Karan Chadha

Research Interests

I am interested in answering questions in **Trustworthy AI**, **Differential Privacy**, **Uncertainty Quantification**, and **Federated Learning**, for which I use theoretical tools from statistics and optimization, complemented with rigorous experimentation. In the near future, I am interested in the following directions:

- Studying the empirical privacy leakage for modern ML models (Foundation Models/LLMs) in realistic attack scenarios and moving beyond DP to find application-relevant definitions to evaluate models on privacy, robustness, fairness and copyright.
- Developing mitigation strategies and training algorithms that ensure foundation models adhere to trustworthy behavior.
- Building better (trustworthy) algorithms and systems for practically relevant ML and data analytics tasks like recommendations, ranking, frequency estimation, etc.
- Leveraging the capabilities of foundation models for private synthetic data generation to help with private training in low resource tasks.

EDUCATION

Stanford University

Ph.D. in Electrical Engineering, GPA: 4.00/4.00 Advised by Prof. John Duchi

Indian Institute of Technology Bombay

Dual Degree (B. Tech. + M. Tech.) in Electrical Engineering, GPA: 9.68/10 Advised by Prof. Ankur Kulkarni, Prof. Jayakrishnan Nair and Prof. Vivek Borkar.

INDUSTRY EXPERIENCE

Student Researcher, Google Deepmind

Worked with Matthew Jagielski and Nicolas Papernot on auditing private prediction.

Machine Learning Intern, Apple

Worked with *Omid Javidbakht*, *Audra McMillan*, *Vitaly Feldman* and *Kunal Talwar* on learning histograms in the unknown dictionary setting with aggregate differential privacy.

Preprints

- Auditing Private Prediction [PDF] <u>K. Chadha</u>, M. Jagielski, N. Papernot, C. Choquette-Choo, and M. Nasr [Arxiv:2402.09403]
- Resampling methods for private statistical inference [PDF] <u>K. Chadha</u>, J. C. Duchi and R. Kuditipudi [Arxiv:2402.07131]

2019–Present

2014 - 2019

Summer 2022

Summer 2023

PUBLICATIONS

- Differentially Private Heavy Hitter Detection using Federated Analytics [PDF] <u>K. Chadha</u>, J. Chen, J. C. Duchi, V. Feldman, H. Hashemi, O. Javidbakht, A. McMillan, and K. Talwar *IEEE SaTML 24*
- Federated Asymptotics: A model for evaluating federated learning algorithms [PDF] <u>K. Chadha</u>*, G. Cheng*, and J. C. Duchi, *AISTATS 23*
- Private optimization in the interpolation regime: faster rates and hardness results [PDF] <u>K. Chadha</u>*, H. Asi*, G. Cheng*, and J. C. Duchi *ICML 22 (Spotlight)*
- Accelerated, optimal, and parallel: Some results on model-based stochastic optimization [PDF] <u>K. Chadha</u>*, G. Cheng*, and J. C. Duchi *ICML 22*
- Minibatch stochastic approximate proximal pointmethods [PDF] <u>K. Chadha</u>*, H. Asi*, G. Cheng*, and J. C. Duchi *Neurips 2020 (Spotlight)*
- Efficiency fairness tradeoff in battery sharing [PDF] <u>K. Chadha</u>, A. A. Kulkarni and J. Nair *Operations Research Letters*, 2021
- Aggregate play and welfare in strategic interactions on networks [PDF] <u>K. Chadha</u> and A. A. Kulkarni *Journal of Mathematical Economics, 2020*
- On independent cliques and linear complementarity problems [PDF] <u>K. Chadha</u> and A. A. Kulkarni *IJPAM, 2022*
- A reinforcement learning algorithm for restless bandits [PDF] V.S. Borkar and <u>K. Chadha</u> Indian Control Conference, 2018

* denotes equal contribution

Scholarships and Awards

•	NVIDIA-TSMC Graduate Fellowship, Stanford University	2019
•	Sharad Maloo Gold Medal (for outstanding academic and extra-curricular achievements)	2019
•	Bhavesh Gandhi Memorial Prize (for standing 1st in the Masters Programme)	2019
•	Honda YES Award	2016
•	Institute Academic Prize	2017, 2018

Skills & Courses

- Courses: Asymptotic Statistics, Information Theory and Statistics, Convex Optimization
- Programming Languages & Frameworks: Python, Numpy, JAX, Pytorch, Tensorflow

ACADEMIC SERVICE

- Reviewer for NeurIPS, ICLR, AISTATS, ICML, SaTML, TMLR
- Organizer, ML Lunch, Stanford, Fall 2020
- Organizer, Workshop on Games and Networks, IIT Bombay, 2019