

Christopher Musco

370 Jay Street, Floor 11 – Brooklyn, NY 11201 – Office 1105

☎ (401) 578 2541 • ✉ cmusco@nyu.edu • 🌐 www.chrismusco.com

Academic Positions

New York University **New York, NY**
Assistant Professor, Tandon School of Engineering
Department of Computer Science and Engineering *2019 - present*

Microsoft Research **Redmond, WA**
Visiting Researcher, Machine Learning and Optimization Group *Summer 2020*

Princeton University **Princeton, NJ**
Research Instructor, Department of Computer Science *2018 – 2019*

Massachusetts Institute of Technology **Cambridge, MA**
Research and Teaching Assistant, Theory of Computation Group *2013 – 2018*

Yahoo Labs **New York, NY**
Research Intern, Scalable Machine Learning Group *Summer 2015*

Research Areas

theory of algorithms • randomized methods • numerical linear algebra • sketching and streaming • foundations of data science • network science • machine learning • active learning and function approximation

Education

Massachusetts Institute of Technology **Cambridge, MA**
Ph.D. in Computer Science *2013 – 2018*

Advisor: Jonathan Kelner

Ph.D. Thesis: Faster Linear Algebra for Data Analysis and Machine Learning

M.S. Thesis: Dimensionality Reduction for Sparse and Structured Matrices

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

Publications

*Authors appear in alphabetical order, in the tradition of mathematics and theoretical computer science. Any exceptions are marked with a *, which indicates that authors are ordered by contribution.*

Navigable Graphs for High-Dimensional Nearest Neighbor Search: Constructions and Limits. Haya Diwan, Jinrui Gou, Cameron Musco, Christopher Musco, Torsten Suel. *Conference on Neural Information Processing Systems (NeurIPS) 2024.*

Near-Optimality Guarantees for Approximating Rational Matrix Functions by the Lanczos Method. Noah Amsel, Tyler Chen, Anne Greenbaum, Cameron Musco, Christopher Musco. *Conference on Neural Information Processing Systems (NeurIPS) 2024.*

[Benchmarking Estimators for Natural Experiments: A Novel Dataset and a Doubly Robust Algorithm](#). R. Teal Witter, Christopher Musco. Preprint 2024.

[Sharper Bounds for Chebyshev Moment Matching with Applications to Differential Privacy and Beyond](#). Cameron Musco, Christopher Musco, Lucas Rosenblatt, Apoorv Vikram Singh. Preprint 2024.

[Coupling without Communication and Drafter-Invariant Speculative Decoding](#). Majid Daliri, Christopher Musco, Ananda Theertha Suresh. Preprint 2024.

[Near-optimal Hierarchical Matrix Approximation from Matrix-vector Products](#). Tyler Chen, Diana Halikias, Feyza Duman Keles, Cameron Musco, Christopher Musco, David Persson. Preprint 2024.

[Faster Linear Systems and Matrix Norm Approximation via Multi-level Sketched Preconditioning](#). Michał Dereziński, Christopher Musco, Jiaming Yang. Preprint 2024.

[Faster Spectral Density Estimation and Sparsification in the Nuclear Norm](#). Yujia Jin, Ishani Karmarkar, Christopher Musco, Aaron Sidford, Apoorv Vikram Singh. *Conference on Learning Theory (COLT)* 2024.

[Agnostic Active Learning of Single Index Models with Linear Sample Complexity*](#). Aarshvi Gajjar, Wai Ming Tai, Xingyu Xu, Chinmay Hegde, Christopher Musco, Yi Li. *Conference on Learning Theory (COLT)* 2024.

[Sampling Methods for Inner Product Sketching](#). Majid Daliri, Juliana Freire, Christopher Musco, Aécio Santos, Haoxiang Zhang. *Proceedings of the VLDB Endowment (VLDB)*. 2024.

[Fixed-Sparsity Matrix Approximation from Matrix-Vector Products](#). Noah Amsel, Tyler Chen, Diana Halikias, Feyza Duman Keles, Cameron Musco, Christopher Musco. Preprint 2024.

