Abhin Shah

abhin-shah.github.io abhin@mit.edu

Education

Massachusetts Institute of Technology (MIT) *2018—2024*
Ph. D., in Electrical Engineering and Computer Science *CGPA: 5.00/5.00*
Advisors: *Prof. Devavrat Shah, Prof. Gregory W. Wornell, & Prof. Alberto Abadie*

Indian Institute of Technology, Bombay (IIT Bombay) *2014—2018*
Bachelor of Technology in Electrical Engineering with Honors *CGPA: 9.61/10.00*
Minor in Computer Science

Research Interests

My research develops theories and algorithms for causal inference with the goal of personalized, data-driven decision-making. Specifically, I build methods to answer individual-level rather than population-level as well as distribution-related rather than mean-related causal questions. I use tools and frameworks from statistics and machine learning, such as exponential family and matrix completion. I also develop fair algorithms for supervised machine learning. In the long term, I am enthusiastic about developing causal inference methods that are robust and fair.

Selected Awards And Achievements

Best Presentation, Laboratory of Information & Decision Systems (LIDS) Student Conference, MIT *2024*
Presidential Graduate Fellowship, MIT *2018—19*
Undergraduate Research Award for exemplary contribution towards research, IIT Bombay *2017*
Institute Academic Prize for excellent academic performance in the junior year, IIT Bombay *2017*
Best Project for Social Cause at Technical, Research, and Development Exposition, IIT Bombay *2015*
All India Rank 126 out of 1.4 million, IIT Joint Entrance Exam Advanced *2014*
Bronze medal, 32nd Annual Mathematics Olympiad, IIT Bombay *2013*
National Talent Search Scholarship, Government of India *2010—18*

News Coverage

Research on "Fair selective regression" featured in MIT News article, "A technique to improve both fairness and accuracy in artificial intelligence" *July 2022*

Publications

($\star$ denotes equal contribution; title is hyperlinked to the online pdf of the paper)

1. **Abhin Shah**, Karthikeyan Shanmugam, Murat Kocaoglu, "Front-door adjustment beyond Markov equivalence with limited graph knowledge", *Conference on Neural Information Processing Systems (NeurIPS)* *2023*
2. **Abhin Shah***, Yuheng Bu*, Joshua Ka-Wing Lee, Subhro Das, Rameswar Panda, Prasanna Sattigeri, Gregory W. Wornell, "Selective regression under fairness criteria", *International Conference on Machine Learning (ICML)*, link *2022*
3. **Abhin Shah**, Wei-Ning Chen, Johannes Ballé, Peter Kairouz, Lucas Theis, "Optimal compression of locally differentially private mechanisms", *International Conference on Artificial Intelligence and Statistics (AISTATS)*, link *2022*
4. **Abhin Shah**, Karthikeyan Shanmugam, Kartik Ahuja, "Finding valid adjustments under non-ignorability with minimal DAG knowledge", *International Conference on Artificial Intelligence and Statistics (AISTATS)*, link *2022*
5. **Abhin Shah**, Devavrat Shah, Gregory W. Wornell, "A computationally efficient method for learning exponential family distributions", *Conference on Neural Information Processing Systems (NeurIPS)* *2021*
6. **Abhin Shah**, Kartik Ahuja, Karthikeyan Shanmugam, Dennis Wei, Kush Varshney, Amit Dhurandhar, "Treatment effect estimation using invariant risk minimization", *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, link *2021*

7. **Abhin Shah**, Devavrat Shah, Gregory W. Wornell, "On learning continuous Markov random fields", *International Conference on Artificial Intelligence and Statistics (AISTATS)*, Oral Presentation *2021*
8. **Abhin Shah**, Nikhil Karamchandani, Suhas Diggavi, "Coded caching: Global vs local content popularity", *Canadian Workshop on Information Theory (CWIT)* *2019*
9. **Abhin Shah**, Sai Vinjanampathy, Bhaskaran Muralidharan, "Classical information driven quantum dot thermal machines", *Annals of Physics* *2018*
10. Su Li, **Abhin Shah**, Michael Gastpar, "Cooperative data exchange with weighted cost based on basis construction", *Annual Allerton Conference on Communication, Control, and Computing (Allerton)* *2017*

