

Ashudeep Singh

Applied Research Scientist, Pinterest, Inc.

(607)-379-7806
✉ mail@ashudeepsingh.com
📄 www.ashudeepsingh.com

Research Interests

Machine Learning · Responsible AI · Fairness in ML · Recommender Systems · Information Retrieval

Education

2015-2021 **Ph.D. Computer Science**, *Cornell University*, Ithaca, NY.

Advisor: Thorsten Joachims

Thesis Committee: Solon Barocas, Karthik Sridharan, David Mimno.

Title: Fairness of Exposure for Ranking Systems.

Relevant Coursework: Machine Learning Theory, Advanced topics in Machine Learning, Design and Analysis of Algorithms, Causality and Learning for Intelligent Decision Making, Ethics and Policy in Data Science.

Grade Point Average (GPA)– 4.0

2010–2015 **B.Tech.-M.Tech. Dual Degree**, *Indian Institute of Technology (IIT) Kanpur*, India.

Major: Computer Science and Engineering.

M.Tech. GPA– 10.0/10.0, BTech. GPA– 9.6/10.0 (Academic Excellence Award for all years)

Industry Experience

August 2021-present **Applied Research Scientist, Advanced Technologies Group**, *Pinterest, Inc.*, Palo Alto, CA.

- **Responsible and Inclusive AI:** Drive responsible ML and inclusive AI frameworks and practices to promote algorithmic fairness, diversity, and inclusive system design in Pinterest's production systems for personalized discovery as well as applications involving large language models and generative AI. This work has led to multiple successful product launches with user engagement wins, research papers at ACM FAccT 2023 and WISE 2023, and a tutorial at NeurIPS 2022.
- **ML for Interactive systems:** Applying and exploring various state-of-the-art machine learning algorithms that learn from sequential human feedback data to enhance Pinterest's search and recommendation systems, spanning Graph Learning, Reinforcement Learning, and sequential models such as Transformers.

January–May 2020 **Research Intern**, *Google Brain*, New York, NY.

Safe Reinforcement Learning for Sequential Recommender Systems

Research Internship project mentored by Alex Beutel (Google Brain).

Formulated and developed a sequential recommendation framework that considers the long-term well-being of users and proposed a novel policy gradient algorithm based on Safe Reinforcement Learning (Safe RL) that provides risk guarantees for the worst-case users. Presented the work at FAccTRec Workshop at ACM RecSys 2020.

May–August 2019 **Research Intern**, *Microsoft Research*, Montreal, QC, Canada.

Feedback Loops and Producer-side Fairness in Recommender Systems

Research Internship project working with Fernando Diaz (FATE Group).

Studied the intertwined phenomenon of *selection bias* and *exposure unfairness* for producers in a recommender system feedback loop. Theoretically characterized the conditions under which such effects amplify, and developed algorithms to mitigate unfairness.

May–August 2017 **Research Intern**, *Facebook*, Menlo Park, CA.

Active Learning for Multilabel Classification on Newsfeed

Research internship project working with Khalid El-Arini (Facebook Newsfeed).

Developed an active learning approach to optimize the trade-off between human labeling cost and model accuracy for a large-scale multilabel classification problem for Facebook Newsfeed, improving labeling efficiency by up to 30%.

May–August **Research Intern**, *Microsoft Research Lab*, New York City, NY.

2016 **Contextual Bandits for Personalization of Notifications in Microsoft Health App**

*Research internship project working with **John Langford** (MSR NYC) and **Ryen White** (Microsoft Health and MSR Redmond).*

Developed a *Contextual Bandits*-based approach to personalize reminder and notification messages on the Microsoft Health App to optimize users' long-term health and fitness.

