

Mikhail Belkin — Curriculum Vitae

Research Interests

- Mathematical, statistical and computational foundations of deep learning.
- Spectral, manifold and kernel methods for machine learning.
- Optimization and scalable implementations.
- Semi-supervised learning, clustering.

Academic Experience

- University of California, San Diego, Professor,
Halicioğlu Data Science Institute,
Computer Science and Engineering (affiliate appointment), 2020 – present.
- Ohio State University, Asst/Assoc/Full Professor,
Department of Computer Science and Engineering,
Department of Statistics (courtesy appointment), 2005 – 2020.
- Simons Institute for the Theory of Computing, Berkeley,
Visiting Faculty, Jan 2017 – May 2017, May 2019 – Aug 2019.
- The Institute of Science and Technology (IST Austria),
Visiting Professor, Sept 2012 – May 2013.
- Statistical and Applied Mathematical Sciences Institute (SAMSI),
Research Fellow, Jan 2007 – March 2007.
- University of Chicago, Department of Computer Science,
Postdoctoral Researcher, Dec 2003 – Aug 2005.
- Institute for Pure and Applied Mathematics, UCLA,
Multiscale Geometry and Analysis in High Dimensions, Fall 2004.
- University of California, Berkeley, Department of Statistics,
Visiting Research Fellow, Feb – Mar 2004.
- Max Planck Institute for Biological Cybernetics, Tübingen,
Research Scientist, Aug – Dec 2003.

Degrees and Education

- The University of Chicago, Ph.D. in Mathematics, 2003.
Thesis: Problems of Learning on Manifolds.
Thesis adviser: Partha Niyogi.
- The University of Chicago, M.Sc. in Mathematics, 1997.
- The University of Toronto, Hon.B.Sc. with High Distinction,
Major in Mathematics, 1995.

Industry Experience

- Amazon Scholar, 2022 – present.
- University Community Healthcare, Chicago.
– Senior Software Engineer, 1998 – 2000.

Grants, Awards and Scholarships

- 2023 ACM Fellow.
- NSF-Simons Collaboration on the Theoretical Foundations of Deep Learning (PI), 2020-2025.
- Google Research Award, 2019.
- NSF: RI: Small: Learning discrete structure from continuous spaces (PI), 2018-2021.
- NSF: Collaborative research: NCS-FO: Learning efficient visual representations from realistic environments across time scales (co-I), 2016-2020.
- NSF IIS:EAGER:The exploration of geometric and non-geometric structure in data (PI), 2015-2017.
- NSF CCF: Small: Geometry and High-dimensional Inference (co-PI), 2014-2017.
- NSF RI: Small: Algebraic and Spectral Structure of Data in High Dimension (single PI), 2011-2015.
- Lumley Research Award, College of Engineering, Ohio State University, 2011.
- NSF EAGER, Integration of Computational Geometry and Statistical Learning for Modern Data Analysis (co-PI), 2010 – 2012.
- AFOSR, Networks of Memories (co-PI), 2009 – 2012.
- NSF Conference Grant, 2009 Machine Summer School on Computational Learning (single PI), 2009.
- NSF Early Career Award: Geometry and High-dimensional Inference, 2007 – 2012.
- NSF/DARPA Computational and Algorithmic Representation of Geometric Objects incubationary grant (Senior Personnel), 2003 – 2004.
- University Fellowship, University of Chicago, 1996 – 1997.
- National Science and Eng. Research Council of Canada Postgraduate Scholarship, 1995 – 1997.
- Several merit-based undergraduate scholarships and awards, including Galois Award in Mathematics, Ted Mossman Scholarship, Alfred T. Delury Scholarship in Mathematics, Samuel Beatty Award.
- Moscow Mathematical Olympiad, second prize, 1990.

