Cameron Musco

☎ (401) 578 7385
✉ cmusco@cs.umass.edu
www.cameronmusco.com

Academic Employment

University of Massachusetts Amherst — **Amherst, MA**
Assistant Professor, College of Information and Computer Sciences — *2019 – Present*

Microsoft Research – New England — **Cambridge, MA**
Postdoctoral Researcher — *2018 – 2019*

IBM Research – Almaden — **San Jose, CA**
Research Intern — *2016*

Education

Massachusetts Institute of Technology — **Cambridge, MA**
Ph.D. Computer Science (Advisor: Nancy Lynch) — *2013 – 2018*

Yale University — **New Haven, CT**
B.S. Computer Science, B.S. Applied Mathematics — *2008 – 2012*

Research Interests

I study algorithms, working at the intersection of theoretical computer science, numerical linear algebra, and machine learning. I am especially interested in randomized methods for linear algebraic computation and data applications.

Publications

Universal Matrix Sparsifiers and Fast Deterministic Algorithms for Linear Algebra. Rajarshi Bhattacharjee, Gregory Dexter, Cameron Musco, Archan Ray, Sushant Sachdeva, David P. Woodruff. *Innovations in Theoretical Computer Science (ITCS)* 2024.

On the Unreasonable Effectiveness of Single Vector Krylov Methods for Low-Rank Approximation. Raphael Meyer, Cameron Musco, Christopher Musco. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2024.

Sublinear Time Low-Rank Approximation of Toeplitz Matrices. Cameron Musco, Kshiteej Sheth. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2024.

No-regret Algorithms for Fair Resource Allocation. Abhishek Sinha, Ativ Joshi, Rajarshi Bhattacharjee, Cameron Musco, Mohammad Hajiesmaili. *Neural Information Processing Systems (NeurIPS)* 2023.

Exact Representation of Sparse Networks with Symmetric Nonnegative Embeddings. Sudhanshu Chanpuriya, Ryan A. Rossi, Anup B. Rao, Tung Mai, Nedim Lipka, Zhao Song, Cameron Musco. *Neural Information Processing Systems (NeurIPS)* 2023.

Finite Population Regression Adjustment and Non-asymptotic Guarantees for Treatment Effect Estimation. Mehrdad Ghadiri, David Arbour, Tung Mai, Cameron Musco, Anup B. Rao. *Neural Information Processing Systems (NeurIPS)* 2023.

Latent Random Steps as Relaxations of Max-Cut, Min-Cut, and More. Sudhanshu Chanpuriya, Cameron Musco. *Differentiable Almost Everything Workshop at ICML* 2023.

Sublinear Time Eigenvalue Approximation via Random Sampling. Rajarshi Bhattacharjee, Gregory Dexter, Petros Drineas, Cameron Musco, Archan Ray. *International Colloquium on Automata, Languages,*

and Programming (ICALP) 2023. Extended version in *Algorithmica*, 2024.

Low-Memory Krylov Subspace Methods for Optimal Rational Matrix Function Approximation. Tyler Chen, Anne Greenbaum, Cameron Musco, Christopher Musco. *SIAM Journal on Matrix Analysis and Applications (SIMAX)* 2023.

Weighted Minwise Hashing Beats Linear Sketching for Inner Product Estimation. Aline Bessa, Majid Daliri, Juliana Freire, Cameron Musco, Christopher Musco, Aécio Santos, Haoxiang Zhang. *Symposium on Principles of Database Systems (PODS)* 2023.

Direct Embedding of Temporal Network Edges via Time-Decayed Line Graphs. Sudhanshu Chanpuriya, Ryan A. Rossi, Sungchul Kim, Tong Yu, Jane Hoffswell, Nedim Lipka, Shunan Guo, Cameron Musco. *International Conference on Learning Representations (ICLR)* 2023.

