

ANANYA KUMAR

ananya@cs.stanford.edu \diamond ananyakumar.wordpress.com

EDUCATION

Stanford University 2018-Now
PhD in Machine Learning, SGF Scholar

- Doing research on distribution shifts and uncertainty estimation in the Statistical Machine Learning group, advised by Percy Liang and Tengyu Ma.

Carnegie Mellon University 2013-2017
B.S. in Computer Science, Minor in Math, 4.00/4.00 GPA, PBK and ACS Scholar

- Selected Courses: Machine Learning (MS), Convex Optimization (PhD), Intermediate Statistics (PhD), Advanced Mobile Robot Design (PhD), Theorist's Toolkit (PhD), Deep Reinforcement Learning (PhD), Operating Systems, Honors Analysis I & II

PUBLICATIONS (* EQUAL CONTRIBUTION)

[9] Sang Michael Xie*, [Ananya Kumar*](#), Robbie Jones*, Fereshte Khani, Tengyu Ma, Percy Liang. In-N-Out: Pre-Training and Self-Training using Auxiliary Information for Out-of-Distribution Robustness. International Conference on Learning Representations (ICLR) 2021.

[8] Selective classification can magnify disparities across groups. Erik Jones*, Shiori Sagawa*, Pang Wei Koh*, [Ananya Kumar](#), and Percy Liang. International Conference on Learning Representations (ICLR) 2021.

[7] Self-Training Avoids Using Spurious Features Under Domain Shift. Yining Chen*, Colin Wei*, [Ananya Kumar](#), Tengyu Ma. Neural Information Processing Systems (NeurIPS) 2020.

[6] Understanding Self-Training for Gradual Domain Adaptation. [Ananya Kumar](#), Tengyu Ma, Percy Liang. International Conference on Machine Learning (ICML) 2020.

[5] Verified Uncertainty Calibration. [Ananya Kumar](#), Percy Liang, Tengyu Ma. Neural Information Processing Systems (NeurIPS) 2019, **Spotlight** (3.0% acceptance rate).

[4] Rigorous Agent Evaluation: An Adversarial Approach to Uncover Catastrophic Failures. Jonathan Uesato*, [Ananya Kumar*](#), Csaba Szepesvari*, Tom Erez, Avraham Ruderman, Keith Anderson, Krishnamurthy (Dj) Dvijotham, Nicolas Heess, Pushmeet Kohli. International Conference on Learning Representations (ICLR) 2019.

[3] Consistent Generative Query Networks. [Ananya Kumar](#), S. M. Ali Eslami, Danilo Rezende, Marta Garnelo, Fabio Viola, Edward Lockhart, Murray Shanahan. NeurIPS Workshop on Bayesian Deep Learning 2018 and NeurIPS Physics Workshop 2018.

[2] Approximate Convex Hull of Data Streams. Avrim Blum*, Vladimir Braverman*, [Ananya Kumar*](#), Harry Lang*, Lin Yang* (Alphabetical Order). International Colloquium on Automata, Languages, and Programming (ICALP) 2018.

[1] Parallel Functional Arrays. [Ananya Kumar](#), Guy E. Blelloch and Robert Harper. Principles of Programming Languages (POPL) 2017.

WORK EXPERIENCE

Research Engineer III at Alphabet DeepMind, London, 2017 - 2018

- Developed a deep generative model for video and scene prediction.
- Developed methods for estimating failure rates of RL agents in safety critical environments.

Intern at Uber ATG, Self Driving Car Team, Pittsburgh, Summer 2016

- Interned in the motion planning team.

AWARDS

2019 Top 400 reviewer for NeurIPS 2019

2018 Stanford Graduate Fellowship

2017 Alumni Award in CS and 2nd place in Math Competition for Senior Thesis Research

2017 CMU's Senior Leadership Recognition

2016 Putnam Competition - Rank 109.5 (Top 200 in North America)

2016 Mark Stehlik SCS Alumni Undergraduate Impact Scholarship

2015 Top 10 finalists at HackMIT 2015 (out of 195 projects)

2013 ACM Regional Contest - 7th/126

2012 Intel International Science and Engineering Fair - 2nd Award

2012 2nd place at Singapore National Olympiad of Informatics

SERVICE, TEACHING, MENTORSHIP

Reviewing NeurIPS 2019, ICLR 2021, TPAMI Journal, Safe ML Workshop at ICLR 2019, UDL Workshop at ICML 2020.

Mentorship Advised undergraduates Robbie Jones and Fahim Tajwar for summer CURIS research projects in 2020, and mentor for SAIL mentorship program in 2019 and 2020.

Teaching Teaching Assistant for Convex Optimization (EE364A) at Stanford University. Teaching Assistant for "Great Theoretical Ideas in Computer Science" at Carnegie Mellon University.

TECHNICAL SKILLS

Languages Python, C/C++, Matlab, Java, SML, Go, Javascript, SQL, x86 assembly
Technologies TensorFlow, PyTorch, Google Cloud Platform, Unix, MPI, GIT, L^AT_EX

REFERENCES

Percy Liang, Associate Professor at Stanford University, pliang@cs.stanford.edu

Tengyu Ma, Assistant Professor at Stanford University, tengyuma@stanford.edu

Avrim Blum, Professor and Chief Academic Officer at TTI-Chicago, avrim@ttic.edu