

WORK EXPERIENCE

Research Scientist, Meta July 2024 - Present

Working as a Research Scientist in Meta's Central Applied Science team, focusing primarily on identifying, addressing, and mitigating privacy risks associated with foundation models.

- Leading the methodological framework for systematically evaluating memorization in foundation models, particularly addressing critical privacy and compliance risks such as PII leakage and copyright infringement.
- Experimenting with novel approaches to mitigate unintended memorization in large language models.
- Proposed and prototyped a masked language modeling-based PII detection system which significantly outperformed the legacy detection system, resulting in widespread adoption across multiple product teams beginning early 2025.

Student Researcher, Google Deepmind Summer 2023

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

Machine Learning Intern, Apple Summer 2022

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

EDUCATION

Stanford University 2019–Present

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

Indian Institute of Technology Bombay 2014–2019

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.

PUBLICATIONS & PREPRINTS

- **Resampling methods for private statistical inference** [\[PDF\]](#)
K. Chadha, J. C. Duchi and R. Kuditipudi
[Arxiv:2402.07131]
- **Auditing Private Prediction** [\[PDF\]](#)
K. Chadha, M. Jagielski, N. Papernot, C. Choquette-Choo, and M. Nasr
ICML 24
- **Differentially Private Heavy Hitter Detection using Federated Analytics** [\[PDF\]](#)
K. Chadha, 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** [\[PDF\]](#)
K. Chadha*, G. Cheng*, and J. C. Duchi,
AISTATS 23

- **Private optimization in the interpolation regime: faster rates and hardness results** [\[PDF\]](#)
K. Chadha*, H. Asi*, G. Cheng*, and J. C. Duchi
ICML 22 (Spotlight)
- **Accelerated, optimal, and parallel: Some results on model-based stochastic optimization** [\[PDF\]](#)
K. Chadha*, G. Cheng*, and J. C. Duchi
ICML 22
- **Minibatch stochastic approximate proximal pointmethods** [\[PDF\]](#)
K. Chadha*, H. Asi*, G. Cheng*, and J. C. Duchi
Neurips 2020 (Spotlight)
- **Efficiency fairness tradeoff in battery sharing** [\[PDF\]](#)
K. Chadha, A. A. Kulkarni and J. Nair
Operations Research Letters, 2021
- **Aggregate play and welfare in strategic interactions on networks** [\[PDF\]](#)
K. Chadha and A. A. Kulkarni
Journal of Mathematical Economics, 2020
- **On independent cliques and linear complementarity problems** [\[PDF\]](#)
K. Chadha and A. A. Kulkarni
IJPAM, 2022
- **A reinforcement learning algorithm for restless bandits** [\[PDF\]](#)
V.S. Borkar and K. Chadha
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

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