

Arpit Agarwal

CONTACT INFORMATION	Data Science Institute Columbia University E-mail: arpit.agarwal@columbia.edu WWW: http://www.columbia.edu/aa4931/	
RESEARCH INTERESTS	My research is in online machine learning problems that surface in applications such as recommendation systems, information retrieval and online marketing. My recent interest is understanding feedback loops between recommendation algorithms and user preferences, and their societal implications.	
WORK EXPERIENCE	Postdoctoral Research Fellow Data Science Institute, Columbia University Mentors: Prof. Tim Roughgarden and Prof. Yashodhan Kanoria	2021 – present New York, USA
EDUCATION	Ph.D. in Computer & Information Science University of Pennsylvania Advisor: Prof. Shivani Agarwal	2016 – 2021 Philadelphia, USA
	M.E. in Computer Science & Engineering (Awarded <i>CSI medal for Best Student in the outgoing batch</i>) Indian Institute of Science	2012 – 2014 Bangalore, India
	B.Tech. (Hons.) in Computer Science & Engineering Kamla Nehru Institute of Technology	2008 – 2012 Sultanpur, India
RESEARCH INTERNSHIPS & VISITS	<ul style="list-style-type: none">• Google Research, Mountain View Mentor: Ashwinkumar Badanidiyuru Varadaraja• Harvard University, Cambridge Mentor: Prof. David Parkes	(June, 2019 – Sept, 2019) (Sept, 2015 – Dec, 2015)
PUBLICATIONS	<ul style="list-style-type: none">• Agarwal, A., Brown, W. <i>Diversified Recommendations for Agents with Adaptive Preferences.</i> To Appear in Neural Information Processing Systems (NeurIPS), 2022.• Agarwal, A., Khanna, S., Li, H., Patil, P. <i>Sublinear Algorithms for Hierarchical Clustering.</i> To Appear in Neural Information Processing Systems (NeurIPS), 2022.• Agarwal, A., Ghuge, R., Nagarajan, V. <i>An Asymptotically Optimal Batched Algorithm for the Dueling Bandit Problem.</i> To Appear in Neural Information Processing Systems (NeurIPS), 2022.• Agarwal, A., Khanna, S., Patil, P. <i>A Sharp Memory-Regret Trade-off for Multi-Pass Streaming Bandits.</i> Conference on Learning Theory (COLT), 2022.	

- Agarwal, A., Ghuge, R., Nagarajan, V.
Batched Dueling Bandits.
International Conference on Machine Learning (**ICML**), 2022.
Long presentation (top 2% of all submissions).
- Agarwal, A., Khanna, S., Patil, P.
PAC Top-k Identification under SST in Limited Rounds.
In Artificial Intelligence and Statistics (**AISTATS**), 2022.
- Agarwal, A., Patil, P., and Agarwal, S.,
Stochastic Dueling Bandits with Adversarial Corruption.
In Algorithmic Learning Theory (**ALT**), 2021.
- Agarwal, A., Mandal, D., Parkes, D., and Shah, N.,
Peer Prediction with Heterogeneous Users.
In ACM Transactions on Economics and Computation (**TEAC**), 2020.
Full version of the EC-17 paper below.
- Agarwal, A., Johnson, N., and Agarwal, S.,
Choice Bandits.
In Neural Information Processing Systems (**NeurIPS**), 2020.
- Agarwal, A., Agarwal, S., Khanna, S., and Patil, P.,
Rank Aggregation from Pairwise Comparisons in the Presence of Adversarial Corruption.
In International Conference on Machine Learning (**ICML**), 2020.
- Agarwal, A., Assadi, S., and Khanna, S.,
Stochastic Submodular Covering with Limited Adaptivity.
In ACM-SIAM Symposium on Discrete Algorithms (**SODA**), 2019
- Agarwal, A., Patil, P., and Agarwal, S.,
Accelerated Spectral Ranking.
In International Conference on Machine Learning (**ICML**), 2018.
- Agarwal, A., Agarwal, S., Assadi, S., and Khanna, S.,
Learning with Limited Rounds of Adaptivity: Coin Tossing, Multi-Armed Bandits, and Ranking from Pairwise Comparisons.
In Conference on Learning Theory (**COLT**), 2017.
- Agarwal, A., Mandal, D., Parkes, D., and Shah, N.,
Peer Prediction with Heterogeneous Users.
In ACM Conference on Economics and Computation (**EC**), 2017.
Invited to **TEAC special issue** for EC 2017 papers.
- Shnayder, V., Agarwal, A., Frongillo, R. and Parkes D.C.,
Informed Truthfulness in Multi-Task Peer Prediction.
In ACM Conference on Economics and Computation (**EC**), 2016. A shorter version appeared in **HCOMP** Workshop on Mathematical Foundations of Human Computation, 2016.
- Agarwal, A. and Agarwal, S.,
On Consistent Surrogate Risk Minimization and Property Elicitation.
In Conference on Learning Theory (**COLT**), 2015.
- Agarwal, A., Narasimhan, H., Kalyanakrishnan, S. and Agarwal, S.,
GEV-Canonical Regression for Accurate Binary Class Probability Estimation when One

