

INDERJIT S. DHILLON

Gottesman Family Centennial Professor
Director, Center for Big Data Analytics
Department of Computer Science, The University of Texas at Austin
2317 Speedway, Suite 2.302, Stop D9500, Austin, TX 78712-1757
Office: GDC 4.704
Phone: 512-471-9725
E-mail: inderjit@cs.utexas.edu
WWW: <http://www.cs.utexas.edu/users/inderjit>

EDUCATION

Ph.D. University of California at Berkeley - May 1997.

Major: Computer Science
Thesis: *A New $O(n^2)$ Algorithm for the Symmetric Tridiagonal Eigenvalue/Eigenvector Problem*
Advisors: Profs. Beresford N. Parlett and James W. Demmel
Minors: Mathematics and Theoretical Computer Science.

B. Tech. Indian Institute of Technology, Bombay, India - April 1989.

Major: Computer Science and Engineering
Thesis: *Parallel Architectures for Sparse Matrix Computations*
Advisors: Prof. S. Biswas and Dr. N. K. Karmarkar

RESEARCH INTERESTS

Large Language Models, Pre-Training & Post-Training Gemini Models, Deep Learning, Machine learning, Big Data, Statistical Pattern Recognition, Data Mining, Numerical Linear Algebra, Social Network Analysis, Numerical Optimization, Bioinformatics.

RESEARCH EXPERIENCE

05/25–: Vice President and Google Fellow, Google, Mountain View, CA.
09/09–: Professor, Departments of Computer Science & Mathematics, The University of Texas, Austin.
12/21-04/25 Distinguished Scientist and Area Tech Lead (ATL), Google, Mountain View, CA.
09/20-12/21: Vice President and Distinguished Scientist, Amazon, Berkeley & Palo Alto, CA.
06/18-08/20: Amazon Fellow and Head, Amazon Research Lab, Berkeley, CA.
02/17-06/18: Amazon Fellow, A9/Amazon, Berkeley & Palo Alto, CA.
01/16-02/17: Principal Member of Research Staff, Voleon Capital Management, Berkeley, CA.
09/99-present: Member, Institute for Comp. Engg. & Sciences (ICES), The University of Texas, Austin.
09/05-08/09: Associate Professor, Department of Computer Science, The University of Texas, Austin.
09/07-12/07: Senior Research Fellow, Institute of Pure & Applied Mathematics (IPAM), UCLA.
09/99-08/05: Assistant Professor, Department of Computer Science, The University of Texas, Austin.
11/97-08/99: Researcher, IBM Almaden Research Center, San Jose, CA.
05/97-10/97: Post-Doctoral Scholar, EECS Department, University of California at Berkeley, AND
08/91-04/97: Graduate Student Researcher, EECS Department, University of California at Berkeley.
9/89-8/91: Member of Technical Staff, Math Sciences Research Center, AT&T Bell Labs, Murray Hill, NJ.

HONORS & AWARDS

2022: **AAAI Fellow** for “significant contributions to machine learning and large-scale data analysis”.
2021: **AAAI Best Paper Runner-up Award** for the paper “Learning from eXtreme Bandit Feedback” (with R. Lopez and Mike Jordan) at the Thirty-Fifth AAAI Conference.
2020: **AI 2000 Machine Learning Most Influential Scholars** Honorable Mention for “top most cited scholars from the top publication venues in Machine Learning over the past 10 years (2009–2019)”.

2016: AAAS Fellow for “contributions to large-scale data analysis and computational mathematics”.

2015: Alcalde’s Texas 10, “selected in annual list of UT’s most inspiring professors as nominated by alumni”.

2014: ACM Fellow for “contributions to large-scale data analysis, machine learning and computational mathematics”.

2014: Gottesman Family Centennial Professor, The University of Texas at Austin.

2014: SIAM Fellow for “contributions to numerical linear algebra, data analysis, and machine learning”.

2014: IEEE Fellow for “contributions to large-scale data analysis and computational mathematics”.

2013: ICES Distinguished Research Award, The University of Texas at Austin.

2012: Best Paper Award at the IEEE Int’l Conference on Data Mining for the paper “Scalable Coordinate Descent Approaches to Parallel Matrix Factorization for Recommender Systems”.

2011: SIAM Outstanding Paper Prize for the journal paper “The Metric Nearness Problem”. The prize is for “outstanding papers published in SIAM journals during the 3 years prior to year of award.”

2010-2011: Moncrief Grand Challenge Award, The University of Texas at Austin.

2014, 2011, 2008, 2005, 2002 & 1999: Plenary talks at the XVIII, XVII, XVI, XV and XIV Householder Symposia on Numerical Linear Algebra (Spa in Belgium, Tahoe City in California, Zeuthen in Germany, Campion in Pennsylvania, Peebles in Scotland & Whistler in Canada).

2002-2010: Faculty Fellowship, Dept of Computer Science, The University of Texas at Austin.

2006: SIAM Linear Algebra Prize for the journal paper “Orthogonal Eigenvectors and Relative Gaps”. The award is for “the most outstanding paper on a topic in applicable linear algebra published in English in a peer-reviewed journal in the three calendar years preceding the year of the award.”

Spring 2006: Dean’s Fellowship, The University of Texas at Austin.

2005: University Cooperative Society’s Research Excellence Award for Best Research Paper for “Clustering with Bregman Divergences”.

2007 & 2005: Best Paper Awards at ICML at the 24th Int’l Conference on Machine Learning for the paper “Information-Theoretic Metric Learning”, and at the 22nd Int’l Conference on Machine Learning for the paper “Semi-Supervised Graph-Based Clustering: A Kernel Approach”.

2004: Best Paper Award at the Third SIAM Int’l Conference on Data Mining for the paper “Clustering with Bregman Divergences”.

2001: NSF CAREER Award for the period 2001-2006.

1999: Householder Award for the *Best Dissertation in Numerical Linear Algebra* for the period 1996-1998, Honorable Mention.

Fall 1996-Spring 1997: Graduate Research Fellowship from Pacific Northwest National Laboratory (PNNL).

1985-1989: Ranked 2nd (in a class of over 300) at Indian Institute of Technology, Bombay.

PHD STUDENTS & POSTDOCS

Current Ph.D. Students: Sai Surya Duvvuri, Nilesh Gupta, Devvrit Khatri, and Nirmal Patel.

Former Postdocs: Abolfazl Hashemi, 2020-21 (now Assistant Professor, Purdue University),
 Nikhil Rao, 2014-2016 (Principal Researcher, Google),
 Piyush Rai, 2012-2013 (Professor, IIT Kanpur),
 Ambuj Tewari, 2010-2012 (Professor, University of Michigan),
 Zhengdong Lu, 2008-2010 (Researcher, Huawei Noah’s Ark Lab, Hong Kong),
 Berkant Savas, 2009-2011 (Linköping Univ., Sweden).

Former Ph.D. Students: Joel Tropp in 2004 (now Professor, Caltech, Pasadena),
Yuqiang Guan in 2005 (Google, LA),
Suvrit Sra in 2007 (Associate Professor, MIT, Boston),
Jason Davis in 2008 (former Director of Data & Search, Etsy Inc.),
Hyuk Cho in 2008 (Professor, Sam Houston State Univ.),
Brian Kulis in 2008 (Associate Professor, Boston University),
Prateek Jain in 2009 (Distinguished Research Scientist, Google Deepmind, Bangalore),
Mátyás Sustik in 2013 (Researcher, Walmart Labs, California),
Cho-Jui Hsieh in 2015 (Associate Professor, UCLA),
Nagarajan Natarajan in 2015 (Researcher, Microsoft Research, Bangalore),
Joyce Whang in 2015 (Assistant Professor, KAIST, Korea),
Si Si in 2016 (Staff Research Scientist, Google Deepmind),
Hsiang-Fu Yu in 2016 (Senior Staff Research Scientist, Google),
Kai-Yang Chiang in 2017 (Senior Software Engineer, Google),
Donghyuk Shin in 2017 (Associate Professor, Arizona State University),
David Inouye in 2017 (Assistant Professor, Purdue University),
Ian Yen in 2018 (Researcher, Snap, Inc.),
Kai Zhong in 2018 (Senior Applied Scientist, Amazon),
Jiong Zhang in 2020 (Staff Research Scientist, Google),
Qi Lei in 2020 (Assistant Professor, NYU),
Rudrajit Das in 2024 (Researcher, Google Research, NY),
Anish Acharya in 2025 (Senior Applied Scientist, Amazon, Austin),
Luwen Huang in 2025.

REPRESENTATIVE ACTIVITIES

- Neural Information Processing Systems Conference (NeurIPS) — Senior Area Chair: 2026 (Sydney, Australia), 2025 (San Diego, CA), 2024 (Vancouver, Canada), 2023 & 2022 (New Orleans), 2018 (Montreal, Canada), 2017 (Long Beach, CA), Area Chair: 2021 (Virtual), 2020 (Virtual), 2015 (Montreal, Canada), Reviewer: 2014 (Montreal, Canada), 2011 (Granada, Spain), 2010, 2009, 2008, 2007 & 2006 (Vancouver, Canada).
- Int'l Conference on Machine Learning (ICML) — Senior Area Chair: 2026 (Seoul, South Korea), 2025 (Vancouver, Canada), 2024 (Vienna, Austria), 2022 (Baltimore, USA), 2021 (Virtual), Area Chair: 2018 (Stockholm, Sweden), 2012 (Edinburgh, Scotland), 2010 (Haifa, Israel), Program Committee(PC) Member: 2011 (Bellevue, Washington) & 2007 (Corvallis, Oregon).
- Int'l Conference on Learned Representations (ICLR) — Senior Area Chair: 2026 (Rio de Janeiro, Brazil), 2025 (Singapore), 2024 (Vienna, Austria), Reviewer: 2018 (Vancouver, Canada).
- Organizing Committee Member, SIAM Conference on Applied Linear Algebra (ALA2027), 2026-2027.
- SIAM Fellows Selection Committee Member, 2022-24.
- General Chair, The Conference on Machine Learning and Systems (MLSys), Austin, TX, 2020.
- Board Member, The Machine Learning and Systems Foundation, 2019–now.
- National Advisory Committee Member, The Statistical and Applied Mathematical Sciences Institute (SAMSI), 2016-2022.
- Action Editor, Journal of Machine Learning (JMLR), 2008-2024.
- Associate Editor, IEEE Transactions of Pattern Analysis and Machine Intelligence (TPAMI), 2011-2017.
- Associate Editor, Foundations and Trends in Machine Learning, 2007-present.
- Member of Editorial Board, Machine Learning Journal, 2008-2020.
- Program Committee Co-Chair, KDD (ACM Int'l Conference on Knowledge Discovery & Data Mining), Chicago, 2013.

- Guest Editor, Mathematical Programming Series B, special issue on “Optimization and Machine Learning”, 2008.
- Associate Editor, SIAM Journal for Matrix Analysis and Applications, 2002-2007.
- Householder Prize Committee Member, 2011-2022.
- Organizing Committee Member, IMA (Institute for Mathematics and its Applications) Annual Program on the Mathematics of Information, Sept 2011 - June 2012.
- Selection Committee Member, SIAM Linear Algebra Prize, 2009.
- Served on 2009 NSF Committee of Visitors(COV) in the Formal and Mathematical Foundations Cluster (FMF), 2015 & 2013 NSF SBIR/STTR panels, 2011 NSF panel in the Division of Information and Intelligent Systems (IIS), 2011 NSF panel in Cyber-enabled Discovery and Innovations (CDI), 2010 NSF panel in FMF, 2006 NSF panel in IIS, 2004 NSF panel in FMF, 2001 NSF panel in the Division of Advanced Computational Research (ACR), and 1999 NSF panel in IIS.
- Organizing Committee Member, NSF Workshop on Algorithms in the Field, May 2011.
- Program Chair, IMA Workshop on “Machine Learning: Theory and Computation”, March 2012 (Minneapolis, MA).
- Organizer, Mysore Park Workshop on Machine Learning, August 2012 (Mysore, India).
- AAAI Conference on Artificial Intelligence — Area Chair: 2022 (Vancouver, CA), 2021 (Virtual), Senior Program Committee Member: 2020 (New York, New York).
- SysML Conference (SysML) — PC Member: 2018 (Stanford, CA)
- ACM Int’l Conference on Knowledge Discovery & Data Mining (KDD) — Senior PC Member: 2015 (Sydney, Australia), 2014 (New York), 2012 (Beijing, China), 2009 (Paris, France), 2007 (San Jose, CA), PC Member: 2011 (San Diego), 2008 (Las Vegas, NV), 2006 (Philadelphia, PA), 2005 (Chicago, IL), 2004 (Seattle, WA), 2000 (Boston, MA).
- SIAM Int’l Conference on Data Mining (SDM) — Area Chair: 2010 (Columbus, OH), 2008 (Atlanta, GA), PC Member: 2011 (Mesa, AZ), 2007 (Minneapolis, MN), 2006 (Bethesda, MD), 2005 (Newport Beach, CA), 2004 (Orlando, FL), 2003 (San Francisco, CA), 2002 (Arlington, VA), 2001 (Chicago, IL).
- IEEE Int’l Conference on Data Mining (ICDM) — PC Vice-Chair: 2005 (New Orleans, LA), PC Member: 2004 (Brighton, UK), 2003 (Melbourne, FL).
- Committee Member: 5-year Review of Manoa College of Natural Sciences, University of Hawaii, Oct 2021.
- Committee Member: 5-year Review of Academic Excellence for Department of Electronic Engineering, City University of Hong Kong, June 2016.
- PC Member for the IKDD Conference on Data Science (IKDD CODS): 2016 (Pune, India), 22nd Annual Conference on Learning Theory (COLT): 2009 (Montreal), SIAM Conference on Applied Linear Algebra: 2009 (Seaside, California), World Wide Web Conference (WWW): 2008 (Beijing, China), ACM Conference on Information & Knowledge Management (CIKM): 2011 (Glasgow, Scotland), 2006 (Arlington, VA).
- Program Co-Chair: NIPS workshops on “Multiresolution Methods for Large Scale Learning”, 2015 (Montreal, Canada), “Numerical Mathematics in Machine Learning”, 2010 (Whistler, Canada), ICML workshop on “Covariance Selection & Graphical Model Structure Learning”, 2014 (Beijing, China), Workshops on “Clustering High-Dimensional Data and its Applications” at SIAM Int’l Conference on Data Mining (SDM): 2005 (Newport Beach, CA), 2004 (Orlando, FL), 2003 (San Francisco, CA), 2002 (Arlington, VA), Workshop on “Clustering High-Dimensional Data and its Applications” at IEEE Int’l Conference on Data Mining (ICDM): 2003 (Melbourne, FL), Workshop on “Text Mining” at the Second SIAM Int’l Conference on Data Mining (SDM): 2002 (Arlington, VA).
- Organized invited minisymposium on “Mathematical Methods in Data Mining”, SIAM Annual Meeting, San Diego, CA, July 2008, and invited minisymposium on “Linear Algebra in Data Mining and Information Retrieval”, SIAM Conference on Applied Linear Algebra, Williamsburg, VA, July 2003.

