

Brian Kulis

Department of Electrical and Computer Engineering
Boston University
8 Saint Mary's Street
Boston, MA 02215

bkulis@bu.edu
http://people.bu.edu/bkulis

Research Interests

Machine learning; large-scale data analysis; statistics; numerical optimization; applications to computer vision, audio, and other domains

Education

- **Ph.D. Computer Science** **University of Texas at Austin**
August, 2003 - November, 2008
Thesis: *Scalable Kernel Methods for Machine Learning*
Research Advisor: Inderjit S. Dhillon
Graduate GPA: 4.0 / 4.0
- **B.A. Computer Science and Mathematics** **Cornell University**
August, 1999 - May, 2003
Graduated With Honors

Professional Experience

- **Associate Professor** **Boston University**
Department of Electrical and Computer Engineering
Department of Computer Science (affiliated)
Division of Systems Engineering (core member)
5/2018 - Boston, MA
- **Amazon Scholar** **Amazon Inc.**
Amazon Alexa
Cambridge, MA
7/2019 -
- **Peter J. Levine Career Development Assistant Professor** **Boston University**
Department of Electrical and Computer Engineering
7/2015 - 5/2018
Department of Computer Science (affiliated)
7/2015 - 5/2018
Division of Systems Engineering (core member)
9/2016 - 5/2018
Boston, MA

- **Assistant Professor**
 Department of Computer Science and Engineering
 Department of Statistics
 1/2012 -6/2015

Ohio State University
Columbus, OH
- **Postdoctoral Fellow**
 EECS Department & International Computer Science Institute
 1/2009 - 12/2011

UC Berkeley
Berkeley, CA
- **Research Fellow**
 Institute for Pure & Applied Mathematics
 9/2007 - 12/2007

UCLA
Los Angeles, CA
- **Research Intern**
 6/2006 - 9/2006

Microsoft Research
Redmond, WA
- **Research Assistant**
 Department of Computer Sciences
 1/2004 - 12/2008

University of Texas at Austin
Austin, TX
- **Research Assistant**
 Department of Computer Science
 9/2001 - 6/2003

Cornell University
Ithaca, NY

Teaching Experience

- **Instructor**
 1/2019 - 5/2019
 Introduction to Machine Learning (Undergraduate Level)

Boston University
- **Instructor**
 8/2020-12-2020
 8/2018-12/2018
 1/2017-5/2017
 Deep Learning (Graduate Level)

Boston University
- **Instructor**
 1/2018 - 5/2018
 1/2017 - 5/2017
 8/2016 - 12/2016
 Advanced Data Structures and Algorithms (Graduate Level)

Boston University
- **Instructor**
 8/2014-12/2014
 1/2014-4/2014
 Survey of Artificial Intelligence II (Undergraduate and Graduate Level)

Ohio State University
- **Instructor**
 1/2013-4/2013
 Machine Learning (Undergraduate and Graduate Level)

Ohio State University
- **Instructor**
 8/2012-12/2012
 Probabilistic Graphical Models (Graduate Level)

Ohio State University
- **Instructor**
 1/2015-5/2015
 3/2012-6/2012
 Bayesian Modeling and Inference (Graduate Level)

Ohio State University

- **Teaching Assistant** University of Texas at Austin
 1/2007 - 5/2007
 Data Mining: A Statistical Learning Perspective (Graduate Level)
 Supervising professor: Inderjit Dhillon.
- **Teaching Assistant** University of Texas at Austin
 9/2006 - 12/2006
 Machine Learning (Graduate Level)
 Supervising professor: Raymond Mooney.
- **Teaching Assistant** University of Texas at Austin
 9/2004 - 12/2004
 Large-scale Data Mining (Graduate Level)
 Supervising professor: Inderjit Dhillon.
- **Teaching Assistant** Cornell University
 1/2003 - 5/2003
 Introduction to Algorithms (Undergraduate Level)
 Supervising professors: Jon Kleinberg and Eva Tardos.
- **Teaching Assistant** Cornell University
 9/2002 - 12/2002
 Introduction to the Theory of Computing (Undergraduate Level)
 Supervising professor: John Hopcroft.

