Ankur Moitra

CONTACT INFORMATION

Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139 USA

Cell: (607) 227-6851
Offices: 2-472 and 32-G594
E-mail: moitra@mit.edu
http://people.csail.mit.edu/moitra/

CITIZENSHIP

USA

RESEARCH INTERESTS

theoretical computer science, machine learning

APPOINTMENTS

Massachusetts Institute of Technology, Department of Mathematics

Norbert Wiener Professor, Summer 2021 – present
Associate Professor with Tenure, Summer 2019 – Summer 2021
Associate Professor, Summer 2017 – Summer 2019
Assistant Professor, Fall 2013 – Summer 2017

University of California, Berkeley, Electrical Engineering and Computer Science

Chancellor's Professor, Fall 2024

Statistics and Data Science Center (SDSC)

Director, Summer 2021 – present

Computer Science and Artificial Intelligence Lab (CSAIL)

Principal Investigator, Fall 2013 – present

Institute for Advanced Study, School of Mathematics

NSF Computing and Innovation Fellow, Fall 2011 – Summer 2013
Also: *Senior Postdoc* at Princeton University, Summer 2012 – Summer 2013

EDUCATION

Massachusetts Institute of Technology, Electrical Engineering and Computer Science

Ph.D. in Computer Science, June 2011
Thesis: *Vertex Sparsification and Universal Rounding Algorithms*
Advisor: F. Thomson Leighton
Received the George M. Sprowls Thesis Award (best thesis)

Massachusetts Institute of Technology, Electrical Engineering and Computer Science

S.M. in Computer Science, May 2009
Thesis: *A Solution to the Papadimitriou-Ratajczak Conjecture*
Advisor: F. Thomson Leighton
Received the William A. Martin Memorial Thesis Award (best thesis)

Cornell University, Electrical and Computer Engineering

B.S. *Summa Cum Laude*, May 2007
Class Rank in School of Engineering: 1^{st} of 714
Minor in Applied Mathematics

HONORS AND AWARDS

QIP Invited Plenary, 2025
Common Ground Award for Excellence in Teaching, 2024
QIP Invited Plenary, 2024
AMS von Neumann Lecture, 2024
Norbert Wiener Chair, 2021 – present

Invited Paper ("Robustness Meets Algorithms") in Communications of the ACM, Research Highlights, 2021
School of Science Prize for Excellence in Graduate Teaching, 2018
ONR Young Investigator Award, 2018
Rockwell International Chair, 2016 – 2019
David and Lucile Packard Foundation Fellow, 2016
Alfred P. Sloan Research Fellow, 2016
Edmund F. Kelly Research Award, 2015 – 2018
NSF CAREER Award, 2015
Invited Paper ("Learning Topic Models — Provably and Efficiently") in Communications of the ACM, Research Highlights, 2015
Google Research Award, 2014
NSF Computing and Innovation Fellow, 2011 – 2013
George M. Sprowls Thesis Award, 2011
Invited Paper ("Disentangling Gaussians") in Communications of the ACM, Research Highlights, 2011
William A. Martin Memorial Thesis Award, 2009
Siebel Scholar, Class of 2009
Fannie and John Hertz Foundation Fellow, 2008 – 2011
MIT Presidential Fellowship, 2007 – 2008
Barry M. Goldwater Scholar, 2006 – 2007

Publications

A. Liu, A. Moitra. Model Stealing for Any Low-Rank Language Model. *Proceedings of the 57th Annual ACM Symposium on Theory of Computing* (STOC 2025), to appear.

J. Gaitonde, A. Moitra, E. Mossel. Bypassing the Noisy Parity Barrier: Learning Higher-Order Markov Random Fields from Dynamics. *Proceedings of the 57th Annual ACM Symposium on Theory of Computing* (STOC 2025), to appear.

A. Moitra, A. Wein. Precise Error Rates for Computationally Efficient Testing. *Annals of Statistics*, to appear.

N. Golowich, A. Moitra. Edit Distance Robust Watermarks for Language Models. *Advances in Neural Information Processing Systems 37* (NeurIPS 2024), to appear.

A. Bakshi, A. Liu, A. Moitra, E. Tang. Structure Learning of Hamiltonians from Real-time Evolution. *Proceedings of the 65th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2024), to appear.

A. Bakshi, A. Liu, A. Moitra, E. Tang. High-Temperature Gibbs States are Unentangled and Efficiently Preparable. *Proceedings of the 65th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2024), to appear.

N. Golowich, A. Moitra, D. Rohatgi. Exploration is Harder than Prediction: Cryptographically Separating Reinforcement Learning from Supervised Learning. *Proceedings of the 65th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2024), to appear.

N. Golowich, A. Moitra. The Role of Inherent Bellman Error in Offline Reinforcement Learning with Linear Function Approximation. *Proceedings of the 1st Annual Reinforcement Learning Conference* (RLC 2024), to appear.

N. Golowich, A. Moitra. Linear Bellman Completeness Suffices for Efficient Online Reinforcement Learning with Few Actions. *Proceedings of the 37th Annual Conference on Learning Theory* (COLT 2024),pages 1939–1981.

B. Chin, A. Moitra. E. Mossel, C. Sandon. he Power of an Adversary in Glauber Dynamics. *Proceedings of the 37th Annual Conference on Learning Theory* (COLT 2024), pages 1102–1124.

N. Golowich, A. Moitra, D. Rohatgi. Exploring and Learning in Sparse Linear MDPs without Computationally Intractable Oracles.*Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (STOC 2024), pages 183–193.

A. Bakshi, A. Liu, A. Moitra, E. Tang. Learning Quantum Hamiltonians at Any Temperature in Polynomial Time. *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (STOC 2024), pages 1470–1477. **Invited to the SIAM Journal on Computing Special Issue**.

