

# NOAH AMSEL

noah.amsel@nyu.edu  $\diamond$  <https://NoahAmsel.github.io/>

*Last updated February 23, 2025*

## EDUCATION

---

### **Courant Institute, New York University**

Ph.D. in Computer Science

Advised by Joan Bruna and Chris Musco.

August 2022 - Present

New York, NY

### **Yale University**

B.S. in Computer Science & Mathematics, with distinction.

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

August 2016 - May 2020

New Haven, CT

## EXPERIENCE

---

### **Polymathic AI**

*Summer Research Intern*

June 2024 - August 2024

New York, NY

- Contributed to deep learning research on large foundation models for scientific data from heterogeneous domains.
- Led project studying transfer capabilities of models for fluid dynamics problems described by differential equations.

### **Qualcomm Technologies, Inc.**

*Engineer, Corporate Research & Development*

November 2021 - July 2022

New York, NY

- Stayed on after Reservoir Labs was acquired by Qualcomm and turned into a new R&D division.
- Worked on software package for global non-convex optimization.

### **Reservoir Labs**

*Research Engineer*

June 2020 - November 2021

New York, NY

- Created algorithm for simulating network performance that achieved order-of-magnitude speed up.
- Deployed our network modeling tool to DOE's Energy Sciences Network.
- Developed framework for designing provably efficient data center networks. Published it at ACM SIGCOMM.

### **Weizmann Institute of Science**

*Kupcinet-Getz International Summer School, Research Fellow*

June - August 2019

Rehovot, Israel

- Developed novel spectral method for fitting latent tree graphical models of DNA data.
- Proved sufficient condition for the algorithm to succeed and finite sample guarantee.
- Wrote open source implementation that scales to to very large problems with high accuracy.

### **Facebook, Inc.**

*Software Engineering Intern*

May - August 2018

New York, NY

- Built Bandwidth Estimation model for Adaptive Bitrate Streaming to improve mobile video quality.
- Implemented and tested the model in C++ and Java; refactored ABR code.
- Validated new model in production; found it reduced error by more than half (RMSE).
- Received return offer.

### **Off the Hook, LLC**

*Data Analyst / Software Developer*

May - July 2017

New York, NY

- Developed computer gambling software in Python based on statistical analysis of horse racing data.
- Evaluated wager opportunities using machine learning trained on data from hundreds of thousands of past races.

## PUBLICATIONS

---

- Noah Amsel, Gilad Yehudai, and Joan Bruna. The low-rank bottleneck in attention. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=y9Xp9NozPR>.
- Noah Amsel, Tyler Chen, Feyza Duman Keles, Diana Halikias, Cameron Musco, and Christopher Musco. Fixed-sparsity matrix approximation from matrix-vector products, 2024a. URL <https://arxiv.org/abs/2402.09379>.
- Noah Amsel, Tyler Chen, Anne Greenbaum, Cameron N 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*, 2024b. URL <https://openreview.net/forum?id=3s8V8QP9XV>.
- 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. *Information and Inference: A Journal of the IMA*, 12(3):iaad032, 08 2023. ISSN 2049-8772. doi: 10.1093/imaiai/iaad032. URL <https://doi.org/10.1093/imaiai/iaad032>.
- Jordi Ros-Giralt, Noah Amsel, Sruthi Yellamraju, James Ezick, Richard Lethin, Yuang Jiang, Aosong Feng, Leandros Tassioulas, Zhenguo Wu, Min Yee Teh, and Keren Bergman. Designing data center networks using bottleneck structures. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference, SIGCOMM '21*, pages 319–348, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383837. doi: 10.1145/3452296.3472898. URL <https://doi.org/10.1145/3452296.3472898>.
- Ariel Jaffe, Noah Amsel, Yariv Aizenbud, Boaz Nadler, Joseph T. Chang, and Yuval Kluger. Spectral neighbor joining for reconstruction of latent tree models. *SIAM Journal on Mathematics of Data Science (SIMODS)*, 3(1):113–141, 2021. doi: 10.1137/20M1365715. URL <https://doi.org/10.1137/20M1365715>.
- 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. *2020 IEEE/ACM Innovating the Network for Data-Intensive Science (INDIS)*, 11 2020. doi: 10.1109/indis51933.2020.00012. URL <https://ieeexplore.ieee.org/document/9307175>.
- 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*, pages 224–232, Florence, Italy, August 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-4823. URL <https://aclanthology.org/W19-4823>.
- 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*, pages 306–315, Brussels, Belgium, November 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-5433. URL <https://aclanthology.org/W18-5433>.

## CONFERENCE TALKS

---

NYAS ML Symposium	The Low-Rank Bottleneck in Attention	2024
SIAM Linear Algebra Temple Num. Anal. Day	Fixed-Sparsity Matrix Approximation from Matrix-Vector Products	2024
SIAM-NNP		
ACM SIGCOMM	Designing Data Center Networks Using Bottleneck Structures	2021
IEEE/ACM INDIS at SC20	Computing Bottleneck Structures at Scale	2020

## POSTER PRESENTATIONS

---

NeurIPS – <i>Spotlight!</i>	Lanczos-FA is Nearly Krylov-Optimal for Rational Functions	2024
ICERM		2023
ASE60 (MIT)		2023
Mihalis Fest (Columbia)		2023
BlackboxNLP @ACL	Finding Hierarchical Structure in Neural Stacks Using Unsupervised Parsing	2019

## TEACHING EXPERIENCE

---

Teaching Assistant	Algorithms for Machine Learning	<i>NYU</i>	2025
Undergraduate Learning Assistant	Algorithms	<i>Yale University</i>	2019
Volunteer Math Tutor		<i>Top Honors</i>	2013–2016

## AWARDS

---

NSF Graduate Student Research Fellowship	2024
Phi Beta Kappa	2020
Intel Science Talent Search Semifinalist	2016