Optimal Sketching Bounds for Sparse Linear Regression. Tung Mai, Alexander Munteanu, Cameron Musco, Anup B. Rao, Chris Schwiegelshohn, David P. Woodruff. *International Conference on Artificial Intelligence and Statistics (AISTATS)* 2023.

Local Edge Dynamics and Opinion Polarization. Nikita Bhalla, Adam Lechowicz, Cameron Musco. *ACM International Conference on Web Search and Data Mining (WSDM)* 2023. **Invited to Special issue of ACM Transactions on Intelligent Systems and Technology.**

Toeplitz Low-Rank Approximation with Sublinear Query Complexity. Michael Kapralov, Hannah Lawrence, Mikhail Makarov, Cameron Musco, Kshiteej Sheth. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2023.

Near-Linear Sample Complexity for L_p Polynomial Regression. Raphael Meyer, Cameron Musco, Christopher Musco, David P. Woodruff, Samson Zhou. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2023.

Simplified Graph Convolution with Heterophily. Sudhanshu Chanpuriya, Cameron Musco. *Neural Information Processing Systems (NeurIPS)* 2022.

Sample Constrained Treatment Effect Estimation. Raghavendra Addanki, David Arbour, Tung Mai, Cameron Musco, Anup B. Rao. *Neural Information Processing Systems (NeurIPS)* 2022.

Kernel Interpolation with Sparse Grids. Mohit Yadav, Daniel Sheldon, Cameron Musco. *Neural Information Processing Systems (NeurIPS)* 2022.

Modeling Transitivity and Cyclicity in Directed Graphs via Binary Code Box Embeddings. Dongxu Zhang, Michael Boratko, Cameron Musco, Andrew McCallum. *Neural Information Processing Systems (NeurIPS)* 2022.

Active Linear Regression for ℓ_p Norms and Beyond. Cameron Musco, Christopher Musco, David P. Woodruff, Taisuke Yasuda. *Foundations of Computer Science (FOCS)* 2022.

Non-Adaptive Edge Counting and Sampling via Bipartite Independent Set Queries. Raghavendra Addanki, Andrew McGregor, Cameron Musco. *European Symposium on Algorithms (ESA)* 2022.

Fast Regression for Structured Inputs. Raphael Meyer, Cameron Musco, Christopher Musco, David P. Woodruff, Samson Zhou. *International Conference on Learning Representations (ICLR)* 2022.

Sublinear Time Approximation of Text Similarity Matrices. Archan Ray, Nicholas Monath, Andrew McCallum, Cameron Musco. *AAAI Conference on Artificial Intelligence (AAAI)* 2022.

Error Bounds for Lanczos-Based Matrix Function Approximation. Tyler Chen, Anne Greenbaum, Cameron Musco, Christopher Musco. *SIAM Journal on Matrix Analysis and Applications (SIMAX)* 2022.

Coresets for Classification - Simplified and Strengthened. Tung Mai, Cameron Musco, Anup B. Rao. *Neural Information Processing Systems (NeurIPS)* 2021.

On the Power of Edge Independent Graph Models. Sudhanshu Chanpuriya, Cameron Musco, Konstantinos Sotiropoulos, Charalampos E. Tsourakakis. *Neural Information Processing Systems (NeurIPS)* 2021.

DeepWalking Backwards: From Embeddings Back to Graphs. Sudhanshu Chanpuriya, Cameron Musco, Konstantinos Sotiropoulos, Charalampos E. Tsourakakis. *International Conference on Machine Learning (ICML)* 2021.

Faster Kernel Matrix Algebra via Density Estimation. Arturs Backurs, Piotr Indyk, Cameron Musco, Tal Wagner. *International Conference on Machine Learning (ICML)* 2021.

Faster Kernel Interpolation for Gaussian Processes. Mohit Yadav, Dan Sheldon, Cameron Musco. *International Conference on Artificial Intelligence and Statistics (AISTATS)* 2021. **Oral presentation.**

Subspace Embeddings Under Nonlinear Transformations. Aarshvi Gajjar, Cameron Musco. *Algorithmic Learning Theory (ALT)* 2021.

