

Dr. Ayesha Rahman

✉ a.rahman@cs.cambridge.ac.uk | 📞 +44 7700 900123 | 🌐 ayesharahman.cs.cam.ac.uk | 
0000-0002-1234-5678 | 🎓 Google Scholar

Research Interests

Machine learning for scientific discovery; probabilistic graphical models; natural language processing; fairness, accountability and transparency in AI; computational social science; privacy-preserving federated learning.

Education

Ph.D. in Computer Science

University of Cambridge, Cambridge, UK

Advisor: Prof. Samuel J. Fletcher

Thesis: “Scalable Bayesian Inference for Large-Scale Structured Prediction Problems” 2017–2022

M.Sc. in Advanced Computer Science (*with Distinction*)

University of Edinburgh, Edinburgh, UK

Thesis: “Efficient Approximate Inference in Deep Latent Variable Models” 2015–2016

B.Sc. (Hons) in Computer Science and Mathematics

University of Dhaka, Dhaka, Bangladesh

First Class Honours; University Gold Medal 2011–2015

Publications

Journal Articles

- [1] **Rahman, A.**, Fletcher, S. J., & Nguyen, T. L. (2024). Variational Riemannian Flows for Posterior Approximation in High-Dimensional Graphical Models. *Journal of Machine Learning Research*, 25(1), 1–48. doi:10.5555/example.2024.001
- [2] **Rahman, A.**, & Osei-Bonsu, K. (2023). FairFed: A Federated Learning Framework for Demographic Parity under Distribution Shift. *IEEE Transactions on Neural Networks and Learning Systems*, 34(8), 4102–4115. doi:10.1109/TNNLS.2023.0012345
- [3] Park, J., **Rahman, A.**, & Whitfield, C. (2022). Scalable Structured Prediction via Conjugate Energy Descent. *Artificial Intelligence*, 312, 103780. doi:10.1016/j.artint.2022.103780
- [4] **Rahman, A.**, Blackwell, R., & Singh, P. (2021). Causal Discovery in Non-Stationary Time Series with Latent Confounders. *Transactions on Machine Learning Research*. openreview.net/example2021

Conference Papers

- [5] **Rahman, A.**, Leclerc, M., & Fletcher, S. J. (2024). Amortized Causal Discovery with Score-Based Diffusion Priors. In *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, Vienna, Austria. PMLR 235, pp. 18 432–18 451.
- [6] **Rahman, A.**, & Huang, Y. (2023). Privacy-Aware Representation Learning via Information-Theoretic Regularisation. In *Advances in Neural Information Processing Systems (NeurIPS)*

2023), Vol. 36, pp. 57 120–57 135.

- [7] Osei-Bonsu, K., **Rahman, A.**, & Patel, N. (2022). Robust Federated Aggregation under Byzantine and Sybil Attacks. In *Proceedings of the 39th International Conference on Machine Learning (ICML 2022)*, Baltimore, MD. PMLR 162, pp. 16 890–16 905.
- [8] **Rahman, A.**, Whitfield, C., & Park, J. (2021). Spectral Normalisation for Structured Prediction with Neural CRFs. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL 2021)*, pp. 3 441–3 454.

Research Experience

Assistant Professor (Lecturer), Department of Computer Science

University of Cambridge, Cambridge, UK

Principal Investigator of the *Statistical AI Lab*

Jan 2023–present

Postdoctoral Research Associate

Max Planck Institute for Intelligent Systems, Tübingen, Germany

Supervisor: Prof. Bernhard Schölkopf

Sep 2022–Dec 2022

Research Intern, Google DeepMind, London, UK

Project: Efficient large-scale amortised inference with neural ODEs

Jun 2020–Sep 2020

Graduate Research Assistant

Machine Intelligence Laboratory, University of Cambridge

Oct 2017–Aug 2022

Teaching Experience

Lecturer, *Machine Learning* (Part II CST)

University of Cambridge — ~180 students

2023–present

Supervisor, *Probabilistic Machine Learning* (Part II CST)

University of Cambridge

2018–2022

Teaching Assistant, *Natural Language Processing*

University of Cambridge — Designed coursework and held weekly Q&A sessions

2019–2021

Instructor, *Introduction to Data Science* (summer school)

African Institute for Mathematical Sciences (AIMS), Accra, Ghana

Jul 2019

Grants, Awards & Honors

UKRI Future Leaders Fellowship (FLF-MR/V021967/1) — £1.2M

“Causally-Aware Machine Learning for Robust Scientific Discovery”

2023–2027

Royal Society Research Grant (RG\R1\231042) — £19 800

2023

Best Paper Award, ICML 2024 (top 0.4% of submissions)

2024

Cambridge International Scholarship, Cambridge Trust

2017–2021

University Gold Medal for highest aggregate, University of Dhaka

2015

Google PhD Fellowship in Machine Learning

2019–2021

Invited Talks

“Towards Causally-Grounded Representation Learning,” <i>invited seminar</i> Oxford University Machine Learning Research Group, Oxford, UK	Mar 2025
“Scalable Bayesian Inference for Structured Prediction,” <i>tutorial</i> European Conference on Machine Learning (ECML-PKDD), Turin, Italy	Sep 2024
“FairFed: Federated Learning with Fairness Guarantees,” <i>keynote</i> Workshop on Socially Responsible Machine Learning, NeurIPS 2023	Dec 2023
“Causal Discovery in Non-Stationary Systems,” <i>invited talk</i> Alan Turing Institute Data Science Seminar, London, UK	Jun 2022

Professional Service

Area Chair / Senior Program Committee

- NeurIPS 2025; ICML 2024, 2025; ICLR 2024, 2025

Program Committee / Reviewer

- AAAI, ACL, EMNLP, UAI (2020–present)
- *Journal of Machine Learning Research*; *Artificial Intelligence*; *IEEE TNNLS*; *Nature Machine Intelligence*

Organizing Roles

- Workshop Co-Organizer, “Causal Representation Learning,” ICML 2024
- Doctoral Consortium Co-Chair, ECML-PKDD 2023
- Co-Founder, Cambridge Women in AI Network (CWAIN), 2019–present

Technical Skills

Programming	Python, Julia, R, C++, Bash
ML Frameworks	PyTorch, JAX, TensorFlow, Hugging Face Transformers
Probabilistic	Pyro, NumPyro, Stan, GPyTorch
Tools	Git, Docker, Slurm/HPC, $\text{\LaTeX}$, Jupyter
Languages	English (fluent), Bengali (native), French (intermediate)

References

Available upon request.