Noah Amsel

NoahAmsel.github.io | noah.amsel@nyu.edu

Education

Courant Institute, New York University August 2022 – Present
Ph.D. in Computer Science New York, NY

- Advised by Joan Bruna and Chris Musco. Supported by NSF GRFP.

Yale University August 2016 – May 2020
B.S. in Computer Science & Mathematics, with distinction. New Haven, CT

- *Magna cum laude*. Thesis: "Online Vector Balancing in Practice," advised by Dan Spielman.

Research Interests

Algorithms and deep learning, specifically related to matrix functions, structured matrices, approximation theory and expressibility.

Experience

Adobe, Inc. Summer 2025
Research Scientist Intern San Jose, CA

- Led research project on evaluation of auxiliary long-term memory systems for AI agents.

Polymathic AI Summer 2024
Summer Research Intern New York, NY

- Led experiments studying OOD transfer capabilities of neural networks for fluid dynamics PDEs.

Qualcomm Technologies, Inc. November 2021 – July 2022
Engineer, Corporate Research & Development New York, NY

- Stayed on when Reservoir Labs was acquired by Qualcomm and turned into a new R&D division.

Reservoir Labs July 2020 – November 2021
Research Engineer New York, NY

- Developed mathematical model of congestion control for designing efficient networks and managing traffic. Deployed it to DOE's Energy Sciences Network and published at ACM SIGCOMM.

Weizmann Institute of Science Summer 2019
Kupcinet-Getz International Summer School, Research Fellow Rehovot, Israel

- Developed new method for phylogenetic inference via spectral analysis. Proved finite sample guarantees.

Facebook, Inc. Summer 2018
Software Engineering Intern New York, NY

- Built Bandwidth Estimation model for Adaptive Bitrate Streaming to improve mobile video quality.

Off the Hook, LLC Summer 2017
Data Analyst / Software Developer New York, NY

- Developed gambling software based on machine learning analysis of horse racing data.

Publications

[14] **Noah Amsel***, Pratyush Avi, Tyler Chen, Feyza Duman Keles, Chinmay Hegde, Cameron Musco, Christopher Musco, and David Persson. *Query Efficient Structured Matrix Learning*. 2025.

[13] **Noah Amsel***, Tyler Chen, Feyza Duman Keles, Diana Halikias, Cameron Musco, Christopher Musco, and David Persson. *Quasi-optimal hierarchically semi-separable matrix approximation*. 2025.

[12] **Noah Amsel**, David Persson, Christopher Musco, and Robert M. Gower. *The Polar Express: Optimal Matrix Sign Methods and Their Application to the Muon Algorithm*. 2025.

[11] **Noah Amsel**, Gilad Yehudai, and Joan Bruna. "Quality over Quantity in Attention Layers: When Adding More Heads Hurts". In: *International Conference on Representation Learning (ICLR)*. 2025.

[10] Yilun Kuang, **Noah Amsel**, Sanae Lotfi, Shikai Qiu, Andres Potapczynski, and Andrew Gordon Wilson. "Customizing the Inductive Biases of Softmax Attention using Structured Matrices". In: *Proceedings of the 42nd International Conference on Machine Learning (ICML)*. 2025.

[9] Gilad Yehudai, **Noah Amsel**, and Joan Bruna. *Compositional Reasoning with Transformers, RNNs, and Chain of Thought*. 2025.

[8] **Noah Amsel**, Tyler Chen, Anne Greenbaum, Cameron Musco, and Christopher Musco. "Nearly Optimal Approximation of Matrix Functions by the Lanczos Method". In: *The Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*. 2024.

[7] **Noah Amsel***, Tyler Chen, Feyza Duman Keles, Diana Halikias, Cameron Musco, and Christopher Musco. *Fixed-sparsity matrix approximation from matrix-vector products*. 2024.