[Improved Active Learning via Dependent Leverage Score Sampling*](#). Atsushi Shimizu, Xiaoou Cheng, Christopher Musco, Jonathan Weare. *International Conference on Learning Representations (ICLR)* 2024. Invited for full oral presentation (1.2% of submissions).

[Improved Bounds for Agnostic Active Learning of Single Index Models](#). Aarshvi Gajjar, Xingyu Xu, Christopher Musco, Chinmay Hegde. *NeurIPS Workshop on Adaptive Experimental Design and Active Learning in the Real World* 2023.

[A Simple and Practical Method for Reducing the Disparate Impact of Differential Privacy*](#). Lucas Rosenblatt, Julia Stoyanovich, Christopher Musco. *AAAI Conference on Artificial Intelligence (AAAI)* 2024.

[Algorithm-agnostic low-rank approximation of operator monotone matrix functions*](#). David Persson, Raphael A. Meyer, Christopher Musco. *SIAM Journal on Matrix Analysis and Applications (SIMAX)* 2024.

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

[Simple Anaysis of Priority Sampling](#). Majid Daliri, Juliana Freire, Christopher Musco, Aécio Santos, Haoxiang Zhang. *SIAM Symposium on Simplicity in Algorithms (SOSA)* 2024.

[Structured Semidefinite Programming for Recovering Structured Preconditioners](#). Arun Jambulapati, Jerry Li, Christopher Musco, Aaron Sidford, Kevin Tian. *Conference on Neural Information Processing Systems (NeurIPS)* 2023.

[Efficient Block Approximate Matrix Multiplication*](#). Chuhan Yang, Christopher Musco. *European Symposium on Algorithms (ESA)* 2023.

[Moments, Random Walks, and Limits for Spectrum Approximation](#). Yujia Jin, Christopher Musco, Apoorv Singh, Aaron Sidford. *Conference on Learning Theory (COLT)* 2023.

[Dimensionality Reduction for General KDE Mode Finding](#). Xinyu Luo, Christopher Musco, Cas Widdershoven. *International Conference on Machine Learning (ICML)* 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. *ACM Symposium*

on *Principles of Database Systems (PODS)* 2023.

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.

Active Learning for Single Neuron Models with Lipschitz Non-Linearities. Aarshvi Gajjar, Chinmay Hegde, Christopher Musco. *International Conference on Artificial Intelligence and Statistics (AISTATS)* 2023. Preliminary version at NeurIPS 2022 Workshop on The Symbiosis of Deep Learning and Differential Equations II (DLDE 2022).

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

A Tight Analysis of Hutchinson's Diagonal Estimator. Prathamesh Dharangutte, Christopher Musco. *SIAM Symposium on Simplicity in Algorithms (SOSA)* 2023.

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

How to Quantify Polarization in Models of Opinion Dynamics. Christopher Musco, Indu Ramesh, Johan Ugander, R. Teal Witter. *International Workshop on Mining and Learning with Graphs (at KDD)* 2022.

Chemically-informed data-driven optimization (ChIDDO): Leveraging physical models and Bayesian learning to accelerate chemical research*. Daniel Frey, Juhee Shin, Christopher Musco, Miguel A. Modestino. *Reaction Chemistry & Engineering* 2022.

Streaming Approach to In Situ Selection of Key Time Steps for Time-Varying Volume Data*. Mengxi Wu, Yi-Jen Chiang, Christopher Musco. *EuroVis* 2022.

Sublinear Time Spectral Density Estimation. Vladimir Braverman, Aditya Krishnan, Christopher Musco. *ACM Symposium on Theory of Computing (STOC)* 2022.

A Sketch-based Index for Correlated Dataset Search*. Aécio Santos, Aline Bessa, Christopher Musco, Juliana Freire. *IEEE 38th International Conference on Data Engineering (ICDE)* 2022.

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

Active Linear Regression in the Online Setting via Leverage Score Sampling. Harvineet Singh, Christopher Musco, Rumi Chunara. *ICML Workshop on Adaptive Experimental Design and Active Learning in the Real World* 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.

Dynamic Trace Estimation. Prathamesh Dharangutte, Christopher Musco. *Conference on Neural Information Processing Systems (NeurIPS)* 2021.

