John C. Duchi

Stanford University
126 Sequoia Hall
390 Serra Mall
Stanford, CA 94305

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

Academic and Research Employment

Assistant Professor of Statistics and Electrical Engineering and (by courtesy) Computer Science, Stanford University, 2014–present.

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

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.

National Science Foundation Graduate Fellowship, Honorable Mention, 2008 & 2009.

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

Phi Beta Kappa, inducted as a junior, 2005.

Teaching Experience

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

Guest lecturer, Statistical Learning Theory, Fall 2012 (taught by Martin Wainwright). University of California, Berkeley.

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

TA, Artificial Intelligence, Spring 2011 (taught by Pieter Abbeel). University of California, Berkeley.

TA, Introduction to Convex Optimization, Spring 2009 (taught by Laurent El Ghaoui). University of California, Berkeley.

TA, Reasoning Methods in Artificial Intelligence, Spring 2006 & Spring 2007 (taught by P. Pandurang Nayak). Stanford University.

TA, Probabilistic Models in Artificial Intelligence, Winter 2007 (taught by Daphne Koller). Stanford University.

Publications

Refereed Journal Publications

1. J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Privacy aware learning. *Journal of the Association for Computing Machinery*, page To appear, 2014. URL `http://arxiv.org/abs/1210.2085`.

2. J. C. Duchi, L. Mackey, and M. I. Jordan. The asymptotics of ranking algorithms. *Annals of Statistics*, 41(5):2292–2323, 2013. URL `http://arxiv.org/abs/1204.1688`.

3. Y. Zhang, J. C. Duchi, and M. J. Wainwright. Communication-efficient algorithms for statistical optimization. *Journal of Machine Learning Research*, 14:3321–3363, 2013.

4. 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.

5. J. C. Duchi, A. Agarwal, M. Johansson, and M. I. Jordan. Ergodic mirror descent. *SIAM Journal on Optimization*, 22(4):1549–1578, 2012.

6. J. C. Duchi, P. L. Bartlett, and M. J. Wainwright. Randomized smoothing for stochastic optimization. *SIAM Journal on Optimization*, 22(2):674–701, 2012.

7. 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.

8. 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.

9. 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, M. I. Jordan, and H. B. McMahan. Estimation, optimization, and parallelism when data is sparse. In *Advances in Neural Information Processing Systems 26*, 2013.

2. 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 26*, 2013.

3. 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 26*, 2013.

4. 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.

5. 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.

6. J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Privacy aware learning. In *Advances in Neural Information Processing Systems 25*, 2012.

7. 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 25*, 2012.

8. Y. Zhang, J. C. Duchi, and M. J. Wainwright. Communication-efficient algorithms for statistical optimization. In *Advances in Neural Information Processing Systems 25*, 2012.

9. 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.

10. A. Agarwal and J. C. Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems 24*, 2011.

11. 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.

12. 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.

13. J. C. Duchi, A. Agarwal, and M. J. Wainwright. Distributed dual averaging in networks. In *Advances in Neural Information Processing Systems 23*, 2010.

14. 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.

15. 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.

16. 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.

17. J. C. Duchi and Y. Singer. Efficient learning using forward-backward splitting. In *Advances in Neural Information Processing Systems 22*, 2009.

18. J. C. Duchi and Y. Singer. Boosting with structural sparsity. In *Proceedings of the 26th International Conference on Machine Learning*, 2009.

19. 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.

20. 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.

21. 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.

22. J. Duchi, D. Tarlow, G. Elidan, and D. Koller. Using combinatorial optimization within max-product belief propagation. In *Advances in Neural Information Processing Systems 19*, 2006.

Preprints and papers under review

1. J. C. Duchi, M. I. Jordan, M. J. Wainwright, and A. Wibisono. Optimal rates for zero-order optimization: the power of two function evaluations. *arXiv:1312.2139 [math.OC]*, 2013. URL `http://arxiv.org/abs/1312.2139`
2. J. C. Duchi and M. J. Wainwright. Distance-based and continuum Fano inequalities with applications to statistical estimation. *arXiv:1311.2669 [cs.IT]*, 2013. URL `http://arxiv.org/abs/1311.2669`.
3. J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy and statistical minimax rates. *arXiv:1302.3203 [math.ST]*, 2013. URL `http://arxiv.org/abs/1302.3203`.
4. Y. Zhang, J. C. Duchi, and M. J. Wainwright. Divide and conquer kernel ridge regression: A distributed algorithm with minimax optimal rates. *arXiv:1305.5029 [math.ST]*, 2013. URL `http://arxiv.org/abs/1305.5029`.
5. A. Agarwal, P. L. Bartlett, and J. Duchi. Oracle inequalities for computationally adaptive model selection. *arXiv:1208.0129 [stat.ML]*, 2012. URL `http://arxiv.org/abs/1208.0129`.

Invited Conference Publications

1. 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.
2. 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.
3. J. C. Duchi, P. L. Bartlett, and M. J. Wainwright. Randomized smoothing for stochastic optimization. In *51st IEEE Conference on Decisions and Control*, 2012.
4. Y. Zhang, J. C. Duchi, and M. J. Wainwright. Communication-efficient algorithms for statistical optimization. In *51st IEEE Conference on Decisions and Control*, 2012.
5. 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.

Talks

Invited Talks

2014	University of Chicago Scientific and Statistical Computing Seminar
2014	Summer School on Machine Learning with Constraints, T.U. Dortmund
2013	Yale University Statistics Department Seminar
2013, 2011	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	University of Wisconsin SILO Seminar
2011	Google New York

Conference Talks

2013 Neural Information Processing Systems (Lake Tahoe, Nevada USA)
2013 Allerton Conference on Communications, Control, and Computing (Allerton, Illinois USA)
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)
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

Reviewing

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).

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*.

Workshop and symposium organization

Co-organizer, *Biglearn 2012* workshop at *Neural Information Processing Systems (NIPS)* 2012 with S. Singh, J. Gonzalez, and Y. Low.

Co-organizer, *Learning on Cores, Clusters and Clouds* workshop at *Neural Information Processing Systems (NIPS)* 2010 with A. Agarwal, O. Dekel and J. Langford.

Outside Interests

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