Journal papers

- *A universal trade-off between the model size, test loss, and training loss of linear predictors*, Nikhil Ghosh, Mikhail Belkin, SIAM Journal on Mathematics of Data Science (SIMODS), 5 (4), 2023.
- *Kernel ridgeless regression is inconsistent for low dimensions*, Daniel Beaglehole, Mikhail Belkin, Parthe Pandit, SIAM Journal on Mathematics of Data Science (SIMODS), 5(4), 2023.
- *Wide and deep neural networks achieve consistency for classification*, Adityanarayanan Radhakrishnan, Mikhail Belkin, Caroline Uhler, Proceedings of the National Academy of Sciences (PNAS), 120 (14), 2023.
- *Loss landscapes and optimization in over-parameterized non-linear systems and neural networks*, Chaoyue Liu, Libin Zhu, Mikhail Belkin, Applied and Computational Harmonic Analysis (ACHA), 59, 85-116, 2022.
- *Simple, fast, and flexible framework for matrix completion with infinite width neural networks*, Adityanarayanan Radhakrishnan, George Stefanakis, Mikhail Belkin, Caroline Uhler, Proceedings of the National Academy of Sciences (PNAS), 119 (16), 2022.
- *Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation* (by invitation), Mikhail Belkin, Acta Numerica, Volume 30, May 2021, pp. 203 - 248.
- *Classification vs regression in overparameterized regimes: Does the loss function matter?* Vidya Muthukumar, Adhyyan Narang, Vignesh Subramanian, Mikhail Belkin, Daniel Hsu, Anant Sahai, The Journal of Machine Learning Research (JMLR), 22 (1), 10104-10172.
- *Overparameterized Neural Networks Can Implement Associative Memory*, Adityanarayanan Radhakrishnan, Mikhail Belkin, Caroline Uhler, Proceedings of the National Academy of Sciences (PNAS), 2020, 117 (44).
- *Two models of double descent for weak features*, Mikhail Belkin, Daniel Hsu, Ji Xu, SIAM Journal on Mathematics of Data Science, 2(4), 1167–1180.
- *Reconciling modern machine-learning practice and the classical bias–variance trade-off*, Mikhail Belkin, Daniel Hsu, Siyuan Ma, Soumik Mandal. Proceedings of the National Academy of Sciences (PNAS), 116 (32) 15849-15854, 2019.
- *Back to the Future: Radial Basis Networks Revisited*, Qichao Que, Mikhail Belkin, IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 2019.
- *Eigenvectors of orthogonally decomposable functions*, Mikhail Belkin, Luis Rademacher, James Voss, SIAM Journal on Computing (SICOMP), 2018.
- *Polynomial learning of distribution families*, Mikhail Belkin, Kaushik Sinha, SIAM Journal on Computing (SICOMP), 44(4), 889-911, 2015.
- *Robust features for the automatic identification of autism spectrum disorder in children*, Justin Eldridge, Alison Lane, Mikhail Belkin, Simon Dennis, Journal of Neurodevelopmental Disorders, 2014.

- *The Geometry and Dynamics of Lifelogs: Discovering the Organizational Principles of Human Experience*, Vishnu Sreekumar, Simon Dennis, Isidoros Doxas, Yuwen Zhuang, Mikhail Belkin, PLOS One, 2014.
- *Heat flow and a faster algorithm to compute the surface area of a convex body*, M. Belkin, H. Narayanan, P. Niyogi, Random Structures & Algorithms, 43: 407–428, 2013.
- *Laplacian Support Vector Machines Trained in the Primal*, S. Melacci, M. Belkin, The Journal of Machine Learning Research, 12:1149–1184, 2011.
- *On Learning with Integral Operators*, L. Rosasco, M. Belkin, E. de Vito, The Journal of Machine Learning Research, 11(Feb):905-934, 2010.
- *Data Spectroscopy: Eigenspaces of Convolution Operators and Clustering*, Tao Shi, Mikhail Belkin, Bin Yu, The Annals of Statistics, vol. 37, Number 6B (2009), 3960-3984.
- *Towards a Theoretical Foundation for Laplacian-Based Manifold Methods*, M. Belkin, P. Niyogi, Journal of Computer and System Sciences, Volume 74 , Issue 8 (December 2008), pp. 1289-1308 Special Issue on Learning Theory, invited.
- *Consistency of Spectral Clustering*, U. von Luxburg, M. Belkin, O. Bousquet, The Annals of Statistics, 2008, Vol. 36, No. 2, 555-586.
- *Manifold Regularization: a Geometric Framework for Learning from Examples*, M. Belkin, P. Niyogi, V. Sindhwani, Journal of Machine Learning Research, 7(Nov):2399–2434, 2006.
- *Semi-supervised Learning on Riemannian Manifolds*, M. Belkin, P. Niyogi, Machine Learning, 56, 209-239, 2004. Special issue on clustering, invited.
- *Laplacian Eigenmaps for Dimensionality Reduction and Data Representation*, M. Belkin, P. Niyogi, Neural Computation, June 2003; 15 (6):1373-1396.