C. Pabbaraju, D. Rohatgi, A. Sevekari, H. Lee, A. Moitra, A. Risteski. Provable Benefits of Score Matching. *Advances in Neural Information Processing Systems 36* (NeurIPS 2023). **Spotlight**.

Z. Chen, K. Liu, N. Mani, A. Moitra. Strong Spatial Mixing for Colorings on Trees and its Algorithmic Applications. *Journal of the ACM*, to appear. *Proceedings of the 64th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2023), pages 810–845.

A. Bakshi, A. Liu, A. Moitra, M. Yau. Tensor Decompositions Meet Control Theory: Learning General Mixtures of Linear Dynamical Systems. *Proceedings of the 40th International Conference on Machine Learning* (ICML 2023), pages 1549–1563.

A. Bakshi, A. Liu, A. Moitra, M. Yau. A New Approach to Learning Linear Dynamical Systems. *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (STOC 2023), pages 335–348.

N. Golowich, A. Moitra, D. Rohatgi. Planning and Learning in Partially Observable Systems via Filter Stability. *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (STOC 2023), pages 349–362.

S. Jain, H. Lawrence, A. Moitra, A. Madry. Distilling Model Failures as Directions in Latent Space. *Eleventh International Conference on Learning Representations* (ICLR 2023). **Spotlight**.

A. Moitra, D. Rohatgi. Provably Auditing Ordinary Least Squares in Low Dimensions. *Eleventh International Conference on Learning Representations* (ICLR 2023).

Z. Chen, N. Mani, A. Moitra. From Algorithms to Connectivity and Back: Finding a Giant Component in Random k-SAT. *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2023), pages 3437–3470.

A. Liu, A. Moitra. Robust Voting Rules from Algorithmic Robust Statistics. *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2023), pages 3471–3512.

N. Golowich, A. Moitra, D. Rohatgi. Learning in Observable POMDPs, without Computationally Intractable Oracles. *Advances in Neural Information Processing Systems 35* (NeurIPS 2022).

J. Li, A. Liu, A. Moitra. Robust Model Selection and Nearly-Proper Learning for GMMs. *Advances in Neural Information Processing Systems 35* (NeurIPS 2022).

D. Cifuentes, A. Moitra. Polynomial Time Guarantees for the Burer-Monteiro Method. *Advances in Neural Information Processing Systems 35* (NeurIPS 2022).

A. Liu, A. Moitra. Minimax Rates for Robust Community Detection. *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2022), pages 823–831.

N. Golowich, A. Moitra. Can Q-Learning be Improved with Advice? *Proceedings of the 35th Annual Conference on Learning Theory* (COLT 2022), pages 4548–4619.

A. Liu, A. Moitra. Learning GMMs with Nearly Optimal Robustness Guarantees. *Proceedings of the 35th Annual Conference on Learning Theory* (COLT 2022), pages 2815–2895.

S. Chen, F. Koehler, A. Moitra, M. Yau. Kalman Filtering with Adversarial Corruptions. *Proceedings of the 54th Annual ACM Symposium on Theory of Computing* (STOC 2022), pages 832–845.

L. Hamilton, A. Moitra. A No-go Theorem for Acceleration in the Hyperbolic Plane. *Advances in Neural Information Processing Systems 34* (NeurIPS 2021), pages 3914–3924.

S. Chen, F. Koehler, A. Moitra, M. Yau. Online and Distribution-Free Robustness: Regression and Contextual Bandits with Huber Contamination. *Proceedings of the 62nd Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2021), pages 684–695.

A. Moitra, E. Mossel, C. Sandon. Learning to Sample from Censored Markov Random Fields. *Proceedings of the 34th Annual Conference on Learning Theory* (COLT 2021), pages 3419–3451.

I. Diakonikolas, G. Kamath, D. Kane, J. Li, A. Moitra, A. Stewart. Robustness Meets Algorithms. *Communications of the ACM* May 2021, **Research Highlights**, pages 107–115.

A. Liu, A. Moitra. Settling the Robust Learnability of Mixtures of Gaussians. *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing* (STOC 2021), pages 518–531.

S. Chen, A. Moitra. Algorithmic Foundations for the Diffraction Limit. *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing* (STOC 2021), pages 490–503.

S. Chen, F. Koehler, A. Moitra, M. Yau. Classification Under Misspecification: Halfspaces, Generalized Linear Models, and Connections to Evolvability. *Advances in Neural Information Processing Systems 33* (NeurIPS 2020). **Spotlight**.

A. Liu, A. Moitra. Tensor Completion Made Practical. *Advances in Neural Information Processing Systems 33* (NeurIPS 2020).

S. Chen, J. Li, A. Moitra. Learning Structured Distributions From Untrusted Batches: Faster and Simpler. *Advances in Neural Information Processing Systems 33* (NeurIPS 2020).

J. Kelner, F. Koehler, R. Meka, A. Moitra. Learning Some Popular Gaussian Graphical Models without Condition Number Bounds. *Advances in Neural Information Processing Systems 33* (NeurIPS 2020). **Spotlight**.

C. Franks, A. Moitra. Rigorous Guarantees for Tyler's M-estimator via Quantum Expansion. *Proceedings of the 33rd Annual Conference on Learning Theory* (COLT 2020), pages 1601–1632.

A. Moitra, E. Mossel, C. Sandon. The Circuit Complexity of Inference. *Proceedings of the 33rd Annual Conference on Learning Theory* (COLT 2020), pages 2910–2946.

A. Liu, A. Moitra. Better Algorithms for Estimating Non-Parametric Models in Crowd-Sourcing and Rank Aggregation. *Proceedings of the 33rd Annual Conference on Learning Theory* (COLT 2020), pages 2780–2829.

