Moses Charikar

Office: Gates Computer Science Bldg
353 Jane Stanford Way
Stanford, CA 94305

Email: moses@cs.stanford.edu
URL: https://profiles.stanford.edu/moses-charikar

RESEARCH INTERESTS

Efficient algorithmic techniques for processing, searching and indexing massive high-dimensional data sets; efficient algorithms for computational problems in high-dimensional statistics and optimization problems in machine learning; approximation algorithms for discrete optimization problems with provable guarantees; convex optimization approaches for non-convex combinatorial optimization problems; low-distortion embeddings of finite metric spaces.

EDUCATION

1995 - 2000 Stanford University, Stanford, CA.
Ph.D. Computer Science.
Advisor: Rajeev Motwani

1991 - 1995 Indian Institute of Technology (IIT), Bombay, India.
Bachelor of Technology, Computer Science and Engineering.

PROFESSIONAL EXPERIENCE

Dec '17 onwards
Donald E. Knuth Professor of Computer Science and Professor, by courtesy, of Mathematics, Stanford University.

August '15 onwards
Professor, Computer Science, Stanford University.

July '11 - July '15
Professor, Computer Science, Princeton University.

July '07 - June '11
Associate Professor, Computer Science, Princeton University.

Sept '01 - June '07
Assistant Professor, Computer Science, Princeton University.

Sept '00 - Aug '01
Research Scientist, Google Inc.

PROFESSIONAL SERVICE

- ACM-SIAM SODA steering committee (2010-2012), SIGACT Committee for the Advancement of Theoretical Computer Science (2011-2017), Workshops chair for STOC 2013, FOCS 2013, Scientific Advisory Board member, Simons Institute for the Theory of Computing (2015-2018), TheoryFest co-chair (2018), TheoryFest chair (2019).

- Program Committee member for APPROX (2001), FOCS (2001), SODA (2003), ESA (2003), FSTTCS (2003), STOC (2004), APPROX (2005), FOCS (2006), SODA (2009), ICS (2011), FOCS (2012), FOCS (2014), FOCS (2016), APPROX (2017), HALG (2018), COLT (2019), SOSA (2020). Program Committee Chair for APPROX (2007), SODA (2010).
- Director, Center for Computational Intractability, 2012-14.
- Co-organizer of Aladdin workshops on "Integrated Logistics" at Princeton (Oct-Nov '02) and CMU (March '03), workshop on "Discrete Metric Spaces and their Applications" at Princeton (Aug '03), Aladdin workshop on "Flexible Network Design" at Princeton (Nov '05), Concentration week on "Metric Geometry and Geometric Embeddings of Discrete Metric Spaces" at Texas A&M (July '06), Mini-course on "Additive Combinatorics" at Princeton (Aug '07), "Women in Theory" workshop at Princeton (June '08, June '10, June '12), workshop on "Geometry and Algorithms" at Princeton (Oct '08), DIMACS Tutorial on Limits of Approximation Algorithms: PCPs and Unique Games (July '09), Co-Chair of Organizing Committee for DIMACS Special Focus on Intractability (2008-2010), Co-organizer of Barriers in Computational Complexity II workshop at Princeton (Aug '10), workshop on Approximation Algorithms: The Last Decade and the Next at Princeton (June '11), Summer School in Theoretical Computer Science at Princeton for high school students (June-Aug '11), STOC 2012 workshop on Recent results regarding the Unique Games Conjecture (May '12), workshop on Provable Bounds in Machine Learning (Aug '12), minicourse on Spectral Methods (June '13), Rising Stars in EECS workshop (2017).

RESEARCH GRANTS

- Amazon Research Award for *Efficient Algorithms for High-Dimensional Statistics* $80,000 (2019).
- Google Research Award for *New Estimators via Locality Sensitive Hashing* $70,000 (2018).
- NSF CCF award 1617577 for *New Perspectives on Mathematical Programming Relaxations*, $450,000 (July '16 - June '19).
- Simons Investigator award, $1,320,000 (August '14 - July '24)
- NSF CCF award 1302518 for *Towards Provable Bounds for Machine Learning* (co-PI Sanjeev Arora), $900,000 (Sep '13 - Aug '17).
- NSF CCF award 1218687 for *Approximation Techniques for Combinatorial Optimization*, $400,000 (Aug '12 - July '16).
- Google Research Award for *Online Bipartite Matching* $70,110 (2011).
- NSF CCF award 0916218 for *Mathematical Programming Methods in Approximation*, $499,996 (Aug '09 - July '12).
- NSF Expeditions award 0832797 for *Understanding, Coping with, and Benefiting from Intractibility* (co-PI's Sanjeev Arora, Bernard Chazelle, Bob Tarjan, Boaz Barak, Avi Wigderson, Russell Impagliazzo, Eric Allender, Mike Saks, Mario Szegedy, Subhash Khot, Assaf Naor), approx $10,000,000 (Aug '08 - July '13).