Selected preprints

- *Feature learning in neural networks and kernel machines that recursively learn features*, A Radhakrishnan, D Beaglehole, P Pandit, M Belkin, arXiv:2212.13881.
- *Mechanism of feature learning in convolutional neural networks*, D Beaglehole, A Radhakrishnan, P Pandit, M Belkin, arXiv:2309.00570.
- *Catapults in SGD: spikes in the training loss and their impact on generalization through feature learning*, L Zhu, C Liu, A Radhakrishnan, M Belkin, arXiv:2306.04815.

- *Linear Recursive Feature Machines provably recover low-rank matrices*,
A Radhakrishnan, M Belkin, D Drusvyatskiy, arXiv:2401.04553

Refereed Conference Proceedings

- *More is Better in Modern Machine Learning: when Infinite Overparameterization is Optimal and Overfitting is Obligatory*,
JB Simon, D Karkada, N Ghosh, M Belkin, ICLR 2024.
- *Quadratic models for understanding neural network dynamics*,
L Zhu, C Liu, A Radhakrishnan, M Belkin, ICLR 2024.
- *On the Nystrom Approximation for Preconditioning in Kernel Machines* A Abedsoltan, M Belkin, P Pandit, L Rademacher, AI&Stats 2024.
- *Aiming towards the minimizers: fast convergence of SGD for overparametrized problems*,
C Liu, D Drusvyatskiy, M Belkin, D Davis, YA Ma, NeurIPS 2023.
- *Cut Your Losses with Squentropy*,
L Hui, M Belkin, S Wright, ICML 2023.
- *Toward large kernel models*,
A Abedsoltan, M Belkin, P Pandit, ICML 2023.
- *Neural tangent kernel at initialization: linear width suffices*,
A Banerjee, P Cisneros-Velarde, L Zhu, M Belkin, UAI 2023.
- *Restricted strong convexity of deep learning models with smooth activations*,
A Banerjee, P Cisneros-Velarde, L Zhu, M Belkin, ICLR 2023.
- *Benign Overfitting in Two-layer Convolutional Neural Networks*, Yuan Cao, Zixiang Chen, Mikhail Belkin, Quanquan Gu, NeurIPS 2022.
- *Benign, Tempered, or Catastrophic: A Taxonomy of Overfitting*, N Mallinar, JB Simon, A. Abedsoltan, P Pandit, M Belkin, P Nakkiran, NeurIPS 2022.
- *Transition to Linearity of General Neural Networks with Directed Acyclic Graph Architecture*, L Zhu, C Liu, M Belkin, NeurIPS 2022.
- *Transition to Linearity of Wide Neural Networks is an Emerging Property of Assembling Weak Models (spotlight)*, Chaoyue Liu, Libin Zhu, Misha Belkin, ICLR 2022.
- *Evaluation of neural architectures trained with square loss vs cross-entropy in classification tasks*,
Like Hui, Mikhail Belkin, ICLR 2021.
- *Risk bounds for over-parameterized maximum margin classification on sub-gaussian mixtures*, Yuan Cao, Quanquan Gu, Mikhail Belkin, Neural Inf. Proc. Systems (NeurIPS) 2021.
- *Multiple Descent: Design Your Own Generalization Curve*, Lin Chen, Yifei Min, Mikhail Belkin, Amin Karbasi, Neural Inf. Proc. Systems (NeurIPS) 2021.
- *Evaluation of Neural Architectures Trained with Square Loss vs Cross-Entropy in Classification Tasks*,
Like Hui, Mikhail Belkin, ICLR 2021.