Finding an Approximate Mode of a Kernel Density Estimate. Jasper C.H. Lee, Jerry Li, Christopher Musco, Jeff M. Phillips, Wai Ming Tai. *European Symposium on Algorithms (ESA)* 2021.

Correlation Sketches for Approximate Join-Correlation Queries*. Aécio Santos, Aline Bessa, Fernando Chirigati, Christopher Musco, Juliana Freire. *ACM Special Interest Group on Management of Data Conference (SIGMOD)* 2021.

Public Transport Planning: When Transit Network Connectivity Meets Commuting Demand*. Sheng Wang, Yuan Sun, Christopher Musco, Zhifeng Bao. *ACM Special Interest Group on Management of Data Conference (SIGMOD)* 2021.

Simple Heuristics Yield Provable Algorithms for Masked Low Rank Approximation. Cameron Musco, Christopher Musco, David 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.

Graph Learning for Inverse Landscape Genetics. Prathamesh Dharangutte, Christopher Musco. *AAAI Conference on Artificial Intelligence (AAAI) 2021*. Short version in NeurIPS 2020 AI for Earth Sciences workshop.

Fourier Sparse Leverage Scores and Approximate Kernel Learning. Tamás Erdélyi, Cameron Musco, Christopher Musco. *Conference on Neural Information Processing Systems (NeurIPS)* 2020. Invited for spotlight presentation (4% of submissions).

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.

The Statistical Cost of Robust Kernel Hyperparameter Tuning. Raphael A. Meyer, Christopher Musco. *Conference on Neural Information Processing Systems (NeurIPS)* 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. *ACM-SIAM Symposium on Discrete Algorithms (SODA)* 2020.

Analyzing the Impact of Filter Bubbles on Social Network Polarization. Uthsav Chitra, Christopher Musco. *ACM International Conference on Web Search and Data Mining (WSDM)* 2020. Preliminary version in KDD Workshop on Issues of Sentiment Discovery and Opinion Mining.

Projection-Cost-Preserving Sketches: Proof Strategies and Constructions . Cameron Musco, Christopher Musco. *Technical Report* 2020.

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.

Inferring networks from random walk-based node similarities. Jeremy Hoskins, Cameron Musco, Christopher Musco, Charalampos Tsourakakis. *Conference on 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.

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. *Conference on Neural Information Processing Systems (NeurIPS)* 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.

Determining Tournament Payout Structures for Daily Fantasy Sports. Christopher Musco, Maxim Sviridenko, Justin Thaler. *SIAM Algorithm Engineering and Experiments (ALENEX)* 2017.

Principal Component Projection Without Principal Component Analysis. Roy Frostig, Cameron Musco, Christopher Musco, 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. *Conference on Neural Information Processing Systems (NeurIPS)* 2015. Invited for full oral presentation (1 of 15, < 1% of submissions). Also presented at the 2016 Copper Mountain Conference on Iterative Methods.

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.

Principled Sampling for Anomaly Detection. Brendan Juba, Fan Long, Christopher Musco, Stelios Sidiropoulos-Douskos, Martin Rinard. *Network and Distributed System Security Symposium (NDSS)*. 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, Special Issue SIAM Journal on Computing 2017.

Funding

NSF AF Small

The Query Complexity of Linear Algebra. 2024 - 2027
PI (w/ Cameron Musco, UMass), \$600,000

NSF CAREER Award

Frontiers in Matrix Sketching. 2021 - 2026
PI, \$562,224

NSF III Medium

Dataset Search and Ranking for Data Augmentation and Explanation. 2021 - 2024
Co-PI (w/ Juliana Freire, NYU), \$1,093,195

DOE Office of Science Grant

Leverage Score Sampling for Parametric PDEs. 2021 - 2023
PI (w/ Chinmay Hegde, NYU), \$300,000