S. Chen, J. Li, A. Moitra. Efficiently Learning Structured Distributions from Untrusted Batches. *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing* (STOC 2020), pages 960–973.

A. Moitra, A. Wein. Spectral Methods from Tensor Networks. *Proceedings of the 51st Annual ACM Symposium on Theory of Computing* (STOC 2019), pages 926–937. **Invited to the SIAM Journal on Computing Special Issue**.

G. Bresler, F. Koehler, A. Moitra. Learning Restricted Boltzmann Machines via Influence Maximization. *Proceedings of the 51st Annual ACM Symposium on Theory of Computing* (STOC 2019), pages 828–839.

S. Chen, A. Moitra. Beyond the Low-Degree Algorithm: Mixtures of Subcubes and Their Applications. *Proceedings of the 51st Annual ACM Symposium on Theory of Computing* (STOC 2019), pages 869–880.

Y. Kim, F. Koehler, A. Moitra, E. Mossel, G. Ramnarayan. How Many Subpopulations is Too Many? Exponential Lower Bounds for Inferring Population Histories. *Proceedings of the 23rd International Conference on Research in Computational Molecular Biology* (RECOMB 2019), pages 136–157. **Invited to the Journal of Computational Biology Special Issue**.

L. Hamilton, A. Moitra. The Paulsen Problem Made Simple. *Israel Journal of Mathematics*, to appear. *Proceedings of the 10th Annual Innovations in Theoretical Computer Science* (ITCS 2019), pages 1–6.

S. Chen, M. Delcourt, A. Moitra, G. Perarnau, L. Postle. Improved Bounds for Sampling Colorings via Linear Programming. *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2019), pages 2216–2234.

A. Liu, A. Moitra. Efficiently Learning Mixtures of Mallows Models. *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2018), pages 627–638.

S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu, M. Zhu. Learning Topic Models — Provably and Efficiently. *Communications of the ACM* April 2018, **Research Highlights**, pages 85–93.

A. Perry, A. Wein, A. Bandeira, A. Moitra. Optimality and Suboptimality of PCA I: Spiked Random Matrix Models. *Annals of Statistics* 46(5):2416–2451, 2018.

I. Diakonikolas, G. Kamath, D. Kane, J. Li, A. Moitra, A. Stewart. Robustly Learning a Gaussian: Getting Optimal Error, Efficiently. *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2018), pages 2683–2702.

L. Hamilton, F. Koehler, A. Moitra. Information Theoretic Properties of Markov Random Fields, and Their Algorithmic Applications. *Advances in Neural Information Processing Systems 30* (NIPS 2017), pages 2460–2469.

A. Perry, A. Wein, A. Bandeira, A. Moitra. Message-passing Algorithms for Synchronization Problems over Compact Groups. *Communications on Pure and Applied Mathematics*, 11(71):2275–2322, 2018.

J. Urschel, V. Brunel, A. Moitra, P. Rigollet. Learning Determinantal Point Processes with Moments and Cycles. *Proceedings of the 34th International Conference on Machine Learning* (ICML 2017), pages 3511–3520.

I. Diakonikolas, G. Kamath, D. Kane, J. Li, A. Moitra, A. Stewart. Being Robust (in High Dimensions) Can Be Practical. *Proceedings of the 34th International Conference on Machine Learning* (ICML 2017), pages 999–1008.

V. Brunel, A. Moitra, P. Rigollet, J. Urschel. Rates of Estimation for Determinantal Point Processes. *Proceedings of the 30th Annual Conference on Learning Theory* (COLT 2017), pages 343–345.

A. Moitra. Approximate Counting, the Lovasz Local Lemma and Inference in Graphical Models. *Journal of the ACM*,66(2):1–25, 2019. Preliminary version in *Proceedings of the 49th Annual ACM Symposium on Theory of Computing* (STOC 2017), pages 356–369.

B. Barak, S. Hopkins, J. Kelner, P. Kothari, A. Moitra, A. Potechin. A Nearly Tight Sum-of-Squares Lower Bound for the Planted Clique Problem. *SIAM Journal on Computing* **Special Issue** 48(2): 687-735, 2019. Preliminary version in *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2016), pages 428–437.

I. Diakonikolas, G. Kamath, D. Kane, J. Li, A. Moitra, A. Stewart. Robust Estimators in High Dimensions without the Computational Intractability. *SIAM Journal on Computing* **Special Issue** 48(2): 742-864, 2019. Preliminary version in *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2016), pages 655–664.

S. Arora, R. Ge, F. Koehler, T. Ma, A. Moitra. Provable Algorithms for Inference in Topic Models. *Proceedings of the 33rd International Conference on Machine Learning* (ICML 2016), pages 2859–2867.

A. Moitra, W. Perry, A. Wein. How Robust are Thresholds for Community Detection? *Proceedings of the 48th Annual ACM Symposium on Theory of Computing* (STOC 2016), pages 828–841.

B. Barak, A. Moitra. Noisy Tensor Completion via the Sum-of-Squares Hierarchy. *Proceedings of the 29th Annual Conference on Learning Theory* (COLT 2016), pages 1–29. **Mathematical Programming, Series B Special Issue**, to appear.

B. Barak, A. Moitra, R. O'Donnell, P. Raghavendra, O. Regev, D. Steurer, L. Trevisan, A. Vijayaraghavan, D. Witmer, J. Wright. Beating the Random Assignment on Constraint Satisfaction Problems of Bounded Degree. *19th International Workshop on Approximation, Randomization, and Combinatorial Optimization* (APPROX 2015), pages 110–123.