Funding

- NSF IIS RI:Small, “Hard Clustering via Bayesian Nonparametrics”, Award 1217433, 2012–2015, Total Award: \$440,000
- NSF CAREER, “Rich and Scalable Optimization for Modern Bayesian Nonparametric Learning”, Award 1452903, 2015–2020, Total Award: \$486,000
- Peter J. Levine Career Development Award, 2015-2018, Total Award: \$90,000
- Verisk Analytics, “Generative Adversarial Networks for Discrete Data”, 2017-2018: \$40,000
- Verisk AI Faculty Award, 2019: \$25,000

Advising

- **Ph. D. Students**
 Ke Jiang (2011-2017; after OSU: Data Scientist at Microsoft), Anirban Roychowdhury (2012-2017; after OSU: Research Scientist at Facebook), Robert Finn (2013-2017; after OSU: Assistant Professor, St. Peters University), Andrew Cutler (2016-), Ali Siahkamari (2016-), Kubra Cilingir (2016-), Xide Xia (2016-)
- **Masters Students**
 Siddharth Singh (2011-2013; after OSU: IBM, then Amazon), Ye Liu (2012-2014; after OSU: PhD student at University of Michigan), Jiaxin Zhang (2012-2014; after OSU: Google), Xiangyang Zhou (2013-2016; after OSU: Google), Siva Sankarapandian (2016-2018; after BU: Proscia Inc.), Tayler Pauls (2018-2019; after BU: startup)
- **Undergraduate Students**
 Elizabeth Burl (B.S. 2015; after OSU: Google), Vijay Thakkar (2017-2019; after BU: graduate student at Georgia Tech), Rachel Manzelli (2017-2019; after BU: machine learning engineer at Marlo Inc.), Natalia Frumkin (2019-2020; after BU: PhD student at UT Austin)

Major Honors and Awards

- Peter J. Levine Career Development Professorship 2015
- NSF CAREER Award 2015
- Best Student Paper Award, 2008 IEEE Conference on Computer Vision and Pattern Recognition 2008
- Best Student Paper Award, 2007 International Conference on Machine Learning 2007
- Best Student Paper Award, 2005 International Conference on Machine Learning 2005
- MCD Fellowship from the University of Texas at Austin 2003
- Award of Excellence from College of Natural Sciences at the University of Texas at Austin 2003

Service

- **Senior Program Committee**
 - Area Chair, International Conference on Artificial Intelligence and Statistics (AISTATS) 2021
 - (Senior) Area Chair, AAAI Conference on Artificial Intelligence 2021
 - Area Chair, International Conference on Learning Representations (ICLR) 2021
 - Area Chair, Advances in Neural Information Processing Systems (NeurIPS) 2020
 - Area Chair, International Conference on Machine Learning (ICML) 2020
 - Area Chair, International Conference on Artificial Intelligence and Statistics (AISTATS) 2020
 - (Senior) Area Chair, AAAI Conference on Artificial Intelligence 2020
 - Area Chair, Advances in Neural Information Processing Systems (NeurIPS) 2019
 - Area Chair, International Conference on Machine Learning (ICML) 2019
 - Area Chair, International Conference on Artificial Intelligence and Statistics (AISTATS) 2019
 - Area Chair, Advances in Neural Information Processing Systems (NIPS) 2018
 - Area Chair, International Conference on Machine Learning (ICML) 2018
 - SPC, AAAI Conference on Artificial Intelligence 2018
 - Area Chair, Advances in Neural Information Processing Systems (NIPS) 2017
 - Area Chair, International Conference on Machine Learning (ICML) 2017
 - Area Chair, International Conference on Artificial Intelligence and Statistics (AISTATS) 2017
 - Area Chair, International Conference on Machine Learning (ICML) 2016
 - Area Chair, International Conference on Artificial Intelligence and Statistics (AISTATS) 2016
 - Area Chair, International Conference on Machine Learning (ICML) 2015
 - Area Chair, International Conference on Artificial Intelligence and Statistics (AISTATS) 2015
 - Area Chair, Advances in Neural Information Processing Systems (NIPS) 2014
 - Area Chair, International Conference on Machine Learning (ICML) 2014
 - Local Arrangements Chair, IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2014
 - Area Chair, International Conference on Machine Learning (ICML) 2013
- **Editorial Board**
 - Associate Editor, Computer Vision and Image Understanding (2016-2019)