Intervention Efficient Algorithms for Approximate Learning of Causal Graphs. Raghavendra Addanki, Andrew McGregor, Cameron Musco. *Algorithmic Learning Theory (ALT)* 2021.

Simple Heuristics Yield Provable Algorithms for Masked Low-Rank Approximation. Cameron Musco, Christopher Musco, David P. Woodruff. *Innovations in Theoretical Computer Science (ITCS)* 2021.

Hutch++: Optimal Stochastic Trace Estimation. Raphael A. Meyer, Cameron Musco, Christopher Musco, David P. Woodruff. *SIAM Symposium on Simplicity in Algorithms (SOSA)* 2021.

Fourier Sparse Leverage Scores and Approximate Kernel Learning. Tamás Erdélyi, Cameron Musco, and Christopher Musco. *Neural Information Processing Systems (NeurIPS)* 2020. **Spotlight presentation**.

Node Embeddings and Exact Low-Rank Representations of Complex Networks. Sudhanshu Chanpuriya, Cameron Musco, Konstantinos Sotiropoulos, Charalampos E. Tsourakakis. *Neural Information Processing Systems (NeurIPS)* 2020.

Spiking Neural Networks Through the Lens of Streaming Algorithms. Yael Hitron, Cameron Musco, Merav Parter. *International Symposium on Distributed Computing (DISC)* 2020.

Near Optimal Linear Algebra in the Online and Sliding Window Models. Vladimir Braverman, Petros Drineas, Cameron Musco, Christopher Musco, Jalaj Upadhyay, David P. Woodruff, Samson Zhou. *IEEE Symposium on Foundations of Computer Science (FOCS)* 2020.

Efficient Intervention Design for Causal Discovery with Latents. Raghavendra Addanki, Shiva Prasad Kasiviswanathan, Andrew McGregor, Cameron Musco. *International Conference on Machine Learning (ICML)* 2020.

InfiniteWalk: Deep Network Embeddings as Laplacian Embeddings with a Nonlinearity. Sudhanshu Chanpuriya, Cameron Musco. *Knowledge Discovery and Data Mining (KDD)* 2020.

Low-Rank Toeplitz Matrix Estimation via Random Ultra-Sparse Rulers. Hannah Lawrence, Jerry Li, Cameron Musco, Christopher Musco. *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* 2020.

Sample Efficient Toeplitz Covariance Estimation. Yonina Eldar, Jerry Li, Cameron Musco, Christopher Musco. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2020.

Fast and Space Efficient Spectral Sparsification in Dynamic Streams. Michael Kapralov, Aida Mousavifar, Cameron Musco, Christopher Musco, Navid Nouri, Aaron Sidford, Jakab Tardos. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2020.

Importance Sampling via Local Sensitivity. Anant Raj, Cameron Musco, Lester Mackey. *International Conference on Artificial Intelligence and Statistics (AISTATS)* 2020.

Random Sketching, Clustering, and Short-Term Memory in Spiking Neural Networks. Yale Hiltron, Nancy Lynch, Cameron Musco, Merav Parter. *Innovations in Theoretical Computer Science (ITCS)* 2020.

Toward a Characterization of Loss Functions for Distribution Learning. Nika Haghtalab, Cameron Musco, Bo Waggoner. *Neural Information Processing Systems (NeurIPS)* 2019.

Learning to Prune: Speeding up Repeated Computations. Daniel Alabi, Adam Tauman Kalai, Katrina Ligett, Cameron Musco, Christos Tzamos, Ellen Vitercik. *Conference on Learning Theory (COLT)* 2019.

A Universal Sampling Method for Reconstructing Signals with Simple Fourier Transforms. Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, Amir Zandieh. *ACM Symposium on Theory of Computing (STOC)* 2019.

Learning Networks from Random Walk-Based Node Similarities. Jeremy Hoskins, Cameron Musco, Christopher Musco, Charalampos Tsourakakis. *Neural Information Processing Systems (NeurIPS)* 2018.

