

EDUCATION

Massachusetts Institute of Technology (MIT)*PhD Candidate in EECS, advised by Aleksander Mądry; GPA: 5.0/5.0*

Cambridge, MA

*Sep 2020 – Present***Stanford University***MS in Computer Science, Concentration: AI/Theory; GPA: 3.938/4.0*

Stanford, CA

*Jan 2018 – Jun. 2019**BS in Computer Science with Honors and Distinction, Concentration: Systems; GPA: 3.960/4.0* *Sept. 2014 – Jun. 2018*WORK EXPERIENCE AND RESEARCH

Tesla Computer Vision Scientist*Tesla Autopilot*

Palo Alto, CA

Aug 2019 - Sep 2020

- *Vision Models for Self-Driving*: Designed and trained computer vision models as part of Tesla Autopilot's Full Self-Driving (FSD) team.

Graduate Research Assistant*Advised by Matei Zaharia*

Stanford, CA

Sep 2018 - Jun 2019

- *I/O Lower Bounds on Computation Graphs*: Developed automated lower bounds on the I/O of arbitrary computations using the spectra of the computation graph's Laplacian.

Undergraduate Research Assistant*Research Assistant, advised by Jure Leskovec*

Stanford, CA

Sep 2016 - Sep 2018

- *Motif Aware State Assignment (MASA) in Noisy Time Series Data*: Designed an alternating minimization method to robustly assign states in noisy time series data by leveraging knowledge of recurring state patterns (motifs).

Facebook, Research Engineering Intern*Facebook AI Research (FAIR)*

New York, NY

Jun 2018 - Sep 2018

- *Dialogue models for Text Adventure Games*: Generated crowd-sourcing tasks and trained initial generative models for project to place context-aware dialogue agents in a text adventure game.

Google, Software Engineering Intern*Indexing Performance Team*

Mountain View, CA

Jun 2017 - Sep 2017

- *Tiered Storage Prototype*: Designed a system for the Tiered Storage of indexing records on top of Spanner based on the priority of documents.

Dropbox, Software Engineering Intern*Web Performance Team*

San Francisco, CA

Jun 2016 - Sep 2016

- *Tracking JS Module Load Times*: Developed tools for tracking the load and execution times of Javascript modules when required by the Asynchronous Module Definition framework RequireJs.

PUBLICATIONS (* INDICATES EQUAL CONTRIBUTION)

- **Saachi Jain***, Hannah Lawrence*, Ankur Moitra, Mądry. Distilling Model Failures as Directions in Latent Space. *International Conference on Learning Representations (ICLR) 2023, Spotlight (Top 25%)*. [Paper](#)
- **Saachi Jain***, Hadi Salman*, Alaa Khaddaj*, Eric Wong, Sung Min Park, Aleksander Mądry. A Data-Based Perspective on Transfer Learning. *Conference on Computer Vision and Pattern Recognition (CVPR) 2023*. [Paper](#)
- **Saachi Jain***, Dimitris Tsipras*, Aleksander Mądry. Combining Diverse Feature Priors. *International Conference on Machine Learning (ICML) 2022*. [Paper](#)
- **Saachi Jain***, Hadi Salman*, Eric Wong, Pengchuan Zhang, Vibhav Vineet, Sai Vemprala, Aleksander Mądry. Missingness Bias in Model Debugging. *International Conference on Learning Representations (ICLR) 2022*. [Paper](#)
- Hadi Salman*, **Saachi Jain***, Eric Wong*, Aleksander Mądry. Certified Patch Robustness via Smoothed Vision Transformers. *Conference on Computer Vision and Pattern Recognition (CVPR) 2022*. [Paper](#)

- **Saachi Jain** and Matei Zaharia. Spectral Lower Bounds on the I/O Complexity of Computation Graphs. *Symposium on Parallelism in Algorithms and Architectures (SPAA) 2020*. [Paper](#)
- **Saachi Jain**, David Hallac, Rok Susic, and Jure Leskovec. MASA: Motif-Aware State Assignment in Noisy Time Series Data. *Workshop on Mining and Learning from Time Series (MiLeTS) at SIGKDD 2019*. [Paper](#)
- Jack Urbanek, Angela Fan, Siddharth Karamcheti, **Saachi Jain**, Samuel Humeau, Emily Dinan, Tim Rocktaschel, Douwe Kiela, Arthur Szlam, Jason Weston. Learning to Speak and Act in a Fantasy Text Adventure Game. *Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP 2019)*. [Paper](#)

PREPRINTS (* INDICATES EQUAL CONTRIBUTION)

- Joshua Vendrow*, **Saachi Jain***, Logan Engstrom, Aleksander Mądry. Dataset Interfaces: Diagnosing Model Failures Using Controllable Counterfactual Generation. *2023*. [Paper](#)
- Hadi Salman*, **Saachi Jain***, Andrew Ilyas*, Logan Engstrom*, Eric Wong, Aleksander Mądry. When does Bias Transfer in Transfer Learning? *2022*. [Paper](#)
- **Saachi Jain***, Adityanarayanan Radhakrishnan*, Caroline Uhler. A Mechanism for Producing Aligned Latent Spaces with Autoencoders. *2021*. [Paper](#)

TEACHING EXPERIENCE

- *Head Course Assistant*: Introduction to Computer Networking (CS144), *Winter 2019*
- *Teaching Assistant*: Cybersecurity (IPS 268), *Fall 2018*
- *Course Assistant*: Advanced Networking (CS244), *Spring 2018*
- *Course Assistant*: Operating Systems (CS140), *Winter 2018*
- *Section Leader*: Programming Abstractions (CS106A/B), *Spring 2016 - Spring 2017*

AWARDS

- Two Sigma Diversity PhD Fellowship
- Apple Fellowship