Sahasrajit Sarmasarkar

Email : sahasras@stanford.edu

Research Interests: Preference Learning, Machine Learning theory, Optimization, and Differential Privacy

Education

Stanford University *(Sept '21 -)*
PhD student in the Department of Electrical Engineering GPA 4.0/4.0

Indian Institute of Technology Bombay *(Jul '16 - Jun '21)*
Bachelors and Masters in Electrical Engineering GPA 9.74/10 (**Second** among 74 students)
Minor in Computer Science and Engineering

Select Publications

* denotes co-first authors/alphabetical order

1. Ayush Sawarni*, Sahasrajit Sarmasarkar* and Vasilis Syrgkanis, *'Preference Learning with Response Time'* [accepted at Neurips 2025] (paper)
2. Chirag Pabbaraju*, Sahasrajit Sarmasarkar*, *'A characterisation of list regression'* [accepted in ALT 2025] (paper)
3. Mohak Goyal*, Sahasrajit Sarmasarkar*, *'Metric Distortion under Probabilistic voting'* [accepted in EC 2025] (paper)
4. Ashish Goel*, Zhihao Jiang*, Aleksandra Korolova*, Kamesh Munagala*, Sahasrajit Sarmasarkar*, *'On Differential Privacy with Multiple Selections,'* [accepted in FORC '25] (paper)
5. Mohak Goyal*, Sukolsak Sakshuwong*, Sahasrajit Sarmasarkar*, Ashish Goel, *'Low Sample Complexity Participatory Budgeting,'* International Colloquium on Automata, Languages and Programming(ICALP), 2023 (paper)
6. Mohak Goyal*, Sahasrajit Sarmasarkar*, Ashish Goel, *'A Mechanism for Participatory Budgeting With Funding Constraints and Project Interactions',* Web And InterNet Economics (WINE), 2023 (paper)

Papers under Review

1. Sahasrajit Sarmsarkar, *'Shuffle and joint differential privacy for Generalized Linear Contextual Bandits'* under review at AISTATS '26

Industry Experience

Machine Learning Research Intern *Mountain View, California* *(June 2025 - Ongoing)*
Optimizing data selection hyper-parameters to improve validation metrics in a budgeted setup with only a small number of training runs available.

Quantitative Finance Intern *Morgan Stanley, New York* *(June 2023 – August 2023)*
Studied and implemented adaptive band-based hedging strategies for a portfolio of FX options with different expiries by peeking into the past and empirically showed its improvement over non-adaptive strategies.

Digital Engineering Intern *Texas Instruments, Bengaluru* *(May 2019 – June 2019)*
Devised algorithms for connections of cells and buffer insertion problems in a carry-save adder network to minimize the overall delay and wave-pipeline the whole network to run at higher speeds.

Technical Skills

Courses taken: Convex optimization, Randomized algorithms, Optimization and Algorithmic Paradigm, Machine learning theory, Algorithms in Decentralized Finance, Information Theory and Stats, Number Theory and Cryptography, Theory of Statistics, Statistical Signal Processing, Markov chains and Queueing Systems, Random graphs, Games and Information, Optimal Controls, Nonlinear Dynamical Systems, Matrix Computations, Information Theory, Error Correcting Codes and Advanced Concentration Inequalities, Reinforcement Learning

Courses taught as a CA: Optimization and Algorithmic Paradigm(CS 261), Market Design for Engineers(MS&E 230), Introduction to Probability (MS&E 120) , Probabilistic Analysis (MS&E 220) & Advanced Applied Optimization (MS & E 214)

Programming Languages: Python, R, C/C++, MATLAB

Selected Research Projects

Preference Learning With Response Time [1]

We study how adding response time data to preference learning frameworks can improve reward model learning.

- Developed orthogonal loss functions to leverage drift–diffusion modeling of response times.
- Proved exponential-to-polynomial improvements in sample efficiency for linear reward models.
- Conducted experiments on linear and neural network–based reward models, as well as a 500k text-to-image preference dataset generated from diffusion models to demonstrate the improvements.