Eigenvector Computation and Community Detection in Asynchronous Gossip Models. Frederik Mallmann-Trenn, Cameron Musco, Christopher Musco. *International Colloquium on Automata, Languages, and Programming (ICALP)* 2018.

Minimizing Polarization and Disagreement in Social Networks. Cameron Musco, Christopher Musco, Charalampos Tsourakakis. *The Web Conference (WWW)* 2018.

Spectrum Approximation Beyond Fast Matrix Multiplication: Algorithms and Hardness. Cameron Musco, Praneeth Netrapalli, Aaron Sidford, Shashanka Ubaru, David P. Woodruff. *Innovations in Theoretical Computer Science (ITCS)* 2018.

Stability of the Lanczos Method for Matrix Function Approximation. Cameron Musco, Christopher Musco, Aaron Sidford. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2018.

Recursive Sampling for the Nyström Method. Cameron Musco, Christopher Musco. *Neural Information Processing Systems (NeurIPS)* 2017.

Is Input Sparsity Time Possible for Kernel Low-Rank Approximation? Cameron Musco, David P. Woodruff. *Neural Information Processing Systems (NeurIPS)* 2017.

Sublinear Time Low-Rank Approximation of Positive Semidefinite Matrices. Cameron Musco, David P. Woodruff. *IEEE Symposium on Foundations of Computer Science (FOCS)* 2017.

Neuro-RAM Unit with Applications to Similarity Testing and Compression in Spiking Neural Networks. Nancy Lynch, Cameron Musco, Merav Parter. *International Symposium on Distributed Computing (DISC)* 2017.

Random Fourier Features for Kernel Ridge Regression: Approximation Bounds and Statistical Guarantees. Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, Amir Zandieh. *International Conference on Machine Learning (ICML)* 2017. Expanded version to appear in Journal of Machine Learning Research 2020.

Spiking Neural Networks: An Algorithmic Perspective. Nancy Lynch, Cameron Musco, Merav Parter. *Workshop on Biological Distributed Algorithms (BDA)* 2017.

New Perspectives on Algorithmic Robustness Inspired by Ant Colony House-Hunting. Tsvetomira Radeva, Cameron Musco, Nancy Lynch. *Workshop on Biological Distributed Algorithms (BDA)* 2017.

Input Sparsity Time Low-Rank Approximation via Ridge Leverage Score Sampling. Michael B. Cohen, Cameron Musco, Christopher Musco. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2017.

Computational Tradeoffs in Biological Neural Networks: Self-Stabilizing Winner-Take-All. Nancy Lynch, Cameron Musco, Merav Parter. *Innovations in Theoretical Computer Science (ITCS)* 2017.

Ant-Inspired Density Estimation via Random Walks. Cameron Musco, Hsin-Hao Su, Nancy Lynch. *Proceedings of the National Academy of Sciences (PNAS)* 2017. An extended abstract initially appeared in *ACM Symposium on Principles of Distributed Computing (PODC)* 2016.

Online Row Sampling. Michael B. Cohen, Cameron Musco, Jakub Pachocki. *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)* 2016. **In Special issue of Theory of Computing, 2020.**

Principal Component Projection Without Principal Component Analysis. Roy Frostig, Cameron Musco, Christopher Musco, Aaron Sidford. *International Conference on Machine Learning (ICML)* 2016.

Faster Eigenvector Computation via Shift-and-Invert Preconditioning. Daniel Garber, Elad Hazan, Chi Jin, Sham M. Kakade, Cameron Musco, Praneeth Netrapalli, Aaron Sidford. *International Conference on Machine Learning (ICML)* 2016.

Randomized Block Krylov Methods for Stronger and Faster Approximate Singular Value Decomposition. Cameron Musco, Christopher Musco. *Neural Information Processing Systems (NeurIPS)* 2015. **Oral presentation (1 of 15 out of 403 accepted papers)**.