- Google Research Award for *Efficient Content Based Similarity Search* (co-PI Kai Li) $110,000 (2006), $120,000 (2007).
- Yahoo! Research Alliance Award for *Content-Based Based Similarity Search for Non-Text, Feature-Rich Datasets* (co-PI Kai Li) $100,000 (2006-07).
- NSF MSPA-MCS award 0528414 for *Embeddings of Finite Metric Spaces - A Geometric Approach to Efficient Algorithms* (co-PI's Sanjeev Arora, Bill Johnson, Misha Gromov), $289,998 (Sept '05 - Aug '08).
- NSF CSR-PDOS award 0509447 for *Content-Searchable Storage for Feature-Rich Data* (co-PI's Kai Li, Perry Cook, Olga Troyanskaya), $290,464 (July '05 - June '06).
- NSF IIS award 0414072 for *Constructing an Enhanced version of Wordnet* (co-PI's Christiane Fellbaum, Daniel Osherson, Rob Schapire), $106,000 (Sept '04 - Aug '05)
- NSF CAREER award 0237113 for *Approximation Algorithms: New directions and Techniques*, $400,091 (July '03 - June '09).
- DOE Early Career Principal Investigator Award for *Algorithmic Techniques for Massive Data Sets*, $256,817 (Sept '02 - Aug '05).
- NSF ITR award 0205594 for *New directions in Clustering and Learning* (co-PI's Sanjeev Arora, Amit Sahai and Yoram Singer), $1,530,000 (July '02 - June '07).

AWARDS AND HONORS

- 10 year best paper award, VLDB 2017.
- Best paper award, 30th Annual Conference on Learning Theory (COLT), 2017.
- Distinguished Alumnus Award, IIT Bombay, 2016.
- Simons Investigator in Theoretical Computer Science, 2014.
- ACM Paris Kanellakis Theory and Practice Award, 2012.
- Howard B. Wentz Jr. junior faculty award, 2004.
- Best paper award, 44th IEEE Symposium on Foundations of Computing (FOCS), 2003.
- Alfred P. Sloan Fellowship, 2003.
- Best student paper award, 31st ACM Symposium on Theory of Computing (STOC), 1999.
- Invited speaker: 17th Annual Conference on Learning Theory (July 2004), 14th Annual Fall Workshop on Computational Geometry (Nov 2004), Statistics and Optimization of Clustering Workshop (July 2005), 11th International Conference on Artificial Intelligence and Statistics (March 2007), Rajeev Motwani Distinguished Lecture at Stanford (March 2011), 54th Annual IEEE Symposium on Foundations of Computer Science (October 2013), 35th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (December 2015).

STUDENTS

- **Graduated students:** Bo Brinkman (Associate Professor, Miami University, Ohio), Tony Wirth (Professor, University of Melbourne), Adriana Karagiozova, Konstantin Makarychev (Associate Professor, Northwestern University), Yury Makarychev (Professor, Toyota Technological Institute, Chicago), Mohammad Hossein Bateni (Staff Research Scientist, Google Research), Aditya Bhaskara (Assistant Professor, University of Utah), Aravindan Vijayaraghavan (Assistant Professor, Northwestern University), Shi Li (Assistant Professor, University at Buffalo), Huy Le Nguyen (Assistant Professor, Northeastern University)), Yonatan Naamad (Applied Scientist, Amazon Research), Paris Syminelakis (Research Staff Member, The Voleon Group).

- **Current graduate students:** Ofir Geri, Xian (Carrie) Wu, Kiran Shiragur, Weiyun (Anna) Ma, Neha Gupta, Paul Liu, Joshua Brakensiek, Lunjia Hu, Anna Thomas, June Vuong, Aidan Perrault.

- **Postgraduate students:** Alantha Newman (Fall '04), Eliott Anshelevich ('05-'06), several postdocs at the Intractability Center including Alexandr Andoni ('09-'10), Ofer Neiman ('10-'11), Jelani Nelson ('12-'13), Ankur Moitra ('12-'13), Ravishankar Krishnaswamy ('12-'14), Alina Ene ('13-'14), Pranjal Awasthi ('13-'15), Roy Schwartz ('14-'15), Anand Louis ('14-'16). Shay Solomon ('16-'17), Yuchen Zhang ('16-'18), Rad Niazadeh ('17-'19), Erik Waingarten ('19-'21).

- **Undergraduate students**: Lorenzo Orecchia (Fall '03), Emily Huang (Fall '03), Mike Dinitz (Spring '04), Bryce Liu (Spring '04), Lev Reyzin (Spring '04), Rahul Bhargava (senior thesis, '03-'04), Emily Huang (senior thesis, '04-'05), Matt Stanton (senior thesis, '04-'05), Jon Ullman (Spring '07), Newton Allen (Spring '10), Lavanya Jose (Fall '11), Christian Tessier-Lavigne (Fall '12), Mark Benjamin (Fall '12), Alex Daifotis, Nader Al-Naji (Spring '13), Diogo Adrados (Fall '13), Delaney Granizo-Mackenzie (senior thesis, '13-'14), Aaron Schild (senior thesis, '13'-14), William Kuszmaul (senior honors thesis, '17-'18), Kendall Beache (Autumn 2020), Krithika Iyer (Autumn 2020).

TEACHING

- Courses at Stanford: *Algorithmic Perspective on Machine Learning* (Autumn '15, Autumn '17, Autumn '19), *Advanced Approximation Algorithms* (Winter '16), *Algorithmic Techniques for Big Data* (Spring '16, Spring '18, Spring '20), *Hierarchies of Integer Programming Relaxations* (Spring '17), *Design and Analysis of Algorithms* (Autumn '16, Winter '21), *Optimization and Algorithmic Paradigms* (Winter '18, Winter '20), *Artificial Intelligence: Principles and Techniques* (Spring '19), *Machine Learning* (Autumn '20, Spring '20, Spring '21).