- **Program Committee**

- International Conference on Artificial Intelligence and Statistics (AISTATS) (2013–2014)
- International Conference on Machine Learning (ICML) (2009–2012)
- IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2009–2013)
- International Conference on Computer Vision (ICCV) (2009, 2011)
- European Conference on Computer Vision (ECCV) (2010, 2014)
- ICDM Workshop on Optimization Methods for Emerging Data Mining Problems (OEDM) (2009–2011)

- **Journal Reviewing**

- IEEE Transactions on Pattern Analysis and Machine Intelligence
- Journal of Machine Learning Research
- International Journal on Computer Vision
- IEEE Transactions on Neural Networks
- Machine Learning Journal

- **Conference Reviewing**

- Neural Information Processing Systems (NIPS) (2006–2013, 2015–2016)
- International Conference on Machine Learning (ICML) (2005–2008)
- Uncertainty in Artificial Intelligence (UAI) (2015–2016)
- ACM SIGKDD Conference (2004–2009)
- SIAM Data Mining Conference (SDM) (2004–2007)
- IEEE International Conference on Data Mining (ICDM) (2004–2005)
- World Wide Web Conference (WWW) (2007)

- **Grant Reviewing**

- Panelist for the National Science Foundation (2013–2016)
- External Reviewer for the Swiss National Science Foundation

- **Workshop and Meeting Organization**

- ICML 2019 Workshop: “Joint Workshop on On-Device Machine Learning & Compact Deep Neural Network Representations (ODML-CDNNR)” (with Sujith Ravi, Zornitsa Kozareva, Lixin Fan, Max Welling, Yurong Chen, Werner Bailer, Haoji Hu, Jonathan Dekhtiar, Yingyan Lin, and Diana Marculescu)
- Co-chair, BU Data Science Day 2018
- ICCV 2013 Workshop: “VisDA: 1st International Workshop on Visual Domain Adaptation” (with Ruonan Li, Kate Saenko, and Fei Sha)
- NIPS 2011 Workshop: “Beyond Mahalanobis: Supervised Large-scale Learning of Similarity” (with Dhruv Batra, Greg Shakhnarovich, and Kilian Weinberger)
- ICCV 2011 Workshop: “Kernels and Distances in Computer Vision” (with Peter Gehler and Christoph Lampert)

- **Tutorial Organization**

- ECCV 2010 Tutorial: “Distance Functions and Metric Learning” (with Ofir Pele and Michael Werman)
- ICML 2010 Tutorial: “Metric Learning”