Distributed House-Hunting in Ant Colonies. Mohsen Ghaffari, Cameron Musco, Tsvetomira Radeva, Nancy Lynch. *ACM Symposium on Principles of Distributed Computing (PODC)* 2015.

Dimensionality Reduction for k-Means Clustering and Low Rank Approximation. Michael B. Cohen, Samuel Elder, Cameron Musco, Christopher Musco, Madalina Persu. *ACM Symposium on Theory of Computing (STOC)* 2015.

Uniform Sampling for Matrix Approximation. Michael B. Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, Aaron Sidford. *Innovations in Theoretical Computer Science (ITCS)* 2015.

Single Pass Spectral Sparsification in Dynamic Streams. Michael Kapralov, Yin Tat Lee, Cameron Musco, Christopher Musco, Aaron Sidford. *IEEE Symposium on Foundations of Computer Science (FOCS)* 2014. **In Special Issue of SIAM Journal on Computing, 2017.**

Invited Talks

Instance Optimal Iterative Methods for Matrix Function Approximation
Simons Workshop on Optimization and Algorithm Design *November 2023*

Sublinear Time Eigenvalue Approximation via Random Sampling
Simons Workshop on Sketching and Algorithm Design *October 2023*

Streaming Algorithms for Distinct Elements
Program in Algorithmic and Combinatorial Thinking (PACT) *July 2023*

Universal Sparsifiers and Fast Deterministic Algorithms for Linear Algebra
Foundations of Computational Mathematics (FOCM) *June 2023*

Theoretical Models for Opinion Polarization via Local Edge Dynamics
Integrity Workshop, Web Search and Data Mining (WSDM) *March 2023*

Sample Constrained Treatment Effect Estimation
Adobe-UMass Edge Computing Workshop *November 2022*

Randomized Iterative Methods for Approximate SVD
Guest Lecture, University of Michigan EECS 598 *November 2021, October 2022*

Representation Power and Theoretical Foundations of Modern Node Embeddings
SIAM Mathematics of Data Science *September 2022*

Data Oblivious Low-Rank Approximation for Kernel Methods
SIAM Annual Meeting *July 2022*

Sublinear Time Eigenvalue Approximation via Random Sampling
Algorithms and Foundations for Data Science Workshop, NUS *June 2022*

Randomized Methods for Sublinear Time Low-Rank Matrix Approximation
Oxford Numerical Analysis Seminar *October 2021*

Conditional Lower Bounds for Spectral Sums

SIAM Annual Meeting *July 2021*

Linear Systems in Theoretical Computer Science

Complexity of Matrix Computations Seminar *May 2021*

Hutch++: Optimal Stochastic Trace Estimation

Rutgers University MSIS Seminar *March 2023*

University of Washington Theory Seminar *April 2022*

University of Maryland CATS Seminar *December 2021*

Boston University MiDAS Seminar *December 2021*

Workshop on Algorithms for Large Data (WALDO) *August 2021*

David Harold Blackwell Summer Research Institute *July 2021*

E-NLA Seminar Series *March 2021*

IBM Research – Zurich *January 2021*

The Statistical Leverage Scores

Broderick Group Meeting, MIT EECS *February 2021*

Online Importance Sampling for Fast Linear Algebra

Adobe-Academia Workshop on Real-Time Experience Optimization *October 2020*

Spiking Neural Networks Through the Lens of Streaming Algorithms

MIT Theory of Distributed of Distributed Systems Seminar *September 2020*

Randomized Numerical Linear Algebra Meets Approximation Theory

SIAM Annual Meeting *July 2020*

Sampling Strategies for Structured Covariance Estimation

SIAM Mathematics of Data Science Conference *May 2020*

Random Sketching, Clustering, and Short-Term Memory in Spiking Neural Networks

MIT Neural Algorithms Reading Group *February 2020*

Low-Rank Approximation from Communication Complexity

UMass Amherst Theory Seminar *March 2020*

Information Theory and Applications (ITA) *January 2020*

Sample Efficient Toeplitz Covariance Estimation

WPI Computer Science Colloquium *September 2019*

UMass Amherst Theory Seminar *September 2019*

DIMACS Workshop on RandNLA, Statistics, and Optimization *September 2019*

Cornell CS Theory Seminar *April 2019*

A Universal Sampling Method for Reconstructing Signals with Simple Fourier Transforms

