Publications

Other Privacy Models

1. Kifer, D. and A. Machanavajjhala (2014). Pufferfish: A framework for mathematical privacy definitions. *ACM Trans. Database Syst.* 39(1), 3:1–3:36.
2. Bun, M. and T. Steinke (2016). Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In: *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*. Ed. by M. Hirt and A. D. Smith. Vol. 9985. Lecture Notes in Computer Science, pp.635–658. http://dx.doi.org/10.1007/978-3-662-53641-4_24.

Adaptive Data Analysis via Privacy

1. Dwork, C., V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth (2015). Generalization in Adaptive Data Analysis and Holdout Reuse. In: *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*. Ed. by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, pp.2350–2358. http://papers.nips.cc/paper/5993-generalization-in-adaptive-data-analysis-and-holdout-reuse.
2. Bassily, R., K. Nissim, A. D. Smith, T. Steinke, U. Stemmer, and J. Ullman (2016). Algorithmic stability for adaptive data analysis. In: *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*. Ed. by D. Wichs and Y. Mansour. ACM, pp.1046–1059. http://doi.acm.org/10.1145/2897518.2897566.
3. Rogers, R. M., A. Roth, A. D. Smith, and O. Thakkar (2016). Max-Information, Differential Privacy, and Post-selection Hypothesis Testing. In: *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*. Ed. by I. Dinur. IEEE Computer Society, pp.487–494. http://dx.doi.org/10.1109/FOCS.2016.59.

Game Theory and Privacy

1. McSherry, F. and K. Talwar (2007). Mechanism Design via Differential Privacy. In: *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*. IEEE Computer Society, pp.94–103. http://dx.doi.org/10.1109/FOCS.2007.41.
2. Kearns, M., M. M. Pai, A. Roth, and J. Ullman (2014). Mechanism design in large games: incentives and privacy. In: *Innovations in Theoretical Computer Science, ITCS'14, Princeton, NJ, USA, January 12-14, 2014*. Ed. by M. Naor. ACM, pp.403–410. http://doi.acm.org/10.1145/2554797.2554834.
3. Ghosh, A. and A. Roth (2015). Selling privacy at auction. *Games and Economic Behavior* 91, 334–346.
4. Chen, Y., S. Chong, I. A. Kash, T. Moran, and S. P. Vadhan (2016). Truthful Mechanisms for Agents That Value Privacy. *ACM Trans. Economics and Comput.* 4(3), 13:1–13:30.

Private Machine Learning

1. McSherry, F. and I. Mironov (2009). Differentially Private Recommender Systems: Building Privacy into the Netflix Prize Contenders. In: *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France, June 28 - July 1, 2009*. Ed. by J. F. E. IV, F. Fogelman-Soulié, P. A. Flach, and M. J. Zaki. ACM, pp.627–636. http://doi.acm.org/10.1145/1557019.1557090.
2. Williams, O. and F. McSherry (2010). Probabilistic Inference and Differential Privacy. In: *Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010. Proceedings of a meeting held 6-9 December 2010, Vancouver, British Columbia, Canada*. Ed. by J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta. Curran Associates, Inc., pp.2451–2459. http://papers.nips.cc/paper/3897-probabilistic-inference-and-differential-privacy.
3. Abadi, M., A. Chu, I. J. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang (2016). Deep Learning with Differential Privacy. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*. Ed. by E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi. ACM, pp.308–318. http://doi.acm.org/10.1145/2976749.2978318.

Private Singular Vector Computation

1. Hardt, M. and A. Roth (2013). Beyond worst-case analysis in private singular vector computation. In: *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*. Ed. by D. Boneh, T. Roughgarden, and J. Feigenbaum. ACM, pp.331–340. http://doi.acm.org/10.1145/2488608.2488650.

2. Dwork, C., K. Talwar, A. Thakurta, and L. Zhang (2014). Analyze gauss: optimal bounds for privacy-preserving principal component analysis. In: *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*. Ed. by D. B. Shmoys. ACM, pp.11–20. http://doi.acm.org/10.1145/2591796.2591883.

Private Optimization

1. Bassily, R., A. D. Smith, and A. Thakurta (2014). Private Empirical Risk Minimization: Efficient Algorithms and Tight Error Bounds. In: *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*. IEEE Computer Society, pp.464–473. http://dx.doi.org/10.1109/FOCS.2014.56.
2. Hsu, J., Z. Huang, A. Roth, T. Roughgarden, and Z. S. Wu (2016). Private Matchings and Allocations. *SIAM J. Comput.* 45(6), 1953–1984.
3. Hsu, J., Z. Huang, A. Roth, and Z. S. Wu (2016). Jointly Private Convex Programming. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*. Ed. by R. Krauthgamer. SIAM, pp.580–599. http://dx.doi.org/10.1137/1.9781611974331.ch43.

Privacy on Graphs

1. Blocki, J., A. Blum, A. Datta, and O. Sheffet (2013). Differentially private data analysis of social networks via restricted sensitivity. In: *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*. Ed. by R. D. Kleinberg. ACM, pp.87–96. http://doi.acm.org/10.1145/2422436.2422449.
2. Kasiviswanathan, S. P., K. Nissim, S. Raskhodnikova, and A. D. Smith (2013). Analyzing Graphs with Node Differential Privacy. In: *Theory of Cryptography - 10th Theory of Cryptography Conference, TCC 2013, Tokyo, Japan, March 3-6, 2013. Proceedings*. Ed. by A. Sahai. Vol. 7785. Lecture Notes in Computer Science. Springer, pp.457–476. http://dx.doi.org/10.1007/978-3-642-36594-2_26.
3. Raskhodnikova, S. and A. D. Smith (2016). Lipschitz Extensions for Node-Private Graph Statistics and the Generalized Exponential Mechanism. In: *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*. Ed. by I. Dinur. IEEE Computer Society, pp.495–504. http://dx.doi.org/10.1109/FOCS.2016.60.

Miscellaneous

1. Blocki, J., A. Blum, A. Datta, and O. Sheffet (2012). The Johnson-Lindenstrauss Transform Itself Preserves Differential Privacy. In: *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*. IEEE Computer Society, pp.410–419. http://dx.doi.org/10.1109/FOCS.2012.67.
2. Muthukrishnan, S. and A. Nikolov (2012). Optimal private halfspace counting via discrepancy. In: *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*. Ed. by H. J. Karloff and T. Pitassi. ACM, pp.1285–1292. http://doi.acm.org/10.1145/2213977.2214090.
3. Gaboardi, M., E. J. G. Arias, J. Hsu, A. Roth, and Z. S. Wu (2014). Dual Query: Practical Private Query Release for High Dimensional Data. In: *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*. Vol. 32. JMLR Workshop and Conference Proceedings. JMLR.org, pp.1170–1178. http://jmlr.org/proceedings/papers/v32/gaboardi14.html.