- Courses at Princeton: *Approximation Algorithms* (Fall '01, Spring '03), *Algorithms for Massive Data Sets* (Spring '02), *Discrete Mathematics* (Fall '02, Fall '03, Fall '04, Fall '05), *Reasoning About Computation* (Fall '07, Fall '10, Fall '11, Fall '12, Spring '14), *Advanced Algorithms* (Spring '04, Fall '06, Spring '08, Spring '11, Spring '13), *Algorithms Seminar* (Spring '10, Fall '14), *CS Independent Work* (Fall '04, Spring '05, Fall '05, Spring '06).

- Undergraduate Advising at Princeton: BSE Freshman Advisor ('02-'03, '03-'04, '06-'07, '07-'08), Computer Science '05 AB Advisor.

- Co-taught PhD Summer School on *Finite Metric Spaces and their Algorithmic Applications* at IT University, Copenhagen (Summer '04).

PATENTS

1. *Method and Apparatus for Estimating Similarity*,
United States Patent number 7,158,961, issued January 2, 2007, assigned to Google Inc.

2. *Non-transferable anonymous credentials*,
with R. Canetti, S. Rajagopalan, S. Ravikumar, A. Sahai and A. Tomkins, United States Patent number 7,222,362, issued May 22, 2007, assigned to IBM Corp.

3. *Similarity search system with compact data structures*,
with K. Li and Q. Lv, United States Patent number 7,966,327, issued June 21, 2011, assigned to Princeton University.

4. *Format Identification for Fragmented Data,*
with D. Ramakrishna, United States Patent numbers 9,384,218 (issued July 5, 2016) 9,495,390 (issued Nov 15, 2016), 10,114,839 (issued Oct 30, 2018), assigned to EMC Corp.

5. *Lossless Compression of Fragmented Image Data,*
with D. Ramakrishna, United States Patent numbers 9,558,566 (issued Jan 31, 2017), 9,684,974 (issued June 20, 2017), 10,249,059 (issued April 2, 2019), 10,282,863 (issued May 7, 2019) assigned to EMC Corp.

PUBLICATIONS

Journals

1. *Targeted exploration and analysis of large cross-platform human transcriptomic data compendia* (Brief Communication), with Q. Zhu, A. Krishnan, A. Wong, M. Aure, L. Bongo, V. Kristensen, K. Li, and O. Troyanskaya, in *Nature Methods*, vol. 12, (2015).

2. *Fitting Tree Metrics: Hierarchical Clustering and Phylogeny*,
with N. Ailon, in *SIAM J. Comput.*, vol. 40(5), (2011). (preliminary version in FOCS 2005).

3. *Beating the Random Ordering Is Hard: Every Ordering CSP Is Approximation Resistant*,
with V. Guruswami, J. Hästad, R. Manokaran and P. Raghavendra, in *SIAM J. Comput.*, vol. 40(3), (2011). (preliminary version in CCC 2009).

4. *Improved Approximation Algorithms for Label Cover Problems*,
with M. T. Hajiaghayi and H. J. Karloff, in *Algorithmica* , vol. 61(1), (2011). (special issue for ESA 2009).

5. *Local Global Tradeoffs in Metric Embeddings*,
with K. Makarychev and Y. Makarychev, in *SIAM J. Comput.*, vol. 39(6), (2010). (special issue for FOCS 2007).

6. ℓ_2^2 *Spreading Metrics for Vertex Ordering Problems*,
with M. T. Hajiaghayi, H. J. Karloff and S. Rao, in *Algorithmica*, vol. 56(4), (2010). (preliminary version in SODA '06).

7. *Near-optimal algorithms for maximum constraint satisfaction problems*,
with K. Makarychev and Y. Makarychev, in *ACM Transactions on Algorithms*, vol. 5(3), (2009). (special issue for SODA '07).

8. *Aggregating inconsistent information: Ranking and clustering*,
with N. Ailon and A. Newman, in *Journal of the ACM*, vol. 55(5), (2008). (preliminary version in STOC '05).

9. Special issue for FOCS 2001,
guest editor, *Journal of Computer Systems and Sciences*, vol. 72, issue 5, (2006).

10. *Embedding the Ulam metric into* ℓ_1,
with R. Krauthgamer, *Theory of Computing*, vol. 2, pp. 207-224 (2006).

11. *On the Integrality Ratio for the Asymmetric Traveling Salesman Problem*,
with M. Goemans and H. Karloff, *Math of Operations Research*, vol. 31, pp. 245-252 (2006). (preliminary version in FOCS '04).

12. *On the Impossibility of Dimension Reduction in* ℓ_1,
with B. Brinkman, in *Journal of the ACM*, vol. 52(5), pp. 766-788 (2005). (preliminary version in FOCS '03).

13. *Clustering with Qualitative Information*,
with V. Guruswami and A. Wirth, in *Journal of Computer and System Sciences*, vol. 71(3), pp. 360-383 (2005). (preliminary version in FOCS '03).

14. *The smallest grammar problem*,
with E. Lehman, D. Liu, R. Panigrahy, M. Prabhakaran, A. Sahai and A. Shelat, in *IEEE Transactions on Information Theory*, vol. 51(7), pp. 2554-2576 (2005). (preliminary version in STOC '02).

