

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. Candidate, Computer Science – Theoretical Computer Science

Advisor: Nancy Lynch

S.M. Thesis: Dimensionality Reduction for k-Means Clustering (September 2015)

Cambridge, MA

2013 – present

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 algorithms, focusing on randomized methods for linear algebra, with applications in machine learning and data analysis. I am also interested in understanding randomized computation and algorithmic robustness by studying computational processes in biological systems.

Publications

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.

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 Networks. 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.**

— Talks and Presentations

Sublinear Time Low-Rank Approximation of Positive Semidefinite Matrices

Foundations of Computer Science (FOCS) *October 2017*

Boston University Theory Seminar *October 2017*

UMass Amherst Theory Seminar *October 2017*

Random Fourier Features for Kernel Ridge Regression

International Conference on Machine Learning (ICML) *August 2017*

Spiking Neural Networks: An Algorithmic Perspective

Workshop on Biological Distributed Algorithms (BDA) *July 2017*

Input Sparsity Time Low-Rank Approximation via Ridge Leverage Score Sampling

Symposium on Discrete Algorithms (SODA) *January 2017*

Computational Tradeoffs in Biological Neural Networks: Self-Stabilizing WTA

Innovations in Theoretical Computer Science (ITCS) *January 2017*

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

Invited Talk, NII Shonan Meeting on Recent Advances in RandNLA *July 2016*

Invited Talk, Algorithms for Modern Massive Data Sets (MMDS) *June 2016*

IBM Research – Almaden, intern research presentation *June 2016*

Ant-Inspired Density Estimation via Random Walks

MIT Theory of Distributed Systems Seminar *April 2016*

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	
Invited Talk, 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	
<i>Mentor</i>	<i>Fall 2016</i>
Mentored undergraduate Brandon Benson, on the research project <i>Swarm Agent Controllers for Formation and Herding under Locality Constraints</i> .	
MIT 6.852: Distributed Algorithms	
<i>Teaching Assistant</i>	<i>Fall 2015</i>
Yale CS 202: Mathematical Tools for Computer Science	
<i>Teaching Assistant</i>	<i>Fall 2010</i>

———— Industry Research Experience

IBM Research – Almaden	San Jose, CA
<i>Research Intern (Mentor: David Woodruff)</i>	<i>Summer 2016</i>
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	
<i>Invited participant.</i>	<i>July 2016</i>
SIAM G2S3 Summer School on Randomization in Numerical Linear Algebra	
<i>Attendee.</i>	<i>June 2015</i>

———— Service

Co-Organizer, Theory of Distributed Systems Seminar	<i>Fall 2017</i>
Co-Organizer, Realistic Distributed Algorithms Reading Group	<i>Spring 2017</i>
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: STOC (2017), FOCS (2015, 2017), SODA (2017, 2018), NIPS (2016), COLT (2016, 2017), ITCS (2018), RANDOM (2017), ICALP (2016, 2017), PODC (2016, 2017), DISC (2015), SPAA (2017), BDA (2015, 2017), IPDPS (2017), SIROCCO (2016)

External Journal Reviewer: SIAM Journal on Matrix Analysis and Applications, Journal of Machine Learning Research, Computational and Applied Mathematics, Expert Systems

———— Honors and Awards

National Science Foundation: Graduate Research Fellowship	<i>2014-Present</i>
Yale University: Computer Science Senior Prize	<i>2012</i>
Yale University: Summa Cum Laude, Phi Beta Kappa	<i>2012</i>

———— Professional Experience

Redfin	Seattle, WA
<i>Software Developer, Data Team</i>	<i>2012-2014</i>

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	Cambridge, MA
<i>Summer Technical Litigation Consultant</i>	<i>Summer 2011</i>

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

Amicus	New Haven, CT
<i>Software Developer</i>	<i>2010-2011</i>

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.