- Ph.D. External Committee Member for Brendon Ames (Univ. of Waterloo, Canada) in 2011, Gilles Meyer (University of Liege, Belgium) in 2011.
- Referee for SIAM Review, SIAM Journal for Matrix Analysis and Applications, SIAM Journal on Scientific Computing, SIAM Journal on Numerical Analysis, Linear Algebra and its Applications, BIT, Proceedings of the National Academy of Sciences (PNAS), Journal of the ACM, Journal of Machine Learning Research (JMLR), Journal of Complex Networks, Internet Mathematics, Data Mining and Knowledge Discovery Journal, AI Review, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), IEEE Transactions on Network Science and Engineering (TNSE) IEEE Transactions on Knowledge and Data Engineering (TKDE), IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB), IEEE Transactions on Signal Processing, IEEE Transactions on Image Processing, Information Processing Letters, Decision Support Systems, ACM Transactions on Internet Computing, International Journal of Neural Systems, etc.

TEACHING

Spring 2025, Spring 2023: Instructor for the graduate course CS391L, “*Machine Learning*”.

Spring 2020, Fall 2011 & 2009, Spring 2008 & 2007: Instructor for the graduate course CS391D, “*Data Mining: A Mathematical Perspective*”.

Fall 2014: Instructor for the graduate course CS395T, “*Scalable Machine Learning*”.

Fall 2014, 2013 & 2012: Instructor for the undergraduate course SSC329C, “*Practical Linear Algebra*”.

Fall 2012, 2008, 2005, 2002 & 1999: Instructor for the graduate breadth course CS383C, “*Numerical Analysis: Linear Algebra*”.

Spring 2012, 2010 & 2009, & Fall 2006: Instructor for the undergraduate course CS378, “*Introduction to Data Mining*”.

Fall 2004, 2003, 2001, Spring 2001 & 2000: Instructor for the graduate topics course CS395T, “*Large-Scale Data Mining*”.

Fall 2004, Spring 2004, 2003, 2002 & Fall 2000: Instructor for the undergraduate course CS323E, “*Elements of Scientific Computing*”.

Spring 1993: Teaching Assistant for CS170, “*Efficient Algorithms and Intractable Problems*”. Instructor - Prof. Manuel Blum.

PUBLICATIONS, TALKS, PATENTS & GRANTS

Google Scholar Profile: Number of Citations = 60,000+; h-index = 103, i10-index = 253. Details available at <http://scholar.google.com/citations?hl=en&user=xBv5ZfkAAAAJ>

Publications in Progress

1. N. Gupta, W. Chang, N. Bui, C.-J. Hsieh and I. S. Dhillon, “LLM-guided Hierarchical Retrieval”, 2025.
2. J. Yen, W. Chang, S. Si, N. Bui, I. S. Dhillon and C.-J. Hsieh, “SLIM: Structure-aware Low-rank Inference Model”, 2025.

Journal Publications

1. G. Comanici, E. Bieber, M. Schaeckermann, I. Pasupat, N. Sachdeva, I. S. Dhillon, et al, “Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities”, arXiv preprint arXiv:2507.06261, 2025.
2. M. Reid, et al, “Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context”, arXiv preprint arXiv:2403.05530, 2024.

3. R. Kumar, A. Mittal, N. Gupta, A. Kusupati, I. S. Dhillon and P. Jain, “EHI: End-to-end learning of Hierarchical Index for Efficient Dense Retrieval”, *Transactions on Machine Learning Research(TMLR)*, 2024.
4. H. Li, H.-F. Yu, L. Ying and I. S. Dhillon, “Accelerating Primal-Dual Methods for Regularized Markov Decision Processes”, *SIAM Journal on Optimization(SIOPT)*, vol. 34(1):764-789, 2024.
5. H. Li, S. Gupta, H.-F. Yu, L. Ying and I. S. Dhillon, “Approximate Newton policy gradient algorithms”, *SIAM Journal on Scientific Computing(SISC)*, vol. 45(5): 2585-2609, 2023.
6. H.-F. Yu, K. Zhong, J. Zhang, W. Chang and I. S. Dhillon, “PECOS: Prediction for Enormous and Correlated Output Spaces”, *Journal of Machine Learning Research(JMLR)*, vol. 23(98), pages 1–32, 2022.
7. A. Hashemi, A. Acharya, R. Das, H. Vikalo, S. Sanghavi and I. S. Dhillon, “On the Benefits of Multiple Gossip Steps in Communication-Constrained Decentralized Federated Learning”, *IEEE Transactions on Parallel and Distributed Systems*, vol. 33:11, pages 2727–2739, 2022.
8. J. Whang, Y. Hou, D. Gleich and I. S. Dhillon, “Non-exhaustive, Overlapping Clustering”, *IEEE Transactions on Pattern Analysis and Machine Intelligence(PAMI)*, vol. 41:11, pages 2644–2659, 2019.
9. K. Chiang, C.-J. Hsieh and I. S. Dhillon, “Using Side Information to Reliably Learn Low-Rank Matrices from Missing and Corrupted Observations”, *Journal of Machine Learning Research(JMLR)*, vol. 19, pages 1–35, 2018.
10. N. Natarajan, I. S. Dhillon, P. Ravikumar and A. Tewari, “Cost-Sensitive Learning with Noisy Labels”, *Journal of Machine Learning Research(JMLR)*, vol. 18, pages 1–33, 2018.
11. P. Jain, I. S. Dhillon and A. Tewari, “Partial hard thresholding”. *IEEE Transactions on Information Theory (IT)*, vol. 63:5, pages 3029–3038, 2017.
12. S. Si, C.-J. Hsieh and I. S. Dhillon, “Memory Efficient Kernel Approximation”, *Journal of Machine Learning Research(JMLR)*, vol. 18(20), pages 1–32, 2017.
13. B. Savas and I. S. Dhillon, “Clustered matrix approximation”, *SIAM Journal of Matrix Analysis and Applications(SIMAX)*, vol. 37:4, pages 1531-1555, 2016.
14. A. Vandaele, N. Gillis, Q. Lei, K. Zhong and I. S. Dhillon, “Efficient and Non-Convex Coordinate Descent for Symmetric Nonnegative Matrix Factorization”, *IEEE Transactions on Signal Processing (TSP)*, vol. 64:21, pages 5571–5584, 2016.
15. J. Whang, D. Gleich and I. S. Dhillon, “Overlapping Community Detection Using Neighborhood-Inflated Seed Expansion”, *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, vol. 28:5, pages 1272–1284, 2016.
16. H.-F. Yu, C.-J. Hsieh, H. Yun, S.V.N. Vishwanathan and I. S. Dhillon, “Nomadic Computing for Big Data Analytics”, *IEEE Computer*, vol. 49:4, pages 52–60, 2016.
17. H. Yun, H.-F. Yu, C.-J. Hsieh, S.V.N. Vishwanathan and I. S. Dhillon, “NOMAD: Non-locking, stochastic Multi-machine algorithm for Asynchronous and Decentralized matrix completion”, *Proceedings of the VLDB Endowment*, vol. 7:11, pages 975–986, 2014.
18. N. Natarajan and I. S. Dhillon, “Inductive matrix completion for predicting gene-disease associations”, *Bioinformatics*, vol. 30:12, pages 60–68, 2014.
19. C.-J. Hsieh, M. Sustik, I. S. Dhillon and P. Ravikumar, “QUIC: Quadratic Approximation for Sparse Inverse Covariance Matrix Estimation”, *Journal of Machine Learning Research(JMLR)*, vol. 15, pages 2911–2947, 2014.
20. H. F. Yu, C.-J. Hsieh, S. Si and I. S. Dhillon, “Parallel Matrix Factorization for Recommender Systems”, *Knowledge and Information Systems(KAIS)*, vol. 41:3, pages 793–819, 2014.
21. K. Chiang, C.-J. Hsieh, N. Natarajan, I. S. Dhillon and A. Tewari, “Prediction and Clustering in Signed Networks: A Local to Global Perspective”, *Journal of Machine Learning Research(JMLR)*, vol. 15, pages 1177–1213, 2014.

22. U. Singh Blom, N. Natarajan, A Tewari, J. Woods, I. S. Dhillon and E. M. Marcotte, “Prediction and Validation of Gene-Disease Associations Using Methods Inspired by Social Network Analyses”, *PLOS ONE* 8(5): e58977, 2013.
23. D. Kim, S. Sra and I. S. Dhillon, “A Non-monotonic Method for Large-scale Nonnegative Least Squares”, *Optimization Methods and Software*, vol. 28:5, pages 1012–1039, 2013.
24. M. Sustik and I. S. Dhillon, “On a Zero-Finding Problem involving the Matrix Exponential”, *SIAM Journal of Matrix Analysis and Applications(SIMAX)*, vol. 33:4, pages 1237–1249, 2012.
25. P. Jain, B. Kulis, J. Davis and I. S. Dhillon, “Metric and Kernel Learning using a Linear Transformation”, *Journal of Machine Learning Research(JMLR)*, vol. 13, pages 519–547, 2012.
26. V. Vasuki, N. Natarajan, Z. Lu, B. Savas and I. S. Dhillon, “Scalable affiliation recommendation using auxiliary networks”, *ACM Transactions on Intelligent Systems and Technology(TIST)*, vol. 3, 2011.
27. D. Kim, S. Sra and I. S. Dhillon, “Tackling Box-Constrained Optimization Via a New Projected Quasi-Newton Approach”, *SIAM Journal on Scientific Computing(SISC)*, vol. 32:6, pages 3548–3563, 2010.
28. B. Kulis, M. Sustik and I. S. Dhillon, “Low-Rank Kernel Learning with Bregman Matrix Divergences”, *Journal of Machine Learning Research(JMLR)*, vol. 10, pages 341–376, 2009.
29. B. Kulis, S. Basu, I. S. Dhillon and R. J. Mooney, “Semi-Supervised Graph Clustering: A Kernel Approach”, *Machine Learning*, 74:1, pages 1–22, 2009.
30. P. Jain, R. Meka and I. S. Dhillon, “Simultaneous Unsupervised Learning of Disparate Clusterings”, *Statistical Analysis and Data Mining*, vol. 1:3, pages 195–210, 2008.
31. I. S. Dhillon, R. Heath Jr., T. Strohmer and J. Tropp, “Constructing Packings in Grassmannian Manifolds via Alternating Projection”, *Experimental Mathematics*, vol. 17:1, pages 9–35, 2008.
32. H. Cho and I. S. Dhillon, “Co-clustering of Human Cancer Microarrays using Minimum Sum-Squared Residue Co-clustering”, *IEEE/ACM Transactions on Computational Biology and Bioinformatics(TCBB)*, vol. 5:3, pages 385–400, 2008.
33. J. Brickell, I. S. Dhillon, S. Sra and J. Tropp, “The Metric Nearness Problem”, *SIAM Journal of Matrix Analysis and Applications(SIMAX)*, vol. 30:1, pages 375–396, April 2008 — **2011 SIAM Outstanding Paper Prize for outstanding papers published in SIAM journals in the three year period from 2008–2010.**
34. D. Kim, S. Sra and I. S. Dhillon, “Fast Projection-Based Methods for the Least Squares Nonnegative Matrix Approximation Problem”, *Statistical Analysis and Data Mining*, vol. 1:1, pages 38–51, 2008.
35. I. S. Dhillon and J. Tropp, “Matrix Nearness Problems using Bregman Divergences”, *SIAM Journal of Matrix Analysis and Applications(SIMAX)*, vol. 29:4, pages 1120–1146, 2007.
36. I. S. Dhillon, Y. Guan and B. Kulis, “Weighted Graph Cuts without Eigenvectors: A Multilevel Approach”, *IEEE Transactions on Pattern Analysis and Machine Intelligence(PAMI)*, vol. 29:11, pages 1944–1957, 2007.
37. M. Sustik, J. Tropp, I. S. Dhillon and R. Heath Jr., “On the existence of Equiangular Tight Frames”, *Linear Algebra and its Applications(LAA)*, vol. 426:2–3, pages 619–635, 2007.
38. A. Banerjee, I. S. Dhillon, J. Ghosh, S. Merugu and D. S. Modha, “A Generalized Maximum Entropy Approach to Bregman Co-Clustering and Matrix Approximations”, *Journal of Machine Learning Research(JMLR)*, vol. 8, pages 1919–1986, 2007.
39. I. S. Dhillon, B. N. Parlett and C. Vömel, “The Design and Implementation of the MRRR Algorithm”, *ACM Transactions on Mathematical Software*, vol. 32:4, pages 533–560, 2006.
40. I. S. Dhillon, B. N. Parlett and C. Vömel, “Glued Matrices and the MRRR Algorithm”, *SIAM Journal on Scientific Computing(SISC)*, vol. 27:2, pages 496–510, 2005.
41. A. Banerjee, S. Merugu, I. S. Dhillon and J. Ghosh, “Clustering with Bregman Divergences”, *Journal of Machine Learning Research(JMLR)*, vol. 6, pages 1705–1749, 2005.