15. *Improved Combinatorial Algorithms for Facility Location Problems*,
with S. Guha, in *SIAM Journal on Computing*, vol. 34(4), pp. 803–824 (2005). (preliminary version in FOCS '99).

16. *Incremental Clustering and Dynamic Information Retrieval*,
with C. Chekuri, T. Feder and R. Motwani, in *SIAM Journal on Computing*, vol. 33(6), pp. 1417–1440 (2004). (preliminary version in STOC '97).

17. *Approximating Wire Length in Zero and Bounded Skew Clock Trees*,
with J. Kleinberg, S. Rajagopalan, S. Ravikumar, A. Sahai and A. Tomkins, in *SIAM Journal on Discrete Mathematics*, vol. 17(4), pp. 582–595 (2004). (preliminary version in SODA '99).

18. *Resource Optimization in QoS Multicast Routing of Real-Time Multimedia*,
with J. Naor and B. Schieber, in *IEEE/ACM Transactions on Networking*, vol.12(2), pp. 340–348 (2004). (preliminary version in INFOCOM '00).

19. *Finding Frequent Items in Data Streams*,
with K. Chen and M. Farach-Colton, in *Theoretical Computer Science*, vol. 312(1), pp. 3-15 (2004) (special issue for ICALP 2002).

20. *Clustering to Minimize the Sum of Cluster Diameters*,
with R. Panigrahy, in *Journal of Computer Systems and Sciences*, vol. 68(2), pp. 417-441 (2004) (special issue for STOC '01).

21. *A derandomization using min-wise independent permutations*,
with A. Broder, and M. Mitzenmacher, in *J. Discrete Algorithms*, vol. 1(1), pp. 11–20 (2003).

22. *A Constant Factor Approximation Algorithm for the k-Median Problem*,
with S. Guha, E. Tardos and D. Shmoys, in *Journal of Computer Systems and Sciences*, vol. 65(1), pp. 129-149 (2002) (special issue for STOC '99).

23. *Query Strategies for Priced Information*,
with R. Fagin, V. Guruswami, J. Kleinberg, P. Raghavan and A. Sahai, in *Journal of Computer Systems and Sciences*, vol. 64, pp. 785-819 (2002). (special issue for for STOC '00).

24. *Algorithms for Capacitated Vehicle Routing*,
with S. Khuller and B. Raghavachari, in *SIAM Journal on Computing*, vol. 31(3), pp. 665-682 (2001). (preliminary version in STOC '98).

25. *Delayed Information and Action in On-Line Algorithms*,
with S. Albers and M. Mitzenmacher, in *Information and Computation* vol. 170, pp. 135-152 (2001) (preliminary version in FOCS '98).

26. *Min-Wise Independent Permutations*,
with A. Broder, A. Frieze and M. Mitzenmacher, in *Journal of Computer Systems and Sciences*, vol. 60(3), pp. 630-659 (2000) (special issue for STOC '98).

27. *On-line Load Balancing for Related Machines*,
with P. Berman and M. Karpinski, in *Journal of Algorithms*, vol. 35(1), pp. 108-121 (2000) (preliminary version in WADS '97).

28. *On Page Migration and Other Relaxed Task Systems*,
with Y. Bartal and P. Indyk, in *Theoretical Computer Science*, vol. 268(1), pp. 43-66 (2001) (preliminary version in SODA '97).

29. *Approximation Algorithms for Directed Steiner Tree Problems*,
with C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha and M. Li, in *Journal of Algorithms*, vol. 33, pp. 73-91 (1999). (preliminary version in SODA '98).

Conference Proceedings

30. *Improved Algorithms for Edge Colouring in the W-Streaming Model*, with P. Liu, in *Proceedings of the 4th SIAM Symposium on Simplicity in Algorithms* (SOSA 2021).

31. *A Model for Ant Trail Formation and its Convergence Properties*, with S. Garg, D. M. Gordon and K. Shiragur, in *Proceedings of the 12th Innovations in Theoretical Computer Science Conference* (ITCS 2021).

32. *Instance Based Approximations to Profile Maximum Likelihood*,
with N. Anari, K. Shiragur, A. Sidford, in *Proceedings of the Advances in Neural Information Processing Systems 33* (NeurIPS 2020).

33. *Kernel Density Estimation through Density Constrained Near Neighbor Search*,
with M. Kapralov, N. Nouri and P. Syminelakis, in *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science* (FOCS 2020).

34. *CoopStore: Optimizing Precomputed Summaries for Aggregation*,
with E. Gan and P. Bailis, in *Proceedings of the VLDB Endowment*, Volume 13 (VLDB 2020).

35. *Unconditional Lower Bounds for Adaptive Massively Parallel Computation*,
with W. Ma and L-Y. Tang, in *Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures* (SPAA 2020).

36. *Institutions Share Successes, Failures, and Advice in Moving the Diversity Needle*,
with D. Garcia, E. Hearn, E. Lazowska and J. Reynolds, in *Proceedings of the 51st ACM Technical Symposium on Computer Science Education* (SIGCSE 2020).

37. *Retrieving Top Weighted Triangles in Graphs*,
with R. Kumar, P. Liu and A.R. Benson, in *Proceedings of the Thirteenth ACM International Conference on Web Search and Data Mining* (WSDM 2020).