NSF/CRA

CIFellow Fellow Award for Dr. Aline Bessa. 2021-2022
PI, \$263,250

Current Graduate Student and Postdoc Advising

Tyler Chen. Postdoc, New York University. 2022 - present

Noah Amsel. Ph.D., New York University (with Prof. Joan Bruna). 2022 - present

Majid Daliri. Ph.D., New York University. 2022 - present

Feyza Duman Keles. Ph.D., New York University (with Prof. Chinmay Hedge). 2023 - present

Aarshvi Gajjar. Ph.D., New York University (with Prof. Chinmay Hedge). 2021 - present

Apoorv Singh. Ph.D., New York University. 2020 - present

R. Teal Witter. Ph.D., New York University (with Prof. Lisa Hellerstein). 2020 - present

Past Advising

Raphael A. Meyer. Ph.D., New York University.	2019 - 2024
Chuhan Yang. Ph.D., New York University.	2020 - 2024
Aline Bessa. Postdoc, New York University.	2021
David Persson. Visiting Ph.D. student.	2023
Akash Rao, M.S., New York University.	2023 - 2024
Siddharth Sagar. M.S., New York University.	2023
Yajur Ahuja. M.S., New York University.	2023
Sundeep Chenreddy. M.S., New York University.	2023
Atsushu Shimuzu. M.S., New York University.	2022 - 2023
Danrong Li. M.S., New York University.	2022 - 2023
Indu Ramesh. M.S., New York University.	2020 - 2022
Xinyu Luo. M.S., New York University.	2020 - 2022
Prathamesh Dharangutte. M.S., New York University.	2019 - 2021
John Shin. M.S., New York University.	2019 - 2020

I have advised over a dozen undergraduate guided studies and research projects, as well as several research projects involving high school students.

Ph.D. Thesis Committees

Haoxiang Zhang, NYU	
Bruno Coelho, NYU	
Aécio Santos, NYU	Defended 2024
Ishan Agarwal, NYU	Defended 2023
Yiming Zhang, NYU	Defended 2022
Francis Williams, NYU	Defended 2021

Teaching

Machine Learning, Graduate (NYU CS-GY 6923)	
<i>Course Instructor</i>	Spring 2022, 2023, Fall 2024
Algorithmic Machine Learning and Data Science (NYU CS-GY 6763)	
<i>Course Instructor</i>	Fall 2019 - 2023
Introduction to Machine Learning (NYU CS-UY 4563)	
<i>Course Instructor</i>	Spring 2020
Advanced Algorithm Design (Princeton COS 521)	
<i>Course Instructor</i>	Fall 2018
Advanced Algorithms (MIT 6.854/18.415)	
<i>Teaching Assistant</i>	Spring 2016
Technical Communication Skills for Graduate Students (MIT 6.S977)	
<i>Workshop Leader</i>	Spring 2016

Mentorship and Outreach

SIAM/MathWorks Math Modeling Challenge

Problem Development Committee, Director of Technical Computing Judging 2018 - present
Support SIAM's annual math modeling competition, which gives high school students an opportunity to apply mathematics, statistics, and data science to a real-world problem. Serve on committee in charge of writing practice and contest problems. Lead judging and award presentation for student Technical Computing prizes.

MIT Graduate Communication Lab

Communication Adviser 2015 - 2018
Helped to establish an EECS Communication Lab to provide over 700 graduate students with support in academic writing, oral presentation, job applications, and generally, in communicating effectively. Offered feedback, guidance, and coaching to individual students during weekly office hours.

Invited Talks

Sampling-based Methods for Inner Product Sketching

Wuhan University Aug. 2024

The Lanczos Method, Matrix Functions, and the Quest for Optimality

Johns Hopkins, Applied Math Seminar Feb. 2024

Tutorial on Randomized Numerical Linear Algebra (invited panelist)

Neural Information Processing Systems (NeurIPS) Dec. 2023

On the Effectiveness of Single Vector Krylov Methods for Low-rank Approximation

Foundations of Computational Mathematics (FoCM) June 2023

SIAM Conference on Applied Linear Algebra May 2024

Tutorial on Matrix Sketching

BIRS Workshop on Perspectives on Matrix Computations Mar. 2023

Integrating Education and Research in Early Career Funding Proposals

Joint Mathematics Meetings, SIAM Special Session on Education Jan. 2023

Robust Active Learning via Leverage Score Sampling

Canadian Mathematics Society, Winter Meeting Dec. 2022

One World Mathematics of Information, Data, and Signals (1W-MINDS) Seminar Oct. 2023