S. Arora, R. Ge, T. Ma, A. Moitra. Simple, Efficient and Neural Algorithms for Sparse Coding. *Proceedings of the 28th Annual Conference on Learning Theory* (COLT 2015), pages 113–149.

A. Moitra. Super-resolution, Extremal Functions and the Condition Number of Vandermonde Matrices. *Proceedings of the 47th Annual ACM Symposium on Theory of Computing* (STOC 2015), pages 821–830.

S. Arora, R. Ge, A. Moitra. New Algorithms for Learning Incoherent and Overcomplete Dictionaries. *Proceedings of the 27th Annual Conference on Learning Theory* (COLT 2014), pages 779–806.

A. Bhaskara, M. Charikar, A. Moitra, A. Vijayaraghavan. Smoothed Analysis of Tensor Decompositions. *Proceedings of the 46th Annual ACM Symposium on Theory of Computing* (STOC 2014), pages 594–603.

C. Daskalakis, A. De, I. Diakonikolas, A. Moitra, R. Servedio. A Polynomial Time Approximation Scheme for Fault-tolerant Distributed Storage. *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2014), pages 628–644.

A. Moitra, M. Saks. A Polynomial Time Algorithm for Lossy Population Recovery. *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2013), pages 110–116.

M. Hardt, A. Moitra. Algorithms and Hardness for Robust Subspace Recovery. *Proceedings of the 26th Annual Conference on Learning Theory* (COLT 2013), pages 354–375.

S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu, M. Zhu. A Practical Algorithm for Topic Modeling with Provable Guarantees. *Proceedings of the 30th International Conference on Machine Learning* (ICML 2013), pages 280–288.

M. Braverman, A. Moitra. An Information Complexity Approach to Extended Formulations. *Proceedings of the 45th Annual ACM Symposium on Theory of Computing* (STOC 2013), pages 161–170.

A. Moitra. An Almost Optimal Algorithm for Computing Nonnegative Rank. *SIAM Journal on Computing* 45(1):156–173, 2016. Preliminary version in *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2013), pages 1454–1464.

S. Arora, R. Ge, A. Moitra, S. Sachdeva. Provable ICA with Unknown Gaussian Noise, and Implications for Gaussian Mixtures and Autoencoders. *Algorithmica* **Special Issue** 72(1):215–236, 2015. Preliminary version in *Advances in Neural Information Processing Systems 25* (NIPS 2012), pages 2384–2392.

S. Arora, R. Ge, A. Moitra. Learning Topic Models – Going Beyond SVD. *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2012), pages 1–10.

A. T. Kalai, A. Moitra, G. Valiant. Disentangling Gaussians. *Communications of the ACM* February 2012, **Research Highlights**, pages 113–120.

S. Arora, R. Ge, R. Kannan, A. Moitra. Computing a Nonnegative Matrix Factorization – Provably. *SIAM Journal on Computing* **Special Issue** 45(4): 1582-1611, 2016. Preliminary version in *Proceedings of the 44th Annual ACM Symposium on Theory of Computing* (STOC 2012), pages 145–162.

N. Alon, A. Moitra, B. Sudakov. Nearly Complete Graphs Decomposable into Large Induced Matchings and Their Applications. *Journal of the European Mathematical Society* 15(5): 1575–1596, 2013. Preliminary version in *Proceedings of the 44th Annual ACM Symposium on Theory of Computing* (STOC 2012), pages 1079–1090.

R. Gelles, A. Moitra, A. Sahai. Efficient and Explicit Coding for Interactive Communication. *IEEE Transactions on Information Theory* 60(3): 1899–1913, 2014. Preliminary version in *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2011), pages 768–777.

A. Moitra, R. O'Donnell. Pareto Optimal Solutions for Smoothed Analysts. *SIAM Journal on Computing* **Special Issue** 41(5):1266–1284, 2012. Preliminary version in *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing* (STOC 2011), pages 225–234.

N. Immorlica, A. Kalai, B. Lucier, A. Moitra, A. Postlewaite, and M. Tennenholtz. Dueling Algorithms. *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing* (STOC 2011), pages 215–224.

M. Andrews, M.T. Hajiaghayi, H. Karloff, A. Moitra. Capacitated Metric Labeling. *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2011), pages 976–995.

A. Moitra, G. Valiant. Settling the Polynomial Learnability of Mixtures of Gaussians. *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2010), pages 93-102.

M. Charikar, T. Leighton, S. Li, A. Moitra. Vertex Sparsifiers and Abstract Rounding Algorithms. *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2010), pages 265-274.

A.T. Kalai, A. Moitra, G. Valiant. Efficiently Learning Mixtures of Two Gaussians. *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing* (STOC 2010), pages 553-562.

T. Leighton, A. Moitra. Extensions and Limits to Vertex Sparsification. *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing* (STOC 2010), pages 47-56.

A. Moitra. Vertex Sparsification and Oblivious Reductions. *SIAM Journal on Computing* **Special Issue** 42(6):2400–2423, 2013. Preliminary version in *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2009), pages 3-12.

T. Leighton, A. Moitra. Some Results on Greedy Embeddings in Metric Spaces. *Discrete and Computational Geometry* **Invited** 44(3): 686–705, 2010. Preliminary version in *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2008), pages 337-346.

Manuscripts

P. Kothari, A. Moitra, A. Wein. Overcomplete Tensor Decomposition via Koszul-Young Flattenings. *ArXiv:2411.14344*, 2024.