38. *Adaptive Discrete Phase Retrieval*,
with X. Wu and Y. Ye
in *Proceedings of the 3rd SIAM Symposium on Simplicity in Algorithms* (SOSA 2020).

39. *A General Framework for Symmetric Property Estimation*,
with K. Shiragur and A. Sidford, in *Proceedings of the Advances in Neural Information Processing Systems 32* (NeurIPS 2019).

40. *Multi-resolution Hashing for Fast Pairwise Summations*,
with P. Syminelakis, in *Proceedings of the 60th IEEE Symposium on Foundations of Computer Science* (FOCS 2019).

41. *Hierarchical Clustering for Euclidean Data*,
with V. Chatziafratis, R. Niazadeh and G. Yaroslavtsev, in *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics* (AISTATS 2019).

42. *Recovery Guarantees For Quadratic Tensors With Sparse Observations*,
with H. Zhang, V. Sharan and Y. Liang, in *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics* (AISTATS 2019).

43. *The One-Way Communication Complexity of Dynamic Time Warping Distance*,
with V. Braverman, W. Kuszmaul, D. P. Woodruff and L. F. Yang, in *35th International Symposium on Computational Geometry* (SoCG 2019), invited to special issue.

44. *Rehashing Kernel Evaluation in High Dimensions*,
with P. Siminelakis, K. Rong, P. Bailis and P. Levis, in *Proceedings of the 36th International Conference on Machine Learning* (ICML 2019).

45. *Hierarchical Clustering better than Average-Linkage*,
with V. Chatziafratis and R. Niazadeh, in *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2019).

46. *Efficient profile maximum likelihood for universal symmetric property estimation*,
with K. Shiragur and A. Sidford, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (STOC 2019).

47. *Sampling Methods for Counting Temporal Motifs*,
with P. Liu and A. R. Benson, in *Proceedings of the 12th ACM International Conference on Web Search and Data Mining* (WSDM 2019).

48. *Efficient Density Evaluation for Smooth Kernels*,
with A. Backurs, P. Indyk, and P. Siminelakis, to appear in *Proceedings of the 59th IEEE Symposium on Foundations of Computer Science* (2018).

49. *Multi-Commodity Flow with In-Network Processing*,
with Y. Naamad, R. Rexford and X. Zhou, in *Proceedings of the 4th International Symposium on Algorithmic Aspects of Cloud Computing* (2018).

50. *Fully Dynamic Almost-Maximal Matching: Breaking the Polynomial Worst-Case Time Barrier*,
with S. Solomon, in *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming* (2018).

51. *On Estimating Edit Distance: Alignment, Dimension Reduction, and Embeddings*,
with O. Geri, M. Kim and W. Kuszmaul, in *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming* (2018).

52. *Hierarchical Clustering with Structural Constraints*,
with V. Chatziafratis and R. Niazadeh, in *Proceedings of the 35th International Conference on Machine Learning* (2018).

53. *Local Density Estimation in High Dimensions*,
with X. Wu and V. Natchu, in *Proceedings of the 35th International Conference on Machine Learning* (2018).

54. *Resilience: A Criterion for Learning in the Presence of Arbitrary Outliers*,
with J. Steinhardt and G. Valiant, in *Proceedings of the 9th Innovations in Theoretical Computer Science* (2018).

55. *Hashing-Based-Estimators for Kernel Density in High Dimensions*,
with P. Siminelakis, in *Proceedings of the 58th Symposium on Foundations of Computer Science* (2017).

56. *Min-Cost Bipartite Perfect Matching with Delays*,
with I. Ashlagi, Y. Azar, A. Chiplunkar, O. Geri, H. Kaplan, R. Makhijani, Y. Wang, R. Wattenhofer, in *Proceedings of the 20th APPROX* (2017).

57. *Local Guarantees in Graph Cuts and Clustering*,
with N. Gupta and R. Schwartz, in *Proceedings of the 19th International Conference on Integer Programming and Combinatorial Optimization* (2017).

58. *A Hitting Time Analysis of Stochastic Gradient Langevin Dynamics*,
with Y .Zhang and P. Liang, in *Proceedings of the 30th Conference on Learning Theory* (COLT 2017), **best paper award**.

59. *Learning with Untrusted Data*,
with J. Steinhardt and G. Valiant, in *Proceedings of the 49th ACM Symposium on Theory of Computing* (2017).

60. *Approximate Hierarchical Clustering via Sparsest Cut and Spreading Metrics*,
with V. Chatziafratis, in *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms* (2017).

61. *Avoiding Imposters and Delinquents: Adversarial Crowdsourcing and Peer Prediction*,
with J. Steinhardt and G. Valiant, in *Proceedings of the 30th Annual Conference on Neural Information Processing Systems* (2016).

62. *On Approximating Target Set Selection*,
with Y. Naamad and A. Wirth, in *Proceedings of the 24th Annual European Symposium on Algorithms* (2016).

63. *Spectral Embedding of k-Cliques, Graph Partitioning and k-Means*,
with P. Awasthi, R. Krishnaswami, and A. Sinop, in *Proceedings of the 7th Innovations in Theoretical Computer Science conference* (2016).

64. *The Hardness of Approximation of Euclidean k-Means*,
with P. Awasthi, R. Krishnaswami, and A. Sinop, in *Proceedings of the 31st International Symposium on Computational Geometry* (2015).