42. A. Banerjee, I. S. Dhillon, J. Ghosh and S. Sra, “Clustering on the Unit Hypersphere using von Mises-Fisher distributions”, *Journal of Machine Learning Research(JMLR)*, vol. 6, pages 1345–1382, 2005.
43. P. Bientinesi, I. S. Dhillon and R. van de Geijn, “A Parallel Eigensolver for Dense Symmetric Matrices Based on Multiple Relatively Robust Representations”, *SIAM Journal on Scientific Computing(SISC)*, vol. 27:1, pages 43–66, 2005.
44. I. S. Dhillon, R. Heath Jr., M. Sustik and J. Tropp, “Generalized finite algorithms for constructing Hermitian matrices with prescribed diagonal and spectrum”, *SIAM Journal of Matrix Analysis and Applications(SIMAX)*, vol. 27:1, pages 61–71, 2005 (a longer version appears as UT CS Technical Report # TR-03-49, Nov 2003).
45. J. Tropp, I. S. Dhillon, R. Heath Jr. and T. Strohmer, “Designing Structured Tight Frames via an Alternating Projection Method”, *IEEE Transactions on Information Theory (IT)*, vol. 51:1, pages 188–209, 2005.
46. J. Tropp, I. S. Dhillon and R. Heath Jr., “Finite-step algorithms for constructing optimal CDMA signature sequences”, *IEEE Transactions on Information Theory (IT)*, vol. 50:11, pages 2916–2921, 2004.
47. I. S. Dhillon and B. N. Parlett, “Multiple Representations to Compute Orthogonal Eigenvectors of Symmetric Tridiagonal Matrices”, *Linear Algebra and its Applications(LAA)*, vol. 387, pages 1–28, 2004.
48. I. S. Dhillon and B. N. Parlett, “Orthogonal Eigenvectors and Relative Gaps”, *SIAM Journal of Matrix Analysis and Applications(SIMAX)*, vol. 25:3, pages 858–899, 2004 — **SIAM Linear Algebra Prize for the best journal paper in applied linear algebra in the three year period from 2003–2005.**
49. I. S. Dhillon, E. M. Marcotte and U. Roshan, “Diametrical Clustering for identifying anti-correlated gene clusters”, *Bioinformatics*, vol. 19:13, pages 1612–1619, 2003.
50. I. S. Dhillon, S. Mallela and R. Kumar, “A Divisive Information-Theoretic Feature Clustering Algorithm for Text Classification”, *Journal of Machine Learning Research(JMLR)*, vol. 3, pages 1265–1287, 2003.
51. I. S. Dhillon and A. Malyshev, “Inner deflation for symmetric tridiagonal matrices”, *Linear Algebra and its Applications(LAA)*, vol. 358:1-3, pages 139–144, 2003.
52. I. S. Dhillon, D. S. Modha and W. S. Spangler, “Class Visualization of High-Dimensional Data with Applications”, *Computational Statistics & Data Analysis (special issue on Matrix Computations & Statistics)*, vol. 4:1, pages 59–90, 2002.
53. I. S. Dhillon and D. S. Modha, “Concept Decompositions for Large Sparse Text Data using Clustering”, *Machine Learning*, 42:1, pages 143–175, 2001.
54. B. N. Parlett and I. S. Dhillon, “Relatively Robust Representations of Symmetric Tridiagonals”, *Linear Algebra and its Applications(LAA)*, vol. 309, pages 121–151, 2000.
55. I. S. Dhillon, “Current inverse iteration software can fail”, *BIT Numerical Mathematics*, 38:4, pages 685–704, 1998.
56. I. S. Dhillon, “Reliable computation of the condition number of a tridiagonal matrix in $O(n)$ time”, *SIAM Journal of Matrix Analysis and Applications(SIMAX)*, 19:3, pages 776–796, 1998.
57. B. N. Parlett and I. S. Dhillon, “Fernando’s solution to Wilkinson’s problem: an application of Double Factorization”, *Linear Algebra and its Applications(LAA)*, vol. 267, pages 247–279, 1997.
58. L. Blackford, A. Cleary, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, A. Petitet, H. Ren, K. Stanley and R. Whaley, “Practical Experience in the Numerical Dangers of Heterogeneous Computing”, *ACM Transactions on Mathematical Software*, vol. 23, no. 2, pages 133–147, 1997.
59. J. Choi, J. Demmel, I. Dhillon, J. Dongarra, S. Ostrouchov, A. Petitet, K. Stanley, D. Walker and R. Whaley, “ScaLAPACK: A Portable Linear Algebra Library for Distributed Memory Computers - Design Issues and Performance”, *Computer Physics Communications*, vol. 97, pages 1–15, 1996.

60. J. W. Demmel, I. S. Dhillon and H. Ren, “On the correctness of some bisection-like eigenvalue algorithms in floating point arithmetic”, *Electronic Transactions of Numerical Analysis*, vol. 3, pages 116–140, 1995.

Conference Publications

1. S. Duvvuri, N. Patel, N. Gupta and I. S. Dhillon, “LUCID: Attention with Preconditioned Representations”, *Proceedings of the Forty-Third International Conference on Machine Learning (ICML)*, 2026.
2. O. Susladkar, T. Prakash, G. Deshmukh, K. Nguyen, J. Zhang, A. Juvekar, T. Bao, L. Chai, S. Mittal, I. S. Dhillon and I. Lourentzou, “Best of Both Worlds: Multimodal Reasoning and Generation via Unified Discrete Flow Matching”, *Proceedings of the Forty-Third International Conference on Machine Learning (ICML)*, 2026.
3. O. Susladkar, T. Prakash, A. Juvekar, K. Nguyen, D. Jang, I. S. Dhillon and I. Lourentzou, “PyraTok: Language-Aligned Pyramidal Tokenizer for Video Understanding and Generation”, *Proceedings of The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026.
4. L. Madaan, R. Tiwari, R. Bansal, S. Duvvuri, M. Zaheer, I. S. Dhillon, D. Brandfonbrener and R. Agarwal, “The Art of Scaling Reinforcement Learning Compute for LLMs”, *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*, 2026 – **oral presentation**.
5. L. Huang, I.S. Dhillon and K. Willcox, “Modeling Longitudinal Student Pathways with Explainable Generative Models”, *Proceedings of the 16th International Conference on Learning Analytics & Knowledge (LAK)*, 2026.
6. T. Prakash, O. Susladkar, S. Mittal and I. S. Dhillon, “Pyramidal Spectrum: Frequency-based Hierarchically Vector Quantized VAE for Videos”, *Proceedings of The IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2026.
7. N. Gupta, C. You, S. Bhojanapalli, S. Kumar, I. S. Dhillon and F. Yu, “ReTuning: Towards Scalable In-context Retrieval”, *Proceedings of the Neural Information Processing Systems Conference (NeurIPS)*, 2025.
8. C. Verma, A. Timmaraju, C.-J. Hsieh, S. Damle, N. Bui, Y. Zhang, W. Chen, X. Liu, P. Jain and I. S. Dhillon, “Matryoshka Model Learning for Improved Elastic Student Models”, *Proceedings of the 31st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2025 – **oral presentation**.
9. S. Duvvuri and I. S. Dhillon, “LASER: Attention with Exponential Transformation”, *Proceedings of the Forty-Second International Conference on Machine Learning (ICML)*, 2025.
10. R. Das, I. S. Dhillon, A. Epasto, A. Javanmard, J. Mao, V. Mirrokni, S. Sanghavi and P. Zhong, “Retraining with Predicted Hard Labels Provably Increases Model Accuracy”, *Proceedings of the Forty-Second International Conference on Machine Learning (ICML)*, 2025.
11. A. Acharya, S. Sanghavi, A. Dimakis and I. S. Dhillon, “Geometric Median (GM) Matching for Robust k -Subset Selection from Noisy Data”, *Proceedings of the Forty-Second International Conference on Machine Learning (ICML)*, 2025.
12. J. Yen, S. Si, Z. Meng, F. Yu, S. Duvvuri, I. S. Dhillon, C.-J. Hsieh and S. Kumar, “LoRA Done RITE: Robust Invariant Transformation Equilibration for LoRA Optimization”, *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*, 2025 – **oral presentation**.
13. R. Wang, S. Si, F. Yu, D. Rothuizen, C.-J. Hsieh and I. S. Dhillon, “Large Language Models are Interpretable Learners”, *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*, 2025.
14. A. Akula, G. Pruthi, I. S. Dhillon, P. Narayana, S. Basu and V. Jampani, “PRISM: A New Lens for Improved Color Understanding”, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1659–1670, 2024.

15. F. Devvrit, S. Kudugunta, A. Kusupati, T. Dettmers, K. Chen, I. S. Dhillon, Y. Tsvetkov, H. Hajishirzi, S. Kakade, A. Farhadi and P. Jain, “MatFormer: Nested Transformer for Elastic Inference”, *Proceedings of the Neural Information Processing Systems Conference(NeurIPS)*, 2024.
16. C.-J. Hsieh, S. Si, F. Yu and I. S. Dhillon, “Automatic Engineering of Long Prompts”, *Findings of the Association for Computational Linguistics (ACL)* pages 10672–10685, 2024.
17. N. Gupta, F. Devvrit, A. Rawat, S. Bhojanapalli, P. Jain and I. S. Dhillon, “Dual-Encoders for Extreme Multi-label Classification”, *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*, 2024.
18. S. Duvvuri, F. Devvrit, R. Anil, C.-J. Hsieh and I. S. Dhillon, “CASPR: Combining Axes Preconditioners through Kronecker Approximation for Deep Learning”, *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*, 2024.
19. Y. Wang, Si Si, D. Li, M. Lukasik, F. Yu, C. J. Hsieh, I. S. Dhillon and S. Kumar, “Two-stage LLM Fine-tuning with Less Specialization and More Generalization”, *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*, 2024.
20. F. Devvrit, S. Duvvuri, R. Anil, V. Gupta, C.-J. Hsieh and I. S. Dhillon, “A Computationally Efficient Sparsified Online Newton Method”, *Proceedings of the Neural Information Processing Systems Conference(NeurIPS)*, 2023.
21. J. Yen, S. Duvvuri, I. S. Dhillon and C.-J. Hsieh, “Block Low-Rank Preconditioner with Shared Basis for Stochastic Optimization”, *Proceedings of the Neural Information Processing Systems Conference(NeurIPS)*, 2023.
22. P. Chen, W.-Chang, J. Jiang H.-F. Yu, I. S. Dhillon and C.-J. Hsieh, “FINGER: Fast Inference for Graph-based Approximate Nearest Neighbor Search”, *Proceedings of the 32nd International World Wide Web Conference(WWW)*, pages 3225–3235, 2023.
23. S. Yang, Y. Dong, R. Ward, I. S. Dhillon, S. Sanghavi and Q. Lei, “Sample Efficiency of Data Augmentation Consistency Regularization”, *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 3825–3853, 2023. 2023.
24. N. Gupta, P. Chen, H.-F. Yu, C.-J. Hsieh and I. S. Dhillon, “ELIAS: End-to-End Learning to Index and Search in Large Output Spaces”, *Proceedings of the Neural Information Processing Systems Conference (NeurIPS)*, pages 19798–19809, 2022.
25. F. Devvrit, A. Sinha, I. S. Dhillon and P. Jain, “S3GC: Scalable Self-Supervised Graph Clustering”, *Proceedings of the Neural Information Processing Systems Conference (NeurIPS)*, pages 3248–3261, 2022.
26. S. Yang, T. Ren, S. Shakkottai, E. Price, I. S. Dhillon and S. Sanghavi, “Linear Bandit Algorithms with Sublinear Time Complexity”, *Proceedings of the 39th International Conference on Machine Learning (ICML)*, pages 25241–25260, 2022.
27. R. Das, A. Acharya, A. Hashemi, S. Sanghavi, I. S. Dhillon and U. Topcu, “Faster Non-Convex Federated Learning via Global and Local Momentum”, *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence (UAI 2022)*, pages 496–506, 2022.
28. Y. Xiong, W.-Chang, C.-J. Hsieh, H.-F. Yu and I. S. Dhillon, “Extreme Zero-Shot Learning for Extreme Text Classification”, *Proceedings of the 2022 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 5455–5468, 2022.
29. M. Cheng, Q. Lei, P.-Y. Chen, I. S. Dhillon and C.-J. Hsieh, “CAT: Customized Adversarial Training for Improved Robustness”, *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 673–679, 2022.
30. A. Block, R. Kidambi, D. Hill, T. Joachims and I. S. Dhillon, “Counterfactual Learning To Rank for Utility-Maximizing Query Autocompletion”, *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pages 791–802, 2022.