Book

A. Moitra. *Algorithmic Aspects of Machine Learning.* Cambridge University Press, New York, 2018.

Talks

Model Stealing for Any Low-Rank Language Model

UCSD Encore Workshop, March 2025

Phase Transitions in Quantum Spin Systems

Berkeley Theory Lunch, February 2025

Learning Markov Random Fields from Dynamics

Stanford Theory Lunch, February 2025
Los Almos National Lab, October 2024
Simons Institute, September 2024

How to do Theoretical Research

TCS-For-All Talk, October 2024
Learning Theory Alliance Mentoring Workshop, June 2024

Learning From Dynamics

IFML/Simons Symposium, November 2024
UC Davis Data Science Seminar, November 2024
Google Research, October 2024
AMS John von Neumann Lecture, January 2024
MIT LIDS Colloquium, December 2023
Harvard CS Theory Seminar, October 2023
Institute for Advanced Study CSDM Seminar, October 2023
Harvard Big Data Conference, August 2023

Reinforcement Learning without Intractable Oracles?

CSAIL Computing the Future, June 2023
ONR Meeting, April 2023
Princeton Machine Learning Seminar June 2022

Rethinking Robustness: From Classification to Contextual Bandits

Boston University Algorithms and Theory Seminar, April 2022
UAI Keynote, July 2021

Tensor Methods and Their Applications

IPAM Tutorial, March 2021

Tensor Completion and Refuting Random CSPs

IPAM Tutorial, March 2021

Algorithmic Foundations for the Diffraction Limit

Institute for Advanced Study CSDM Seminar, October 2023
One World Mathematics of Information, Data and Signals Seminar, February 2022
Georgia Tech ARC Colloquium, April 2021
UT Austin Machine Learning Lab Inaugural Lecture, November 2020

Learning with Massart Noise, and Connections to Fairness

USC Probability and Statistics Seminar, October 2021
Princeton ORFE Colloquium, September 2021
Northwestern Quarterly Theory Workshop, June 2020
CSAIL-MSR Trustworthy and Robust AI Workshop, June 2020

Efficient Federated Learning

Microsoft Trustworthy AI Workshop, November 2019

Robustly Learning a Signal, Under a Group Action

Stanford Theory Seminar, October 2019
ONR Program Review, October 2019

Learning Gaussian Graphical Models without Condition Number Bounds

ICERM Workshop on Data Science, May 2019
Oberwolfach Workshop on High-Dimensional Statistics, May 2019

Learning Restricted Boltzmann Machines

Simons Institute, March 2022
Institute for Mathematics and its Applications, September 2020
University of Texas, Austin Theory Seminar, March 2019
MIT MIFODS Workshop, January 2019

Sum-of-Squares, With a View Towards Average-case Complexity

Kavli Institute Tutorial, January 2019

The Paulsen Problem Made Simple

ICERM Workshop on Real Algebraic Geometry and Optimization, October 2018

Approximate Counting and the Lovasz Local Lemma

MIT Probability Seminar, October 2018
Harvard Probability and Random Matrix Theory, December 2017
Princeton University Theory Lunch, December 2017
Toyota Technological Institute CS Colloquium, October 2017
STOC, June 2017
Institute for Advanced Study CSDM Seminar, March 2017

Robustness Meets Algorithms

UPenn Theory. Seminar, November 2019
Stanford Statistics Seminar, October 2019
NUS Distinguished Lecture, January 2019
MIT Applied Math Colloquium, September 2018
SWAT Plenary Talk, June 2018
HALG Tutorial, June 2018
Princeton Center for Theoretical Science, May 2018
NYU Math and Data Seminar, April 2018
Simons Math+X Symposium, January 2018
JASON Fall Meeting, November 2017
DARPA Safe Machine Learning Workshop, October 2017
Yale University Statistics Seminar, October 2017
Packard Fellows Retreat, September 2017
ICML Tutorial, July 2017
EPFL Summer Research Institute, June 2017
Northeastern Theory Seminar, April 2017
MIT Stochastics and Statistics Seminar, March 2017
Simons Institute, November 2016
Southern California Theory Day, November 2016
Georgia Tech ARC Colloquium, October 2016

Sum-of-Squares Proofs, with Pablo Parrilo

Simons Institute Tutorial, August 2017

How Robust are Reconstruction Threshold for Community Detection?

International Bayes Club Seminar, March 2021
Institute for Mathematics and Applications, June 2019
MIT MIFODS Workshop, June 2018
Northwestern Quarterly Theory Workshop, May 2017
MIT Statistics and Data Science Day, April 2017

Planted Clique, Sum-of-Squares and Pseudo-Calibration

Institute for Mathematics and Applications, May 2016
Simons Institute, May 2016

Beyond Matrix Completion

Princeton PACM Colloquium, March 2016
FSTTCS Plenary Talk, December 2015
Harvard Big Data Conference, August 2015
European Meeting of Statisticians, July 2015
New England Machine Learning Day, May 2015
MIT Stochastics and Statistics Seminar, April 2015
NYU Theory Seminar, April 2015
Harvard Probability and Random Matrix Theory, March 2015
UChicago Theory Seminar, March 2015

Simple, Efficient and Neural Algorithms for Sparse Coding

Simons Institute, March 2015
ITA, February 2015
Workshop on Algorithmic Challenges in Machine Learning (UCSD), January 2015
Symposium on Learning, Algorithms and Complexity (IISc), January 2015
Workshop on Learning Theory (FOCM), December 2014

Super-resolution, Extremal Fctns and the Condition Number of Vandermonde Matrices

BIRS Workshop on Analytic Techniques in TCS, August 2018
STOC, June 2015
MSR/MIT Reading Group, November 2014
Workshop on Sparse Fourier Transform (FOCS), October 2014

Tensor Decompositions and Their Applications

Learning at Scale (MADALGO), August 2014

New Algorithms for Dictionary Learning

Learning at Scale (MADALGO) August 2014
Curves and Surfaces, June 2014
Mathematical Foundations of Learning Theory, June 2014
COLT, June 2014
Workshop on Overcoming Intractability in Learning (STOC), May 2014