65. *Label optimal regret bounds for online local learning*,
with P. Awasthi, K. Lai, and A. Risteski, in *Proceedings of the 28th Conference on Learning Theory* (2015).

66. *Relax, no need to round: integrality of clustering formulations*,
with P. Awasthi, A. Bandeira, R. Krishnaswami, S. Villar and R. Ward, in *Proceedings of the 6th Innovations in Theoretical Computer Science conference* (2015).

67. *Online Bipartite Matching with Decomposable Weights*,
with M. Henzinger and H. Nguyen, in *Proceedings of the 22nd Annual European Symposium on Algorithms* (2014).

68. *Uniqueness of Tensor Decompositions with Applications to Polynomial Identifiability*,
with A. Bhaskara and A. Vijayaraghavan, in *Proceedings of the 27th Annual Conference on Learning Theory* (2014).

69. *Smoothed Analysis of Tensor Decompositions,*
with A. Bhaskara, A. Moitra and A. Vijayaraghavan, in *Proceedings of the 46th ACM Symposium on Theory of Computing* (2014).

70. *Multireference Alignment using Semidefinite Programming,*
with A. Bandeira, A. Singer and A. Zhu, in *Proceedings of the 5th Innovations in Theoretical Computer Science* (2014).

71. *Better Algorithms and Hardness for Broadcast Scheduling via a Discrepancy Approach*,
with N. Bansal, R. Krishnaswamy and S. Li, in *Proceedings of the 25nd Annual ACM-SIAM Symposium on Discrete Algorithms* (2014).

72. *A Dependent LP-rounding Approach for the k-Median Problem*,
with S. Li, in *Proceedings of the 39th International Colloquium on Automata, Languages and Programming* (2012).

73. *On Quadratic Programming with a Ratio Objective*,
with A. Bhaskara, and A. Vijayaraghavan, in *Proceedings of the 39th International Colloquium on Automata, Languages and Programming* (2012).

74. *High-confidence near-duplicate image detection*,
with W. Dong, Z. Wang and K. Li, in *International Conference on Multimedia Retrieval* (2012).

75. *Polynomial integrality gaps for strong SDP relaxations of Densest k-subgraph*,
with A. Bhaskara, V. Guruswami, A. Vijayaraghavan, and Y. Zhou, in *Proceedings of the 23nd Annual ACM-SIAM Symposium on Discrete Algorithms* (2012).

76. *Near Linear Lower Bound for Dimension Reduction in ℓ_1*,
with A. Andoni, O. Neiman and H. Nguyen, in *Proceedings of the 52nd Annual IEEE Conference on Foundations of Computer Science* (2011).

77. *Efficient k-nearest neighbor graph construction for generic similarity measures*,
with W. Dong and K. Li, in *Proceedings of the 20th International Conference on World Wide Web WWW* (2011).

78. *Tight Hardness Results for Minimizing Discrepancy*,
with A. Newman and A. Nikolov, in *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms* (2011).

79. *Vertex Sparsifiers and Abstract Rounding Algorithms*,
with T. Leighton, S. Li and A. Moitra, in *Proceedings of the 51st Annual IEEE Conference on Foundations of Computer Science* (2010).

80. *Detecting high log-densities: an $O(n^{1/4})$ approximation for densest k-subgraph*,
with A. Bhaskara, E. Chlamtac, U. Feige and A. Vijayaraghavan, in *Proceedings of the 42nd ACM Symposium on Theory of Computing* (2010).

81. *Improved Approximation Algorithms for Label Cover Problems*,
with M. Hajiaghayi and H. Karloff, in *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, pp. 23-34 (2009).

82. *Every Permutation CSP of arity 3 is Approximation Resistant*,
with V. Guruswami and R. Manokaran, in *Proceedings of the 24th IEEE Conference on Computational Complexity (CCC)*, pp. 62-73 (2009).

83. *Integrality Gaps for Sherali-Adams Relaxations*,
with K. Makarychev and Y. Makarchev, in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*,pp. 283-292 (2009).

84. *MaxMin Allocation via Degree Lower-Bounded Arborescences*,
with M. H. Bateni and V. Guruswami, in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*,pp. 543-552 (2009).

85. *Efficiently matching sets of features with random histograms*,
with W. Dong, Z. Wang, and K. Li, in *Proceedings of the 16th International ACM Conference on Multimedia*, pp. 179-188 (2008).

86. *Modeling LSH for performance tuning* ,
with W. Dong, Z. Wang, W. Josephson, and K. Li, in *Proceedings of the 17th ACM Conference on Information and Knowledge Management (CIKM)* pp. 669-678 (2008).

87. *Asymmetric distance estimation with sketches for similarity search in high-dimensional spaces*,
with W. Dong and K. Li, in *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 123-130, (2008).

88. *Online multicast with egalitarian cost sharing*,
with H. J. Karloff, C. Mathieu, J. Naor, M. E. Saks, in *Proceedings of the 20th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 70-76, (2008).

89. *Local Global Tradeoffs in Metric Embeddings.*
with K. Makarychev and Y. Makarychev, in *Proceedings of the 48th Annual IEEE Conference on Foundations of Computer Science* (2007).

90. *On the Advantage over Random for Maximum Acyclic Subgraph.*
with K. Makarychev and Y. Makarychev, in *Proceedings of the 48th Annual IEEE Conference on Foundations of Computer Science* (2007).