Cornell Scientific Computing and Numerics Seminar *April 2019*

Importance Sampling for Infinite Dimensional Optimization

INFORMS Annual Meeting *November 2018*

Recent Advances in Positive Semidefinite Matrix Approximation

Simons Workshop on Randomized Numerical Linear Algebra *September 2018*

Low-Rank Approximation of Positive Semidefinite Matrices

Guest Lecture, Harvard CS 226/MIT 6.889, Sketching for Big Data *November 2017*

Low-Rank Approximation and Clustering Via Sketching
Guest Lecture, Harvard CS 226/MIT 6.889, Sketching for Big Data *November 2017*

Sublinear Time Low-Rank Approximation of Positive Semidefinite Matrices
Boston University Theory Seminar *October 2017*
UMass Amherst Theory Seminar *October 2017*

Computational Tradeoffs in Biological Neural Networks: Self-Stabilizing WTA
MIT Theory of Distributed Systems Seminar *October 2016*

Dimensionality Reduction and Linear Sketching for Large Scale Data Analysis
BigData@CSAIL Annual Meeting, Poster Session *November 2016*
CSAIL Industry Alliance Program Annual Meeting, Poster Session *June 2015*

Fast Low-Rank Approximation and PCA: Beyond Sketching
NII Shonan Meeting on Recent Advances in RandNLA *July 2016*
Algorithms for Modern Massive Data Sets (MMDS) *June 2016*

Ant-Inspired Density Estimation via Random Walks
MIT Theory of Distributed Systems Seminar *April 2016*
MIT Theoretical Computer Science Group Theory Lunch *Febuary 2016*

Randomized Block Krylov Methods for Stronger and Faster Approximate SVD
Copper Mountain Conference on Iterative Methods *March 2016*
University of Utah Data Group Meeting *January 2016*
Neural Information Processing Systems (NeurIPS) Oral Presentation *December 2015*
MIT Theoretical Computer Science Group Theory Lunch *August 2015*

Chebyshev Polynomials and Approximation Theory in Theoretical Computer Science
MIT Danny Lewin Theory Retreat *October 2015*

Distributed House-Hunting in Ant Colonies
University of Arizona Social Insect Lab *June 2015*

Dimensionality Reduction for k-Means Clustering
MIT Algorithms and Complexity Seminar *April 2015*

Single Pass Spectral Sparsification in Dynamic Streams
CSoI NSF Site Visit, Purdue University, Poster Session *December 2015*

Uniform Sampling for Matrix Approximation
MIT Algorithms and Complexity Seminar *November 2014*

Linear Sketching and Applications to Distributed Computation
MIT Theory of Distributed Systems Seminar *November 2014*

Sparse Recovery Based Sketching for Streaming and Distributed Graph Algorithms
MIT Theoretical Computer Science Group Theory Lunch *June 2014*

Teaching

UMass Amherst CS 614: Randomized Algorithms w/ Applications to Data Science
Course Instructor *Spring '24*

UMass Amherst CS 690RA: Randomized Algorithms
Course Instructor *Spring '22*

UMass Amherst CS 514: Algorithms for Data Science
Course Instructor *Fall '19-'23; Spring '20*

UMass Amherst CS 891M: Theory Seminar
Course Instructor *Fall '20, '21, '22*

MIT 6.852: Distributed Algorithms
Teaching Assistant *Fall 2015*

Yale CS 202: Mathematical Tools for Computer Science
Teaching Assistant *Fall 2010*