Publications

• Journal Publications

1. T. Campbell, B. Kulis, and J. How. *Dynamic Clustering Algorithms via Small-Variance Analysis of Markov Chain Mixture Models*. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 41, No. 6, pp. 1338–1352, 2019.
2. J. Hoffman, E. Rodner, J. Donahue, B. Kulis, and K. Saenko. *Asymmetric and Category Invariant Feature Transformations for Domain Adaptation*. International Journal of Computer Vision (IJCV), Vol. 109, Nos. 1-2, pp. 28–41, 2014.
3. B. Kulis. *Metric Learning: A Survey*. Foundations and Trends in Machine Learning, Vol. 5, No. 4, pp. 287–364, 2012.
4. B. Kulis and K. Grauman. *Kernelized Locality-Sensitive Hashing*. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 34, No. 6, pp. 1092–1104, 2012.
5. P. Jain, B. Kulis, J. Davis, and I. Dhillon. *Metric and Kernel Learning using a Linear Transformation*. Journal of Machine Learning Research, Vol. 13, pp. 519–547, 2012.
6. B. Kulis, P. Jain, and K. Grauman. *Fast Similarity Search for Learned Metrics*. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 31, No. 12, pp. 2143–2157, 2009. (Spotlight Paper, Dec. 2009 Issue)
7. B. Kulis, M. Sustik, and I. Dhillon. *Low-Rank Kernel Learning with Bregman Matrix Divergences*. Journal of Machine Learning Research, Vol. 10, pp. 341–376, 2009.
8. B. Kulis, S. Basu, I. Dhillon, and R. Mooney. *Semi-supervised Graph Clustering: A Kernel Approach*. Machine Learning, Vol. 74, No. 1, pp. 1–22, 2009. (Invited Paper)
9. I. Dhillon, Y. Guan, and B. Kulis. *Weighted Graph Cuts without Eigenvectors: A Multilevel Approach*. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 29, No. 11, pp. 1944–1957, 2007.
10. J. Hopcroft, O. Khan, B. Kulis, and B. Selman. *Tracking Evolving Communities in Large Linked Networks*. Proceedings of the National Academy of Sciences, Vol. 101, pp. 5249–5253, 2004.

• Conference Publications

1. R. Kumar, M. Rodehorst, J. Wang, J. Gu, and B. Kulis. *Building a Robust Word-Level Wakeword Verification Network*. Proc. 21st INTERSPEECH Conference, 2020.
2. J. Wang, R. Kumar, M. Rodehorst, B. Kulis, and S. Vitaladevuni. *An Audio-Based Wakeword-Independent Verification System*. Proc. 21st INTERSPEECH Conference, 2020.
3. H. Liu, A. Abhyankar, Y. Mishchenko, T. Senechal, G. Fu, B. Kulis, N. Stein, A. Shah, and S. Vitaladevuni. *Metadata-Aware End-to-End Keyword Spotting*. Proc. 21st INTERSPEECH Conference, 2020.
4. X. Xia, M. Zhang, T. Xue, Z. Sun, H. Fang, B. Kulis, and J. Chen. *Joint Bilateral Learning for Real-time Universal Photorealistic Style Transfer*. Proc. 16th European Conference on Computer Vision (ECCV), 2020.
5. K. Cilingir, R. Manzelli, and B. Kulis. *Deep Divergence Learning*. Proc. 37th International Conference on Machine Learning (ICML), 2020.
6. A. Siahkamari, A. Gangrade, B. Kulis, and V. Saligrama. *Piecewise Linear Regression via a Difference of Convex Functions*. Proc. 37th International Conference on Machine Learning (ICML), 2020.
7. X. Wang, S. Wang, P. Chen, Y. Wang, B. Kulis, X. Lin, and P. Chin. *Protecting Neural Networks with Hierarchical Random Switching: Towards Better Robustness-Accuracy Trade-off for Stochastic Defenses*. Proc. 28th International Joint Conference on Artificial Intelligence (IJCAI), 2019.
8. K. He, F. Cakir, X. Xia, B. Kulis, and S. Sclaroff. *Deep Metric Learning to Rank*. Proc. Computer Vision and Pattern Recognition Conference (CVPR), 2019.