31. E. Chien, W. Chang, C.-J. Hsieh, H.-F. Yu, J. Zhang, O. Milenkovic and I. S. Dhillon, “Node Feature Extraction by Self-Supervised Multi-scale Neighborhood Prediction”, *Proceedings of the Tenth International Conference on Learning Representations (ICLR)*, 2022.
32. A. Acharya, A. Hashemi, P. Jain, S. Sanghavi, I. S. Dhillon and U. Topcu, “Robust Training in High Dimensions via Block Coordinate Geometric Median Descent”, *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 11145–11168, 2022.
33. P. Etter, K. Zhong, H.-F. Yu, L. Ying and I. S. Dhillon, “Enterprise-Scale Search: Accelerating Inference for Sparse Extreme Multi-Label Ranking Trees”, *Proceedings of the 31st International World Wide Web Conference (WWW)*, pages 452–461, 2022.
34. J. Zhang, W. Chang, H.-F. Yu and I. S. Dhillon, “Fast Multi-Resolution Transformer Fine-tuning for Extreme Multi-label Text Classification”, *Proceedings of the Neural Information Processing Systems Conference (NeurIPS)*, 2021.
35. P. Chen, H.-F. Yu, I. S. Dhillon and C.-J. Hsieh, “DRONE: Data-aware Low-rank Compression for Large NLP Models”, *Proceedings of the Neural Information Processing Systems Conference (NeurIPS)*, 2021.
36. X. Liu, W. Chang, H.-F. Yu, C.-J. Hsieh and I. S. Dhillon, “Label Disentanglement in Partition-based Extreme Multilabel Classification”, *Proceedings of the Neural Information Processing Systems Conference (NeurIPS)*, 2021.
37. T. Baharav, D. Jiang, K. Kolluri, S. Sanghavi, I. S. Dhillon, “Enabling Efficiency-Precision Trade-offs for Label Trees in Extreme Classification”, *Proceedings of the 2021 ACM Conference on Information and Knowledge Management (CIKM)*, pages 3717–3726, 2021.
38. R. Sen, A. Rakhlin, L. Ying, R. Kidambi, D. Foster, D. Hill and I. S. Dhillon, “Top-k eXtreme Contextual Bandits with Arm Hierarchy”, *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pages 9422–9433, 2021.
39. W. Chang, D. Jiang, H.-F. Yu, C.-H. Teo, J. Zhang, K. Zhong, K. Kolluri, Q. Hu, N. Shandilya, V. Ievgrafoc, J. Singh and I. S. Dhillon, “Extreme Multi-label Learning for Semantic Matching in Product Search”, *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 2643–2651, 2021.
40. N. Yadav, R. Sen, D. Hill, A. Mazumdar and I. S. Dhillon “Session-Aware Query Auto-completion using Extreme Multi-label Ranking”, *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 3835–3844, 2021.
41. R. Lopez, I.S. Dhillon and M. Jordan, “Learning from eXtreme Bandit Feedback”, *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI)*, pages 8732–8740, 2021.
42. W. Chang, H.-F. Yu, K. Zhong, Y. Yang and I. S. Dhillon, “Taming Pretrained Transformers for eXtreme Multi-label Text Classification”, *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 3163–3171, 2020.
43. J. Whang, Y. Jung, S. Kang, D. Yoo and I. S. Dhillon, “Scalable Anti-TrustRank with Qualified Site-level Seeds for Link-based Web Spam Detection”, *Proceedings of the WWW Conference (Companion Volume)*, pages 593–602, 2020.
44. X. Liu, H.-F. Yu, I. S. Dhillon and C.-J. Hsieh, “Learning to Encode Position for Transformer with Continuous Dynamical Model”, *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 6327–6335, 2020.
45. Y. Shen, H.-F. Yu, S. Sanghavi and I. S. Dhillon, “Extreme Multi-label Classification from Aggregated Labels”, *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 8752–8762, 2020.
46. R. Sen, H. F. Yu and I. S. Dhillon, “Think Globally, Act Locally: A Deep Neural Network Approach to High-Dimensional Time Series Forecasting”, *Proceedings of the Neural Information Processing Systems Conference (NeurIPS)*, pages 4838–4847, 2019.

47. J. Zhang, H. F. Yu and I. S. Dhillon, “AutoAssist: A Framework to Accelerate Training of Deep Neural Networks”, *Proceedings of the Neural Information Processing Systems Conference(NeurIPS)*, pages 5996–6006, 2019.
48. K. Zhong, Z. Song, P. Jain and I. S. Dhillon, “Provable Non-linear Inductive Matrix Completion”, *Proceedings of the Neural Information Processing Systems Conference(NeurIPS)*, pages 11435–11445, 2019.
49. Q. Lei, J. Zhuo, C. Caramanis, I. S. Dhillon and A. Dimakis, “Primal-Dual Block Generalized Frank-Wolfe”, *Proceedings of the Neural Information Processing Systems Conference(NeurIPS)*, pages 13866–13875, 2019.
50. Q. Lei, A. Jalal, I. S. Dhillon and A. Dimakis, “Inverting Deep Generative models, One layer at a time”, *Proceedings of the Neural Information Processing Systems Conference(NeurIPS)*, pages 13910–13919, 2019.
51. Q. Lei, J. Yi, R. Vaculin, L. Wu and I. S. Dhillon. “Similarity Preserving Representation Learning for Time Series Analysis”, *28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.
52. H. Zhang, H. Chen, Z. Song, D. Boning, I.S. Dhillon and C.-J. Hsieh, “The Limitations of Adversarial Training and the Blind-Spot Attack”, *Seventh International Conference on Learning Representations (ICLR)*, 2019.
53. H.F. Yu, C.-J. Hsieh and I. S. Dhillon, “Parallel Asynchronous Stochastic Coordinate Descent with Auxiliary Variables”, *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 2641–2649, 2019.
54. J. Zhang, P. Raman, S. Ji, H.F. Yu, S.V.N. Vishwanathan and I. S. Dhillon, “Extreme Stochastic Variational Inference: Distributed Inference for Large Scale Mixture Models”, *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 935–943, 2019.
55. Q. Lei, L. Wu, P.-Y. Chen, A. Dimakis, I. S. Dhillon and M. Witbrock, “Discrete Adversarial Attacks and Submodular Optimization with Applications to Text Classification”, *The Conference on Systems and Machine Learning (SysML)*, 2019.
56. A. Acharya, R. Goel, A. Metallinou and I. S. Dhillon, “Online Embedding Compression for Text Classification using Low Rank Matrix Factorization”, *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence(AAAI)*, pages 6196–6203, 2019.
57. J. Zhang, Q. Lei and I. S. Dhillon, “Stabilizing Gradients for Deep Neural Networks via Efficient SVD Parameterization”, *Proceedings of the 35th International Conference on Machine Learning(ICML)*, pages 5801–5809, 2018.
58. J. Zhang, Y. Lin, Z. Song and I. S. Dhillon, “Learning Long Term Dependencies via Fourier Recurrent Units”, *Proceedings of the 35th International Conference on Machine Learning(ICML)*, pages 5810–5818, 2018.
59. T.-W. Weng, H. Zhang, H. Chen, Z. Song, C.-J. Hsieh, D. Boning, I. S. Dhillon, L. Daniel, “Towards Fast Computation of Certified Robustness for ReLU Networks”, *Proceedings of the 35th International Conference on Machine Learning(ICML)*, pages 5273–5282, 2018.
60. P. Wang, H. Zhang, V. Mohan, I. S. Dhillon and Z. Kolter. “Realtime query completion via deep language models”, SIGIR Workshop on eCommerce, 2018.
61. H. F. Yu, C.-J. Hsieh, Q. Lei and I. S. Dhillon, “A Greedy Approach for Budgeted Maximum Inner Product Search”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, 2017.
62. J. Whang and I. S. Dhillon, “Non-Exhaustive, Overlapping Co-Clustering”, *Proceedings of the 2017 ACM Conference on Information and Knowledge Management(CIKM)*, pages 2367–2370, 2017.
63. I. Yen, X. Huang, W. Dai, P. Ravikumar, I. S. Dhillon and E. Xing, “PPDspare: A Parallel Primal-Dual Sparse Method for Extreme Classification”, *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 545–553, 2017.

64. C.-J. Hsieh, S. Si and I. S. Dhillon, “Communication-Efficient Distributed Block Minimization for Nonlinear Kernel Machines”, *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 245–254, 2017.
65. K. Zhong, Z. Song, P. Jain, P. Bartlett and I. S. Dhillon, “Recovery Guarantees for One-hidden-layer Neural Networks”, *Proceedings of the 34th International Conference on Machine Learning(ICML)*, pages 4140–4149, 2017.
66. S. Si, H. Zhang, S. Keerthi, D. Mahajan, I. S. Dhillon and C.-J. Hsieh, “Gradient Boosted Decision Trees for High Dimensional Sparse Output”, *Proceedings of the 34th International Conference on Machine Learning(ICML)*, pages 3182–3190, 2017.
67. Q. Lei, I. Yen, C.-Y. Wu, P. Ravikumar and I. S. Dhillon, “Doubly Greedy Primal-Dual Coordinate Methods for Sparse Empirical Risk Minimization”, *Proceedings of the 34th International Conference on Machine Learning(ICML)*, pages 2034–2042, 2017.
68. K. Zhong, R. Guo, S. Kumar, B. Yan, D. Simcha and I. S. Dhillon, “Fast Classification with Binary Prototypes”, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR: W&CP 54, pages 1255–1263, 2017.
69. K. Chiang, C.-J. Hsieh and I. S. Dhillon, “Rank Aggregation and Prediction with Item Features”, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR: W&CP 54, pages 748–756, 2017.
70. X. Huang, I. Yen, R. Zhang, Q. Huang, P. Ravikumar and I. S. Dhillon. “Greedy Direction Method of Multiplier for MAP Inference of Large Output Domain”, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR: W&CP 54, pages 1550–1559, 2017.
71. J. Zhang, I. Yen, P. Ravikumar and I. S. Dhillon. “Scalable Convex Multiple Sequence Alignment via Entropy-Regularized Dual Decomposition”, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR: W&CP 54, pages 1514–1522, 2017.
72. H. F. Yu, H. Y. Huang, I. S. Dhillon and C. J. Lin. “A Unified Algorithm for One-class Structured Matrix Factorization with Side Information”, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence(AAAI)*, JMLR: W&CP 54, pages 2845–2851, 2017.
73. H. F. Yu, N. Rao and I. S. Dhillon. “Temporal Regularized Matrix Factorization for High-dimensional Time Series Prediction”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 847–855, 2016.
74. P. Jain, N. Rao and I. S. Dhillon. “Structured Sparse Regression via Greedy Hard Thresholding”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 1516–1524, 2016.
75. Q. Lei, K. Zhong and I. S. Dhillon. “Coordinate-wise Power Method”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2056–2064, 2016.
76. K. Zhong, P. Jain and I. S. Dhillon. “Mixed Linear Regression with Multiple Components”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2190–2198, 2016.
77. Y. You, X. Lian, C.-J. Hsieh, J. Liu, H.-F. Yu, I. S. Dhillon and J. Demmel, “Asynchronous Parallel Greedy Coordinate Descent”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 4682–4690, 2016.
78. I. Yen, X. Huang, K. Zhong, Z. Ruohan, P. Ravikumar and I. S. Dhillon, “Dual Decomposed Learning with Factorwise Oracle for Structural SVM of Large Output Domain”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 5024–5032, 2016.
79. S. Si, K. Chiang, C.-J. Hsieh, N. Rao and I. S. Dhillon, “Goal-Directed Inductive Matrix Completion”, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 1165–1174, 2016.
80. N. Natarajan, O. Koyejo, P. Ravikumar and I. S. Dhillon, “Optimal Classification with Multivariate Losses”, *Proceedings of the 33rd International Conference on Machine Learning(ICML)*, pages 1530–1538, 2016.

81. S. Si, C.-J. Hsieh and I. S. Dhillon, “Computationally Efficient Nyström Approximation using Fast Transforms”, *Proceedings of the 33rd International Conference on Machine Learning(ICML)*, pages 2655–2663, 2016.
82. K. Chiang, C.-J. Hsieh and I. S. Dhillon, “Robust Principal Component Analysis with Side Information”, *Proceedings of the 33rd International Conference on Machine Learning(ICML)*, pages 2291–2299, 2016.
83. I. Yen, X. Huang, K. Zhong, P. Ravikumar and I. S. Dhillon, “PD-Sparse: A Primal and Dual Sparse Approach to Extreme Classification”, *Proceedings of the 33rd International Conference on Machine Learning(ICML)*, pages 3069–3077, 2016.
84. I. Yen, X. Lin, J. Zhang, P. Ravikumar and I. S. Dhillon, “A Convex Atomic-Norm Approach to Multiple Sequence Alignment and Motif Discovery”, *Proceedings of the 33rd International Conference on Machine Learning(ICML)*, pages 2272–2280, 2016.
85. D. Inouye, P. Ravikumar and I. S. Dhillon. “Square Root Graphical Models: Multivariate Generalizations of Univariate Exponential Families that Permit Positive Dependencies”, *Proceedings of the 33rd International Conference on Machine Learning(ICML)*, pages 2445–2453, 2016.
86. Y. Hou, J. Whang, D. Gleich and I. S. Dhillon, “Fast Multiplier Methods to Optimize Non-exhaustive, Overlapping Clustering”, *Proceedings of the 2016 SIAM International Conference on Data Mining(SDM)*, pages 297–305, 2016.
87. K. Howard, N. Natarajan, A. Ramakarishnan, K. Ponds, W. Layton, M. Cowperthwaite and I.S. Dhillon “A Risk Score For All-Cause Readmissions Using Get With The Guidelines-Stroke Data Elements”, *International Stroke Conference(Stroke)*, 2016.
88. K. Chiang, C.-J. Hsieh and I. S. Dhillon. “Matrix Completion with Noisy Side Information”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 3447–3455, 2015 — *spotlight presentation*.
89. O. Koyejo, N. Natarajan, P. Ravikumar and I. S. Dhillon. “Consistent Multilabel Classification”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 3321–3329, 2015.
90. D. Inouye, P. Ravikumar and I. S. Dhillon. “Fixed-Length Poisson MRF: Adding Dependencies to the Multinomial”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 3213–3221, 2015.
91. N. Rao, H. F. Yu, P. Ravikumar and I. S. Dhillon. “Collaborative Filtering with Graph Information: Consistency and Scalable Methods”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2107–2115, 2015.
92. I. Yen, K. Zhong, C. J. Hsieh, P. Ravikumar and I. S. Dhillon, “Sparse Linear Programming via Primal and Dual Augmented Coordinate Descent”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2368–2376, 2015.
93. N. Natarajan, N. Rao and I. S. Dhillon. “PU Matrix Completion with Graph Information”, *Computational Advances in Multi-Sensor Adaptive Processing(CAMSAP)*, pages 37–40, 2015.
94. D. Shin, S. Cetintas, K. Lee and I. S. Dhillon, “Tumblr Blog Recommendation with Boosted Inductive Matrix Completion”, *Proceedings of the 24th ACM Conference on Information and Knowledge Management(CIKM)*, pages 203–212, 2015.
95. K. Zhong, P. Jain and I. S. Dhillon, “Efficient Matrix Sensing Using Rank-1 Gaussian Measurements”, *Proceedings of the 26th International Conference on Algorithmic Learning Theory(ALT)*, pages 3–18, 2015.
96. Y. Hou, J. Whang, D. Gleich and I. S. Dhillon, “Non-exhaustive, Overlapping Clustering via Low-Rank Semidefinite Programming”, *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 427–436, 2015.
97. J. Whang, A. Lenharth, I. S. Dhillon and K. Pingali, “Scalable Data-driven PageRank: Algorithms, System Issues & Lessons Learned”, *International European Conference on Parallel and Distributed Computing(Euro-Par)*, pages 438–450, 2015.