Skills and Technologies

- Programming Languages: Python, C/C++, Java
- ML Frameworks: TensorFlow, PyTorch, Keras, Scikit-Learn
- Other Tools: Pandas, Jupyter, Colab, TensorBoard, Ray, MLflow

Publications

Conference Publications

Pedro Silva, Bhawna Juneja, Shloka Desai, Ashudeep Singh, Nadia Fawaz. **“Representation Online Matters: Practical End-to-End Diversification in Search and Recommender Systems”**. In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT), 2023. [↗](#)

Sahil Verma, Ashudeep Singh, Varich Boonsanong, John P. Dickerson, Chirag Shah. **“RecRec: Algorithmic Recourse for Recommender Systems”**. In Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, 2023. [↗](#)

Ashudeep Singh, David Kempe, Thorsten Joachims. **“Fairness in Ranking under Uncertainty”**. In Proceedings of Advances in Neural Information Processing Systems (NeurIPS) 2021. [↗](#)

Marco Morik*, Ashudeep Singh*, Jessica Hong, Thorsten Joachims. **“Controlling Fairness and Bias in Dynamic Learning-to-Rank”**. In Proceedings of 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval 2020. (* equal contribution) [↗](#) **[Best Paper Award]**

Ashudeep Singh and Thorsten Joachims. **“Policy Learning for Fairness in Ranking”**. In Proceedings of Advances in Neural Information Processing Systems (NeurIPS) 2019, Vancouver, BC, Canada. [↗](#)

Ashudeep Singh and Thorsten Joachims. **“Fairness of Exposure in Rankings”**. In ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD) 2018, London, United Kingdom. [↗](#)

Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, Thorsten Joachims. **“Recommendations as Treatments: Debiasing Learning and Evaluation”** In Proceedings of The International Conference on Machine Learning (ICML) 2016, New York, NY, USA. [↗](#)

David Adamson, Akash Bharadwaj, Ashudeep Singh, Colin Ashe, David Yaron, Carolyn P. Rosé. **“Predicting Student Learning from Conversational Cues”**. In Proceedings of 12th International Conference of Intelligent Tutoring Systems (ITS) 2014, Honolulu, HI, USA. [↗](#)

David Adamson, Divyanshu Bhartiya, Biman Gujral, Radhika Kedia, Ashudeep Singh, Carolyn P. Rosé. **“Automatically Generating Discussion Questions”**. In Proceedings of 16th International Conference of Artificial Intelligence in Education (AIED) 2013, Memphis, TN, USA. [↗](#)

Workshop Papers

Ashudeep Singh, Yoni Halpern, Nithum Thain, Konstantina Christakopoulou, Ed H. Chi, Jilin Chen, Alex Beutel. **“Building Healthy Recommendation Sequences for Everyone: A Safe Reinforcement Learning Approach”**. At FAccTRec Workshop at ACM RecSys, 2020. [↗](#)

Ashudeep Singh, Thorsten Joachims. “**Equality of Opportunity in Rankings**”. At Workshop on Prioritising Online Content at NeurIPS 2017. [↗](#)

Ashudeep Singh, Thorsten Joachims. “**Learning item embeddings using biased feedback**”. At Causal Inference and Machine Learning for Intelligent Decision Making Workshop at NeurIPS 2017. [↗](#)

Awards and Achievements

- 2020 Awarded the **Best Paper Award** at ACM SIGIR 2020.
- 2019 Outstanding Teaching Assistant Award by the Department of Computer Science for CS6780: Advanced Machine Learning class.
- 2019 Awarded the NeurIPS Travel Award to attend NeurIPS 2019, Vancouver, BC, Canada.
- 2018 Awarded the ACM Student Travel Award to attend SIGKDD 2018, London, UK.
- 2015 **Ranked first** in the M.Tech. class of 108 students graduating in 2015 at IIT Kanpur.
- 2011–2015 Awarded **Academic Excellence Award** for outstanding academic achievements at IIT Kanpur for each academic year.
- 2010–2014 Awarded **CBSE Merit Scholarship for Professional Studies** by Central Board of Secondary Education, India.
- 2012 Recipient of **Summer Undergraduate Research Grant for Excellence (SURGE)**, granted by Dean Resource Planning and Generation, IIT Kanpur.