9. R. Manzelli, V. Thakkar, A. Siahkamari, and B. Kulis. *Conditioning Deep Generative Raw Audio Models for Structured Automatic Music*. Proc. 19th International Society for Music Information Retrieval Conference (ISMIR), 2018.
10. A. Cutler and B. Kulis. *Inferring Human Traits from Facebook Statuses*. Proc. 10th International Conference on Social Informatics (Socinfo), 2018.
11. K. Jiang, S. Sra, and B. Kulis. *Combinatorial Topic Models using Small-Variance Asymptotics*. Proc. 20th International AISTATS Conference, 2017.
12. A. Roychowdhury, B. Kulis, and S. Parthasarathy. *Robust Monte Carlo Sampling using Riemannian Nosé-Poincaré Hamiltonian Dynamics*. Proc. 33rd International Conference on Machine Learning (ICML), 2016.
13. K. Jiang, Q. Que, and B. Kulis. *Revisiting Kernelized Locality-Sensitive Hashing for Improved Large-Scale Image Retrieval*. Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.
14. R. Finn and B. Kulis. *A Sufficient Statistics Construction of Exponential Family Lévy Measure Densities for Nonparametric Conjugate Models*. Proc. 18th International AISTATS Conference, 2015.
15. X. Zhou, J. Zhang, and B. Kulis. *Power-Law Graph Cuts*. Proc. 18th International AISTATS Conference, 2015.
16. A. Roychowdhury and B. Kulis. *Gamma Processes, Stick-Breaking, and Variational Inference*. Proc. 18th International AISTATS Conference, 2015.
17. A. Roychowdhury, K. Jiang, and B. Kulis. *Small-Variance Asymptotics for Hidden Markov Models*. Neural Information Processing Systems (NIPS), 2013. (Spotlight Presentation).
18. T. Campbell, M. Liu, B. Kulis, and J. How. *Dynamic Clustering via Asymptotics of the Dependent Dirichlet Process*. Neural Information Processing Systems (NIPS), 2013.
19. T. Broderick, B. Kulis, and M. I. Jordan. *MAD-Bayes: MAP-based Asymptotic Derivations from Bayes*. Proc. 30th International Conference on Machine Learning (ICML), 2013.
20. K. Jiang, B. Kulis and M. I. Jordan. *Small-Variance Asymptotics for Exponential Family Dirichlet Process Mixture Models*. Neural Information Processing Systems (NIPS), 2012.
21. J. Hoffman, B. Kulis, K. Saenko, and T. Darrell. *Discovering Latent Domains for Multisource Domain Adaptation*. Proc. 12th European Conference on Computer Vision (ECCV), 2012.
22. B. Kulis and M. I. Jordan. *Revisiting k-means: New Algorithms via Bayesian Nonparametrics*. Proc. 29th International Conference on Machine Learning (ICML), 2012.
23. B. Kulis, K. Saenko, and T. Darrell. *What You Saw Is Not What You Get: Domain Adaptation Using Asymmetric Kernel Transforms*. Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2011. **(Oral Presentation: 3.5% Acceptance Rate)**
24. M. E. Taylor, B. Kulis, and F. Sha. *Metric Learning for Reinforcement Learning Agents*. 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2011.
25. P. Jain, B. Kulis, and I. Dhillon. *Inductive Regularized Learning of Kernel Functions*. Neural Information Processing Systems (NIPS), 2010. (Spotlight Presentation).
26. K. Saenko, B. Kulis, M. Fritz, and T. Darrell. *Adapting Visual Category Models to New Domains*. Proc. 11th European Conference on Computer Vision (ECCV), 2010.
27. B. Kulis and P. Bartlett. *Implicit Online Learning*. Proc. 27th International Conference on Machine Learning (ICML), 2010.
28. B. Kulis and T. Darrell. *Learning to Hash with Binary Reconstructive Embeddings*. Neural Information Processing Systems (NIPS), 2009. (Spotlight Presentation)
29. B. Kulis and K. Grauman. *Kernelized Locality-Sensitive Hashing for Scalable Image Search*. Proc. 12th IEEE International Conference on Computer Vision (ICCV), 2009.
30. B. Kulis, S. Sra, and I. Dhillon. *Convex Perturbations for Scalable Semidefinite Programming*. Proc. 12th International AISTATS Conference, 2009.