Iterative Sampling Methods for Matrix Approximation

University of Michigan, Randomized Numerical Linear Algebra Seminar Oct. 2022

Noisy Implicit Methods for Matrix Functions

SIAM Annual Meeting July 2022

Sublinear Time Spectral Density Estimation

Queens College CUNY Computer Science Colloquium Mar. 2023

National University of Singapore, Workshop on Algorithms and Foundations for Data Science June 2022

Oxford University, Computational Mathematics and Applications Seminar Feb. 2022

University of British Columbia, Applied Math Seminar Jan. 2022

Workshop on Algorithms for Large Data Aug. 2021

Cornell University, Scientific Computing and Numerics Seminar Mar. 2021

Optimal Stochastic Trace Estimation

New York University, Numerical Analysis and Scientific Computing Seminar Apr. 2022

University of Michigan, Randomized Numerical Linear Algebra Seminar Nov. 2021

Stanford University, Operations Research Seminar	<i>Oct. 2020</i>
A TCS Perspective on Iterative Eigensolvers	
Complexity of Matrix Computations Seminar	<i>Oct. 2021</i>
Mathematical Modeling and Coding in the Classroom	
Mu Alpha Theta National Convention	<i>July 2021</i>
Randomized Algorithms for Probing the Spectrum of a Matrix	
George Mason University, AI Tea Series	<i>Apr. 2021</i>
Importance Sampling for Approximating High-Dimensional Kernel Matrices	
Emory University, Numerical Analysis and Scientific Computing Seminar	<i>Feb. 2021</i>
Structured Covariance Estimation	
Oxford Statistics and Machine Learning in Finance Seminar	<i>June 2022</i>
Yahoo! Research, Seminar	<i>Apr. 2020</i>
Johns Hopkins University, Theory Seminar	<i>Oct. 2019</i>
From Randomized Linear Algebra to Randomized Functional Analysis	
SIAM Mathematics of Data Science Conference	<i>May 2020</i>
University of Utah, Data Science Seminar	<i>Jan. 2020</i>
Rutgers University, DIMACS RandNLA Workshop	<i>Sept. 2019</i>
Learning Signals with Simple Fourier Transforms	
Columbia University, Theory Seminar	<i>Oct. 2018</i>
UC Berkeley, Simons Institute Workshop	<i>Sept. 2018</i>
Princeton University, Theory Seminar	<i>Sept. 2018</i>
The Lanczos Method in Data Science: Challenges and Opportunities	
SIAM Conference on Parallel Processing for Scientific Computing, Minisymposium	<i>Feb. 2020</i>
Georgetown University, CS Seminar	<i>Feb. 2018</i>
Unifying Random Features and Leverage Scores for Kernel Matrix Approximation	
Institute for Advanced Study, Computer Science and Discrete Mathematics Seminar	<i>Feb. 2018</i>
Harvard University, guest lecture in "Sketching Algorithms for Big Data"	<i>Nov. 2017</i>
Microsoft Research, Redmond, Research Seminar	<i>Sept. 2017</i>
Sketching and Sampling Methods for Large Scale Linear Algebra	
Twitter, Cambridge, Research Seminar	<i>July 2017</i>
Recursive Sampling for the Nyström Method	
New England Machine Learning Day	<i>May 2017</i>
Introduction to Linear Sketching	
MIT, Theory Retreat	<i>Sept. 2016</i>
Ridge Leverage Score Sampling	
National Institute of Informatics, Shonan Meeting	<i>July 2016</i>
MIT, Algorithms & Complexity Seminar	<i>July 2016</i>
University of Utah, Data Group Meeting	<i>Jan. 2016</i>
Large Scale Column Subset Selection	
Yahoo Labs, Science Week	<i>July 2015</i>
Dimensionality Reduction for k-Means Clustering	

