

# PREETUM NAKKIRAN

preetum@ucsd.edu | [preetum.nakkiran.org](http://preetum.nakkiran.org)  
Postdoctoral Scholar, University of California San Diego

**Research Objective:** I try to do ML research that will remain relevant in 10+ years. Currently I work on understanding when and why deep learning works (*generalization, representation, pretraining, etc.*) and building conceptual tools for reasoning about learning more broadly.

## ACADEMIC POSITIONS

---

*University of California, San Diego* 2021-Present  
Postdoctoral Scholar. Hosted by Prof. Mikhail Belkin.  
Supported by NSF/Simons Collaboration on the Theoretical Foundations of Deep Learning  
(<https://deepfoundations.ai/>)

## EDUCATION

---

*University of California, Berkeley* 2012-2016  
B.S. in Electrical Engineering and Computer Science (GPA 4.0/4.0)

*Harvard University* 2016-2021  
PhD in Computer Science  
Thesis: *Towards an Empirical Theory of Deep Learning* (<https://preetum.nakkiran.org/thesis.pdf>)  
Advisors: Boaz Barak and Madhu Sudan.

## INDUSTRY EXPERIENCE & ACADEMIC VISITS

---

**Research Intern: Google Brain** Summer 2020  
Independent research with Hanie Sedghi and Behnam Neyshabur at Google Brain, on generalization in deep learning (one ICLR paper).

**Research Intern: OpenAI** Summer 2019  
Worked with Ilya Sutskever at OpenAI, on independent research in deep learning. Worked on projects involving adversarial examples, the effect of learning-rate on generalization, and “double descent” (ICLR, Distill papers).

**Visiting Student: UC Berkeley** Spring 2020  
Visited Jacob Steinhardt at UC Berkeley and interacted with the Berkeley ML research groups.

**Research Intern: Google Speech Team** Summer 2014  
Worked on deep-learning for detecting “ok google” audio on a phone. (two research publications and two patents).

**Software Intern: Google Ads Team** Summer 2013  
On the Ads team, worked on regression models for forecasting trends in ads-metrics.

## SELECTED INVITED TALKS

---

- CMU ML Faculty Seminar Talk 2021
- Simons Institute Talk (“Is Overfitting Actually Benign?”) 2021
- Deep Learning Classics and Trends 2021
- UToronto, Roger Grosse group (“The Deep Bootstrap”) 2021
- Foundations of Deep Learning: Simons Collaboration seminar (“The Deep Bootstrap”) 2021
- UCLA Big Data and Machine Learning seminar (“Distributional Generalization”) 2021
- Google X Blueshift (“Distributional Generalization”) 2020
- Max Planck Institute: Math ML Seminar (“Distributional Generalization”) 2020
- IAS Theory of Deep Learning Workshop (“Deep Double Descent”) 2019
- China Theory Week (“Algorithmic Polarization”) 2018

## PUBLICATIONS AND MANUSCRIPTS

---

Publications in: NeurIPS, ICLR, Distill, STOC, COLT, ITCS, ISIT, APPROX, INTERSPEECH, ICASSP.

- **“What You See is What You Get: Distributional Generalization for Algorithm Design in Deep Learning”**  
Bogdan Kulynych\*, Yao-Yuan Yang\*, Yaodong Yu, Jarosław Błasiok, Preetum Nakkiran.  
*In submission to ICML 2022.*
- **“Deconstructing Distributions: A Pointwise Framework of Learning”**  
Gal Kaplun\*, Nikhil Ghosh\*, Saurabh Garg, Boaz Barak, Preetum Nakkiran.

- **“Limitations of Neural Collapse for Understanding Generalization in Deep Learning”**  
Like Hui, Mikhail Belkin, Preetum Nakkiran.  
*In submission to ICML 2022.*
- **“Deep Double Descent: Where Bigger Models and More Data Hurt”**  
Preetum Nakkiran, Gal Kaplun\*, Yamini Bansal\*, Tristan Yang, Boaz Barak, Ilya Sutskever.  
*ICLR 2020.*  
<https://arxiv.org/pdf/1912.02292.pdf>
- **“The Deep Bootstrap: Good Online Learners are Good Offline Generalizers”**  
Preetum Nakkiran, Behnam Neyshabur, Hanie Sedghi.  
*ICLR 2021.*  
<https://arxiv.org/pdf/2010.08127.pdf>
- **“Distributional Generalization: A New Kind of Generalization”**  
Preetum Nakkiran\*, Yamini Bansal\*.  
*ICML OPPO 2021.*  
<https://arxiv.org/pdf/2009.08092.pdf>
- **“Learning Rate Annealing Can Provably Help Generalization, Even for Convex Problems”**  
Preetum Nakkiran.  
*ICML OPT 2020: Best Student Paper.*  
<https://arxiv.org/pdf/2005.07360.pdf>
- **“Optimal Regularization Can Mitigate Double Descent”**  
Preetum Nakkiran, Prayaag Venkat, Sham Kakade, Tengyu Ma.  
*ICLR 2021.*  
<https://arxiv.org/pdf/2003.01897.pdf>
- **“SGD on Neural Networks Learns Functions of Increasing Complexity”**  
Preetum Nakkiran, Gal Kaplun, Dimitris Kalimeris, Tristan Yang, Benjamin L Edelman, Fred Zhang, Boaz Barak.  
*NeurIPS 2019 Spotlight.*  
<https://arxiv.org/pdf/1905.11604>
- **“Computational Limitations in Robust Classification and Win-Win Results”**  
Akshay Degwekar, Preetum Nakkiran, Vinod Vaikuntanathan.  
*COLT 2019.*  
<https://arxiv.org/pdf/1902.01086>
- **“Adversarial Examples are Just Bugs, Too”**  
Preetum Nakkiran.  
*Distill 2019.*  
<https://distill.pub/2019/advex-bugs-discussion/response-5/>
- **“The Generic Holdout: Preventing False-Positives in Adaptive Data Science”**  
Jarosław Błasiok and Preetum Nakkiran.  
*Preprint.*  
<https://arxiv.org/pdf/1809.05596.pdf>
- **“Algorithmic Polarization for Hidden Markov Models”**  
Venkatesan Guruswami, Preetum Nakkiran, Madhu Sudan.  
*ITCS 2019.*  
<https://arxiv.org/pdf/1810.01969.pdf>
- **“Differentially Private Simultaneous Mechanisms: A New Model and Mechanisms”**  
Rohit Agrawal, Christina Ilvento, Preetum Nakkiran.  
*TPDP 2019.*
- **“Tracking the  $\ell_2$  Norm with Constant Update Time”**  
Chi-Ning Chou, Zhixian Lei, Preetum Nakkiran.