Professional Service

- **Area Chair/Meta-Reviewer** for ICML 2022, NeurIPS 2023.
- **Senior Program Committee** member for ACM EAAMO 2022.
- **Program Committee** (PC member)
 - ACM FAccT Conference 2021-2023
 - FAccTRec workshop at ACM RecSys 2020
 - FACTS-IR Workshop at SIGIR 2019
 - Repl4NLP Workshop at ACL 2018
- **Reviewer**
 - NeurIPS 2019-2022
 - ICML 2019-2023
 - ICLR 2021
 - Recsys 2021
 - AAAI 2020
- **Ethics Reviewer** for NeurIPS 2022-2023 and **Datasets & Benchmarks Reviewer** for NeurIPS 2022.
- **Invited Talks**
 - “Fairness of Exposure in Ranking” at:
 - Mechanism Design for Social Good (MD4SG) (March 2019)
 - Amazon (Feb 2021)
 - Google (March 2021)
 - Meta (March 2021)
 - Spotify (March 2021)
 - “Building Healthy Recommendation Sequences for Everyone: A Safe Reinforcement Learning Approach” at Virtual Workshop on Responsible Recommender Systems at Meta (November 2022)
 - “Responsible ML for Real-World Search and Recommender Systems: A Multistakeholder Perspective” Guest Lecture at USC Marshall School of Business for the Operations Management class (BUAD 311) (October 2023)

Teaching and Mentorship

- **Tutorial** at Neural Information Processing Systems (NeurIPS) 2022 Conference (2022)
 - Title: Fair and Socially Responsible ML for Recommendations [↗](#)
 - Co-taught with Manish Raghavan (MIT) and Hannah Korevaar (Meta).
- **Teaching Assistant** for

- Advanced Machine Learning (CS6780, Cornell) (Spring 2019)
 - Awarded *Outstanding TA Award* by the Department of Computer Science.
 - Machine Learning for Data Science (CS4786, Cornell) (Spring 2016)
 - Machine Learning for Intelligent Systems (CS4780/5780, Cornell) (Fall 2015)
 - Machine Learning for Vision (CS679, IIT Kanpur) (Spring 2015)
 - Fundamentals of Computing (Graduate Student Instructor, ESC101, IIT Kanpur) (Fall 2014)
- Mentored and collaborated with several undergraduate and masters students at Cornell University. (2017-2021)
 - Zili Zhou (Cornell)
 - Ibiyemi Abiodun (Cornell)
 - Hansol (Hannah) Lee (Cornell, now Stanford)
 - Andrei Kozyrev (Cornell, now Spotify)
 - Marco Morik (TU Berlin)
 - Jessica Hong (Cornell, now Google)
 - Vianca Hurtado (Cornell)
 - Kenta Takatsu (Cornell, now UMass Amherst)
 - Xiang (Felix) Fu (Cornell, now MIT)
 - Junyoung (Jared) Lim (Cornell, now Google)
 - Pian Pawakapan (Cornell)

Positions of Responsibility and Extra Curricular Activities

- Served on the Ph.D. Admissions Committee for the Department of Computer Science, Cornell University. (2019)
- Co-developed **ViCoRecS: Virtual Conference Recommender System** to provide attendees at KDD 2020 with relevant Networking and Paper recommendations, which was used by ~1000 users. [↗](#) (2020)
 - Published the outcome of a research experiment conducted as a part of deploying this system at NeurIPS 2021.
- Organized the **Machine Learning Discussion Group** at Cornell University. [↗](#) (2016-18)
- Student Guide, Academic Mentor, and Link Student for Counselling Service, IIT Kanpur. (2011–13)

References

- Thorsten Joachims (Ph.D. Advisor)
 Professor, Department of Computer Science
 Cornell University, Ithaca, NY
tj@cs.cornell.edu
- Fernando Diaz
 Research Scientist, Google Research, Montreal, QC, Canada
 Associate Professor, Carnegie Mellon University, Pittsburgh, PA (starting Fall 2023)
diazf@acm.org
- Alex Beutel
 Member of Technical Staff, Research Scientist
 Open AI, New York, NY
alexbe@openai.com