- *On the linearity of large non-linear models: when and why is the tangent kernel constant?* Chaoyue Liu, Libin Zhu, Mikhail Belkin, NeurIPS 2020 (spotlight presentation).
- *Accelerating Stochastic Training for Over-parametrized Learning*, Chaoyue Liu, Mikhail Belkin, ICLR 2020.
- *Kernel Machines Beat Deep Neural Networks on Mask-based Single-channel Speech Enhancement*, Like Hui, Siyuan Ma, Mikhail Belkin, INTERSPEECH 2019.
- *Learning kernels that adapt to GPU*, Siyuan Ma, Mikhail Belkin, SysML 2019.
- *Does data interpolation contradict statistical optimality?*, Mikhail Belkin, Alexander Rakhlin, Alexandre B. Tsybakov, AI Statistics 2019.
- *Overfitting or perfect fitting? Risk bounds for classification and regression rules that interpolate*, Mikhail Belkin, Daniel Hsu, Partha Mitra, Neural Information Processing Systems (NeurIPS), 2018.
- *To understand deep learning we need to understand kernel learning*, Mikhail Belkin, Siyuan Ma, Soumik Mandal, ICML 2018.
- *The power of interpolation: effectiveness of SGD in modern over-parametrized learning*, Siyuan Ma, Raef Bassily, Mikhail Belkin, ICML 2018.
- *Approximation beats concentration? An approximation view on inference with smooth kernels*, Mikhail Belkin, COLT 2018.
- *Unperturbed: spectral analysis beyond Davis-Kahan*, Justin Eldridge, Mikhail Belkin, Yusu Wang, ALT 2018.
- *Diving into the shallows: a computational perspective on large-scale shallow learning*, Siyuan Ma, Mikhail Belkin, Neural Information Processing Systems (NIPS), 2017 (spotlight presentation, 5% of submissions).
- *Graphons, mergeons, and so on!* Justin Eldridge, Mikhail Belkin, Yusu Wang, Neural Information Processing Systems (NIPS), 2016 (oral presentation, 2% of submissions).
- *Clustering with Bregman Divergences: an Asymptotic Analysis*, Chaoyue Liu, Mikhail Belkin, Neural Information Processing Systems (NIPS), 2016.
- *Learning a Hidden Basis Through Imperfect Measurements: An Algorithmic Primitive*, M. Belkin, L. Rademacher, J. Voss, The 29th Conference on Learning Theory (COLT), 2016.
- *Learning Privately from Multiparty Data*, J. Hamm, P. Cao, and M. Belkin, International Conference on Machine Learning (ICML), 2016

- *Back to the Future: Radial Basis Networks Revisited*,
Qichao Que, Mikhail Belkin,
The 19th International Conference on Artificial Intelligence and Statistics (AISTATS), 2016.
- *The Hidden Convexity of Spectral Clustering*,
James Voss, Mikhail Belkin, Luis Rademacher,
Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16), 2016, oral presentation.
- *Beyond Hartigan Consistency: Merge Distortion Metric for Hierarchical Clustering*,
Justin Eldridge, Mikhail Belkin, Yusu Wang,
The 28th Conference on Learning Theory (COLT), 2015, **Mark Fulk award** (best student paper).
- *A Pseudo-Euclidean Iteration for Optimal Recovery in Noisy ICA*,
James Voss, Mikhail Belkin, Luis Rademacher,
Neural Information Processing Systems (NIPS), 2015.
- *Crowd-ML: A Learning Framework For A Crowd of Smart Devices*,
J. Hamm, A. Champion, G. Chen, M. Belkin, and D. Xuan,
IEEE International Conference on Distributed Computing Systems (ICDCS), 2015.
- *Learning with Fredholm Kernels*,
Qichao Que, Mikhail Belkin, Yusu Wang
Neural Information Processing Systems (NIPS), 2014.
- *The More, the Merrier: the Blessing of Dimensionality for Learning Large Gaussian Mixtures*,
Joseph Anderson, Mikhail Belkin, Navin Goyal, Luis Rademacher, James Voss,
The 27th Conference on Learning Theory (COLT 2014).
- *Inverse Density as an Inverse Problem: The Fredholm Equation Approach*,
Qichao Que, Mikhail Belkin,
Neural Information Processing Systems (NIPS), spotlight presentation, 2013.
- *Fast Algorithms for Gaussian Noise Invariant Independent Component Analysis*,
James Voss, Luis Rademacher, Mikhail Belkin,
Neural Information Processing Systems (NIPS), 2013.
- *Blind Signal Separation in the Presence of Gaussian Noise*,
Mikhail Belkin, Luis Rademacher, James Voss,
The 26th Conference on Learning Theory (COLT 2013).
- *Toward understanding complex spaces: graph Laplacians on manifolds with singularities and boundaries*,
M. Belkin, Q. Que, Y. Wang, X. Zhou,
The 25th Conference on Learning Theory (COLT 2012).
- *Recognizing Daily Contexts from Multisensory Data: A Codebook Approach*,
J. Hamm, B. Stone, M. Belkin, S. Dennis,
Fourth Int. Conference on Mobile Computing, Applications and Services (MOBICASE), 2012
- *Data Skeletonization via Reeb Graphs*,
X. Ge, I. Safa, M. Belkin, Y. Wang,
Twenty-Fifth Annual Conference on Neural Information Processing Systems (NIPS), 2011.