IBM T.J. Watson, Mathematical Sciences Research Seminar	<i>Aug. 2015</i>
MIT, Theory Lunch	<i>June 2014</i>
Uniform Sampling for Matrix Approximation	
Yahoo Labs, Scalable Machine Learning Research Seminar	<i>June 2015</i>
Single Pass Spectral Sparsification in Dynamic Streams	
MIT, Annual Sublinear Algorithms Day	<i>Apr. 2015</i>
Harvard University, Theory Seminar	<i>Nov. 2014</i>

Service

Organization

NYU Theoretical Computer Science Seminar	<i>2022 - present</i>
NYU Computer Science Theory "Pandemic Presentations" Workshop	<i>2022</i>
SIAM Annual Meeting Minisymposium Organizer	<i>2021</i>
SIAM Mathematics of Data Science Minisymposium Organizer	<i>2020, 2024</i>

Department Service

Faculty Search Committee	<i>2019, 2021-2024</i>
Chair, Ph.D. Admissions Committee	<i>2021 - 2023</i>
Department Chair Search Committee	<i>2021 - 2023</i>
Online Faculty Oversight Committee	<i>2021 - 2023</i>

Conference Programming

SOSA Program Committee Member	<i>2025</i>
ESA Track S Program Committee Member	<i>2024</i>
STOC Program Committee Member	<i>2023</i>
FOCS Program Committee Member	<i>2022</i>
ICML Area Chair	<i>2023</i>
NeurIPS Area Chair	<i>2019, 2021-2024</i>
AAAI Senior Area Chair	<i>2022</i>

Journal Editing

Editor, SICOMP Special Issue for STOC	<i>2023</i>
---------------------------------------	-------------

Conference Reviewing

STOC, FOCS, SODA, ICALP, SoCG, ITCS, ICML, NeurIPS, COLT, AAI, AISTATS, SOCG, ESA, APPROX, RANDOM, STACS, IWOCA, WALCOM, NDSS

Journal Reviewing

Journal of the ACM, SIAM Journal on Computing, Communications on Pure and Applied Mathematics, SIAM Journal on Matrix Analysis and Applications, SIAM Journal on Scientific Computing, ACM Transactions on Algorithms, Theoretical Computer Science, Journal of Machine Learning Research, Linear Algebra and Its Applications, SIAM Journal on Mathematics of Data Science, Science Advances, Mathematical Programming, IEEE Transactions on Signal Processing, IEEE Transactions on Knowledge and Data Engineering, International Journal of Foundations of Computer Science, IEEE Journal on Selected Areas in Information Theory, BIT Numerical Mathematics, AMS Notices

Grants Panels

National Science Foundation Panel	<i>2019, 2022</i>
-----------------------------------	-------------------

NSF Graduate Research Fellowship Reviewer 2022
Department of Energy External Reviewer 2022, 2023

Notables.....

ICML 2018 Outstanding Reviewer, ICML Expert Review 2021

Invited Workshops and Events

Banff International Research Station (BIRS)

Invited participant 2023
Workshop on Perspectives on Matrix Computations: Theoretical Computer Science Meets Numerical Analysis.

Rutgers, DIMACS

Invited participant 2019
Workshop on Randomized Numerical Linear Algebra, Statistics, and Optimization.

UC Berkeley, Simons Institute for the Theory of Computing

Invited participant 2018
Workshop on Randomized Numerical Linear Algebra and Applications.

Japan National Institute of Informatics

Invited participant 2016
Shonan meeting on Recent Advances in Randomized Numerical Linear Algebra.

Honors and Awards

NYU Tandon

Junior Faculty Research Award 2023

NYU Tandon

Goddard Junior Faculty Fellowship Award 2023

NYU Tandon

Nominated for Jacobs Excellence in Education Innovation Award 2022

National Science Foundation

Graduate Research Fellowship, recipient 2014-2017

Yale University

Cum Laude, Distinction in Both Majors 2012

Other Professional Experience

Redfin

Software Engineer, Data Team **Seattle, WA**
2012 – 2014

Elysium Digital

Summer Consultant **Cambridge, MA**
2011

Amicus

Software Developer **New Haven, CT**
2010 – 2011