

# Harvineet Singh

📞 (+1) 702 849 8182

✉️ [hs3673@nyu.edu](mailto:hs3673@nyu.edu)

📄 [harvineet.github.io](https://harvineet.github.io)

🐙 Github: [harvineet](https://github.com/harvineet)

## Education

- 2018–  
(Expected) **New York University**, Center for Data Science,  
*Ph.D. in Data Science.*  
May 2023 Research keywords: causal inference, algorithmic fairness, interactive machine learning
- 2010–2015 **Indian Institute of Technology (IIT) Delhi**, Department of Mathematics,  
*Integrated Master of Technology in Mathematics and Computing.*

## Work Experience

- Jun–Oct 2022 **Research Intern**, AMAZON SCIENCE, Tübingen, Germany.  
Mentors: Dr. Chris Russell and Dr. Matthäus Kleindessner.
  - Investigated **minmax-fairness** and a method for enforcing this notion of fairness in classification.
  - Benchmarked fairness methods on multiple tabular datasets.
- Jun–Aug 2021 **Research Intern**, MICROSOFT RESEARCH, Redmond, US.  
Mentors: Dr. Emre Kiciman and Dr. Adith Swaminathan.
  - Developed methods to improve **bandit models** through feature selection for a causal target metric.
  - Applied to ad ranking pipeline and tested ways to estimate the metrics from observational data.
- Jun–Aug 2020 **Summer Research Fellow**, CRCS, Harvard University, US.  
Mentor: Dr. Hima Lakkaraju.
  - Investigated methods for **robust evaluation** of decision policies in dynamic environments.
  - Developed an approach that combines adversarial and causal approaches to robust learning.
- Jul 2015–Aug 2018 **Research Engineer**, ADOBE RESEARCH, India.  
Member of Big Data Experience Lab.
  - Devised machine learning models for **user behavior prediction** and recommender system tasks.
  - Developed prototypes and worked with product engineering teams to transfer technologies to Adobe's digital marketing solutions.

## Publications

- Elita Lobo, Harvineet Singh, Marek Petrik, Cynthia Rudin, and Himabindu Lakkaraju.  
**Data Poisoning Attacks on Off-Policy Policy Evaluation Methods.**  
*Conference on Uncertainty in Artificial Intelligence (UAI) 2022. Oral Presentation.* [Link].
- Harvineet Singh, Shalmali Joshi, Finale Doshi-Velez, and Himabindu Lakkaraju.  
**Towards Robust Off-Policy Evaluation via Human Inputs.**  
*AAAI/ACM Conference on AI, Ethics, and Society (AIES) 2022.* [Link].
- Harvineet Singh, Vishwali Mhasawade, and Rumi Chunara.  
**Generalizability challenges of mortality risk prediction models: A retrospective analysis on a multi-center database.**  
*PLOS Digital Health 2022.* [Link].
- Harvineet Singh, Rina Singh, Vishwali Mhasawade, and Rumi Chunara.  
**Fairness Violations and Mitigation under Covariate Shift.**  
*ACM Conference on Fairness, Accountability, and Transparency (FAccT) 2021.* [Link].  
Earlier version in *Workshop on Fair ML for Health at NeurIPS 2019. Oral Presentation.*

- Gaurush Hiranandani\*, Harvineet Singh\*, Prakhar Gupta\*, Iftikhar A. Burhanuddin, Zheng Wen, and Branislav Kveton.  
**Cascading Linear Submodular Bandits: Accounting for Position Bias and Diversity in Online Learning to Rank.**  
*Conference on Uncertainty in Artificial Intelligence, (UAI)* 2019. **Oral Presentation.** [Link].
- Moumita Sinha, Vishwa Vinay, and Harvineet Singh.  
**Modeling Time to Open of Emails with a Latent State for User Engagement Level.**  
*ACM International Conference on Web Search and Data Mining, (WSDM)* 2018. [Link].
- Sumit Shekhar, Dhruv Singal, Harvineet Singh, Manav Kedia, and Akhil Shetty.  
**Show and Recall: Learning What Makes Videos Memorable.** [Link].  
*Workshop on MBCC at IEEE International Conference on Computer Vision (ICCV)* 2017.

## Patents

- Sumit Shekhar, Dhruv Singal, Harvineet Singh and Atanu Sinha. 'Summarizing Video Content based on Memorability of the Video Content'. U.S. Patent No. 10/311,913.
- Ritwick Chaudhry, Harvineet Singh, Shiv K. Saini and Pradeep Dogga. 'Personalized E-Learning Using A Deep-Learning-Based Knowledge Tracing And Hint-Taking Propensity Model'. U.S. Patent No. 10,943,497.
- Moumita Sinha, Kandarp S. Khandwala, Harvineet Singh and D. P. Tejas. 'Predicting Unsubscription of Subscribing Users'. U.S. Patent No. 16/192,517.
- Prakhar Gupta, Iftikhar Ahamath Burhanuddin, Harvineet Singh and Atanu Sinha. 'Intelligent Analytics Interface'. U.S. Patent No. 10,546,003.

## Achievements

- 2022 Selected to present at **Future Leaders Summit** at University of Michigan, a symposium comprising talks from graduate students and postdocs on Responsible Data Science. [Link]
- 2019 Awarded **Travel Grant** to attend UAI 2019, FairML4H@NeurIPS 2019.
- 2014 Awarded **Scholarship** by Ministry of Human Resource Development for academic excellence.

## Professional Service

- 2022 **Tutorial Co-chair** at the Conference on Health, Inference, and Learning 2022 [Link]
- 2022 **Student Co-organizer** Trustworthy AI Day at KDD 2022. [Link]
- 2022 **Co-organizer** NYU AI School, a week-long school for undergraduates interested in AI. [Link]
- 2021 **Student Co-organizer** Deep Learning Day at KDD 2021. [Link]
- 2020–2021 **Co-lead** Causal inference and fairness group in Mechanism Design for Social Good (MD4SG).
- 2020 **Mentor** ARISE program to provide STEM research exposure to high school students.
- Paper Reviewer** FAccT 2022, 2021; NeurIPS 2022, 2021, 2020; ICLR 2022 (Highlighted Reviewer); ICML 2021, 2022; AISTATS 2022, 2021; AIES 2022; EAAMO 2021; MLHC 2021.

## Press Coverage of Work

- May 17, 2022 'Finding the Fairness in AI'. **Communications of ACM.** [Link]
- May 11, 2022 'Responsible Data Science and AI: Accounting for Dataset Shifts and Designing Fair Machine Learning Systems'. **NYU Center for Data Science Blog.** [Link]
- Feb 25, 2020 'Researchers Ask: Will You Remember that Video Tomorrow?'. **Adobe Research Blog.** [Link]
- Aug 29, 2017 'Adobe wants to bring its AI smarts to email marketing campaigns'. **TechCrunch.** [Link]
- Mar 25, 2016 'Adobe Previews In-Development Features at Summit 2016'. **Techvibes.** [Link]