Sashank J. Reddi *Carnegie Mellon University*

Email sjakkamr@cs.cmu.edu

Phone +1-580-878-0019

Webpage `www.cs.cmu.edu/~sjakkamr`

Department Machine Learning Department, School of Computer Science

Mailing Address Gates Hillman Center, Room 8223, Carnegie Mellon University,
5000 Forbes Avenue, Pittsburgh, PA 15213, USA.

Research Interest My research interests broadly lie in machine learning and statistics. I am, particularly, interested in high dimension statistics and developing large-scale convex and nonconvex optimization methods for machine learning.

Education

Carnegie Mellon University Aug 2011 - Present
PhD Student in Machine Learning Department
Advisors: Alex Smola and Barnabás Póczos
GPA: 4.17/4

Indian Institute of Technology Bombay, India Jul 2005 - Apr 2010
Bachelor of Science & Master of Science in Computer Science
GPA: 9.11/10.00

Work Experience

Microsoft Research Silicon Valley, USA May 2013 - Aug 2013
Role: Research Intern
Large-scale optimization techniques for kernel methods.

Google, India July 2010 - July 2011
Role: Software Engineer
Developed algorithms for calendar sync and large-scale intent-based search.

University of Waterloo, Canada May 2008 - July 2008
Role: Research Intern
Designed algorithms for optimal allocation of resources in multiple virtual operating systems.

* = equal contribution, the following list excludes workshops and short papers.

List of Publications

Sashank J. Reddi, Sunita Sarawagi, Sundar Vishwanathan. *Map Estimation in Binary MRFs via Bipartite Multicuts*. Proc. of 20th Neural Information Processing Systems (NIPS), pp. 955-963, 2010. **(Awarded Honorable Mention for Outstanding Student Paper Award)**

Ariel Procaccia, Sashank J. Reddi, Nisarg Shah. *A Maximum Likelihood Approach For Selecting Sets of Alternatives*. Proc. of 28th Conference on Uncertainty in Artificial Intelligence (UAI), pp. 695-704, 2012.

Sashank J. Reddi, Emma Brunskill. *Incentive Decision Processes*. Proc. of 28th Conference on Uncertainty in Artificial Intelligence (UAI), pp. 418-427, 2012.

Sashank J. Reddi, Barnabás Póczos. *Scale Invariant Conditional Dependence Measures*. Proc. of 30th International Conference on Machine Learning (ICML), pp. 1355-1363, 2013.

Sashank J. Reddi, Barnabás Póczos. *k-NN Regression on Functional Data with Incomplete Observations*. Proc. of 30th Conference on Uncertainty in Artificial Intelligence (UAI), pp. 692-701, 2014.

Sashank J. Reddi, Barnabás Póczos, Alex Smola. *Doubly Robust Covariate Shift Correction.* Proc. of 29th Association for the Advancement of Artificial Intelligence Conference on Artifical Intelligence (AAAI), pp. 2949-2955, 2015.

Sashank J. Reddi*, Aaditya Ramdas*, Barnabás Póczos, Aarti Singh, Larry Wasserman. *On the Decreasing Power of Kernel and Distance based Nonparametric Hypothesis Tests in High Dimensions.* Proc. of 29th Association for the Advancement of Artificial Intelligence Conference on Artifical Intelligence (AAAI), pp. 3571-3577, 2015.

Sashank J. Reddi*, Aaditya Ramdas*, Barnabás Póczos, Aarti Singh, Larry Wasserman. *On the High-dimensional Power of Linear-time Kernel Two-Sample Testing under Mean-difference Alternatives*, Proc. of International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 772-780, 2015.

Sashank J. Reddi*, Ahmed Hefny*, Avinava Dubey, Carlton Downey, Suvrit Sra. *Large-scale Randomized Coordinate Descent Methods with Non-separable linear constraints.* Proc. of 31st Conference on Uncertainty in Artificial Intelligence (UAI), pp. 762-771, 2015.

Sashank J. Reddi, Barnabás Póczos, Alex Smola. *Communication Efficient Coresets for Empirical Loss Minimization.* Proc. of 31st Conference on Uncertainty in Artificial Intelligence (UAI), pp. 752-761, 2015.

Sashank J. Reddi, Ahmed Hefny, Suvrit Sra, Barnabás Póczos, Alex Smola. *On Variance Reduction in Stochastic Gradient Descent and its Asynchronous Variants.* Advances in Neural Information Processing Systems (NIPS), pp. 2647-2655, 2015.

Sashank J. Reddi, Ahmed Hefny, Suvrit Sra, Barnabás Póczos, Alex Smola. *Stochastic Variance Reduction for Nonconvex Optimization*, Proc. of 33rd International Conference on Machine Learning (ICML), pp. 314-323, 2016.

Sashank J. Reddi, Suvrit Sra, Barnabás Póczos, Alex Smola. *Fast Incremental Method for Nonconvex Optimization*, Proc. of 55th IEEE Conference on Decision and Control (CDC), 2016.

Sashank J. Reddi, Suvrit Sra, Barnabás Póczos, Alex Smola. *Stochastic Frank-Wolfe Methods for Nonconvex Optimization*, Proc. of 54th Annual Allerton Conference on Communication, Control and Computing, 2016.

Sashank J. Reddi*, Avinava Dubey*, Sinead Williamson, Barnabás Póczos, Alex Smola, Eric Xing. *Variance Reduction in Stochastic Gradient Langevin Dynamics.* Advances in Neural Information Processing Systems (NIPS), 2016.

Hongyi Zhang, Sashank J. Reddi, Suvrit Sra. *Fast Stochastic Optimization on Riemannian Manifolds.* Advances in Neural Information Processing Systems (NIPS), 2016.

Sashank J. Reddi, Suvrit Sra, Barnabás Póczos, Alex Smola. *Fast Stochastic Methods for Nonsmooth Nonconvex Optimization.* Advances in Neural Information Processing Systems (NIPS), 2016.

Preprints

Aaditya Ramdas, Sashank J. Reddi, Barnabás Póczos, Aarti Singh, Larry Wasserman. *Adaptivity and Computation-Statistics Tradeoffs for Kernel and Distance based High Dimensional Two Sample Testing* , Submitted to Annals of Statistics. arXiv:1508.00655, 2015.

Sashank J. Reddi, Jakub Konečný, Peter Richtárik, Barnabás Póczos, Alex Smola. *AIDE: Fast and Communication Efficient Distributed Optimization*, 2016.

Major Academic Achievements

Facebook Fellowship Finalist in Machine Learning, 2015, Facebook Inc.
Awarded Honorable Mention for **Outstanding Student Paper Award** at **NIPS** 2010.
Awarded CMU Machine Learning Graduate Fellowship.

Awarded Microsoft Alumni Scholarship for academic performance in 2009.
Ranked 3rd at Regional Mathematics Olympiad 2002.

Relevant Coursework

Machine Learning, Intermediate Statistics, Machine Learning Theory, Graduate Algorithms, Statistical Machine Learning, Databases and Data Mining, Convex Optimization, Probabilistic Graphical Models, Advanced Statistical Theory, Advanced Optimization and Randomized Methods.

References

ALEXANDER J. SMOLA (Professor, Carnegie Mellon University, Pittsburgh).
BARNABÁS PÓCZOS (Assistant Professor, Carnegie Mellon University, Pittsburgh).
SUVRIT SRA (Principal Research Scientist, Massachusetts Institute for Technology).
LARRY WASSERMAN (Professor, Carnegie Mellon University, Pittsburgh).