31. P. Jain, B. Kulis, I. Dhillon, and K. Grauman. *Online Metric Learning and Fast Similarity Search*. Neural Information Processing Systems (NIPS), 2008. **(Oral Presentation: 2.7% Acceptance Rate)**
32. P. Jain, B. Kulis, and K. Grauman. *Fast Image Search for Learned Metrics*. Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2008. **(CVPR 2008 Best Student Paper Award)**
33. J. Davis, B. Kulis, P. Jain, S. Sra, and I. Dhillon. *Information-theoretic Metric Learning*. Proc. 24th International Conference on Machine Learning (ICML), 2007. **(ICML 2007 Best Student Paper Award)**
34. B. Kulis, A. Surendran, and J. Platt. *Fast Low-rank Semidefinite Programming for Embedding and Clustering*. Proc. 11th International AISTATS Conference, 2007.
35. B. Kulis, M. Sustik, and I. Dhillon. *Learning Low-Rank Kernel Matrices*. Proc. 23rd International Conference on Machine Learning (ICML), 2006.
36. I. Dhillon, Y. Guan, and B. Kulis. *A Fast Kernel-based Multilevel Algorithm for Graph Clustering*. Proc. 11th ACM SIGKDD Conference, 2005.
37. B. Kulis, S. Basu, I. Dhillon, and R. Mooney. *Semi-supervised Graph Clustering: A Kernel Approach*. Proc. 22nd International Conference on Machine Learning (ICML), 2005. **(ICML 2005 Best Student Paper Award)**
38. I. Dhillon, Y. Guan, and B. Kulis. *Kernel k -means, Spectral Clustering and Normalized Cuts*. Proc. 10th ACM SIGKDD Conference, 2004.
39. J. Hopcroft, O. Khan, B. Kulis, and B. Selman. *Natural Communities in Large Linked Networks*. Proc. 9th ACM SIGKDD Conference, 2003.

- **Preprints, Workshop Papers, and Technical Reports**

1. A. Siahkamari, V. Saligrama, D. Castanon, and B. Kulis. *Learning Bregman Divergences*. Arxiv: 1905.11545, 2019.
2. S. Sankarapandian, A. Kag, R. Manzelli, and B. Kulis. *Learning Compact Networks via Adaptive Network Regularization*. NIPS 2018 Workshop on Compact Deep Neural Networks with Industrial Applications, 2018.
3. R. Manzelli, V. Thakkar, A. Siahkamari, and B. Kulis. *An End-to-End Model for Automatic Music Generation: Combining Deep Raw and Symbolic Audio Networks*. 6th International Workshop on Musical Metacreation, 2018.
4. A. Cutler and B. Kulis. *Inferring Human Traits from Facebook Statuses*. Arxiv: 1805.08718, 2018.
5. X. Xia and B. Kulis. *W-Net: A Deep Model for Fully Unsupervised Image Segmentation*. Arxiv: 1711.08506, 2017.
6. B. Usman, K. Saenko, and B. Kulis. *Stable Distribution Alignment using the Dual of the Adversarial Distance*. International Conference on Learning Representations (ICLR) 2018 Workshop Track, 2018.
7. T. Campbell, B. Kulis, and J. How. *Dynamic Clustering Algorithms via Small-Variance Asymptotics of Markov Chain Mixture Models*. Arxiv: 1707.08493, 2017.
8. K. Jiang, S. Sra, and B. Kulis. *Combinatorial Topic Models using Small-Variance Asymptotics*. Arxiv: 1604.02027, 2016.
9. R. Finn and B. Kulis. *A Sufficient Statistics Construction of Bayesian Nonparametric Exponential Family Conjugate Models*. Arxiv: 1601.02257, 2016.
10. K. Jiang, Q. Que, and B. Kulis. *Revisiting Kernelized Locality-Sensitive Hashing for Improved Large-Scale Image Retrieval*. Arxiv: 1411.4199, 2014.
11. X. Zhou, J. Zhang, and B. Kulis. *Power-Law Graph Cuts*. Arxiv: 1411.1971, 2014. (Also presented at the NIPS 2014 Workshop on Networks: From Graphs to Rich Data.)
12. A. Roychowdhury and B. Kulis. *Gamma Processes, Stick-Breaking, and Variational Inference*. Arxiv: 1410.1068, 2014. (Also presented at the NIPS 2014 Workshop on Advances in Variational Inference.)
13. T. Broderick, B. Kulis, and M. I. Jordan. *MAD-Bayes: MAP-based Asymptotic Derivations from Bayes*. Arxiv: 1212.2126, 2012.