- *An Iterated Graph Laplacian Approach for Ranking on Manifolds*,
X. Zhou, M. Belkin, N. Srebro,
17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2011.
- *Semi-supervised Learning by Higher Order Regularization* ,
X. Zhou, M. Belkin,
14th International Conference on Artificial Intelligence and Statistics (AISTATS), 2011.
- *Polynomial Learning of Distribution Families*,
M. Belkin, K. Sinha,
51st Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2010.
- *Toward Learning Gaussian Mixtures with Arbitrary Separation*,
M. Belkin, K. Sinha,
The 23rd Annual Conference on Learning Theory (COLT), 2010.
- *Learning speaker normalization using semisupervised manifold alignment*,
A. Plummer, M. Beckman, M. Belkin, E. Fosler-Lussier, and B. Munson,
Proc. Interspeech, 2010.
- *Semi-supervised Learning Using Sparse Eigenfunction Bases*,
K. Sinha, M. Belkin,
Twenty-Third Annual Conference on Neural Information Processing Systems (NIPS), 2009.
- *A Note on Learning with Integral Operators*,
L. Rosasco, M. Belkin, E. de Vito,
The 22nd Annual Conference on Learning Theory (COLT), 2009.
- *Constructing Laplace Operator from Point Clouds in \mathbf{R}^d*
M. Belkin, J. Sun, Y. Wang,
ACM-SIAM Symposium on Discrete Algorithms (SODA), 2009.
- *Data Spectroscopy: Learning Mixture Models using Eigenspaces of Convolution Operators*,
T. Shi, M. Belkin, B. Yu,
The 25th International Conference on Machine Learning (ICML), 2008.
- *Component Based Shape Retrieval Using Differential Profiles*,
L. Ding, M. Belkin,
ACM International Conference on Multimedia Information Retrieval (MIR), 2008.
- *Probabilistic Mixtures of Differential Profiles for Shape Recognition*,
L. Ding, M. Belkin,
The 19th International Conference on Pattern Recognition (ICPR), 2008.
- *Discrete Laplace Operator for Meshed Surfaces*,
M. Belkin, J. Sun, Y. Wang,
The 24th Annual Symposium on Computational Geometry (SOCG), 2008.
- *The Value of Labeled and Unlabeled Examples when the Model is Imperfect*,
K. Sinha, M. Belkin, 21st Annual Conference on Neural Information Processing Systems (NIPS),
spotlight presentation, 2007.

