA)
 2015 Oberwolfach workshop on Probabilistic Techniques in Modern Statistics, (Oberwolfach, Germany)
 2015 Workshop on Information Theory and Learning (Simons Institute, UC Berkeley, USA)
 2015 Information Theory and Applications (UC San Diego, USA)
 2015 Workshop on Optimization and Statistical Learning (Les Houches, France)
 2014 Meeting in Mathematical Statistics: New Procedures for New Data (Luminy, France)
 2014 Conference on Decisions and Control, Differential Privacy session (Los Angeles, USA)
 2014 University of Chicago Scientific and Statistical Computing Seminar
 2014 Summer School on Machine Learning with Constraints, T.U. Dortmund
 2014 University of Pennsylvania Wharton Statistics Seminar
 2014 University of Pennsylvania Computer Science Seminar
 2014 University of Washington Statistics Seminar
 2014 Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory
 2014 Princeton Computer Science Seminar
 2014 Princeton Operations Research and Financial Engineering Seminar
 2014 Cornell Computer Science Seminar
 2014 Cornell Operations Research Seminar
 2014 Massachusetts Institute of Technology Operations Research Center
 2014 Carnegie Mellon Statistics Seminar
 2014 Harvard University Statistics Seminar
 2014 Yale University Statistics Seminar
 2013 Neural Information Processing Systems (Lake Tahoe, Nevada USA)
 2013 Allerton Conference on Communications, Control, and Computing (Allerton, Illinois USA)
 2013 Yale University Statistics Seminar
 2013 Stanford University
 2013 University of Southern California EE Seminar
 2013 University of Illinois CSL Seminar
 2013 Microsoft Research New England
 2013 University of Pennsylvania
 2013 University of Michigan AI Seminar
 2012 Neural Information Processing Systems (Lake Tahoe, Nevada USA)
 2012 Allerton Conference on Communications, Control, and Computing (Allerton, Illinois USA)

- 2012 International Symposium on Mathematical Programming (Berlin, Germany)
- 2012 International Conference on Machine Learning (Edinburgh, Scotland)
- 2012 University of Wisconsin SILO Seminar
- 2011 Google New York
- 2011 Stanford University
- 2011 Allerton Conference on Communications, Control, and Computing (Allerton, Illinois USA)
- 2010 International Conference on Machine Learning (Haifa, Israel)
- 2010 Conference on Learning Theory (Haifa, Israel)
- 2010 Workshop on Algorithms for Modern Massive Data Sets (Stanford, California USA)
- 2009 Neural Information Processing Systems (Vancouver, Canada)

Professional Service

Senior Program Committees and Editorial Boards

- 2020– Mathematical Programming (Series A), Associate Editor
- 2021 International Conference on Machine Learning
- 2021 Artificial Intelligence and Statistics
- 2020 International Conference on Machine Learning
- 2020 Conference on Learning Theory
- 2019 Conference on Learning Theory
- 2018 Artificial Intelligence and Statistics
- 2018 IEEE International Symposium on Information Theory
- 2017 Neural Information Processing Systems
- 2015 Conference on Learning Theory
- 2015 International Conference on Machine Learning

Reviewing

Journal Reviewing *Annals of Statistics*, *Journal of Machine Learning Research (JMLR)*, *IEEE Transactions on Information Theory*, *SIAM Journal on Optimization (SIOPT)*, *Machine Learning Journal*, *Foundations and Trends in Machine Learning*, *IEEE Transactions on Automatic Control*, *IEEE Transactions on Control of Networked Systems*, *IMA Journal on Numerical Analysis*, *Electronic Journal of Statistics*, *Foundations of Computational Mathematics*, *Operations Research*, *Proceedings of the National Academy of Sciences*.

Conference Reviewing International Conference on Machine Learning (ICML), Neural Information Processing Systems (NIPS), Artificial Intelligence and Statistics (AISTATS), Conference on Learning Theory (COLT), Conference on Uncertainty in Artificial Intelligence (UAI), Machine Learning and Systems (MLSys).

Outside Interests

Cooking, swimming, cycling, running, triathlons, hiking, backpacking.