14. B. Kulis and M. I. Jordan. *Revisiting k-means: New Algorithms via Bayesian Nonparametrics*. Arxiv: 1111.0352, 2011.
15. K. Saenko, B. Kulis, M. Fritz, and T. Darrell. *Transferring Visual Category Models to New Domains*. University of California at Berkeley Technical Report # UCB/EECS-2010-54, May, 2010.
16. P. Jain, B. Kulis, J. V. Davis, and I. S. Dhillon. *Metric and Kernel Learning using a Linear Transformation*. ArXiv:0910.5932, October, 2009.
17. B. Kulis and T. Darrell. *Learning to Hash with Binary Reconstructive Embeddings*. University of California at Berkeley Technical Report # UCB/EECS-2009-101, July, 2009.
18. P. Jain, B. Kulis, and K. Grauman. *Fast Similarity Search for Learned Metrics*. University of Texas at Austin Technical Report # TR-07-48, September, 2007.
19. B. Kulis, S. Sra, S. Jegelka, and I. Dhillon. *Scalable Semidefinite Programming using Convex Perturbations*. University of Texas at Austin Technical Report # TR-07-47, September, 2007.
20. P. Jain, B. Kulis, and I. Dhillon. *Online Linear Regression using Burg Entropy*. University of Texas at Austin Technical Report # TR-07-08, February, 2007.
21. J. Davis, B. Kulis, S. Sra, and I. Dhillon. *Information-theoretic Metric Learning*. NIPS 2006 Workshop on Learning to Compare Examples, 2006.
22. I. Dhillon, Y. Guan and B. Kulis. *A Unified View of Kernel k-means, Spectral Clustering and Graph Cuts*. University of Texas at Austin Technical Report # TR-04-25, July, 2004.

Talks

- *New Directions in Metric Learning*
 CG Week Workshop on Comp. Aspects of Learning and Processing Metrical Data, Online June, 2020
 UMass Lowell Machine Learning Seminar, Lowell, MA, USA November, 2019
 Harvard Center for Mathematical Science and Applications, Cambridge, MA, USA September, 2019
- *Incorporating Context into Wakeword Cloud-side Verification*
 Amazon Research Days, Cambridge, MA, USA November, 2019
- *Scalable Nonparametric Machine Learning*
 Spotify Research Talk, Somerville, MA, USA March, 2019
 Amazon Research Talk, Cambridge, MA, USA February, 2019
 Boston University Data Science (BUDS) Day, Boston, MA, USA January, 2016
- *Small-Variance Asymptotics for Large-Scale Learning*
 Harvard Center for Mathematical Science and Applications, Cambridge, MA, USA January, 2019
 WPI ECE Graduate Seminar, Worcester, MA, USA April, 2016
 MIT Machine Learning Seminar, Cambridge, MA, USA, October, 2015
 Rutgers Computer Science Seminar, Piscataway, NJ, USA March, 2015
 Boston University ECE Seminar, Boston MA, USA March, 2015
 Northeastern University CCIS Colloquium, Boston, MA, USA February, 2015
 University of Pennsylvania GRASP Seminar, Philadelphia, PA, USA December, 2014
 Ohio State University CSE Focus on Faculty Talk, Columbus, OH, USA October, 2014
 Ohio State University Statistics Department Seminar, Columbus, OH, USA September, 2014
 University of Toronto Computer Science, Toronto, ON, Canada April, 2014
 Brown University Computer Science Seminar, Providence, RI, USA May, 2013
 Cornell University AI Seminar, Ithaca, NY, USA April, 2013
- *Domain Adaptation and Structured Divergence Learning*
 ICML Workshop on Divergences and Divergence Learning, Atlanta, GA, USA June, 2013
- *Small-Variance Asymptotics, Bayesian Nonparametrics, and k-means*
 NIPS Workshop on Graphical Models and Kernel Methods, Lake Tahoe, NV, USA December, 2012
 Mysore Learning Workshop, Mysore, India August, 2012