Pre-Prints

(† denotes alphabetical ordering; title is hyperlinked to the online pdf of the paper)

1. Alberto Abadie†, Anish Agarwal, Raaz Dwivedi, **Abhin Shah**, "Doubly Robust Inference in Causal Latent Factor Models", *under review at ACM Conference on Economics and Computation (EC), full version in preparation for submission to Econometrica*
2. **Abhin Shah**, Raaz Dwivedi, Devavrat Shah, Gregory W. Wornell "On counterfactual inference with unobserved confounding", *NeurIPS 2022 Workshop on Causality for Real-world Impact (CML4Impact), full version under review at Operations Research*
3. **Abhin Shah**, Devavrat Shah, Gregory W. Wornell, "On computationally efficient learning of exponential family distributions", *under review at IEEE Transactions on Information Theory*
4. **Abhin Shah**, Maohao Shen, Jongha Jon Ryu, Subhro Das, Prasanna Sattigeri, Yuheng Bu, Gregory W. Wornell, "Group fairness with uncertainty in sensitive attributes", *ICML 2023 Workshop on Spurious Correlations, Invariance, and Stability, full version under review at IEEE International Symposium on Information Theory (ISIT)*, US Patent Application, 18/503166, November 2023

Working papers

1. "Learning counterfactual distribution under unobserved confounding with exponential family" with Raaz Dwivedi, Devavrat Shah, Gregory W. Wornell
2. "A unified view on learning exponential family distributions via noise-contrastive estimation" with Jongha Jon Ryu, Devavrat Shah, Gregory W. Wornell
3. "Treatment effect estimation beyond sequential ignorability: An application to clinical healthcare" with Shalmali Joshi, Murat Kocaoglu, Karthikeyan Shanmugam

Research Talks

Doubly robust inference for causal latent factor model

- Laboratory of Information & Decision Systems (LIDS) Student Conference, MIT *Jan 2024*
- Econometrics Lunch, MIT *Dec 2023*

On counterfactual inference with unobserved confounding

- Morgan Stanley ML Seminar Series, Virtual *Nov 2023*
- INFORMS Annual Meeting, Phoenix, Arizona *Oct 2023*
- Machine Learning Tea (MLTea) Seminar Series, MIT *Oct 2023*
- Allerton Conference, Monticello, Illinois *Sept 2023*
- DeepMind/UCL Computational Statistics and Machine Learning Seminar, Virtual *June 2023*
- ICSA Applied Statistics Symposium, Ann Arbor, Michigan *June 2023*
- Graphical Models, Statistical Inference, and Algorithms workshop, Harvard University *May 2023*
- CSAIL–LIDS Machine Learning Advances Symposium, MIT *May 2023*
- Machine Learning Foundations and Optimization Group, Google Research, Virtual *April 2023*
- Machine Learning Seminar, IBM Research, Virtual *March 2023*
- Machine Learning Department, Carnegie Mellon University (CMU), Virtual *March 2023*

- Econometrics Lunch, MIT *March 2023*
- Brain Causality Group, Google Research, Virtual *Feb 2023*
- LIDS and Stata tea talk, MIT *Nov 2022*

Group fairness with uncertainty in sensitive attributes

- CMU Artificial Intelligence Seminar Series, Virtual *Oct 2023*
- Simons Collaboration TOC4Fairness Seminar, Virtual *Oct 2023*
- Machine Learning Tea (MLTea) Seminar Series, MIT *May 2023*
- LIDS and Stata tea talk, MIT *March 2023*

Selective regression under fairness criteria

- International Conference on Machine Learning (ICML), Baltimore *July 2022*
- Jane Street Symposium, Virtual *Jan 2022*

Finding valid adjustments under non-ignorability with minimal DAG knowledge

- International Conference on Artificial Intelligence and Statistics (AISTATS), Virtual *March 2022*
- Brain Causality Group, Google Research, Virtual *July 2021*
- Causality and Machine Learning group, Microsoft Research, Virtual *July 2021*