Extended Formulations and Information Complexity

Simons Institute, April 2015
Dagstuhl Tutorial, March 2014

A Polynomial Time Algorithm for Lossy Population Recovery

Simons Institute, September 2014
Duke University Algorithms Seminar, August 2014
ICERM Workshop on Numerical Linear Algebra, May 2014
ITA, February 2014
FOCS, October 2013
TCS+ Seminar, September 2013
Princeton University Discrete Math Seminar, May 2013
Princeton University Theory Lunch, April 2013

Polynomial Methods in Learning and Statistics

UAI Tutorial, July 2013
Institute for Advanced Study Members Seminar, March 2012

An Information Complexity Approach to Extended Formulations

MIT TOC Colloquium, September 2013
STOC, June 2013
NYC Theory Day, May 2013
Columbia University Discrete Math Seminar, April 2013
ARC Theory Day, April 2013
Center for Computational Intractability Meeting, December 2012

New Algorithms for Nonnegative Matrix Factorization and Beyond

Harvard CS Colloquium, October 2014
Learning at Scale (MADALGO), August 2014
University of Washington CS Colloquium, November 2014
University of Texas, Austin CS Colloquium, April 2013
University of California, San Diego CS Colloquium, April 2013
Carnegie Mellon University CS/MLD Colloquium, March 2013
Google Research NYC Theory Seminar, March 2013

Cornell University CS Colloquium, March 2013
Princeton University CS/PACM Colloquim, February 2013
MIT Applied Math Special Seminar, February 2013
Columbia University IEOR/CS Colloquium, February 2013
IBM T.J. Watson Theory Seminar, February 2013
Microsoft Research, Silicon Valley Theory Seminar, February 2013
Caltech CMS Colloquium, January 2013
Microsoft Research, New England Theory Seminar, January 2013

An Almost Optimal Algorithm for Computing Nonnegative Rank

ISSAC Tutorial, July 2015
SODA, January 2013

Learning Topic Models – Going Beyond SVD

FOCS, October 2012

Computing a Nonnegative Matrix Factorization – Provably

Cornell University Theory Seminar, October 2012
IBM T.J. Watson Theory Lunch, July 2012
Google Research NYC Theory Seminar, June 2012
Carnegie Mellon University Theory Seminar, December 2011
University of Texas, Austin Theory Seminar, November 2011

Finding Structure in Big Data (popular talk)

Institute for Advanced Study Board of Trustees Meeting, May 2012

Nearly Complete Graphs Decomposable into Large Induced Matchings

STOC, May 2012
Princeton University Theory Lunch, April 2012

Vertex Sparsification: An Introduction, Connections and Applications

WorKer Tutorial, June 2015
Part I: Institute for Advanced Study CSDM Seminar, November 8th 2011
Part II: Institute for Advanced Study CSDM Seminar, November 15th 2011

Pareto Optimal Solutions for Smoothed Analysts

Princeton University Discrete Math Seminar, March 2012
Rutgers University Discrete Math Seminar, February 2012
University of California, Berkeley Theory Seminar, September 2011
Workshop on Beyond Worst Case Analysis (Stanford), September 2011
STOC, May 2011

Efficiently Learning Mixtures of Gaussians

International Workshop on Statistical Learning, June 2013
University of California, San Diego Theory Seminar, April 2013
Center for Computational Intractability Meeting, September 2011
Microsoft Research, New England Theory Seminar, January 2011
Institute for Advanced Study CSDM Seminar, January 2011
Microsoft Research, Redmond Theory Lunch, December 2010
Yale University Statistics Seminar , November 2010
Carnegie Mellon University Theory Seminar, September 2010
Microsoft Research, Silicon Valley Theory Seminar, August 2010
IBM T.J. Watson Theory LUnch, May 2010

Capacitated Metric Labeling

SODA, January 2011

Vertex Sparsification and Oblivious Reductions

Columbia University CS Colloquium, February 2012
Columbia University Discrete Math Seminar, November 2011
DIMACS Theory Seminar, October 2011
University of California, Los Angeles CS Colloquium, March 2011
Toyota Technological Institute CS Colloquium, March 2011
Stanford University CS Colloquium, February 2011
University of Southern California CS Colloquium, February 2011
Georgia Institute of Technology ARC Colloquium, January 2011
Microsoft Research, Silicon Valley Theory Seminar, January 2011
MIT Operations Research Seminar, December 2010
University of Washington Theory Seminar, December 2010

Extensions and Limits to Vertex Sparsification

STOC, May 2010

Approximation Algorithms for Multicommodity Type Problems

Carnegie Mellon University Theory Lunch, September 2010
China Theory Week (Tshingua), September 2010
Cornell University Theory Seminar, February 2010
Princeton University Theory Lunch, December 2009
University of California, Berkeley Theory Seminar, November 2009
Stanford Algorithms Seminar, November 2009
MIT Combinatorics Seminar, September 2009
FOCS, October 2009
ATT Research Theory Seminar, May 2009
Bell Labs Theory Seminar, May 2009

Some Results on Greedy Embeddings in Metric Spaces

Microsoft Research, Redmond, March 2009
FOCS, October 2008

Service

Associate Editor: ACM Transactions on Algorithms (TALG)

Program Chair: COLT 2025 (with Nika Haghtalab)