- *Convergence of Laplacian Eigenmaps*,
M. Belkin, P. Niyogi,
Twentieth Annual Conference on Neural Information Processing Systems (NIPS), 2006.
- *On the Relation Between Low Density Separation, Spectral Clustering and Graph Cuts*,
H. Narayanan, M. Belkin, P. Niyogi,
Twentieth Annual Conference on Neural Information Processing Systems (NIPS), 2006.
- *Heat Flow and a Faster Algorithm to Compute the Surface Area of a Convex Body*,
M. Belkin, H. Narayanan, P. Niyogi,
47th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2006.
- *Maximum Margin Semi-Supervised Learning for Structured Variables*,
Y. Altun, D. McAllester, M. Belkin,
Nineteenth Annual Conference on Neural Information Processing Systems (NIPS), 2005.
- *Beyond the Point Cloud: from Transductive to Semi-supervised Learning* ,
V. Sindhwani, P. Niyogi, M. Belkin,
The 22nd International Conference on Machine Learning (ICML) 2005.
- *Towards a Theoretical Foundation for Laplacian-based Manifold Methods*
M. Belkin, P. Niyogi, The Eighteenth Annual Conference on Learning Theory (COLT), 2005.
- *On Manifold Regularization*,
M. Belkin, P. Niyogi, V. Sindhwani,
The Tenth International Workshop on Artificial Intelligence and Statistics (AISTATS), 2005.
- *Limits of Spectral Clustering*,
U. von Luxburg, O. Bousquet, M. Belkin,
Neural Information Processing Systems (NIPS), 2004, **outstanding student paper award**.
- *Regularization and Semi-Supervised Learning on Large Graphs*,
M. Belkin, I. Matveeva, P. Niyogi,
The Seventeenth Annual Conference on Learning Theory (COLT), 2004.
- *On the Convergence of Spectral Clustering on Random Samples: the Normalized Case*,
U. von Luxburg, O. Bousquet, M. Belkin,
The Seventeenth Annual Conference on Learning Theory (COLT), 2004.
- *Tikhonov Regularization and Semi-Supervised Learning on Large Graphs (Invited)*,
M. Belkin, I. Matveeva, P. Niyogi,
2004 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2004),
Special Session: Manifolds and Geometry and Signal Processing, 2004.
- *Using Manifold Structure for Partially Labeled Classification*,
M. Belkin, P. Niyogi,
Sixteenth Annual Conference on Neural Information Processing Systems (NIPS), 2002.
- *Using Eigenvectors of the Bigram Graph to Infer Morpheme Identity*,
M. Belkin, J. Goldsmith,
Proceedings of the Morphology/Phonology Learning Workshop of ACL-02, Association for Computational Linguistics, 2002.

- *Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering*, M. Belkin, P. Niyogi, Fifteenth Annual Conference on Neural Information Processing Systems (NIPS), 2001.

Book Chapters

- *Semi-Supervised Learning*, Xueyuan Zhou, Mikhail Belkin, Academic Press Library in Signal Processing: Signal Processing Theory and Machine Learning, 2013.
- *The Geometric Basis of Semi-supervised Learning*, V. Sindhwani, M. Belkin, P. Niyogi, Semi-supervised Learning (Chapelle, Schoelkopf, Zien: editors), MIT Press, 2006.

Ph.D. Thesis

- *Problems of Learning on Manifolds*, University of Chicago, Department of Mathematics, 2003. Thesis adviser: Partha Niyogi.

Selected Recent Editorial and Program Committee Service

- SIAM Journal on Mathematics of Data Science (SIMODS), Editor-in-chief, from 2024.
- COLT 2021 Program Chair (with Samory Kpotufe).
- SIAM Journal on Mathematics of Data Science (SIMODS), Associate Editor, 2020 – 2024.
- The Journal of Machine Learning Research, Action Editor, 2011 – 2018.
- IEEE Transactions on Pattern Recognition and Machine Intelligence, Associate Editor, 2011 – 2016.