Student Advising and Mentorship

Ph.D. Student Supervision

Helia Karisani, co-advised with Mohammad Hajiesmaili *2023-present*
Mohammadreza Daneshvaramoli, co-advised w/ Mohammad Hajiesmaili *2023-present*
Rajarshi Bhattacharjee *2021-present*
Thuy Trang Nguyen *2020-present*
Kyle Doney *2020-present*
Archan Ray *2019-present*
Mohit Yadav, co-advised with Dan Sheldon *2020-2024*
Now Research Scientist at Pythia Labs
Sudhanshu Chanpuriya *2019-2023*
Now Postdoc at UIUC
Raghavendra Addanki, co-advised with Andrew McGregor *2019-2022*
Now Research Scientist at Adobe

Other Student Supervision

Early Research Scholars Program – Mentor *2021-present*
Cuong Than – Ph.D. Synthesis Project, Advisor: Hung Le *2023-Present*
Jacob Gray – Undergraduate Independent Study *2023-present*
Ed Almusalamy – REU/Ind. Study, co-advised with Ramesh Sitaraman *2022-2023*
Manasi Gore – Undergraduate Honors Thesis *2022-2023*
Shib Dasgupta – Ph.D. Synthesis Project, Advisor: Andrew McCallum *2022-2023*
Cooper Sigrist – Ph.D. Synthesis Project, Advisor: Hava Siegelmann *2022-2023*
Concepta Njolima – Visiting Undergraduate Student, Berea College *Fall 2022*
Dongxu Zhang – Ph.D. Synthesis Project, Advisor: Andrew McCallum *2022*
Nathaniel Hansche – Undergraduate Independent Study *2021-2022*
Adam Lechowicz – Undergraduate Honors Thesis *2021-2022*
UMass Rising Researcher Award. Now Ph.D. student at UMass Amherst, DOE CSGF
Johno Pomerat – Undergraduate Independent Study *Fall 2021*
Aarshvi Gajjar – Masters Project *2020-2021*
Now Ph.D. student at NYU
Nikita Bhalla – Masters Independent Study *2020-2021*
Chaitanya Thakkar – Masters Independent Study *2020-2021*
Max Nelson – Masters Project *Fall 2020*
Kanchi Masalia – Masters Independent Study *Spring 2020*
Neeraj Sharma – Masters Independent Study *Spring 2020*

Mohit Yadav – Ph.D. Synthesis Project, Advisor: Dan Sheldon *2019-2020*
Harshul Shukla – Undergraduate Independent Study *Fall 2019*
Hannah Lawrence – Summer Intern at Microsoft Research *Summer 2019*
Now Ph.D. student at MIT, Hertz Fellow

Ph.D. Thesis Committees

Rico Angell, Advisor: Andrew McCallum *In Progress*
Javier Burroni, Advisor: Dan Sheldon *In Progress*
Taisuke Yasuda – CMU, Advisor: David Woodruff *In Progress*
Rik Sengupta, Advisor: Andrew McGregor *In Progress*
Gregory Dexter – Purdue University, Advisor: Petros Drineas *In Progress*
Tongyi Cao, Advisor: Akshay Krishnamurthy *In Progress*
Nishant Yadav, Advisor: Andrew McCallum *December 2023*
Dongxu Zhang, Advisor: Andrew McCallum *February 2023*
Max Nelson – UMass Amherst Linguistics, Advisors: Joe Pater, Gaja Jarosz *June 2022*
Raj Maity, Advisor: Arya Mazumdar *August 2021*
Larkin Flodin, Advisor: Arya Mazumdar *June 2021*
David Tench, Advisor: Andrew McGregor *August 2020*

Service and Outreach

Departmental Service (UMass Amherst)

Teaching Development Committee *2021-2023*
Annual Faculty Review (AFR) Committee (elected) *2022*
Associate Dean of Educational Programs and Teaching, Search Committee *2022*
Executive Committee (elected) *2021-2022*
Theory Hiring Committee *2019 (member), 2021 (member), 2022 (co-chair)*
Undergraduate Awards Committee *2021-2022*
Committee Against Racism and for Equity (CARE) – Student Learning *2020-2021*
Graduate Admissions Committee *2020-2021*
Graduate Awards Committee *2019-2020*