Program Committees: SODA 2025, FOCS 2023, STOC 2023, ITCS 2023, NeurIPS 2022 (area chair), COLT 2022 (area chair), ICLR 2022 (area chair), NeurIPS 2021 (area chair), COLT 2021 (area chair), ICLR 2021 (area chair), COLT 2020 (senior pc), NeurIPS 2019 (area chair), ICML 2019 (area chair), COLT 2019, NeurIPS 2018 (area chair), COLT 2018, ICALP 2018, STOC 2018, RANDOM 2017, ICML 2016, SODA 2015, FOCS 2014, ICML 2013, APPROX 2013

Journal Reviewer: Journal of the ACM, Proceedings of the National Academy of Sciences, Duke Math Journal, Journal of the European Mathematical Society, Foundations of Computational Mathematics, SIAM Journal on Computing, Israel Journal of Math, Annals of Statistics, Mathematical Statistics and Learning, Journal of Machine Learning Research, Mathematical Programming, Machine Learning, Journal of Combinatorial Theory Series B, Mathematics of Operations Research, Computational Complexity, Information and Computation, IEEE Transactions on Information Theory, IEEE Transactions on Computers, SIAM Journal on Matrix Analysis, SIAM Journal on Optimization, ACM Transactions on Algorithms, Algorithmica

Conference Reviewer: FOCS 2022, STOC 2022, ISIT 2022, STOC 2021, SODA 2021, ITCS 2021, FOCS 2020, STOC 2020, SODA 2020, FOCS 2019, STOC 2019, SODA 2019, FOCS 2018, ITCS 2018, SODA 2018, APPROX 2017, FOCS 2017, COLT 2017, ISIT 2017, STOC 2017, ITCS 2017, SODA 2017, FOCS 2016, IPCO 2016, STOC 2016,

SODA 2016, NIPS 2015, ESA 2015, FOCS 2015, COLT 2015, ICALP 2015, STOC 2015, ESA 2014, RANDOM 2014, COLT 2014, STOC 2014, SODA 2014, FOCS 2013, NIPS 2013, COLT 2013, STOC 2013, ICML 2013, STACS 2013, ITCS 2013, SODA 2013, ESA 2012, FOCS 2012, COLT 2012, STOC 2012, SWAT 2012, SODA 2012, ICS 2011, SODA 2011, ESA 2010, COLT 2010, STOC 2010, SODA 2010, SPAA 2009

Advising

Dhruv Rohatgi, EECS, 2021 – present
National Defense Science and Engineering Graduate Fellow

Hannah Lawrence, EECS, 2020 – present
Fannie and John Hertz Foundation Fellow
National Science Foundation Fellow

Allen Liu, EECS, 2020 – present
Fannie and John Hertz Foundation Fellow
National Science Foundation Fellow

Noah Golowich, EECS, 2019 – present
Fannie and John Hertz Foundation Fellow
National Science Foundation Fellow

Linus Hamilton, Math, 2016 – 2022
Now at Two Sigma
Thesis: *Applications and Limits of Convex Optimization*
Fannie and John Hertz Foundation Fellow

Sitan Chen, EECS, 2016 – 2021
Next an **Assistant Professor** in Computer Science at Harvard
Now a **NSF Postdoctoral Fellow** at UC Berkeley
Thesis: *Rethinking Algorithm Design for Modern Challenges in Data Science*
Paul and Daisy Soros Fellow

Frederic Koehler, Math, 2016 – 2021
Next an **Assistant Professor** in Statistics at the University of Chicago
Now a **Rajeev Motwani Fellow** at Stanford
Thesis: *Provable Algorithms for Learning and Variational Inference in Undirected Graphical Models*

Jerry Li, EECS, 2013 – 2018
Next an **Assistant Professor** in Computer Science at University of Washington
Thesis: *Principled Approaches to Robust Machine Learning and Beyond*
George M. Sprowls Thesis Award
National Science Foundation Fellow

Madalina Persu, EECS, 2013 – 2018
Thesis: *Tensors, Sparse Problems and Conditional Hardness*
Now at Two Sigma

Amelia Perry, Math, 2013 – 2018, deceased
Thesis (unsubmitted): *Message Passing Algorithms and Invariant Theory*
Charles M. Johnson Paper Prize

Alex Wein, Math, 2013 – 2018
Now an **Assistant Professor** in Mathematics at UC Davis
Thesis: *Statistical Estimation in the Presence of Group Actions*
Charles M. Johnson Paper Prize
National Defense Science and Engineering Graduate Fellow

POSTDOCS

Ainesh Bakshi, 2022 – present

Kuikui Liu, 2022 – 2023
Now an **Assistant Professor** in Computer Science at MIT

Zongchen Chen, 2021 – 2023
Now an **Assistant Professor** in Computer Science at Georgia Tech

Vatsal Sharan, 2020 – 2021
Now an **Assistant Professor** in Computer Science at USC

Diego Cifuentes, 2018 – 2021
Now an **Assistant Professor** in Industrial Engineering at Georgia Tech

Colin Sandon, 2018 – 2022
Now a **Research Scientist** at EPFL

Tselil Schramm, 2018 – 2020
Now an **Assistant Professor** in Statistics at Stanford

Andrej Risteski, 2017 – 2019
Now an **Assistant Professor** in Machine Learning at CMU

THESIS COMMITTEE

Nitya Mani (Math), Mitchell Harris (Math), Sam Park (EECS), Adam Block (Math), Alex Po-Yen Wu (CSBi), Sinho Chewi (Math), Yuval Dagan (EECS), Chenyang Yuan (EECS), Paxton Turner (Math), Matthew Brennan (EECS), Manolis Zampetakis (EECS), Brandon Tran (Math), Vishesh Jain (Math), Julia Gaudio (Sloan), Sarah Nyquist (CSBi), Jonathan Weed (Math), Haihao Lu (Math/ORC), Francisco Unda (Math), Themistoklis Gouleakis (EECS), Gautam Kamath (EECS), David Rolnick (Math), Cheng Mao (Math), Brian Cleary (CSBi), Adrian Vladu (Math), Ben Yang (Math), Mohammad Bavarian (Math), Hilary Finucane (Math), Hamza Fawzi (EECS), Oren Rippel (Math), Aaron Potechin (Math), Dimiter Ostrev (Math), Mark Lipson (Math), Alan Deckelbaum (Math)