91. *Multi-Probe LSH: Efficient Indexing for High-Dimensional Similarity Search.*
with Q. Lv, W. Josephson, Z. Wang, and K. Li, in *Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB)* (2007), **10 year best paper award**.

92. *Sizing sketches: a rank-based analysis for similarity search.*
with Z. Wang, W. Dong, W. Josephson, Q. Lv, and K. Li, in *ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems* (2007).

93. *Improved approximation for directed cut problems.*
with A. Agarwal and N. Alon, in *Proceedings of the 39th Annual ACM Symposium on Theory of Computing* (2007).

94. *Near-Optimal Algorithms for Maximum Constraint Satisfaction Problems*,
with K. Makarychev and Y. Makarychev, in *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms* (2007).

95. *A Divide and Conquer Algorithm for d-Dimensional Arrangement*,
with K. Makarychev and Y. Makarychev, in *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms* (2007).

96. *Near Optimal Approximation Algorithms for Unique Games*,
with K. Makarychev and Y. Makarychev, in *Proceedings of the 38th Annual ACM Symposium on Theory of Computing* (2006).

97. *New Approximation Guarantees for Chromatic Number*,
with S. Arora and E. Chlamtac, in *Proceedings of the 38th Annual ACM Symposium on Theory of Computing* (2006).

98. *Ferret: A Toolkit for Content-Based Similarity Search*,
with Christine Lv, William Josephson, Zhe Wang and Kai Li, in *Proceedings of ACM SIGOS EuroSys Conference* (2006).

99. *ℓ_2^2 Spreading Metrics for Vertex Ordering Problems*,
with M. Hajiaghayi, H. Karloff and S. Rao, in *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms* (2006).

100. *Directed Metrics and Directed Graph Partitioning Problems*,
with K. Makarychev and Y. Makarychev, in *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms* (2006).

101. *A Robust Maximum Completion Time Measure for Scheduling*,
with S. Khuller, in *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms* (2006).

102. *Fitting tree metrics: Hierarchical clustering and Phylogeny*,
with N. Ailon, in *Proceedings of the 46th Annual IEEE Conference on Foundations of Computer Science* (2005).

103. *On non-uniform multicommodity buy-at-bulk network design*,
with A. Karagiozova. in *Proceedings of the 37th Annual ACM Symposium on Theory of Computing* (2005).

104. *$O(\sqrt{\log n})$ approximation algorithms for min UnCut, min 2CNF deletion, and directed cut problems*,
with A. Agarwal, K. Makarychev and Y. Makarychev, in *Proceedings of the 37th Annual ACM Symposium on Theory of Computing* (2005).

105. *Aggregating inconsistent information: ranking and clustering*,
with N. Ailon and A. Newman, in *Proceedings of the 37th Annual ACM Symposium on Theory of Computing* (2005).

106. *Sampling Bounds for Stochastic Optimization*,
with C. Chekuri and M. Pál, in *Proceedings of the 9th International Workshop on Randomization and Computation RANDOM* (2005).

107. *A tight threshold for metric Ramsey phenomena*,
with A. Karagiozova. in *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms* (2005).

108. *Approximating the average response time in broadcast scheduling*,
with N. Bansal, S. Khanna and J. Naor, in *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms* (2005).

109. *Image Similarity Search with Compact Data Structures*,
with Q. Lv and K. Li, in *Proceedings of the 13th ACM Conference on Information and Knowledge Management* (2004).

110. *On the Advantage of Network Coding for Improving Network Throughput*,
with A. Agarwal, in *Proceedings of IEEE Information Theory Workshop*, (2004).

111. *On the Integrality Ratio for Asymmetric TSP*,
with M. Goemans and H. Karloff. in *Proceedings of the 45th Annual IEEE Conference on Foundations of Computer Science* (2004).

112. *Maximizing Quadratic Programs: Extending Grothendieck's Inequality*,
with A. Wirth, in *Proceedings of the 45th Annual IEEE Conference on Foundations of Computer Science* (2004).

113. *On the Impossibility of Dimension Reduction in ℓ_1*,
with B. Brinkman, in *Proceedings of the 44th Annual IEEE Conference on Foundations of Computer Science* (2003), **best paper award**.

114. *Clustering with Qualitative Information*,
with V. Guruswami and A. Wirth, in *Proceedings of the 44th Annual IEEE Conference on Foundations of Computer Science* (2003).

115. *Better Streaming Algorithms for Clustering Problems*,
with L. O'Callaghan and R. Panigrahy, in *Proceedings of the 35th Annual ACM Symposium on Theory of Computing* (2003).

116. *Dimension Reduction in the ℓ_1 Norm*,
with A. Sahai, in *Proceedings of the 43rd Annual IEEE Conference on Foundations of Computer Science* (2002).

117. *Finding Frequent Items in Data Streams*,
with K. Chen and M. Farach-Colton, in *Proceedings of the 29th International Colloquium on Automata Languages and Programming*, (2002).

118. *New Algorithms for Subset Query, Partial Match, Orthogonal Range Searching and Related Problems*,
with P. Indyk and R. Panigrahy, in *Proceedings of the 29th International Colloquium on Automata Languages and Programming*, (2002).

119. *Similarity Estimation Techniques from Rounding Algorithms*,
in *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, (2002), invited to special issue of Discrete and Computational Geometry.