Organizer

UMass-NYU Quantum Algorithms for Linear Algebra Reading Group *Spring '23*
UMass Amherst Theory Seminar *Fall '20, '21, '22*
SIAM Annual Meeting: *Computational Lower Bounds in Linear Algebra* *July 2021*
SIAM Mathematics of Data Science: *Randomized Functional Analysis* *May 2020*
Neural Algorithms Reading Group (MIT) *Spring 2019*
Machine Learning Ideas Reading Group (MSR New England) *2018-2019*
Realistic Distributed Algorithms Reading Group (MIT) *Spring 2017*

Outreach

Judge: HackUMass Hackathon *2019, 2020, 2022*
Workshop Leader: Hack(H)er413 Hackathon *2020*
Member: MIT CSAIL Algorithms Office Hours *2016-2018*

Program Committees

ACM-SIAM Symposium on Discrete Algorithms (SODA) *2023*

International Colloquium on Automata, Languages and Programming (ICALP) *2022*

ACM Symposium on Theory of Computing (STOC) *2022*

International Conference on Machine Learning (ICML), Area Chair *2023*

Neural Information Processing Systems (NeurIPS), Area Chair *2021, 2022, 2023*

Conference Reviewing

FOCS, STOC, SODA, NeurIPS, ICML, ICLR, COLT, AISTATS, ALT, ITCS, RANDOM, APPROX, ICALP, SOSA, ESA, PODC, DISC, SPAA, BDA, IPDPS, SIROCCO, CSR, ICRA

AISTATS Top Reviewer *2022*

ICML Expert Reviewer *2021*

ICML Top 5% Reviewer *2019*

NeurIPS Top 400 Reviewer *2019*

Journal Reviewing

SIAM Journal on Computing, SIAM Journal on Matrix Analysis and Applications, SIAM Journal on Numerical Analysis, SIAM Journal on Scientific Computing, SIAM Journal on Mathematics of Data Science, Numerical Linear Algebra with Applications, Journal of Machine Learning Research, ACM Transactions on Computation Theory, ACM Transactions on Algorithms, Algorithmica, Theoretical Computer Science, ACM Transactions on Parallel Computing, IEEE Transactions on Information Theory, Proceedings of the National Academy of Sciences, PLOS Computational Biology, Science Advances, Distributed Computing, Advances in Computational Mathematics, Computational and Applied Mathematics, Information Processing Letters

Grant Panels

NSF *2021*

Funding

Google Research Scholar Award (PI)
Random Sketching for Scalable and Data Efficient Learning, $60,000 *2022*

NSF CAREER Award (PI)
Fast Linear Algebra: Algorithms and Fundamental Limits, $571,000 *2021-2026*

Adobe Research Collaboration Grant (PI)
Data Sketching for Real-Time Learning, $45,000 *2021-2022*

Adobe Gift Funds
Query Time Training of ML Models, $85,000 *Ongoing*

NSF RI: Medium (co-PI, PI: Andrew McCallum)
Extreme Clustering, $1,104,000 *2018-2023*

Honors and Awards

UMass Amherst CICS: Dean's Award for Anti-Racism Leadership *2021*

UMass Amherst: Finalist for University Distinguished Teaching Award *2019, 2020*

UMass Amherst CICS: Nominated for College Outstanding Teaching Award *2019, 2022*

National Science Foundation: Graduate Research Fellowship *2014-2018*

Yale University: Computer Science Senior Prize *2012*

Yale University: Summa Cum Laude, Phi Beta Kappa *2012*

Industry Experience

Redfin — **Seattle, WA**
Software Developer, Data Team — *2012-2014*

Elysium Digital — **Cambridge, MA**
Summer Technical Litigation Consultant — *Summer 2011*

Amicus — **New Haven, CT**
Software Developer — *2010-2011*