Optimal compression of locally differentially private mechanisms

- International Conference on Artificial Intelligence and Statistics (AISTATS), Virtual *March 2022*
- Federated Research Group, Google Research, Virtual *Aug 2021*
- Neural Compression Group, Google Research, Virtual *Aug 2021*

Treatment effect estimation using invariant risk minimization

- International Conference on Acoustics, Speech and Signal Processing (ICASSP), Virtual *June 2021*
- IBM Research, Virtual *Aug 2020*

A computationally efficient method for learning exponential family distributions

- Conference on Neural Information Processing Systems (NeurIPS), Virtual *Dec 2021*

On learning continuous Markov random fields

- International Conference on Artificial Intelligence and Statistics (AISTATS), Virtual *March 2021*

Selected Poster Presentations

On counterfactual inference with unobserved confounding

- INFORMS Annual Meeting, Phoenix, Arizona *Oct 2023*
- Young Researchers Workshop, Cornell University, Ithaca *Oct 2023*
- INFORMS Applied Probability Society (APS) Conference, Nancy, France *June 2023*
- ACM SIGMETRICS Workshop on Causal Inference for Engineers, Orlando, Florida *June 2023*
- American Causal Inference Conference, Austin, Texas *May 2023*
- NeurIPS Workshop on Causality for Real-world Impact, New Orleans, Louisiana *Dec 2022*

Front-door adjustment beyond Markov equivalence with limited graph knowledge

- ICML Workshop on Spurious Correlations, Invariance, and Stability *July 2023*

Group fairness with uncertainty in sensitive attributes

- ICML Workshop on Spurious Correlations, Invariance, and Stability *July 2023*
- Social and Ethical Responsibilities of Computing Symposium, MIT *April 2023*

Finding valid adjustments under non-ignorability with minimal DAG knowledge

- American Causal Inference Conference, Austin, Texas *May 2023*

Optimal compression of locally differentially private mechanisms

- Google's Workshop on Federated Learning and Analytics, Virtual *Nov 2021*

On learning continuous Markov random fields

- MIFODS Workshop on Graphical Models, Exchangeable Models and Graphons, MIT *Aug 2021*

Teaching

- Kaufman Teaching Certificate Program, MIT *2023*
- Teaching Assistant, Algorithms for Inference (6.438), MIT *2020*
- Teaching Assistant, Linear Algebra (MA 106), IIT Bombay *2016*
- Teaching Assistant, Quantum Physics and Application (PH 107), IIT Bombay *2015*

Academic Services

Scientific Meetings

- Chair, Social Aspects: Accountability, Transparency and Interpretability Session, International Conference on Machine Learning *2022*

Committees

- Student Search Advisory Group, MIT EECS Faculty Search *2023*

Mentoring Activities

- MIT EECS Graduate Application Assistance Program (GAAP) for *underrepresented applicants* *2022—23*
- MIT EECS for *MEng* students *2022—23*
- IIT Bombay Student Mentoring Program (ISMP) for *incoming undergraduates* *2017—18*
- IIT Bombay Academic Mentoring Program (DAMP) for *sophomores & juniors* *2016—18*

Reviewing Activities

- **Top-10% reviewer** — ICML 2022, AISTATS 2022
- NeuIPS (2023, 2022, 2021), ICML (2023, 2022), AISTATS (2022, 2021), ACIC 2023, ISIT 2023, JSAIT 2020

Internship Experience

Google Research with Johannes Ballé, Lucas Theis, and Peter Kairouz *Summer 2021*

IBM Research with Kush Varshney, Kartik Ahuja, Karthikeyan Shanmugam, Dennis Wei, and Amit Dhurandhar *Summer 2020*

EPFL, Switzerland with Prof. Michael Gastpar *Summer 2017*

References

Devavrat Shah
Professor of EECS, MIT
Ph. D., Advisor
devavrat@mit.edu

Gregory W. Wornell
Professor of EECS, MIT
Ph. D., Advisor
gww@mit.edu

Alberto Abadie
Professor of Economics, MIT
Ph. D., Advisor
abadie@mit.edu