[6] Yariv Aizenbud, Ariel Jaffe, Meng Wang, Amber Hu, **Noah Amsel**, Boaz Nadler, Joseph T Chang, and Yuval Kluger. "Spectral top-down recovery of latent tree models". In: *Information and Inference: A Journal of the IMA* (2023).

[5] Ariel Jaffe, **Noah Amsel**, Yariv Aizenbud, Boaz Nadler, Joseph T. Chang, and Yuval Kluger. "Spectral Neighbor Joining for Reconstruction of Latent Tree Models". In: *SIAM Journal on Mathematics of Data Science (SIMODS)* (2021).

[4] Jordi Ros-Giralt, **Noah Amsel**, Sruthi Yellamraju, James Ezick, Richard Lethin, Yuang Jiang, Aosong Feng, Leandros Tassiulas, Zhenguo Wu, Min Yee Teh, and Keren Bergman. "Designing Data Center Networks Using Bottleneck Structures". In: *Proceedings of the ACM SIGCOMM 2021 Conference*. 2021.

[3] **Noah Amsel**, Jordi Ros-Giralt, Sruthi Yellamraju, James Ezick, Brendan von Hofe, Alison Ryan, and Richard Lethin. "Computing Bottleneck Structures at Scale for High-Precision Network Performance Analysis". In: *2020 IEEE/ACM Innovating the Network for Data-Intensive Science (INDIS)*. 2020.

[2] William Merrill, Lenny Khazan, **Noah Amsel**, Yiding Hao, Simon Mendelsohn, and Robert Frank. "Finding Hierarchical Structure in Neural Stacks Using Unsupervised Parsing". In: *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*. 2019.

[1] Yiding Hao, William Merrill, Dana Angluin, Robert Frank, **Noah Amsel**, Andrew Benz, and Simon Mendelsohn. "Context-Free Transductions with Neural Stacks". In: *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*. 2018.

*Alphabetical author ordering.

Talks

The Low Rank Bottleneck in Attention

- NYC Student Theory Day, *NYU* 2025
- NYAS ML Symposium, *New York Academy of Sciences* 2024
- CCM Machine Learning Seminar, *Flatiron Institute* 2024

Fixed-Sparsity Matrix Approximation from Matrix-Vector Products

- SIAM Linear Algebra, *Sorbonne University* 2024
- Mid Atlantic Numerical Analysis Day, *Temple University* 2024

Near-Optimal Approximation of Matrix Functions by the Lanczos Method

- Rutgers Theory Seminar, *Rutgers University* 2025
- SIAM-NNP, *NJIT* 2023

Designing Data Center Networks Using Bottleneck Structures

- ACM SIGCOMM, *Virtual* 2021

Computing Bottleneck Structures at Scale

- IEEE/ACM INDIS at SC20, *Virtual* 2020

Posters

Quality over Quantity in Attention

- ICML, *Singapore* 2025

Lanczos-FA is Nearly Krylov-Optimal for Rational Functions

- NeurIPS (Spotlight!), *Vancouver* 2024
- Mihalis Fest, *Columbia University* 2023
- ASE60, *MIT* 2023
- Modern Applied and Computational Analysis, *ICERM* 2023

Finding Hierarchical Structure in Neural Stacks Using Unsupervised Parsing

- BlackboxNLP (@ACL), *Florence* 2019

Teaching

Teaching Assistant (Algorithms for Machine Learning), *NYU* 2025
Undergraduate Learning Assistant (Algorithms), *Yale University* 2019
Volunteer Math Tutor, *Top Honors* ..... 2013 – 2016

Awards

NSF Graduate Research Fellowship, *National Science Foundation* 2024
Phi Beta Kappa, *Yale University* 2020
Semifinalist, *Intel Science Talent Search* 2016

Service

Organizing	NYC Machine Learning Speaker Series at the Flatiron Institute (2023 – 2025)
Reviewing	JMLR (2024), IEEE Signal Processing Magazine (2025)
Outreach	NYU AI Summer School Instructor (2024)