- *Revisiting Kernelized Hashing*
ICCV Workshop on Vision for the Web, Florence, Italy October, 2012
- *Revisiting k-means: New Algorithms via Bayesian Nonparametrics*
ICML 2012, Edinburgh, Scotland June, 2012
NYU CBLI Seminar, New York, NY, USA June, 2012
Columbia University Machine Learning Seminar, New York, NY, USA June, 2012
Ohio State University AI Seminar, Columbus, OH, USA April, 2012
Purdue University ML Seminar, West Lafayette, IN, USA March, 2012
TTI Chicago, Chicago, IL, USA March, 2012
- *Learning for Search and Adaptation in Large-Scale Data*
University of Washington Machine Learning Seminar, Seattle, WA, USA May, 2011
University of Rochester Computer Science, Rochester, NY, USA March, 2011
Harvard University Computer Science, Cambridge, MA, USA March, 2011
Boston University ECE, Boston, MA, USA March, 2011
Google Research, Mountain View, CA, USA March, 2011
Imperial College Department of Computing, London, UK February, 2011
Ohio State CSE, Columbus, OH, USA January, 2011
- *Learning Transformations for Adaptation and Search*
UC Berkeley SDH Seminar, Berkeley, CA, USA October, 2010
UC Davis Statistics Seminar, Davis, CA, USA October, 2010
Max Planck Institute, Tübingen, Germany September, 2010
University of Ljubljana, Ljubljana, Slovenia September, 2010
TTI Chicago, Chicago, IL, USA August, 2010
- *Implicit Online Learning*
ICML 2010, Haifa, Israel June, 2010
IPAM Lake Arrowhead Workshop, Lake Arrowhead, CA, USA June, 2010
UC Berkeley Machine Learning Tea, Berkeley, CA, USA October, 2009
- *Hashing Algorithms for Scalable Image Search*
IPAM Lake Arrowhead Workshop, Lake Arrowhead, CA, USA June, 2009
- *Scalable Metric Learning and Fast Similarity Search*
University of Michigan, Ann Arbor, MI, USA March, 2009
- *Scalable Metric Learning with the LogDet Divergence*
Yahoo! Research, Santa Clara, CA, USA February, 2009
- *Fast Image Search for Learned Metrics*
CVPR 2008, Anchorage, AK, USA June, 2008
- *Learning to Compare Objects*
MIT CSAIL Vision Interfaces Group, Cambridge, MA, USA June, 2008
University of North Carolina at Chapel Hill, Chapel Hill, NC, USA March, 2008
University of California at Santa Barbara, Santa Barbara, CA, USA February, 2008
- *Fast Image Search for Learned Metrics*
IPAM Lake Arrowhead Workshop, Lake Arrowhead, CA, USA December, 2007
- *Locality-sensitive Hashing and Fast Image Search*
IPAM Research Seminar at UCLA, Los Angeles, CA, USA October, 2007
- *Information-theoretic Metric Learning*
ICML 2007, Corvallis, OR, USA June, 2007
- *Spectral Clustering and Kernel k-means*
Two Guest Lectures for CS395T Data Mining Course, Austin, TX, USA April, 2007

- *Graph Cuts without Eigenvectors*
Microsoft Research, Redmond, WA, USA July, 2006
- *Learning Low-Rank Kernel Matrices*
ICML 2006, Pittsburgh, PA, USA June, 2006
- *Scalable Kernel Methods for Machine Learning*
Dissertation Defense, Austin, TX, USA November, 2008
Dissertation Proposal, Austin, TX, USA May, 2006
- *Semi-supervised Graph Clustering: A Kernel Approach*
ICML 2005, Bonn, Germany August, 2005