98. C.-J. Hsieh, H. F. Yu and I. S. Dhillon, “PASSCoDe: Parallel ASynchronous Stochastic dual Coordinate Descent”, *Proceedings of the 32nd International Conference on Machine Learning(ICML)*, pages 2370–2379, 2015.
99. C.-J. Hsieh, N. Natarajan and I. S. Dhillon, “PU Learning for Matrix Completion”, *Proceedings of the 32nd International Conference on Machine Learning(ICML)*, pages 2445–2453, 2015.
100. I. Yen, K. Zhong, J. Lin, P. Ravikumar and I. S. Dhillon, “A Convex Exemplar-based Approach to MAD-Bayes Dirichlet Process Mixture Models”, *Proceedings of the 32nd International Conference on Machine Learning(ICML)*, pages 2418–2426, 2015.
101. D. Park, J. Neeman, J. Zhang, S. Sanghavi and I. S. Dhillon, “Preference Completion: Large-scale Collaborative Ranking from Pairwise Comparisons”, *Proceedings of the 32nd International Conference on Machine Learning(ICML)*, pages 1907–1916, 2015.
102. H.-F. Yu, C.-J. Hsieh, H. Yun, S.V.N. Vishwanathan and I. S. Dhillon, “A Scalable Asynchronous Distributed Algorithm for Topic Modeling”, *Proceedings of the 24th International World Wide Web Conference(WWW)*, pages 1340–1350, 2015.
103. J. Whang, I. S. Dhillon and D. Gleich, “Non-exhaustive, Overlapping k -means”, *Proceedings of the 2015 SIAM International Conference on Data Mining(SDM)*, pages 936–944, 2015.
104. A. Jha, S. Ray, B. Seaman and I. S. Dhillon, “Clustering to Forecast Sparse Time-Series Data”, *Proceedings of the International Conference on Data Engineering(ICDE)*, pages 1388–1399, 2015.
105. C.-J. Hsieh, I. S. Dhillon, P. Ravikumar, S. Becker and P. Olsen, “QUIC & DIRTY: A Quadratic Approximation Approach for Dirty Statistical Models”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2006—2014, 2014.
106. C.-J. Hsieh, S. Si and I. S. Dhillon, “Fast Prediction for Large-Scale Kernel Machines”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 3689–3697, 2014.
107. S. Si, D. Shin, I. S. Dhillon and B. Parlett, “Multi-Scale Spectral Decomposition of Massive Graphs”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2798–2806, 2014.
108. N. Natarajan, O. Koyejo, P. Ravikumar and I. S. Dhillon, “Consistent Binary Classification with Generalized Performance Metrics”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2744–2752, 2014.
109. K. Zhong, I. Yen, I. S. Dhillon and P. Ravikumar. “Proximal Quasi-Newton for Computationally Intensive l_1 -regularized M-estimators”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2375–2383, 2014.
110. I. Yen, C.-J. Hsieh, P. Ravikumar and I. S. Dhillon. “Constant Nullspace Strong Convexity and Fast Convergence of Proximal Methods under High-Dimensional Settings”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 1008–1016, 2014.
111. I. Yen, T. Lin, S. Lin, P. Ravikumar and I. S. Dhillon, “Sparse Random Feature Algorithm as Coordinate Descent in Hilbert Space”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2456–2464, 2014.
112. D. Inouye, P. Ravikumar and I. S. Dhillon, “Capturing Semantically Meaningful Word Dependencies with an Admixture of Poisson MRFs”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 3158—3166, 2014.
113. C.-J. Hsieh, S. Si and I. S. Dhillon, “A Divide-and-Conquer Solver for Kernel Support Vector Machines”, *Proceedings of the 31st International Conference on Machine Learning(ICML)*, pages 566–574, 2014.
114. H. F. Yu, P. Jain, P. Kar and I. S. Dhillon, “Large-scale Multi-label Learning with Missing Labels”, *Proceedings of the 31st International Conference on Machine Learning(ICML)*, pages 593–601, 2014.
115. S. Si, C.-J. Hsieh and I. S. Dhillon, “Memory Efficient Kernel Approximation”, *Proceedings of the 31st International Conference on Machine Learning(ICML)*, pages 701–709, 2014.

116. D. Inouye, P. Ravikumar and I. S. Dhillon, “Admixtures of Poisson MRFs: A Topic Model with Word Dependencies”, *Proceedings of the 31st International Conference on Machine Learning(ICML)*, pages 683–691, 2014.
117. C.-J. Hsieh, M. Susik, I. S. Dhillon, P. Ravikumar and R. Poldrack, “BIG & QUIC: Sparse Inverse Covariance Estimation for a Million Variables”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 3165–3173, 2013 – **oral presentation**.
118. N. Natarajan, I. S. Dhillon, P. Ravikumar and A. Tewari, “Learning with Noisy Labels”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 1196–1204, 2013.
119. H. Wang, A. Banerjee, C.-J. Hsieh, P. Ravikumar and I. S. Dhillon, “Large Scale Distributed Sparse Precision Estimation”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 584–592, 2013.
120. J. Whang, P. Rai and I. S. Dhillon, “Stochastic Blockmodel with Cluster Overlap, Relevance Selection, and Similarity-Based Smoothing”, *Proceedings of the IEEE International Conference on Data Mining(ICDM)*, pages 817–826, 2013.
121. J. Whang, D. Gleich and I. S. Dhillon, “Overlapping Community Detection Using Seed Set Expansion”, *Proceedings of the 22nd ACM Conference on Information and Knowledge Management(CIKM)*, pages 2099–2108, 2013.
122. N. Natarajan, D. Shin and I. S. Dhillon, “Which app will you use next? Collaborative Filtering with Interactional Context”, *Proceedings of the 7th ACM Conference on Recommender Systems(RecSys)*, pages 201–208, 2013.
123. C.-J. Hsieh, I. S. Dhillon, P. Ravikumar and A. Banerjee, “A Divide-and-Conquer Method for Sparse Inverse Covariance Estimation”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2339–2347, 2012.
124. H. F. Yu, C.-J. Hsieh, S. Si and I. S. Dhillon, “Scalable Coordinate Descent Approaches to Parallel Matrix Factorization for Recommender Systems”, *Proceedings of the IEEE International Conference on Data Mining(ICDM)*, pages 765–774, 2012 — **Best Paper Award**.
125. J. Whang, X. Sui and I. S. Dhillon, “Scalable and Memory-Efficient Clustering of Large-Scale Social Networks”, *Proceedings of the IEEE International Conference on Data Mining(ICDM)*, pages 705–714, 2012.
126. D. Shin, S. Si and I. S. Dhillon, “Multi-Scale Link Prediction”, *Proceedings of the 21st ACM Conference on Information and Knowledge Management(CIKM)*, pages 215–224, 2012.
127. K. Chiang, J. Whang and I. S. Dhillon, “Scalable Clustering of Signed Networks using Balance Normalized Cut”, *Proceedings of the 21st ACM Conference on Information and Knowledge Management(CIKM)*, pages 615–624, 2012.
128. X. Sui, T. Lee, J. Whang, B. Savas, S. Jain, K. Pingali and I. S. Dhillon, “Parallel Clustered Low-rank Approximation of Graphs and Its Application to Link Prediction”, *Proceedings of the 25th Int’l Workshop on Languages and Compilers for Parallel Computing(LCPC)*, pages 76–95, 2012.
129. C.-J. Hsieh, K. Chiang and I. S. Dhillon, “Low-Rank Modeling of Signed Networks”, *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 507–515, 2012.
130. H. Song, B. Savas, T. Cho, V. Dave, Z. Lu, I. S. Dhillon, Y. Zhang and L. Qiu, “Clustered Embedding of Massive Social Networks”, *Proceedings of ACM SIGMETRICS Int’l Conference on Measurement and Modeling of Computer Systems*, pages 331–342, 2012.
131. P. Jain, A. Tewari and I. S. Dhillon, “Orthogonal Matching Pursuit with Replacement”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 1215–1223, 2011.
132. A. Tewari, P. Ravikumar and I. S. Dhillon, “Greedy Algorithms for Structurally Constrained High Dimensional Problems”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 882–890, 2011.

133. I. S. Dhillon, P. Ravikumar and A. Tewari, “Nearest Neighbor based Greedy Coordinate Descent”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2160–2168, 2011.
134. C.-J. Hsieh, M. Sustik, I. S. Dhillon and P. Ravikumar, “Sparse Inverse Covariance Matrix Estimation using Quadratic Approximation”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 2330–2338, 2011.
135. K. Chiang, N. Natarajan, A. Tewari and I. S. Dhillon, “Exploiting Longer Cycles for Link Prediction in Signed Networks”, *Proceedings of the 20th ACM Conference on Information and Knowledge Management(CIKM)*, pages 1157–1162, 2011.
136. C.-J. Hsieh and I. S. Dhillon, “Fast Coordinate Descent Methods with Variable Selection for Non-negative Matrix Factorization”, *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 1064–1072, 2011.
137. B. Savas and I. S. Dhillon, “Clustered low rank approximation of graphs in information science applications”, *Proceedings of the 2011 SIAM International Conference on Data Mining(SDM)*, pages 164–175, 2011.
138. Z. Lu, B. Savas, W. Tang and I. S. Dhillon, “Supervised Link Prediction using Multiple Sources”, *Proceedings of the IEEE International Conference on Data Mining(ICDM)*, pages 923–928, 2010.
139. R. Meka, P. Jain and I. S. Dhillon, “Guaranteed Rank Minimization via Singular Value Projection”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 937–945, 2010 — *spotlight presentation*.
140. P. Jain, B. Kulis and I. S. Dhillon, “Inductive Regularized Learning of Kernel Functions”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 946–954, 2010 — *spotlight presentation*.
141. V. Vasuki, N. Natarajan, Z. Lu and I. S. Dhillon, “Affiliation recommendation using auxiliary networks”, *Proceedings of the 4th ACM Conference on Recommender Systems(RecSys)*, pages 103–110, 2010.
142. D. Kim, S. Sra and I. S. Dhillon, “A scalable trust-region algorithm with application to mixed-norm regression”, *Proceedings of the 27th International Conference on Machine Learning(ICML)*, pages 519–526, 2010.
143. R. Meka, P. Jain and I. S. Dhillon, “Matrix Completion from Power-Law Distributed Samples”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 1258–1266, 2009.
144. W. Tang, Z. Lu and I. S. Dhillon, “Clustering with Multiple Graphs”, *Proceedings of the IEEE International Conference on Data Mining(ICDM)*, pages 1016–1021, 2009.
145. Z. Lu, D. Agarwal and I. S. Dhillon, “A Spatio-Temporal Approach to Collaborative Filtering”, *Proceedings of the 3rd ACM Conference on Recommender Systems(RecSys)*, pages 13–20, 2009.
146. S. Sra, D. Kim, I. S. Dhillon and B. Schölkopf, “A new non-monotonic algorithm for PET image reconstruction”, *IEEE Medical Imaging Conference (MIC 2009)*, pages 2500–2502, 2009.
147. Z. Lu, P. Jain and I. S. Dhillon, “Geometry-aware Metric Learning”, *Proceedings of the 26th International Conference on Machine Learning(ICML)*, pages 673–680, 2009.
148. M. Deodhar, J. Ghosh, G. Gupta, H. Cho and I. S. Dhillon, “A Scalable Framework for Discovering Coherent Co-clusters in Noisy Data”, *Proceedings of the 26th International Conference on Machine Learning(ICML)*, pages 241–248, 2009 — **Best Student Paper Honorable Mention**.
149. B. Kulis, S. Sra and I. S. Dhillon, “Convex Perturbations for Scalable Semidefinite Programming”, *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR: W&CP 5, pages 296–303, 2009.
150. P. Jain, B. Kulis, I. S. Dhillon and K. Grauman, “Online Metric Learning and Fast Similarity Search”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 761–768, 2008 — **oral presentation**.