TEACHING

Massachusetts Institute of Technology

Lecturer for 18.200: Principles of Discrete Applied Mathematics — Spring 2024
Undergraduate course on discrete math co-taught with Peter Shor

Lecturer for 6.C06/18.C06: Linear Algebra and Optimization — Fall 2023
Undergraduate course on linear algebra co-taught with Pablo Parrilo

Lecturer for 18.408: Algorithmic Aspects of Machine Learning — Spring 2023
Revised graduate course on theoretical foundations for deep learning

Lecturer for 6.C06/18.C06: Linear Algebra and Optimization — Fall 2022
Undergraduate course on linear algebra co-taught with Pablo Parrilo

Lecturer for 18.200: Principles of Discrete Applied Mathematics Spring 2022
Undergraduate course on discrete math co-taught with Michel Goemans

Lecturer for 6.S084/18.S096: Linear Algebra and Optimization Fall 2021
Undergraduate course on linear algebra co-taught with Pablo Parrilo

Lecturer for 18.408: Theoretical Foundations for Deep Learning Spring 2021
New graduate course on theoretical foundations for deep learning

Lecturer for 6.S084/18.S096: Linear Algebra and Optimization Fall 2020
Pilot undergraduate course for College of Computing co-taught with Pablo Parrilo

Lecturer for 6.042/18.062: Mathematics for Computer Science Fall 2019
Undergraduate course on discrete math co-taught with Zachary Abel and Ronitt Rubinfeld

Lecturer for 6.042/18.062: Mathematics for Computer Science Fall 2018
Undergraduate course on discrete math co-taught with Tom Leighton

Lecturer for 18.200: Principles of Discrete Applied Mathematics Spring 2018
Undergraduate course on discrete math co-taught with Michel Goemans

Lecturer for 18.408: Algorithmic Aspects of Machine Learning Fall 2017
Graduate course on recent progress in machine learning

Lecturer for 6.042/18.062: Mathematics for Computer Science Fall 2016
Undergraduate course on discrete math co-taught with Tom Leighton

Lecturer for 6.854/18.415: Advanced Algorithms Spring 2016
Redesigned core graduate course on algorithm design

Lecturer for 18.200: Principles of Discrete Applied Mathematics Fall 2015
Undergraduate course on discrete math co-taught with Peter Shor

Lecturer for 18.409: Algorithmic Aspects of Machine Learning Spring 2015
Updated graduate course on recent progress in machine learning

Lecturer for 6.042/18.062: Mathematics for Computer Science Fall 2014
Undergraduate course on discrete math co-taught with Tom Leighton

Lecturer for 18.434: Topics in Theoretical Computer Science Spring 2014
Undergraduate seminar on polytopes and optimization

Lecturer for 18.S996: Algorithmic Aspects of Machine Learning Fall 2013
New graduate course on modern algorithmic approaches in machine learning

Grants

Microsoft Research Award. *Trustworthy AI.* $250,000, 7/18 – indeterminate.

Office of Naval Research, Young Investigator Program. *An Algorithmic Theory of Robustness.* $510,000, 7/18 – 7/21

David and Lucile Packard Foundation. *Machine Learning with Provable Guarantees.* $875,000, 10/16 – 10/21.

National Science Foundation, Algorithmic Foundations. *Algebraic Proof Systems, Convexity, and Algorithms*, co-PIs Boaz Barak (Harvard), Jonathan Kelner (MIT) and Pablo Parrilo (MIT). $3,000,000 (PI portion: $750,000), 9/16 – 9/21.

Alfred P. Sloan Foundation. *Machine Learning with Provable Guarantees.* $55,000, 9/16 – 9/18.

National Science Foundation, Algorithmic Foundations. *CAREER: Algorithmic Aspects of Machine Learning.* $500,000, 9/15 – 9/20.

NEC Corporation Fund Award. *Algorithms for Discovering Sparse Representations.* $75,000, 7/14 – 1/17,

Google Research Award. *New Theoretical Challenges in Dictionary Learning.* $63,732, 9/14 – 9/16.

Summer Programs

Massachusetts Institute of Technology

Faculty Supervisor for SPUR/RSI — Summers 2015 – 2022

Six week intensive math research experience for 20+ students each year at all levels (high school, undergraduate, graduate) both as advisees and mentors, co-supervised with David Jerison (2015-2017, 2019-2020) and Davesh Maulik (2018)

Project Awards: Fourteen semifinalists and one finalist in the Siemens competition, seven Regeneron STS scholars, three finalists and one third-place overall in the International Science and Engineering Fair and one outstanding presentation award at the 2016 MAA undergraduate poster session.

Rutgers University

Lecturer for The Math Behind the Machine — Summer 2013

Summer course for high school students, aimed at introducing them to theoretical computer science

References

Sanjeev Arora (arora@cs.princeton.edu)
Charles C. Fitzmorris Professor, Princeton University

Robert Kleinberg (rdk@cs.cornell.edu)
Professor, Cornell University

Tom Leighton (ftl@akamai.com)
CEO, Akamai Technologies
Professor, MIT

Madhu Sudan (madhu@cs.harvard.edu)
Gordon McKay Professor, Harvard University

Avi Wigderson (avi@math.ias.edu)
Herbert H. Maass Professor, Institute for Advanced Study