

Cameron Musco

32 Vassar Street, Office G670
Cambridge, MA 02139
☎ (401) 578 7385
✉ cnmusco@mit.edu
🌐 www.cameronmusco.com

Education

Massachusetts Institute of Technology

Ph.D. Computer Science

Advisor: Nancy Lynch

Ph.D. Thesis: The Power of Randomized Algorithms: From Numerical Linear Algebra to Biological Systems

Cambridge, MA

2013 – 2018

Yale University

B.S. Computer Science, B.S. Applied Mathematics

CS Thesis: Graph Construction Through Laplacian Function Optimization (advised by Daniel Spielman)

Applied Math Thesis: Fast Approximation of Maximum Flow Using Electrical Flows (advised by Daniel Spielman)

New Haven, CT

2008 – 2012

Research Interests

I study the algorithmic foundations of data science and machine learning, focusing on simple, randomized methods that adapt to modern computational environments. I am also interested in understanding randomized computation and algorithmic robustness by studying computation in biological systems.

Publications

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 (NIPS)* 2017.

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

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.* **Invited to special issue of Theory of Computing.**

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 (NIPS) 2015.* **Full 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.**

Submitted Manuscripts

Learning Networks from Random Walk-Based Node Similarities. Jeremy Hoskins, Cameron Musco, Christopher Musco, Charalampos Tsourakakis. Submitted for publication, 2018.

Talks and Presentations

Spectrum Approximation Beyond Fast Matrix Multiplication: Algorithms and Hardness
Innovations in Theoretical Computer Science (ITCS) January 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

Foundations of Computer Science (FOCS)	<i>October 2017</i>
Boston University Theory Seminar	<i>October 2017</i>
UMass Amherst Theory Seminar	<i>October 2017</i>
Random Fourier Features for Kernel Ridge Regression	
International Conference on Machine Learning (ICML)	<i>August 2017</i>
Spiking Neural Networks: An Algorithmic Perspective	
Workshop on Biological Distributed Algorithms (BDA)	<i>July 2017</i>
Input Sparsity Time Low-Rank Approximation via Ridge Leverage Score Sampling	
Symposium on Discrete Algorithms (SODA)	<i>January 2017</i>
Computational Tradeoffs in Biological Neural Networks: Self-Stabilizing WTA	
Innovations in Theoretical Computer Science (ITCS)	<i>January 2017</i>
MIT Theory of Distributed Systems Seminar	<i>October 2016</i>
Dimensionality Reduction and Linear Sketching for Large Scale Data Analysis	
BigData@CSAIL Annual Meeting, Poster Session	<i>November 2016</i>
CSAIL Industry Alliance Program Annual Meeting, Poster Session	<i>June 2015</i>
Fast Low-Rank Approximation and PCA: Beyond Sketching	
NII Shonan Meeting on Recent Advances in RandNLA	<i>July 2016</i>
Algorithms for Modern Massive Data Sets (MMDS)	<i>June 2016</i>
IBM Research – Almaden, Intern Research Presentation	<i>June 2016</i>
Ant-Inspired Density Estimation via Random Walks	
MIT Theory of Distributed Systems Seminar	<i>April 2016</i>
MIT Theoretical Computer Science Group Theory Lunch	<i>February 2016</i>
Randomized Block Krylov Methods for Stronger and Faster Approximate SVD	
Copper Mountain Conference on Iterative Methods	<i>March 2016</i>
University of Utah Data Group Meeting	<i>January 2016</i>
Neural Information Processing Systems (NIPS) Oral Presentation	<i>December 2015</i>
MIT Theoretical Computer Science Group Theory Lunch	<i>August 2015</i>
Chebyshev Polynomials and Approximation Theory in Theoretical Computer Science	
MIT Danny Lewin Theory Retreat	<i>October 2015</i>
Distributed House-Hunting in Ant Colonies	
University of Arizona Social Insect Lab	<i>June 2015</i>
Dimensionality Reduction for k-Means Clustering	
MIT Algorithms and Complexity Seminar	<i>April 2015</i>
Single Pass Spectral Sparsification in Dynamic Streams	
CSOI NSF Site Visit, Purdue University, Poster Session	<i>December 2015</i>
Uniform Sampling for Matrix Approximation	
MIT Algorithms and Complexity Seminar	<i>November 2014</i>
Linear Sketching and Applications to Distributed Computation	
MIT Theory of Distributed Systems Seminar	<i>November 2014</i>
Sparse Recovery Based Sketching for Streaming and Distributed Graph Algorithms	
MIT Theoretical Computer Science Group Theory Lunch	<i>June 2014</i>

Teaching and Mentorship Experience

MIT Undergraduate Research Opportunities Program

Mentor

Fall 2016

Mentored undergraduate Brandon Benson, on the research project *Swarm Agent Controllers for Formation and Herding under Locality Constraints*.

MIT 6.852: Distributed Algorithms

Teaching Assistant

Fall 2015

Yale CS 202: Mathematical Tools for Computer Science

Teaching Assistant

Fall 2010

Industry Research Experience

Microsoft Research New England

Postdoctoral Researcher

Cambridge, MA

July 2018-Present

IBM Research – Almaden

Research Intern (Mentor: David Woodruff)

San Jose, CA

Summer 2016

Worked in IBM's Theory Group, studying low-rank approximation of positive semidefinite matrices, computational lower bounds for linear algebra problems, and randomized methods for matrix norm approximation.

Other Research Experience

NII Shonan Meeting on Recent Advances in Randomized Linear Algebra

Invited participant.

July 2016

SIAM G2S3 Summer School on Randomization in Numerical Linear Algebra

Attendee.

June 2015

Service

Co-Organizer, Realistic Distributed Algorithms Reading Group

Spring 2017

Helped organize reading group primarily focused on distributed computation in biological systems.

Member of CSAIL Algorithms Office Hours: Consult researchers in various fields who seek advice in tackling algorithmic problems in their work.

External Conference Reviewer: FOCS (2018, 2017, 2015), STOC (2017), SODA (2018, 2017), NIPS (2018, 2016), COLT (2017, 2016), ITCS (2018), RANDOM (2018, 2017), APPROX (2018), ICALP (2018, 2017, 2016), ESA (2018), PODC (2017, 2016), DISC (2018, 2015), SPAA (2017), BDA (2018, 2017, 2015), IPDPS (2017), SIROCCO (2016)

External Journal Reviewer: SIAM Journal on Matrix Analysis and Applications, Journal of Machine Learning Research, PLOS Computational Biology, Distributed Computing, ACM Transactions on Parallel Computing, Computational and Applied Mathematics, Expert Systems

Honors and Awards

National Science Foundation: Graduate Research Fellowship

2014-Present

Yale University: Computer Science Senior Prize

2012

Yale University: Summa Cum Laude, Phi Beta Kappa

2012

Professional Experience

Redfin

Software Developer, Data Team

Seattle, WA

2012-2014

Helped lead integration of large-scale data analytics with customer facing website. Expanded and improved infrastructure for importing, validating, and analyzing real estate listings, property records, and region data.

Elysium Digital

Summer Technical Litigation Consultant

Cambridge, MA

Summer 2011

Analyzed source code, researched prior art, and provided technical expertise for intellectual property lawsuits.

Amicus

New Haven, CT

Software Developer

2010-2011

Developed first release of fundraising software for political campaigns and non-profits. Focused on general development as well as machine learning algorithms, including the automatic generation of efficient door-to-door canvassing schedules.