151. J. Davis and I. S. Dhillon “Structured Metric Learning for High-Dimensional Problems”, *Proceedings of the Fourteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 195–203, 2008.
152. R. Meka, P. Jain, C. Caramanis and I. S. Dhillon, “Rank Minimization via Online Learning”, *Proceedings of the 25th International Conference on Machine Learning(ICML)*, pages 656–663, July 2008.
153. P. Jain, R. Meka and I. S. Dhillon, “Simultaneous Unsupervised Learning of Disparate Clusterings”, *Proceedings of the Seventh SIAM International Conference on Data Mining*, pages 858–869, April 2008 — **Best paper award, runner-up.**
154. J. Davis, B. Kulis, P. Jain, S. Sra and I. S. Dhillon, “Information-Theoretic Metric Learning”, *Proceedings of the 24th International Conference on Machine Learning(ICML)*, pages 209–216, June 2007 — **Best Student Paper Award.**
155. D. Kim, S. Sra and I. S. Dhillon, “Fast Newton-type Methods for the Least Squares Nonnegative Matrix Approximation Problem”, *Proceedings of the Sixth SIAM International Conference on Data Mining*, pages 343–354, April 2007 — **Best of SDM’07 Award.**
156. J. Davis and I. S. Dhillon, “Differential Entropic Clustering of Multivariate Gaussians”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 337–344, December 2006.
157. J. Davis and I. S. Dhillon “Estimating the Global PageRank of Web Communities”, *Proceedings of the Twelfth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 116–125, August 2006.
158. B. Kulis, M. Sustik and I. S. Dhillon, “Learning Low-Rank Kernel Matrices”, *Proceedings of the 23rd International Conference on Machine Learning(ICML)*, pages 505–512, July 2006.
159. I. S. Dhillon and S. Sra, “Generalized Nonnegative Matrix Approximations with Bregman Divergences”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 283–290, December 2005.
160. B. Kulis, S. Basu, I. S. Dhillon and R. J. Mooney, “Semi-Supervised Graph-Based Clustering: A Kernel Approach”, *Proceedings of the 22nd International Conference on Machine Learning(ICML)*, pages 457–464, July 2005 — **Distinguished Student Paper Award.**
161. I. S. Dhillon, Y. Guan and B. Kulis, “A Fast Kernel-based Multilevel Algorithm for Graph Clustering”, *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 629–634, August 2005.
162. I. S. Dhillon, S. Sra and J. Tropp, “Triangle Fixing Algorithms for the Metric Nearness Problem”, *Proceedings of the Neural Information Processing Systems Conference(NIPS)*, pages 361–368, December 2004.
163. A. Banerjee, I. S. Dhillon, J. Ghosh, S. Merugu and D. S. Modha, “A Generalized Maximum Entropy Approach to Bregman Co-Clustering and Matrix Approximations”, *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 509–514, August 2004 (a longer version appears as UT CS Technical Report # TR-04-24, June 2004).
164. I. S. Dhillon, Y. Guan and B. Kulis, “Kernel k-means, Spectral Clustering and Normalized Cuts”, *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 551–556, August 2004.
165. A. Banerjee, I. S. Dhillon, J. Ghosh and S. Merugu, “An Information Theoretic Analysis of Maximum Likelihood Mixture Estimation for Exponential Families”, *Proceedings of the Twenty-First International Conference on Machine Learning(ICML)*, pages 57–64, July 2004.
166. A. Banerjee, S. Merugu, I. S. Dhillon and J. Ghosh, “Clustering with Bregman Divergences”, *Proceedings of the Third SIAM International Conference on Data Mining*, pages 234–245, April 2004 — **Best Paper Award.**
167. H. Cho, I. S. Dhillon, Y. Guan and S. Sra, “Minimum Sum-Squared Residue Co-clustering of Gene Expression Data”, *Proceedings of the Third SIAM International Conference on Data Mining*, pages 114–125, April 2004.

168. R. Heath Jr., J. Tropp, I. S. Dhillon and T. Strohmer, “Construction of Equiangular Signatures for Synchronous CDMA Systems”, *Proceedings of IEEE International Symposium on Spread Spectrum Techniques and Applications*, pages 708–712, August 2004.
169. J. Tropp, I. S. Dhillon and R. Heath Jr., “Optimal CDMA Signatures: A Finite-Step Approach”, *Proceedings of IEEE International Symposium on Spread Spectrum Techniques and Applications*, pages 335–340, August 2004.
170. J. Tropp, I. S. Dhillon, R. Heath Jr. and T. Strohmer, “CDMA Signature Sequences with Low Peak-to-Average-Power Ratio via Alternating Projection”, *Proceedings of the Thirty-Seventh IEEE Asilomar Conference on Signals, Systems, and Computers*, pages 475–479, November 2003.
171. I. S. Dhillon, S. Mallela and D. S. Modha, “Information-Theoretic Co-clustering”, *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 89–98, August 2003.
172. A. Banerjee, I. S. Dhillon, J. Ghosh and S. Sra, “Generative Model-Based Clustering of Directional Data”, *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 19–28, August 2003.
173. I. S. Dhillon and Y. Guan, “Information-Theoretic Clustering of Sparse Co-occurrence Data”, *Proceedings of the 3rd IEEE International Conference on Data Mining(ICDM)*, pages 517–520, November 2003 (a longer version appears as UT CS Technical Report # TR-03-39, Sept 2003).
174. I. S. Dhillon, Y. Guan and J. Kogan, “Iterative Clustering of High Dimensional Text Data Augmented by Local Search”, *Proceedings of the 2nd IEEE International Conference on Data Mining(ICDM)*, pages 131–138, December 2002.
175. I. S. Dhillon, S. Mallela and R. Kumar, “Enhanced Word Clustering for Hierarchical Text Classification”, *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 191–200, July 2002.
176. I. S. Dhillon, “Co-Clustering Documents and Words Using Bipartite Spectral Graph Partitioning”, *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, pages 269–274, August 2001.
177. I. S. Dhillon, D. S. Modha and W. S. Spangler, “Visualizing Class Structure of Multi-Dimensional Data”, In *Proceedings of the 30th Symposium of the Interface: Computing Science and Statistics*, Interface Foundation of North America, vol. 30, pages 488–493, May 1998.
178. L. Blackford, J. Choi, A. Cleary, E. D’Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker and R. Whaley, “ScaLAPACK: A Linear Algebra Library for Message-Passing Computers”, *Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing*, March 1997.
179. I. S. Dhillon, G. Fann and B. N. Parlett, “Application of a New Algorithm for the Symmetric Eigenproblem to Computational Quantum Chemistry”, *Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing*, March 1997.
180. L. Blackford, J. Choi, A. Cleary, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker and R. Whaley, “ScaLAPACK: A Portable Linear Algebra Library for Distributed Memory Computers - Design Issues and Performance”, *Proceedings of Supercomputing ’96*, pages 95–106, 1996.
181. I. S. Dhillon, N. K. Karmarkar and K. G. Ramakrishnan, “An Overview of the Compilation Process for a New Parallel Architecture”, *Supercomputing Symposium ’91*, pages 471–486, June 1991.

Other Papers

1. N. Rajani, K. McArdle and I. S. Dhillon, “Parallel k nearest neighbor graph construction using tree-based data structures”, 1st High Performance Graph Mining workshop, KDD Conference, 2015.
2. H.-F. Yu, N. Rao and I. S. Dhillon, “Temporal Regularized Matrix Factorization”, *Time Series Workshop at Neural Information Processing Systems Conference(NIPS)*, 2015.

3. M. Deodhar, H. Cho, G. Gupta, J. Ghosh and I. S. Dhillon, "Hunting for Coherent Co-clusters in High-Dimensional and Noisy Datasets", IEEE International Conference on Data Mining(ICDM) (Workshop on Foundations of Data Mining), pages 654–663, December 2008.
4. J. Brickell, I. S. Dhillon and D. Modha, "Adaptive Website Design using Caching Algorithms", Twelfth ACM International Conference on Knowledge Discovery and Data Mining (KDD) (Workshop on Web Mining and Web Usage Analysis (WebKDD-2006)), pages 1–20, August 2006.
5. I. S. Dhillon and Y. Guan, "Clustering Large, Sparse, Co-occurrence Data", 3rd SIAM International Conference on Data Mining (Workshop on Clustering High-Dimensional Data and its Applications), May 2003.
6. I. S. Dhillon, S. Mallela and R. Kumar, "Information-Theoretic Feature Clustering for Text Classification", Nineteenth International Conference on Machine Learning (ICML) (Workshop on Text Learning (TextML-2002)), July 2002.
7. I. S. Dhillon, Y. Guan and J. Kogan, "Refining clusters in high-dimensional text data", 2nd SIAM International Conference on Data Mining (Workshop on Clustering High-Dimensional Data and its Applications), April 2002.

Book Chapters

1. P. Berkhin and I. S. Dhillon, "Clustering", In: *Encyclopedia of Complexity and Systems Science*, Invited Book Chapter, Springer, 2009.
2. A. Banerjee, I. S. Dhillon, J. Ghosh and S. Sra, "Text Clustering with Mixture of von Mises-Fisher Distributions", In: M. Sahami, A. Srivastava(eds), *Text Mining: Classification, Clustering and Applications*, Invited Book Chapter, CRC Press, pages 121–153, 2009.
3. J. Brickell, I. S. Dhillon and D. Modha, "Adaptive Website Design using Caching Algorithms", In: O. Nasraoui, M. Spiliopoulou, J. Srivastava, B. Mobasher, B. Masand(eds), *Advances in Web Mining and Web Usage Analysis*, Invited Book Chapter, Springer Lecture Notes in Computer Science (LNCS/LNAI), vol. 4811, pages 1–20, Sept 2007.
4. A. K. Cline and I. S. Dhillon, "Computation of the Singular Value Decomposition", In: L. Hogben, R. Brualdi, A. Greenbaum and R. Mathias(eds): *Handbook of Linear Algebra*, Invited Book Chapter, CRC Press, pages 45-1–45-13, 2006.
5. M. Teboulle, P. Berkhin, I. S. Dhillon, Y. Guan and J. Kogan, "Clustering with Entropy-like k -means Algorithms", invited book chapter, In: *Grouping Multidimensional Data – Recent Advances in Clustering*, Springer-Verlag, pages 127–160, 2005.
6. I. S. Dhillon, J. Kogan and C. Nicholas, "Feature Selection and Document Clustering", In: Michael Berry(ed): *A Comprehensive Survey of Text Mining*, Springer-Verlag, pages 73–100, 2003.
7. I. S. Dhillon, Y. Guan and J. Fan, "Efficient Clustering of Very Large Document Collections", invited book chapter, In: R. Grossman, C. Kamath, P. Kegelmeyer, V. Kumar and R. Namburu(eds): *Data Mining for Scientific and Engineering Applications*, In *Data Mining for Scientific and Engineering Applications*, Kluwer Academic Publishers, Massive Computing vol. 2, pages 357–381, 2001.
8. I. S. Dhillon and D. S. Modha, "A Data Clustering Algorithm on Distributed Memory Multiprocessors", In: M.Zaki and C.T.Ho(eds): *Large-Scale Parallel Data Mining, Lecture Notes in Artificial Intelligence, vol. 1759*, Springer-Verlag, pages 245–260, March 2000 (also IBM Research Report RJ 10134).

Book

1. L. Blackford, J. Choi, A. Cleary, E. D’Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker and R. Whaley, "ScaLAPACK Users’ Guide", *SIAM*, 1997.

Technical Reports

1. B. Kulis, S. Sra, S. Jegelka and I. S. Dhillon. "Scalable Semidefinite Programming using Convex Perturbations". UT CS Technical Report # TR-07-47, September 2007.
2. P. Jain, B. Kulis and I. S. Dhillon, "Online Linear Regression using Burg Entropy", UT CS Technical Report # TR-07-08, Feb 2007.
3. S. Sra and I. S. Dhillon, "Nonnegative Matrix Approximation: Algorithms and Applications", UT CS Technical Report # TR-06-27, June 2006.
4. I. S. Dhillon and S. Sra, "Generalized Nonnegative Matrix Approximations with Bregman Divergences", UT CS Technical Report # TR-05-31, June 2005.
5. I. S. Dhillon, Y. Guan and B. Kulis, "A Unified View of Kernel k -means, Spectral Clustering and Graph Cuts", UT CS Technical Report # TR-04-25, June 2004.
6. I. S. Dhillon, S. Sra and J. Tropp, "Triangle Fixing Algorithms for the Metric Nearness Problem", UT CS Technical Report # TR-04-22, June 2004.
7. I. S. Dhillon, S. Sra and J. Tropp, "The Metric Nearness Problem with Applications", UT CS Technical Report # TR-03-23, July 2003.
8. A. Banerjee, I. S. Dhillon, J. Ghosh and S. Sra, "Clustering on Hyperspheres using Expectation Maximization", UT CS Technical Report # TR-03-07, February 2003.
9. I. S. Dhillon and S. Sra, "Modeling data using directional distributions", UT CS Technical Report # TR-03-06, January 2003.
10. I. S. Dhillon, "A New $O(n^2)$ Algorithm for the Symmetric Tridiagonal Eigenvalue/Eigenvector Problem", PhD Thesis, University of California, Berkeley, May 1997 (also available as UCB Tech. Report No. UCB//CSD-97-971).
11. M. Gu, J. W. Demmel and I. S. Dhillon, "Efficient Computation of the Singular Value Decomposition with Applications to Least Squares Problems", *Technical Report LBL-36201, Lawrence Berkeley National Laboratory*, 1994 (also available as LAPACK working note no. 88).
12. J. Choi, J. Demmel, I. Dhillon, J. Dongarra, S. Ostrouchov, A. Petitet, K. Stanley, D. Walker and R. Whaley, "Installation Guide for ScaLAPACK", University of Tennessee Computer Science Technical Report, UT-CS-95-280, March 1995 (version 1.0), updated August 31, 2001 (version 1.7) — also available as LAPACK working note no. 93.
13. I. S. Dhillon, N. K. Karmarkar and K. G. Ramakrishnan, "Performance Analysis of a Proposed Parallel Architecture on Matrix Vector Multiply Like Routines", *Technical Memorandum 11216-901004-13TM, AT&T Bell Laboratories*, Murray Hill, NJ, 1990.
14. I. S. Dhillon, "A Parallel Architecture for Sparse Matrix Computations", *B.Tech. Project Report, Indian Institute of Technology, Bombay*, 1989.

PLENARY TALKS AT MAJOR CONFERENCES/WORKSHOPS

- June 2025:** "Kronecker Product Based Optimizers for Deep Learning", Plenary Talk, *Householder XXII Symposium*, Ithaca, NY.
- Aug 2023:** "Multi-Output Prediction: Theory and Practice", Keynote Talk, 15th International Conference on Contemporary Computing (IC3), Noida, India.
- Aug 2022:** "Multiscale Methods for Machine Learning", Keynote Talk, *The 28th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, Washington DC.
- May 2021:** "Multiscale Methods for Machine Learning", Keynote Talk, SIAM Conference on Applied Linear Algebra (LA21), New Orleans, Louisiana.
- Aug 2020:** "Think Globally, Act Locally: A Deep Neural Network Approach to High-Dimensional Time Series Forecasting", Keynote Talk, 6th KDD (Int'l Conference on Knowledge Discovery and Data Mining) Workshop on Mining and Learning from Time Series (MILETS), Virtual Conference.