Class is Rare.

In International Conference on Machine Learning (**ICML**), 2014.

TALKS/
PRESENTATIONS

- A Sharp Memory-Regret Trade-off for Multi-Pass Streaming Bandits
– Conference on Learning Theory (COLT) 2022
- Choice Bandits
– Neural Information Processing Systems (NeurIPS) 2020
- Stochastic Submodular Covering with Limited Adaptivity
– Algorithms seminar, Google Research, Mountain View 2019
– ACM-SIAM Symposium on Discrete Algorithms (SODA), San Diego 2019
- Accelerated Spectral Ranking
– NYC research seminar, Google Research, New York 2018
– International Conference on Machine Learning (ICML), Stockholm 2018
– Summer UG research program, University of Pennsylvania, Philadelphia 2018
- Learning with Limited Rounds of Adaptivity: Coin Tossing, Multi-Armed Bandits, and Ranking from Pairwise Comparisons
– CSA Department Seminar, Indian Institute of Science, Bangalore 2017
– ML seminar, Microsoft Research, Bangalore 2017
– Conference on Learning Theory (COLT), Amsterdam, Netherlands 2017
- On Consistent Surrogate Risk Minimization and Property Elicitation
– ACM IKDD, Pune, India 2016
– EconCS seminar, Harvard University, Cambridge 2015
- GEV-Canonical Regression for Accurate Binary Class Probability Estimation when One Class is Rare
– International Conference on Machine Learning (ICML), Beijing 2014
- Randomization at work: An Introduction to Randomized Algorithms
– CSA Undergraduate Summer School, Indian Institute of Science, Bangalore 2013

TEACHING
EXPERIENCE

- Teaching Assistant, Advanced Topics in Machine Learning (CIS 620) University of Pennsylvania Fall 2018
- Teaching Assistant, Machine Learning (CIS 520) University of Pennsylvania Fall 2017
- Teaching Assistant, Machine Learning (E0 270) Indian Institute of Science Spring 2016

ADVISING &
SERVICE

- **Mentoring undergraduate students**
 - Arnab Sarker, University of Pennsylvania, “A unified framework for quantile elicitation with applications”, 2018-2019.
 - Jane Lee, University of Pennsylvania, “Multiclass classification under asymmetric label noise”, 2018-2019.
- **Organizational activities**

- Volunteer for Penn Research in Machine Learning (PRiML) forum, University of Pennsylvania, 2018-2019.
- Member of Departmental Curriculum Committee, Computer Science & Automation, Indian Institute of Science, 2015-2016.
- Lead volunteer for Big Data Initiative, Indian Institute of Science, 2014-2015.
- **Conference reviewing activities**– ICML 2021, NeurIPS 2020, COLT 2020, STOC 2020, SODA 2020, ICML 2019, AISTATS 2019, N(eur)IPS 2018, COLT 2018, 2017.
- **Journal reviewing activities**– JAIR, JMLR.