Shuffled differential privacy for generalized linear contextual bandits [1]

- Studied (ϵ, δ) differential privacy in generalized linear contextual bandits under both stochastic and adversarial contexts, protecting context vectors and rewards.
- Designed algorithms achieving $\tilde{O}(\sqrt{T/\epsilon})$ regret under shuffle DP and joint DP, matching non-private bounds up to a $\sqrt{d}/\sqrt{\epsilon}$ factor without any distributional assumptions on context vectors.

On Targeted Advertising with Differential Privacy [4]
Worked on the design of a mechanism that takes noisy input from the user (preserving the privacy of the user) and returns multiple selections and a local model to return a single item from multiple selections.

- Proved the optimality of laplace noise addition amongst a class of differentially private mechanisms.
- Evaluated our results empirically on an existing trained deep neural network on movie-lens dataset and constructed a local PCA based model to demonstrate the advantage of multi-selection.

List PAC learning for agnostic and realizable regression [2]

We study the problem of list regression where the aim of the learner is to predict a set of labels for every input point x and the loss is measured with respect to the "best" label.

- We identify two dimensions of the hypothesis class—referred to as the k fat-shattering dimension and the k OIG dimension—whose finiteness is both necessary and sufficient for agnostic and realizable regression respectively.
- These are extensions of the fat-shattering and OIG (One Inclusion Graph) dimension that characterise learnability for PAC learning under agnostic and realizable regression.

Metric Distortion in Probabilistic Voting [3]

We extend the study of metric distortion in social choice to probabilistic voting models, including the widely-used Plackett-Luce model. Our results yield distortion bounds that better align with intuitive expectations about voting behavior.

- In this model we show that Copeland's distortion is at most 2, whereas that of Random Dictator (RD) scales with the square root of the number of candidates.
- This contrasts with classical model where RD beats Copeland with a distortion of 3 versus 5.
- We generalize Borda's classical distortion bound of 2m - 1 to Plackett-Luce models, obtaining results that range from constant to polynomial depending on level of randomness.

Scholastic Achievements

- Scored a Semester Performance Index **10/10** in the seventh and eighth semesters at IIT-Bombay and awarded an institute academic prize for the same.
- Achieved **All India Rank 1** in **Kishore Vaigyanik Protsahan Yojana (KVPY) 2016** conducted by **Indian Institute of Science,Bangalore** out of nearly **100,000** candidates.
- Secured **All India Rank 98 and 49** in **Joint Entrance Examination (Advanced) 2016** and **Joint Entrance Examination (Mains) 2016** respectively among 1.4 million students.
- Was among the top **35** students in **Indian National Physics Olympiad** out of nearly 40,000 students and selected to attend the **OCSC (Orientation-Cum Selection Camp)** for **International Physics Olympiad, 2016**.

Miscellaneous

- I oversaw **maintenance tasks** for the Stanford PB website, including the implementation of a novel **token voting** system, along with various bug fixes and minor updates.
- **Honourable mention** in Stanford **RAISE** (Research, Action, and Impact through Strategic Engagement) Doctoral Fellowship.
- Teaching Assistant for **Probability and Random Processes, Markov Chains and Queueing Systems, Optimization and Algorithmic Paradigm and Advanced Applied Optimization.**
- Led a team of 4 students of the **System Integration** and **Controlled Area Network** (CAN) subsystem in IIT Bombay Racing team as the design engineer.
- Achieved third place in the **Logic GC-2018** competition hosted by the Maths and Physics Club at IITB, and secured third position in the **Electric Jhatka GC-2017** organized by the ERC club, IITB.
- Received the **top** prize for the course project titled "**Universal Sensor Signal Conditioning**" in Analog Circuits Lab (EE 230) at IIT-B.