Alex Damian

✉ adamian@mit.edu   🌐 alex-damian.github.io   ⑧ Google Scholar

RESEARCH INTERESTS

My research is focused on developing a mathematical foundation for deep learning

▶ **Optimization:** non-convex optimization, implicit regularization, adaptive optimizers

▶ **Representation Learning:** how neural networks extract hierarchical representations from raw data

▶ **Computational Limits**: effects of scale/runtime on learning, computational-statistical gaps

▶ **Transformers**: self-attention, scaling laws, memorization, in-context learning

EMPLOYMENT

Massachusetts Institute of Technology   **Cambridge, MA**
Assistant Professor, Mathematics and EECS[AI+D]   *Starting July 2026*

Harvard University – *Kempner Institute*   **Cambridge, MA**
Kempner Research Fellow   *July 2025 – June 2026*

EDUCATION

Princeton University – *Advisor: Jason D. Lee*   **Princeton, NJ**
Ph.D. in Applied and Computational Mathematics   *August 2020 – June 2025*

Duke University – *Angier B. Duke Scholar*   **Durham, NC**
Bachelor of Science in Mathematics   *August 2016 – May 2020*

HONORS

▶ **Jane Street Graduate Research Fellowship** (2024-2025)

▶ **National Science Foundation Graduate Research Fellowship** (2021-2024)

▶ **Julia Dale Award** (2020)
Duke Mathematics' highest award for graduating seniors

▶ **Putnam Honorable Mention** (2019)

▶ **Angier B. Duke Scholarship** (2016 – 2020)
Full-ride merit scholarship

PUBLICATIONS

*NOTE: * denotes equal contributions to the paper*

▶ ''The Generative Leap: Sharp Sample Complexity for Efficiently Learning Gaussian Multi-Index Models''
Alex Damian, Jason D. Lee, Joan Bruna
Advances in Neural Information Processing Systems, 2025, Spotlight Presentation

▶ ''Learning Compositional Functions with Transformers from Easy-to-Hard Data''
Zixuan Wang*, Eshaan Nichani*, Alberto Bietti, **Alex Damian**, Daniel Hsu, Jason D Lee, Denny Wu
Conference on Learning Theory, 2025

- ‣ ''Understanding Optimization in Deep Learning with Central Flows''
  Jeremy Cohen*, **Alex Damian***, Ameet Talwalkar, Zico Kolter, Jason D. Lee
  International Conference on Learning Representations, 2025

- ‣ ''Computational-Statistical Gaps in Gaussian Single-Index Models''
  Alex Damian, Loucas Pillaud-Vivien, Jason D. Lee, Joan Bruna
  Conference on Learning Theory, 2024

- ‣ ''How Transformers Learn Causal Structure with Gradient Descent''
  Eshaan Nichani, **Alex Damian**, Jason D. Lee
  International Conference on Machine Learning, 2024

- ‣ ''Smoothing the Landscape Boosts the Signal for SGD: Optimal Sample Complexity for Learning Single Index Models''
  Alex Damian, Eshaan Nichani, Rong Ge, Jason D. Lee
  Advances in Neural Information Processing Systems, 2023, Oral Presentation

- ‣ ''Fine-Tuning Language Models with Just Forward Passes''
  Sadhika Malladi, Tianyu Gao, Eshaan Nichani, **Alex Damian**, Jason D. Lee, Danqi Chen, Sanjeev Arora
  Advances in Neural Information Processing Systems, 2023, Oral Presentation

- ‣ ''Provable Guarantees for Nonlinear Feature Learning in Three-Layer Neural Networks''
  Eshaan Nichani, **Alex Damian**, Jason D. Lee
  Advances in Neural Information Processing Systems, 2023, Spotlight Presentation

- ‣ ''Self-Stabilization: The Implicit Bias of Gradient Descent at the Edge of Stability''
  Alex Damian*, Eshaan Nichani*, Jason D. Lee
  International Conference on Learning Representations, 2023

- ‣ ''Neural Networks can Learn Representations with Gradient Descent''
  Alex Damian, Jason D. Lee, Mahdi Soltanolkotabi
  Conference on Learning Theory, 2022

- ‣ ''Label Noise SGD Provably Prefers Flat Global Minimizers''
  Alex Damian, Jason D. Lee, Tengyu Ma
  Advances in Neural Information Processing Systems, 2021

- ‣ ''PULSE: Self-Supervised Photo Upsampling via Latent Space Exploration of Generative Models''
  Sachit Menon*, **Alex Damian***, Nikhil Ravi, Shijia Hu, Cynthia Rudin
  IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020

- ‣ ''New Techniques for Preserving Global Structure and Denoising with Low Information Loss in Single-Image Super-Resolution''
  Yijie Bei*, **Alex Damian***, Shijia Hu*, Sachit Menon*, Nikhil Ravi*, Cynthia Rudin
  IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2018

- ‣ ''Squeeze-free Hamiltonian Paths in Grid Graphs''
  Alex Damian, Robin Flatland
  Canadian Conference on Computational Geometry, 2015

INVITED TALKS

▶ ''Understanding Optimization in Deep Learning with Central Flows''
Optimization Unplugged Workshop at EPFL, August 2025

▶ ''Learning From Gaussian Data: Single and Multi-Index Models''
Cargese Workshop, August 2025

▶ ''Learning From Gaussian Data: Single and Multi-Index Models''
Youth in High Dimensions, July 2025

▶ ''How Transformers Learn Causal Structure with Gradient Descent''
Aspen Workshop on Foundation Models, October 2024

▶ ''Computational-Statistical Gaps in Gaussian Single-Index Models''
Leman Workshop, May 2024

▶ ''Computational-Statistical Gaps in Gaussian Single-Index Models''
Cambridge, April 2024

▶ ''The Secret Life of Optimizers: How Optimizers Implicitly Shape the Loss Landscape''
Duke University, October 2023

▶ ''Smoothing the Landscape Boosts the Signal for SGD''
Cargese Workshop, August 2023

▶ ''Recent Advances in the Generalization Theory of Neural Networks''
ICML Tutorial with Tengyu Ma, July 2023

▶ ''Implicit Biases of Stochastic Gradient Descent''
PACM Graduate Student Seminar, October 2022

▶ ''Label Noise SGD Provably Prefers Flat Global Minimizers''
New York University, December 2021

TEACHING

▶ **Teaching Assistant** (ECE 538B, Princeton University)   August 2021 - December 2021

▶ **Teaching Assistant** (MATH 403, Duke University)   January 2020 - May 2020

▶ **Lead Instructor** (MATH 281S, Duke University)   August 2019 - December 2019

Developed the curriculum (lectures and problem sets) and prepared a class of 25 students for undergraduate math competitions including the Putnam and the Virginia Tech Regional Math Contest.

SERVICE

▶ **Reviewing**   February 2022 - Present
Reviewed papers for NeurIPS (2023-2025), COLT (2022-2025), ICML (2024, 2025), and JMLR (2023).

▶ **Mathematics of Modern Machine Learning (M3L) Workshop**   April 2023 - December 2023
Co-organized the Mathematics of Modern Machine Learning (M3L) workshop at NeurIPS 2023.