*APPROX-RANDOM 2019.*

<https://arxiv.org/pdf/1807.06479.pdf>

- **“General Strong Polarization”**  
Jarosław Błasiok, Venkatesan Guruswami, Preetum Nakkiran, Atri Rudra, Madhu Sudan.  
*STOC 2018.*  
<https://arxiv.org/pdf/1802.02718.pdf>
- **“Predicting Positive and Negative Links with Noisy Queries: Theory & Practice”**  
Charalampos E. Tsourakakis, Michael Mitzenmacher, Kasper Green Larsen, Jarosław Błasiok, Ben Lawson, Preetum Nakkiran, Vasileios Nakos.  
*Allerton 2018.*  
<https://arxiv.org/pdf/1709.07308.pdf>
- **“Near-Optimal UGC-hardness of Approximating Max  $k$ -CSP<sub>R</sub>”**  
Pasin Manurangsi, Preetum Nakkiran, and Luca Trevisan.  
*APPROX 2016.*  
<http://arxiv.org/pdf/1511.06558v1>
- **“Optimal Systematic Distributed Storage Codes with Fast Encoding”**  
Preetum Nakkiran, K.V. Rashmi, and Kannan Ramchandran.  
*ISIT 2016.*  
<http://arxiv.org/pdf/1509.01858>
- **“Having Your Cake and Eating It Too: Jointly Optimal Erasure Codes for I/O, Storage, and Network-bandwidth”**  
K.V. Rashmi, Preetum Nakkiran, Jingyan Wang, Nihar B. Shah, Kannan Ramchandran.  
*USENIX Conference on File and Storage Technologies (FAST) 2015.*  
<https://www.usenix.org/system/files/conference/fast15/fast15-paper-rashmi.pdf>
- **“Compressing Deep Neural Networks using a Rank-Constrained Topology”**  
Preetum Nakkiran, Raziel Alvarez, Rohit Prabhavalkar, and Carolina Parada.  
*Conference of the International Speech Communication Association (Interspeech) 2015.*  
<http://research.google.com/pubs/archive/43813.pdf>
- **“Automatic Gain Control and Multi-style Training for Robust Small-Footprint Keyword Spotting with Deep Neural Networks”**  
Rohit Prabhavalkar, Raziel Alvarez, Carolina Parada, Preetum Nakkiran, and Tara Sainath.  
*International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2015.*  
<http://research.google.com/pubs/archive/43289.pdf>
- **“Fundamental limits on communication for oblivious updates in storage networks”**  
Preetum Nakkiran, Nihar B Shah, and K.V. Rashmi.  
*IEEE Global Communications Conference (GLOBECOM) 2014.*  
<http://arxiv.org/pdf/1409.1666>
- **“Iterative Hard Thresholding for Keyword Extraction from Large Text Corpora”**  
Steve Yadlowsky, Preetum Nakkiran, Jingyan Wang, Rishi Sharma, and Laurent El Ghaoui.  
*International Conference on Machine Learning and Applications (ICMLA) 2014.*  
<http://www.eecs.berkeley.edu/~elghaoui/Pubs/IhtSummarizationICMLA14.pdf>

## PATENTS

---

- **“Rank-constrained neural networks”** (US9767410B1)  
Raziel Alvarez Guevara, Preetum Nakkiran
- **“Automatic selective gain control of audio data for speech recognition”** (US9842608B2)  
Raziel Alvarez Guevara, Preetum Nakkiran

## NOTABLE AWARDS & FUNDING

---

- Google PhD Fellowship Recipient 2020
- NSF Graduate Research Fellowship Recipient 2016
- Berkeley Outstanding Graduate Student Instructor Award 2016
- U.S. Physics Team - Top 20 students in US Physics Olympiad 2012

## TEACHING & SERVICE

---

**Co-founder of Harvard ML Foundations Seminar ([mlfoundations.org](http://mlfoundations.org))**

Reviewer: NeurIPS, ICLR, ICML, JMLR, Distill, STOC, CRYPTO, ITCS, IEEE Transactions on Information Theory.

Teaching:

- Teaching Assistant, Berkeley: CS 70 (Discrete Mathematics) Fall 2014
- Head Teaching Assistant, Berkeley: EE 16A (Linear Algebra, Systems, and Circuits) Fall 2015
- Head Teaching Assistant, Berkeley: EE 121 (Coding for Digital Communication) Spring 2016
- Teaching Assistant, Harvard: CS 221 (Computational Complexity) Spring 2018

Received Berkeley's **Outstanding Graduate Student Instructor Award (2016)**  
and Harvard's **Distinction and Excellence in Teaching Award (2018)**.