Conference (co-)organizer

- Oberwolfach workshop on Interpolation and Over-parameterization in Statistics and Machine Learning, Sept 2023.
- MSRI Summer School on Machine Learning, UCSD, 2023.
- Summer Cluster: Deep Learning Theory, Simons Institute for Theory of Computing, 2022.
- Siam Conference on Mathematics of Data Science (MDS), San Diego, 2022.
- TOPML (Workshop on the Theory of Overparameterized Machine Learning) 2021, 2022.
- Midwest ML Symposium, University of Chicago/TTIC, 2017, 2018.
- Information Modeling and Control of Complex Systems Workshop, Ohio State University, 2016, 2017.
- Simons Institute Workshop on Spectral Algorithms: From Theory to Practice (co-chair), 2014.
- Partha Niyogi Memorial Conference, The University of Chicago, 2011.
- 2010 SDM (Siam Int. Conference on Data Mining) Workshop on Clustering: Theory and applications.
- 2009 AAAI Fall Symposium on Manifold Learning.
- 2009 Machine Learning Summer School/Workshop on Theory and Practice of Comput. Learning.
- Workshop on Geometry, Random Matrices, and Statistical Inference, SAMSI, Jan 2007.
- 2005 Chicago Machine Learning Summer School.

Refereed and served on panels/program committees for:

- The Annals of Statistics, Science, IEEE Transactions on Pattern Analysis and Machine Intelligence, Bernoulli, ACM Transactions on Sensor Networks, IEEE Transactions on Image Processing, PNAS, IEEE Transactions on Signal Processing, International Journal of Computer Vision, Machine Learning Journal, Journal of Machine Learning Research, Journal of the American Statistical Association (JASA), Pattern Recognition, FOCS, NIPS, ICML, COLT, AISTATS, AAAI, ICCV, AFOSR, NSF, RGC Hong Kong, US-Israel Binational Science Foundation.

Selected Invited Talks

- Mathematics and Machine Learning 2023, CalTech (plenary), Dec 2023.
- UC Berkeley CLIMB Seminar, Nov 2023.
- C3.AI Workshop on Generative AI, UIUC, Oct 2023.
- Simons Foundation 2023 Mathematical and Scientific Foundations of Deep Learning Annual Meeting, Sept 2023.
- Samsung Global Research Symposium on Mathematical Theory of AI, Seoul, South Korea, Aug 2023.
- Eric and Wendy Schmidt Center, Broad Institute and MIT EECS Joint Colloquium Series, May 2023.
- Distinguished Seminar in Optimization & Data, University of Washington, May 2023.
- Future of Large Scale Machine Learning Workshop, Rice University, April 2023.
- Statistical Machine Learning Symposium, Columbia University, April 2023.
- PSU-Purdue-UMD Joint Seminar on Mathematical Data Science, March 2023.
- NIST Applied and Computational Mathematics Division seminar, Feb 2023.
- Winter School on Deep Learning 2023, Indian Statistical Institute (plenary), Jan 2023.
- Optimization and Statistical Learning, OSL 2023, Les Houches, France, Jan 2023.
- Third Workshop on Seeking Low-Dimensionality in Deep Neural Networks, Abu Dhabi, Jan 2023.
- A Short Course on the Mathematical Aspects of Deep Learning at the Institut Henri Poincare, Oct. 2022.
- Data Science Week 2022, Purdue University (keynote), Nov 2022.
- 8th International Conference on Computational Harmonic Analysis (plenary), Sept 2022.
- SIAM Conference on Mathematics of Data Science, Sept 2022.
- Canadian Applied and Industrial Mathematics Society Annual Meeting (plenary), June 2022.
- University of Chicago Statistics Colloquium, May 2022.
- Oxford ML and Physics Seminar, 2022.
- CalTech H.B. Keller Colloquium, 2022.
- CMU Philosophy Colloquium, 2022.
- Machine Learning Across Disciplines: New Theoretical Developments, The Griffin Applied Economics Incubator Conference, 2022.
- DeepLearn Summer School (tutorial), 2022.
- MIT IDSS Distinguished Seminar, 2022.
- University of Minnesota Machine Learning Seminar, Sept. 2021.

