

# VAISHNAVH NAGARAJAN

[www.cs.cmu.edu/~vaishnan](http://www.cs.cmu.edu/~vaishnan)

[vaishnavh@cs.cmu.edu](mailto:vaishnavh@cs.cmu.edu)

## Education

- **Ph.D. in Computer Science** *Fall 2015 - present*  
Carnegie Mellon University (CMU)  
**Advisor:** *J. Zico Kolter*  
**Research Interests:** Theoretical foundations of AI, Machine Learning, Deep Learning; Learning Theory & Optimization.
- **Bachelors in Technology** *2011 - 2015*  
Indian Institute of Technology (IIT) Madras GPA: 9.88/10.0 (Rank 2)  
**Advisor:** *Balaraman Ravindran*  
**Thesis:** KWIK Inverse Reinforcement Learning

## Conference Publications

- [1] **Understanding the failure modes of out-of-distribution generalization.** [\[arxiv\]](#)  
Vaishnavh Nagarajan, Anders Andreassen and Behnam Neyshabur  
*International Conference on Learning Representations 2021 (ICLR 2021)*
- [2] **A learning theoretic perspective on local explainability.** [\[arxiv\]](#)  
Jeffrey Li\*, Vaishnavh Nagarajan\*, Gregory Plumb and Ameet Talwalkar  
*International Conference on Learning Representations 2021 (ICLR 2021)*
- [3] **Provably safe PAC-MDP exploration using analogies.** [\[arxiv\]](#)  
Melrose Roderick, Vaishnavh Nagarajan and J. Zico Kolter  
*In Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS 2021)*
- [4] **Uniform convergence may be unable to explain generalization in deep learning.** [\[arxiv\]](#)  
Vaishnavh Nagarajan and J. Zico Kolter.  
*In Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*  
**Oral** paper (0.55% acceptance) and winner of **The Outstanding New Directions Paper Award**
- [5] **Deterministic PAC-Bayesian generalization bounds for deep networks via generalizing noise-resilience.** [\[arxiv\]](#)  
Vaishnavh Nagarajan and J. Zico Kolter.  
*International Conference on Learning Representations 2019 (ICLR 2019)*
- [6] **Revisiting adversarial risk.** [\[arxiv\]](#)  
Arun Sai Suggala, Adarsh Prasad, Vaishnavh Nagarajan and Pradeep Ravikumar  
*In Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS 2019)*

- [7] **Gradient descent GAN optimization is locally stable** [\[arxiv\]](#)  
Vaishnavh Nagarajan and J. Zico Kolter.  
*In Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*  
**Oral** paper (1.2% acceptance)
- [8] **Lifelong learning in costly feature spaces.** [\[arxiv\]](#)  
with Avrim Blum and Maria-Florina Balcan.  
*In Proceedings of the 28th International Conference in Algorithmic Learning Theory (ALT 2017)*
- [9] **Learning-theoretic foundations of algorithm configuration for combinatorial partitioning problems.** [\[arxiv\]](#)  
with Maria-Florina Balcan, Ellen Vitercik and Colin White.  
*In Proceedings of the 30th Annual Conference on Learning Theory (COLT 2017)*
- [10] **Every team deserves a second chance: Identifying when things go wrong.** [\[PDF\]](#)  
Vaishnavh Nagarajan\*, Leandro S. Marcolino\* and Milind Tambe.  
*In Proceedings of the 14th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2015)*

## Journal Publications

- [11] **Lifelong learning in costly feature spaces.**  
with Avrim Blum and Maria-Florina Balcan.  
*In Theoretical Computer Science (invited) (TCS 2019)*
- [12] **Every team deserves a second chance: An extended study on predicting team performance.**  
Leandro S. Marcolino, Aravind Lakshminarayanan, Vaishnavh Nagarajan and Milind Tambe.  
*In Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS 2016)*

## Workshops/Short papers

- [13] **Uniform convergence may be unable to explain generalization in deep learning.**  
Vaishnavh Nagarajan and J. Zico Kolter.
- *In Workshop on Understanding and Improving Generalization in Deep Learning. (ICML '19; spotlight talk in workshop)*
  - *In IAS/Princeton Workshop on Theory of Deep Learning 2019 (spotlight talk)*
- [14] **Theoretical Insights into Memorization in GANs.** [\[PDF\]](#)  
Vaishnavh Nagarajan, Colin Raffel and Ian Goodfellow.  
*In Workshop on Integration of Deep Learning Theories (NeurIPS 2018)*
- [15] **Generalization in Deep Learning: The Role of Distance from Initialization.** [\[arxiv\]](#)  
Vaishnavh Nagarajan and J. Zico Kolter.  
*In Workshop on Deep Learning: Bridging Theory and Practice (NeurIPS 2017)*  
**spotlight** talk in workshop)

- [16] **A reinforcement learning approach to online learning of decision trees.** [\[arxiv\]](#)  
 Abhinav Garlapati, Aditi Raghunathan, Vaishnavh Nagarajan and Balaraman Ravindran.  
*In Proceedings of the 12th European Workshop on Reinforcement Learning, International Conference on Machine Learning (EWRL-ICML 2015)*
- [17] **KWIK inverse reinforcement learning.** [\[PDF\]](#)  
 Vaishnavh Nagarajan and Balaraman Ravindran.  
*The Multi-disciplinary Conference on Reinforcement Learning and Decision Making. (RLDM 2015)*