- Jul 2020:** “Multi-Output Prediction: Theory and Practice”, Invited Talk, ICML (Int’l Conference on Machine Learning) Workshop on Extreme Classification, Virtual Conference.
- May 2020:** “Multi-Output Prediction: Theory and Practice”, Keynote Talk, 24th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), Singapore.
- Aug 2019:** “Stable Recurrent Unit with Efficient Singular Value Gating”, Keynote Talk, 11th International Conference on Contemporary Computing (IC3) , Noida, India.
- Oct 2018:** “Stabilizing Gradients for Deep Neural Networks”, Keynote Talk, Harvard Data Science Initiative Conference (HDSI), Boston, MA.
- Sep 2018:** “Stabilizing Gradients for Deep Neural Networks”, Keynote Talk, Amazon Research Day, Bangalore, India.
- Dec 2017:** “Stabilizing Gradients for Deep Neural Networks”, Keynote Talk, ICMLDS (International Conference on Machine Learning and Data Science), Delhi, India.
- Dec 2017:** “Stabilizing Gradients for Deep Neural Networks with Applications to Extreme Classification”, Invited Talk, NIPS (Neural Information Processing Systems) Workshop on Extreme Classification, Long Beach, CA.
- Sept 2017:** “Multi-Target Prediction Using Low-Rank Embeddings: Theory & Practice”, Keynote Talk, ECML PKDD (European Conference on Machine Learning), Skopje, Macedonia.
- Dec 2016:** “A Primal and Dual Sparse Approach to Extreme Classification”, Invited Talk, NIPS (Neural Information Processing Systems) Workshop on Extreme Classification, Barcelona, Spain.
- Dec 2016:** “Temporal Regularized Matrix Factorization for High-dimensional Time Series Prediction”, Invited Talk, NIPS (Neural Information Processing Systems) Time Series Workshop, Barcelona, Spain.
- Oct 2015:** “Bilinear Prediction using Low-Rank Models”, Keynote Talk, 26th International Conference on Algorithmic Learning Theory (ALT), Banff, Canada.
- June 2015:** “Proximal Newton Methods for Large-Scale Machine Learning”, Distinguished Talk, ShanghaiTech Symposium on Data Science, Shanghai, China.
- Dec 2014:** “NOMAD: A Distributed Framework for Latent Variable Models”, Invited Talk, NIPS (Neural Information Processing Systems) Workshop on Distributed Machine Learning and Matrix computations, Montreal, Canada.
- Dec 2014:** “Divide-and-Conquer Methods for Big Data”, Keynote Talk, ICMLA (13th International Conference on Machine Learning and Applications), Detroit, Michigan.
- June 2014:** “Parallel Asynchronous Matrix Completion”, Plenary Talk, *Householder XVIII Symposium*, Spa, Belgium.
- Mar 2014:** “Divide-and-Conquer Methods for Big Data”, Keynote Talk, CoDS (1st iKDD Conference on Data Sciences), New Delhi, India.
- Dec 2013:** “Scalable Network Analysis”, Keynote Talk, COMAD (19th Int’l Conference on Management of Data), Ahmedabad, India.
- July 2013:** “BIG & QUIC: Sparse Inverse Covariance Estimation for a Million Variables”, Plenary Talk, SPARS (Signal Processing with Adaptive Sparse Structured Representations), EPFL, Lausanne, Switzerland.
- Dec 2011:** “Fast and Memory-Efficient Low-rank Approximation of Massive Graphs”, Plenary Talk, NIPS (Neural Information Processing Systems) Workshop on Low-rank Methods for Large-Scale Machine Learning, Sierra Nevada, Spain.
- June 2011:** “Social Network Analysis: Fast and Memory-Efficient Low-Rank Approximation of Massive Graphs”, Plenary Talk, *Householder XVII Symposium*, Tahoe City, California.
- June 2009:** “Matrix Computations in Machine Learning”, Plenary Talk, ICML (Int’l Conference on Machine Learning) Workshop on Numerical Methods in Machine Learning, McGill University, Montreal, Canada.

- June 2008:** “The Log-Determinant Divergence and its Applications”, Plenary Talk, *Householder XVII Symposium*, Zeuthen, Germany.
- May 2008:** “Machine Learning with Bregman Divergences”, Plenary Talk, *EurOPT-2008*, Neringa, Lithuania.
- July 2006:** “Orthogonal Eigenvectors and Relative Gaps”, SIAM Linear Algebra Prize Talk, *Ninth SIAM Conference on Applied Linear Algebra*, Dusseldorf, Germany.
- July 2006:** “From Shannon to von Neumann: New Distance Measures for Matrix Nearness Problems”, Plenary Talk, *Ninth SIAM Conference on Applied Linear Algebra*, Dusseldorf, Germany.
- May 2005:** “Matrix Nearness Problems using Bregman Divergences”, Plenary Talk, *Householder XVI Symposium*, Seven Springs, Pennsylvania.
- Aug 2002:** “Fast and Accurate Eigenvector Computation in Finite Precision Arithmetic”, Semi-plenary Talk, *The Fourth Foundations of Computational Mathematics Conference (FoCM)*, Minneapolis, Minnesota.
- June 2002:** “Matrix Problems in Data Mining”, Plenary Talk, *Householder XV Symposium*, Peebles, Scotland.
- June 1999:** “Orthogonal Eigenvectors through Relatively Robust Representations”, Plenary Talk, *Householder XIV Symposium*, Whistler, Canada.

INVITED TALKS

- Apr 2024:** “MatFormer: Nested Transformer for Elastic Inference”, Distinguished Speaker Series, Data Science Institute, University of Chicago, Chicago, IL.
- Nov 2023:** “Learning in Large Output Spaces: An End-to-End Approach”, Invited Colloquium Talk, Department of Computer Sciences, Cornell University, Ithaca, New York.
- June 2022:** “Sparse Inverse Covariance Estimation for a Million Variables”, Invited Talk, *Householder XXI Symposium*, Fasano, Italy.
- Feb 2022:** “Multiscale Methods for Machine Learning”, Invited minisymposium talk, *SIAM Conference on Parallel Processing for Scientific Computing*, Seattle, WA.
- May 2021:** “Multiscale Methods for Machine Learning”, Invited Talk, Information System Laboratory (ISL) Colloquium, Stanford University, CA.
- Aug 2020:** “Multi-Output Prediction: Theory and Practice”, Invited Talk, Institute of Advanced Studies (IAS), Special Year on Optimization, Statistics and Theoretical Machine Learning, Princeton, NJ.
- Oct 2019:** “AutoAssist: A Framework to Accelerate Training of Deep Neural Networks”, Invited Talk, Industrial and Systems Engineering (ISyE) Departmental Seminar, Georgia Tech, Atlanta.
- Dec 2018:** “Multi-Target Prediction Using Low-Rank Embeddings: Theory & Practice”, Invited Talk, ICMLDS (International Conference on Machine Learning and Data Science), Hyderabad, India.
- Sep 2018:** “Stabilizing Gradients for Deep Neural Networks”, Invited Talk, International Center for Theoretical Studies (ICTS), Bangalore, India.
- May 2018:** “Stabilizing Gradients for Deep Neural Networks”, Invited Talk, Linear Algebra and Optimization Seminar, Stanford University, CA.
- Nov 2017:** “Stabilizing Gradients for Deep Neural Networks”, Invited Talk, Dept of Computer Science (LAPACK Seminar), UC Berkeley, CA.
- Oct 2017:** “Multi-Target Prediction Using Low-Rank Embeddings: Theory & Practice”, Keynote Talk, University of Louisiana Research Day, Lafayette, Louisiana.
- April 2017:** “Sparse Inverse Covariance Estimation for a Million Variables”, Invited Talk, Linear Algebra and Optimization Seminar, Stanford University, CA.
- March 2017:** “Sparse Inverse Covariance Estimation for a Million Variables”, Invited Talk, Dept of Computer Science (LAPACK Seminar), UC Berkeley, CA.

- June 2015:** “Bilinear Prediction using Low-Rank Models”, Invited Talk, Workshop on Low-rank Optimization and Applications, Hausdorff Center for Mathematics(HCM), Bonn, Germany.
- May 2015:** “Proximal Newton Methods for Large-Scale Machine Learning”, Invited Colloquium Talk, Mathematics Department, UT Austin.
- Nov 2014:** “NOMAD: A Distributed Framework for Latent Variable Models”, Invited Talk, Intel Research Labs, Santa Clara, CA, 2014.
- Nov 2014:** “Divide-and-Conquer Methods for Large-Scale Data Analysis”, Distinguished Lecture, School of Computational Science & Engg, Georgia Tech, Atlanta.
- July 2014:** “Sparse Inverse Covariance Estimation for a Million Variables”, Invited Talk, International Conference on Signal Processing and Communications (SPCOM), Bangalore, India.
- June 2014:** “Scalable Network Analysis”, Invited Talk, Adobe Data Science Symposium, San Jose, CA.
- May 2014:** “Informatics in Computational Medicine”, Invited Talk, ICES Computational Medicine Day, UT Austin.
- May 2014:** “Scalable Network Analysis”, Invited Talk, IBM TJ Watson Research Labs, New York.
- Dec 2013:** “Divide & Conquer Methods for Big Data Analytics”, Keynote Talk, Workshop on Distributed Computing for Machine Learning & Optimization, Mysore Park, Mysore, India.
- Sept 2013:** “Sparse Inverse Covariance Estimation for a Million Variables”, Invited Talk: SAMSI Program on Low-Dimensional Structure in High-Dimensional Systems: Opening Workshop, Raleigh, North Carolina.
- Apr 2013:** “Sparse Inverse Covariance Estimation using Quadratic Approximation”, Invited talk, Johns Hopkins University, Baltimore, Maryland.
- Sept 2012:** “Sparse Inverse Covariance Estimation using Quadratic Approximation”, Invited Talks: SAMSI Massive Datasets Opening Workshop, Raleigh, North Carolina, and MLSLP Symposium, Portland, Oregon.
- Aug 2012:** “Orthogonal Matching Pursuit with Replacement”, Mathematics Department, TU Berlin, Germany.
- Feb 2012:** “Orthogonal Matching Pursuit with Replacement”, Invited talk, Information Theory & Applications (ITA) Workshop, San Diego, California.
- Jan 2012:** “Fast and accurate low-rank approximation of massive graphs”, Invited talk, Statistics Yahoo! Seminar, Purdue University, Indiana.
- June 2011:** “Fast and accurate low-rank approximation of massive graphs”, Plenary talk, Summer Workshop on Optimization in Machine Learning, UT Austin.
- Aug 2010:** “Fast and accurate low-rank approximation of massive graphs”, Invited talk, Workshop on Advanced Topics in Humanities Network Analysis, IPAM, UCLA, California.
- Feb 2010:** “Guaranteed Rank Minimization via Singular Value Projection”, Invited talk, Information Theory & Applications (ITA) Workshop, San Diego, California.
- Dec 2009:** “Guaranteed Rank Minimization via Singular Value Projection”, Plenary talk, Workshop on Algorithms for processing Massive Data Sets, IIT Kanpur, India.
- Apr 2009:** “Metric and Kernel Learning”, Invited Colloquium Talk, Computer Science and Engineering Department, Penn State University, State College, Pennsylvania.
- Feb 2009:** “Newton-type methods for Nonnegative Tensor Approximation”, Invited talk, NSF Workshop on Future Directions in Tensor-Based Computation and Modeling, NSF, Arlington, Virginia.
- June 2008:** “Rank Minimization via Online Learning”, Invited talk, *Workshop on Algorithms for Modern Massive Data Sets*, Stanford University, California.
- Mar 2008:** “The Symmetric Tridiagonal Eigenproblem”, Invited talk, Bay Area Scientific Computing Day, Mathematical Sciences Research Institute (MSRI), Berkeley, California.