- IBM ML Seminar, Oct 2021.
- First MODE Workshop on Differentiable Programming for Experiment Design (keynote), Sept 2021.
- Princeton Summer School on Deep Learning (5-hour tutorial), July 2021.
- ICML workshop on opportunities and pitfalls of over-parameterization, July 2021.
- DeepMind ML seminar, July 2021.
- GdR ISIS workshop on deep learning theory, June 2021.
- CMU/Duolingo ML seminar, June 2021.
- Google ML Seminar, March 2021.
- ML4H Seminar, Broad Institute, March 2021.
- Seminar in Mathematics, Physics and Machine Learning, Lisbon, Apr 2021.
- One World Mathematics of Information, Data, and Signals (MINDS) Seminar, Feb 2021.
- DeepMath 2020 conference, Nov 2020.
- IEEE DeepTech.AI 2020 (keynote), Bangalore, India, Oct 2020.
- Stanford Statistics Seminar, Oct 2020.
- Neymann Seminar, Statistics, University of California Berkeley, Oct 2020.
- National Academy of Sciences Webinar. Deep Learning and Neural Networks: Illustrating the Impact of the Mathematical Sciences, Sept 2020.
- 14th Annual Machine Learning Symposium at the New York Academy of Sciences (keynote), March 2020.
- Eastern European Machine Learning Summer School, July 2020.
- Yale Institute for Network Science, Distinguished seminar, Yale, Jan 2020.
- Duke University CS colloquium, Duke, Jan 2020.
- TTIC colloquium, Chicago, Nov 2019.
- Cornell Statistics and Data Science Seminar, Ithaca, Oct 2019.
- The Center for Brains, Minds and Machines, MIT, Oct 2019.
- Institute for Advanced Study, Workshop on Theory of Deep Learning: Where next, Oct 2019.
- Humboldt University, Workshop on High Dimensional Covariance Operators, Berlin, Sept 2019.
- International Conference on Continuous Optimization, Berlin, Aug 2019.
- Simons Institute, Workshop on Frontiers of Deep Learning, Berkeley, July 2019.
- OpenAI colloquium, July 2019.
- ICML Workshop on Understanding and Improving Generalization in Deep Learning (keynote), Los Angeles, July 2019.
- Conference on Geometric Data Analysis, University of Chicago, May 2019.
- University of California San Diego, HDSI colloquium, Apr 2019.
- University College London, Gatsby unit, Apr 2019.
- Machine Learning Advances and Applications Seminar, Fields Institute, Univ. of Toronto, Nov 2018.
- SILO Seminar, UW Madison, Nov 2018.
- Statistics/Machine Learning Seminar, CMU, Oct 2018.

- Workshop on Approximation Theory and Machine Learning Workshop, Purdue University, Sept 2018.
- Workshop on Optimization and Machine Learning, Lehigh University, Aug 2018.
- The Center for Brains, Minds and Machines, MIT, Apr 2018.
- ORFE colloquium, Princeton University, Feb 2018.
- Information Theory and Application Workshop, San Diego, Feb 2018.
- Workshop on Inverse Problems in Machine Learning, Caltech, Jan 2018.
- Workshop on Computational Brain Research at IIT Madras, Chennai, India, Jan 2018.
- Simons Institute, Workshop on Optimization Statistics and Uncertainty, Berkeley, Nov 2017.
- University of Chicago, Computer Science, Nov 2017.
- NCState, Workshop on Data Science, Nov 2017.
- NYU, CS, New York, Oct 2017.
- APS INFORMS, Chicago, July 2017.
- ENSAE, CREST, Paris, France, June 2017.
- Workshop on statistical aspects of geometric and topological data anal., Banyuls, France, June 2017.
- Neymann Seminar, Statistics, University of California Berkeley, May 2017.
- UCLA CS Colloquium, Apr 2017.
- Oberwolfach workshop on Stat. Recovery of Discrete, Geometric and Invariant Structures, Apr 2017.
- Information Theory and Application Workshop, San Diego, Feb 2017.
- Workshop on Computational Brain Research at IIT Madras, Chennai, India, Jan 2017.
- IDSS Statistics and Stochastics Seminar, MIT, Oct 2016.
- Data Science Initiative (DSI) Colloquium, Boston University, 2016.
- ICML 2016 Workshop on Geometry in Machine Learning, June 2016.
- Simons Institute for Data Analysis, NYC, June 2016.
- Wilks Statistics Seminar, ORFE, Princeton University, March 2016.
- Shape Analysis and Learning by Geometry and Machine, IPAM, UCLA, Feb 2016.