## Talks

- Uniform convergence may be unable to explain generalization in deep learning.
  - Google Research (New York) Learning Theory (**invited**) *Fall 2020*
  - Google Brain (Mountain View) Deep Learning Phenomena (**invited**) *Summer 2020*
  - Center for Human Compatible AI, UC Berkeley (**invited**) *Summer 2020*
  - NeurIPS 2019 Oral presentation *Winter 2019*
  - CMU AI Lunch *Fall 2019*
  - IAS/Princeton University Workshop on Theory of Deep Learning: Where next? *Fall 2019*
  - ICML Workshop: Understanding and Improving Generalization in Deep Learning *Summer 2019*
- Gradient Descent GAN optimization is locally stable.
  - NeurIPS 2017 Oral presentation *Winter 2017*
  - CMU AI lunch *Fall 2017*
- Lifelong learning in costly feature spaces.
  - ALT 2017 *Fall 2017*
- Learning the best algorithm for max-cut, clustering, and other partitioning problems.
  - Learning, Algorithm Design & Beyond Worst-Case Analysis, Simons Institute, Berkeley. (**invited**) *Fall 2016*
  - CMU Theory Lunch *Fall 2016*

## Service

- Reviewer for ALT 2021; ICLR 2021 (**reviewer award**); NeurIPS 2020 (top 10% reviewer), 2019 (top 50% reviewer), 2018 (top 30% reviewer); ICML 2021 (Expert reviewer), 2020, 2019 (**top 5%** reviewer), COLT 2019, AISTATS 2019.
- Member of the admissions committee for CMU's MS in Computer Science program for 2018.

- Representative of the Computer Science Department in the SCS4ALL PhD Committee, a student advisory council for the CMU School of Computer Science. *Fall 2017 - Fall 2018*
- Organized the Learning Theory Reading group in CMU. *Fall 2016*

## Teaching

- Teaching assistant, 10-715: Advanced Introduction to Machine Learning *Fall 2016*
- Teaching assistant, 15-780: Graduate Artificial Intelligence *Spring 2018*

## Internships

- **PhD Research Internship** *Summer 2020*  
 Google X  
**Host:** Behnam Neyshabur  
 Theoretically explained when and why machine learning models fail to generalize under test-time distribution shifts.
- **PhD Research Internship** *Summer 2019*  
 Bosch Center for AI  
**Host:** David Reeb  
 Developed generalization bounds for high-dimensional linear models that circumvent limitations of uniform convergence bounds.
- **PhD Research Internship** *Summer 2018*  
 Google Brain  
**Host:** Colin Raffel, Ian Goodfellow  
 Explained why Generative Adversarial Networks (GANs) counterintuitively do not memorize their training data. Explored metrics for measuring diversity of GAN samples and developed a theoretically-grounded technique for improving sample diversity.
- **Undergraduate Research Internship** *Summer 2014*  
 University of Southern California (USC)  
**Advisor:** Milind Tambe  
 Identified that a machine learning model can predict the success/failure of an artificial multi-agent team playing Computer Go.
- **Undergraduate Internship** *Summer 2013*  
[Report Bee](#)  
**Advisor:** Madhavan Mukund, Chennai Mathematical Institute  
 Designed an index that quantifies learning experiences of schoolchildren. Implemented the model within Report Bee's web application.

## Scholastic Achievements

- Among **national top 1%** candidates in national level olympiads (2011) in **five different fields**, namely, Informatics, Maths, Physics, Chemistry and Astronomy.
- Directly qualified for the Indian National Math Olympiad (INMO) 2011 based on outstanding performance in the Regional Math Olympiad 2010 and INMO 2010.
- One of the 35 students that qualified further for the national selection camp for International Chemistry Olympiad.
- Secured **All India Rank 70** (out of 0.5 million candidates) and State Rank 3 in IIT Joint Entrance Examination 2011, **All India Rank 56** (out of 1.1 million candidates) in All India Engineering Entrance Examination 2011 and **All India Rank 16** (out of 0.1 million candidates) in Indian Institute of Space Science and Technology Admission Test 2011

## Honors and Awards

- **Reviewer Award** at ICLR 2021 for being an outstanding reviewer.
- Winner of the **Outstanding New Directions Paper Award** at NeurIPS 2019 (given to only one out of  $\sim 1400$  accepted papers).
- Awarded the ACM-India/IARCS student grant to attend AAMAS 2015 in Istanbul, Turkey.
- One of  $\sim 30$  Viterbi-India scholars selected by Viterbi School of Engineering (USC) and Indo-US Science and Technology Forum for a fully funded research internship in Summer 2014.
- Awarded the prestigious **KVPY** Fellowship 2009 by the **Government of India** to attract highly motivated students for pursuing a research career in science.
- Invited participant in the Council of Scientific and Industrial Research Programme on Youth for Leadership in Science 2009.

## Other Activities

- Board member of CMU Indian Graduate Student Association (IGSA). *Dec 2015 - Dec 2018*
- National Service Scheme Volunteer involved in Scientific Toys & Assistive Technology. *2011-12*
- Taught basic maths to underprivileged primary school children in villages in India, in association with the NGO, *AID India*. *Dec 2011*
- Scribe for the students of **Vidya Sagar** (formerly, the Spastics Society of India). *2008-09*