- Feb 2008:** “Metric and Kernel Learning”, Invited Colloquium Talk, ORFE (Operations Research & Financial Engineering) Department, Princeton University, Princeton, New Jersey.
- Nov 2007:** “Metric and Kernel Learning”, Invited Colloquium Talk, Department of Computer Sciences, Cornell University, Ithaca, New York.
- Sept, Oct & Nov 2007:** “Clustering Tutorial”, “Metric and Kernel Learning” & “Multilevel Graph Clustering”, Special program on *Mathematics of Knowledge and Search Engines*, Institute of Pure & Applied Mathematics (IPAM), UCLA, California.
- June 2007:** “Machine Learning and Optimization”, Invited Panel Speaker, *A-C-N-W Optimization Tutorials*, Chicago, Illinois.
- Feb 2007:** “Fast Newton-type Methods for Nonnegative Matrix Approximation”, Invited talk, *NISS Workshop on Non-negative Matrix Factorization*, Raleigh, N. Carolina.
- Jan 2007:** “Machine Learning with Bregman Divergences”, Invited talk, *BIRS Seminar on Mathematical Programming in Data Mining and Machine Learning*, organized by M.Jordan, J.Peng, T.Poggio, K.Scheinberg, D.Schuermans and T.Terlaky, Banff, Canada.
- June 2006:** “Kernel Learning with Bregman Matrix Divergences” Invited talk, *Workshop on Algorithms for Modern Massive Data Sets*, Stanford University and Yahoo! Research, California.
- May 2006:** “Spectral Measures for Nearness Problems”, Plenary talk, *Sixth International Workshop on Accurate Solution of Eigenvalue Problems*, Pennsylvania State University, University Park, Pennsylvania.
- Dec 2005:** “Co-Clustering, Matrix Approximations and Bregman Divergences”, Invited Colloquium Talk, Department of Computer Sciences, Cornell University, Ithaca, New York.
- Nov 2005:** “Co-Clustering, Matrix Approximations and Bregman Divergences”, Invited talk, McMaster University, Hamilton, Canada.
- Mar 2004:** “Information-Theoretic Clustering, Co-clustering and Matrix Approximations”, Invited talk, IBM TJ Watson Research Center, New York.
- Feb 2004:** “Fast Eigenvalue/Eigenvector Computation for Dense Symmetric Matrices”, Invited talk, University of Illinois, Urbana-Champaign.
- Nov 2003:** “Inverse Eigenvalue Problems in Wireless Communications”, *BIRS Seminar on Theory and Numerics of Matrix Eigenvalue Problems*, organized by J.Demmel, N.Higham and P.Lancaster, Banff, Canada.
- Oct 2003:** “Inverse Eigenvalue Problems in Wireless Communications”, *Dagstuhl Seminar on Theoretical and Computational Aspects of Matrix Algorithms*, organized by N.Higham, V.Mehrmann, S.Rump and D.Szyld, Wadern, Germany.
- Aug 2003:** “Accurate Computation of Eigenvalues and Eigenvectors of Dense Symmetric Matrices”, Sandia National Laboratories, Albuquerque.
- May 2003:** “Information-Theoretic Clustering, Co-clustering and Matrix Approximations”, *IMA Workshop on Data Analysis and Optimization*, organized by R.Kannan, J.Kleinberg, C.Papadimitriou and P.Ragahavan, Minneapolis, Minnesota.
- Dec 2001:** “Clustering High-Dimensional Data and Data Approximation”, Invited Colloquium Talk, University of Minnesota, Minneapolis.
- Oct 2000:** “Concept Decompositions for Large-Scale Information Retrieval”, Invited talk, *Computational Information Retrieval Workshop*, Raleigh, Carolina.
- Apr 2000:** “Matrix Approximations for Large Sparse Text Data using Clustering”, Invited talk, *IMA Workshop on Text Mining*, Minneapolis, Minnesota.
- Sept 1999:** “Class Visualization of High-Dimensional Data”, Invited talk, *Workshop on Mining Scientific Datasets*, Minneapolis, Minnesota.
- Mar & Apr 1999:** “Eigenvectors and Concept Vectors”, Invited talks at UC Santa Barbara, Stanford, UT Austin, Yale, UW Madison and Caltech.

OTHER MAJOR TALKS

- Aug 2012:** “Sparse Inverse Covariance Estimation using Quadratic Approximation”, ISMP 2012, Berlin.
- Feb, Apr & May 2012:** “Link Prediction for Large-Scale Social Networks”, Tech talks, LinkedIn (Mountain View, CA), Amazon (Seattle, WA) and Google (Mountain View, CA).
- Feb 2012:** “Parallel Clustered Low-rank Approximation of Social Network Graphs”, Invited minisymposium talk, *SIAM Conference on Parallel Processing for Scientific Computing*, Savannah, Georgia.
- July 2011:** “Fast and Memory-Efficient Low-Rank Approximation of Massive Graphs”, Invited minisymposium talk, *ICIAM*, Vancouver, Canada.
- July 2010:** “Guaranteed Rank Minimization via Singular Value Projection”, Invited minisymposium talk, *SIAM Annual Meeting*, Pittsburgh, Pennsylvania.
- June 2008:** “On some modified root finding problems”, Invited minisymposium talk, *15th Conference of the International Linear Algebra Society (ILAS)*, Cancun, Mexico.
- July 2007:** “Fast Newton-type Methods for Nonnegative Tensor Approximation”, Invited minisymposium talk, *Sixth Int'l Conference on Industrial and Applied Mathematics (ICIAM)*, Zurich, Switzerland.
- July 2005:** “Glued Matrices and the MRRR Algorithm”, Invited minisymposium talk, *SIAM Annual Meeting*, New Orleans, Louisiana.
- Feb 2004:** “A Parallel Eigensolver for Dense Symmetric Matrices based on Multiple Relatively Robust Representations”, Invited minisymposium talk, *SIAM Conference on Parallel Processing for Scientific Computing*, San Francisco, California.
- Aug 2003 – June 2004:** “Information Theoretic Clustering, Co-Clustering and Matrix Approximations”, PARC, Yahoo!, Verity, Univ. of Wisconsin-Madison.
- Aug 2003:** “Information-Theoretic Co-clustering”, *The Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, Washington DC.
- July 2003:** “Data Clustering using Generalized Distortion Measures”, Invited minisymposium talk, *SIAM Conference on Applied Linear Algebra*, Williamsburg, Virginia.
- July 2003:** “Matrix Nearness Problems in Data Mining”, Invited minisymposium talk, *SIAM Conference on Applied Linear Algebra*, Williamsburg, Virginia.
- July 2002:** “Enhanced Word Clustering for Hierarchical Text Classification”, *The Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining(KDD)*, Edmonton, Canada.
- June 2002:** “Accurate Computation of Eigenvalues and Eigenvectors of Tridiagonal Matrices”, *Fourth International Workshop on Accurate Eigensolving and Applications*, Split, Croatia.
- Apr 2002:** “Refining Clusters in High-Dimensional Text Data”, *2nd SIAM International Conference on Data Mining*, Arlington, Virginia.
- July 2001:** “Large-Scale Data Mining”, Invited minisymposium talk, *SIAM Annual Meeting*, San Diego, California.
- Oct 2000:** “Multiple Representations for Orthogonal Eigenvectors”, Invited minisymposium talk, *SIAM Conference on Applied Linear Algebra*, Raleigh, Carolina.
- Dec 1999:** “Class Visualization of Multidimensional Data with Applications”, Invited minisymposium on Data Mining, *6th International Conference on High Performance Computing (HiPC '99)*, Calcutta, India.
- May 1999:** “Concept-Revealing Subspaces for Large Text Collections”, Invited minisymposium talk, *SIAM Annual Meeting*, Atlanta, Georgia.
- Aug 1998:** “Concept Identification in Large Text Collections”, *5th International Symposium, IRREGULAR'98*, Berkeley, California.
- Apr 1998:** “Orthogonal Eigenvectors without Gram-Schmidt”, *Ph.D. Dissertation Talk*, Berkeley, California.

- Oct 1997:** “When are Factors of Indefinite Matrices Relatively Robust?”, *Sixth SIAM Conference on Applied Linear Algebra*, Snowbird, Utah.
- July 1997:** “Perfect Shifts and Twisted Q Factorizations”, *SIAM 45th Anniversary and Annual Meeting*, Stanford, California.
- Mar 1997:** “A New Algorithm for the Symmetric Eigenproblem Applied to Computational Quantum Chemistry”, *Eighth SIAM Conference on Parallel Processing for Scientific Computing*, Minneapolis, Minnesota.
- Aug 1996:** “Accuracy and Orthogonality”, *First International Workshop on Accurate Eigensolving and Applications*, Split, Croatia.
- July 1996:** “A New Approach to the Symmetric Tridiagonal Eigenproblem”, *Householder XIII Symposium*, Pontresina, Switzerland.
- June 1991:** “Compilation Process for a New Parallel Architecture based on Finite Geometries”, *Supercomputing Symposium '91*, New Brunswick, Canada.

PATENTS

1. US20240176993A1: “Scalable Self-Supervised Graph Clustering”, P. Jain, I. S. Dhillon, F. Devvrit and A. Sinha, 2024.
2. US11250039: “Extreme multi-label classification”, W. Chang, H.-F. Yu, and I. S. Dhillon, 2022.
3. US11126660: “High dimensional time series forecasting”, R. Sen, H.-F. Yu, and I. S. Dhillon, 2021.
4. US11106690: “Neural query auto-correction and completion”, I. S. Dhillon, V. Mohan and P. Wang, 2021.
5. US10872601: “Natural language processing”, A. Acharya, A. Metallinou, R. Goel and I. S. Dhillon, 2020.
6. US6269376: “Method and system for clustering data in parallel on a distributed memory multiprocessor system”, I. S. Dhillon and D. Modha, 2001.
7. US6560597: “Concept decomposition using clustering”, I. S. Dhillon and D. Modha, 2003.

GRANTS

1. “ Collaborative Research: EAGER: End-to-end Neural Training for Very Large Output Spaces”, National Science Foundation, IIS-2439754, \$100,000, 08/15/24-07/31/26.
2. “Efficient Training Algorithms for Large-Scale Distributed Training”, Meta Inc., \$100,000, 06/01/21-01/31/23.
3. “Predoctoral Training in Biomedical Big Data Science”, National Institutes of Health (NIH), \$1,106,990, 03/01/16-02/28/21.
4. “AF: Medium: Dropping Convexity: New Algorithms, Statistical Guarantees and Scalable Software for Non-convex Matrix Estimation”, National Science Foundation, IIS-1564000, \$902,415, 09/01/16-08/31/22.
5. “BIGDATA:Collaborative Research:F:Nomadic Algorithms for Machine Learning in the Cloud”, National Science Foundation, IIS-1546452, \$1,206,758, 01/01/16-9/30/21.
6. “Co-morbidity prediction using patient healthcare data”, Xerox Foundation, \$90,000, 06/01/15-05/31/17.
7. “Multi-modal Recommendation via Inductive Tensor Completion”, Adobe, \$50,000, 06/01/15-05/31/17.
8. “I-Corps: Faster than Light Big Data Analytics”, National Science Foundation, IIP-1507631, \$50,000, 01/01/15-06/30/15.

9. "AF: Divide-and-Conquer Numerical Methods for Analysis of Massive Data Sets", National Science Foundation, CCF-1320746, \$491,044, 09/01/13-08/31/16.
10. "AF: Fast and Memory-Efficient Dimensionality Reduction for Massive Networks", National Science Foundation, CCF-1117055, \$360,000, 09/01/11-08/31/14.
11. "Disease Modeling via Large-Scale Network Analysis", Army Research Office, \$359,004, 09/01/10-12/31/13.
12. "Scalable mining of SKT cell phone data", UT Austin, \$31,938, 06/01/10-01/15/11.
13. "Mining smartphone data", Motorola Mobility, \$50,000, 1/01/11-12/31/11.
14. "RI: Matrix-structured statistical inference", National Science Foundation, IIS-1018426, \$157,331, 08/15/10-07/31/11.
15. "Spatial-Temporal Approach to Scalable Dynamical Collaborative Filtering", Yahoo! Research, \$15,000, 10/01/09-12/31/10.
16. "NetSE: Multi-Resolution Analysis of Network Matrices", National Science Foundation, CCF-0916309, \$499,996, 09/01/09-08/31/13.
17. "Link Prediction and Missing Value Imputation on Multiple Data Sets", Sandia National Laboratories, \$45,614, 01/01/09-08/31/09.
18. "Graph Data Mining", Sandia National Laboratories, \$24,360, 06/01/08-08/31/08.
19. "Non-Negative Matrix and Tensor Approximations: Algorithms, Software and Applications", National Science Foundation, CCF-0728879, \$250,000, 10/01/07-09/30/12.
20. "III-COR: Versatile Co-clustering Analysis for Bi-modal and Multi-modal Data", National Science Foundation, IIS-0713142, \$430,000, 09/01/07-08/31/10.
21. "Sparse Data Estimation through Hierarchical Aggregation", Yahoo! Research, \$25,000, 01/01/08-12/31/08.
22. Supplemental Award for National Science Foundation Grant "Novel Matrix Problems in Modern Applications", CCF-0431257, \$30,000, 08/15/05-07/31/09.
23. REU Supplemental Award for National Science Foundation Grant "Scalable Algorithms for Large-Scale Data Mining", ACI-0093404, \$6,000, 06/01/05-05/31/08.
24. "Novel Matrix Problems in Modern Applications", National Science Foundation, CCF-0431257, \$200,000, 08/15/04-07/31/09.
25. "Web data mining and algorithms", Sabre Holdings, Inc., \$72,000, 09/01/04-08/31/07.
26. "RI: Mastodon: A Large-Memory, High-Throughput Simulation Infrastructure", 18 co-PIs, National Science Foundation, CISE Research Infrastructure, extramural funding: \$1,418,231, UT match: \$508,000, total grant: \$1,927,031, 2003-2008.
27. "ITR: Feedback from Multi-Source Data Mining to Experimentation for Gene Network Discovery", with R. Mooney, D. Miranker, V. Iyer and E. Marcotte, NSF Information Technology Research Award, IIS-0325116, \$1,700,000, 11/03/03-10/31/08.
28. "Reconstructing Gene Networks by Mining Expression, Genomic and Literature Data", with E. Marcotte, Texas Higher Education Coordinating Board, Advanced Research Program (TARP), Award # 003658-0431-2001, \$147,000, 01/01/02-08/31/04.
29. "Scalable Algorithms for Large-Scale Data Mining", NSF CAREER Award, ACI-0093404, \$478,305, 06/01/01-05/31/08.
30. "Data Mining Seminar Series", Tivoli Inc., \$5,000, 06/01/01-05/31/03.
31. "Linear Algebra for Text Classification", Lawrence Berkeley National Laboratories, \$11,570.13, 08/01/00-05/31/01.
32. "An Interdisciplinary Practicum Course in Data Mining", Tivoli Inc., \$20,000, 06/01/00-05/31/03.
33. "Data Mining Seminar Series", Tivoli Inc., \$5,000, 06/01/00-05/31/01.

34. "Research in Web Mining", Neonyoyo Inc., \$8,520, 06/01/00-08/31/00.
35. "Studies in Numerical Linear Algebra and Applications," University of Texas, Summer Research Award, \$15,112, 06/01/00-07/31/00.
36. "SCOUT: Scientific Computing Cluster of UT (CISE Research Instrumentation)," with D. Burger, S. Keckler, H. Vin and T. Warnow, National Science Foundation, CISE-9985991, extramural funding: \$139,481, UT match: \$70,000, total grant: \$209,481, 03/15/00-03/14/03.