120. *Approximating The Smallest Grammar: Kolmogorov Complexity in Natural Models*,
with E. Lehman, D. Liu, R. Panigrahy, M. Prabhakaran, A. Rasala, A. Sahai, and A. Shelat, in *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, (2002).

121. *On Semidefinite Programming Relaxations for Graph Coloring and Vertex Cover*,
in *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms* (2002).

122. *Clustering to Minimize the Sum of Cluster Diameters*,
with R. Panigrahy, in *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, (2001).

123. *Approximating Min-Sum k-Clustering in Metric Spaces*,
with Y. Bartal and D. Raz, in *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, (2001).

124. *Algorithms for Facility Location with Outliers*,
with S. Khuller, D. Mount and G. Narasimhan, in *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms* (2001).

125. *Combinatorial Feature Selection Problems*,
with V. Guruswami, S. Rajagopalan, S. Ravikumar and A. Sahai, in *Proceedings of the 41st Annual IEEE Conference on Foundations of Computer Science* (2000).

126. *Greedy Approximation Algorithms for Finding Dense Components in Graphs*,
in *Proceedings of APPROX* (2000).

127. *Minimum Outage Transmission over Fading Channels with Delay Constraint*,
with R. Negi and J. Cioffi, in *Proceedings of the IEEE Intnl. Conference on Communication* (2000).

128. *Towards Estimation Error Guarantees for Distinct Values* with S. Chaudhuri, R. Motwani and V. Narasayya, in *Proceedings of the 19th ACM Symposium on Principles of Database Systems* (2000).

129. *Query Strategies for Priced Information*,
with R. Fagin, V. Guruswami, J. Kleinberg, P. Raghavan and A. Sahai, in *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing* (2000).

130. *Resource Optimization in QoS Multicast Routing of Real-Time Multimedia*,
with J. Naor and B. Schieber, in *Proceedings of the 19th Annual IEEE INFOCOM* (2000).

131. *Improved Combinatorial Algorithms for the Facility Location and k-Median Problems*,
with S. Guha, in *Proceedings of the 40th Annual IEEE Conference on Foundations of Computer Science* (1999).

132. *A Constant Factor Approximation Algorithm for the k-Median Problem*,
with S. Guha, E. Tardos and D. Shmoys, in *Proceedings of the 31st Annual ACM Symposium on Theory of Computing* (1999), student submission received **best student paper award**.

133. *On Targeting Markov Segments*,
with P. Raghavan, S. Rajagopalan, S. Ravikumar, and A. Tomkins, in *Proceedings of the 31st Annual ACM Symposium on Theory of Computing* (1999).

134. *Approximating Wire Length in Zero and Bounded Skew Clock Trees*,
with J. Kleinberg, S. Rajagopalan, S. Ravikumar, A. Sahai and A. Tomkins, in *Proceedings of the 10th Annual ACM-SIAM Symp. on Discrete Algorithms* (1999).

135. *A Derandomization Using Min-Wise Independent Permutations*,
with A. Broder and M. Mitzenmacher, in *Proceedings of the 2nd International Workshop on Randomization and Approximation Techniques in Computer Science* (1998).

136. *Approximating a Finite Metric by a Small Number of Tree Metrics*,
with C. Chekuri, A. Goel, S. Guha and S. Plotkin, in *Proceedings of the 39th Annual IEEE Conference on Foundations of Computer Science* (1998).

137. *Delayed Information and Action in On-Line Algorithms*,
with S. Albers and M. Mitzenmacher, in *Proceedings of the 39th Annual IEEE Conference on Foundations of Computer Science* (1998).

138. *The Finite Capacity Dial-A-Ride Problem*,
with B. Raghavachari, in *Proceedings of the 39th Annual IEEE Conference on Foundations of Computer Science* (1998).

139. *Min-Wise Independent Permutations*,
with A. Broder, A. Frieze and M. Mitzenmacher, in *Proceedings of the 30th Annual ACM Symposium on Theory of Computing* (1998).

140. *Rounding Via Trees: Deterministic Approximation Algorithms for Group Steiner Trees and k-Median*,
with C. Chekuri, A. Goel and S. Guha, in *Proceedings of the 30th Annual ACM Symposium on Theory of Computing* (1998).

141. *Algorithms for Capacitated Vehicle Routing*,
with S. Khuller and B. Raghavachari, in *Proceedings of the 30th Annual ACM Symposium on Theory of Computing* (1998).

142. *Approximation Algorithms for Directed Steiner Tree Problems*,
with C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha and M. Li, in *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms* (1998).

143. *The Dynamic Servers Problem*,
with D. Halperin and R. Motwani, in *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms* (1998).

144. *On-line Load Balancing for Related Machines*,
with P. Berman and M. Karpinski, in *Proceedings of Workshop on Algorithms and Data Structures (WADS)*, (1997).

145. *Constrained TSP and Low Power Computing*,
with C. Silverstein, R. Motwani and P. Raghavan, in *Proceedings of Workshop on Algorithms and Data Structures (WADS)*, (1997).

146. *Incremental Clustering and Dynamic Information Retrieval*,
with C. Chekuri, T. Feder and R. Motwani, in *Proceedings of the 29th Annual ACM Symposium on Theory of Computing* (1997).

147. *On Page Migration and Other Relaxed Task Systems*,
with Y. Bartal and P. Indyk, in *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms* (1997).