

RESEARCH SUMMARY

I explore how Bayesian and causal principles can serve as inductive biases in scalable models, enabling reliable, uncertainty-aware predictions under distribution shifts. My work spans theory, novel algorithm design, and applied benchmarks, achieving state-of-the-art results across diverse domains. Ultimately, I aim to accelerate scientific discovery and build safe, robust intelligent agents.

SELECTED PUBLICATIONS

- Anish Dhir, Cristiana Diaconu, Valentinian Mihai Lungu, James Requeima, Richard E. Turner, and Mark van der Wilk. “Estimating Interventional Distributions with Uncertain Causal Graphs through Meta-Learning.”, ICML 2025 Workshop on Scaling Up Intervention Models, 2025.
- Anish Dhir, Ruby Sedgwick, Avinash Kori, Ben Glocker, and Mark van der Wilk. “Continuous Bayesian Model Selection for Multivariate Causal Discovery.”, In Forty-first International Conference on Machine Learning (ICML), 2025.
- Anish Dhir, Matthew Ashman, James Requeima, and Mark van der Wilk. “A Meta-Learning Approach to Bayesian Causal Discovery.”, The Thirteenth International Conference on Learning Representations (ICLR), 2025.
- Anish Dhir, Samuel Power, and Mark van der Wilk. “Bivariate Causal Discovery using Bayesian Model Selection.” In Forty-first International Conference on Machine Learning (ICML), 2024.
- Anish Dhir, and Mark van der Wilk. “Causal Discovery using Marginal Likelihood.” In NeurIPS 2022 Workshop on Causality for Real-world Impact, 2022.

EXPERIENCE

Amazon

Tuebingen

*Applied Science Intern - AGI Foundations**May. 2024 - Nov. 2024*

- Research on scaling neural networks (LLMs) while preserving feature learning behavior as well as stability (using tensor programs).
- Analysed router collapse with increasing width in mixture of expert models.
- Implemented fixes on models such as Mixtral 8x7B.

Babylon Health

London

*Research Scientist**Dec. 2018 - Oct. 2021*

- Research and applications focused on Causal Discovery (finding causal relations in data as opposed to correlations) and Simulation (using causal relations to compute causal effects and counterfactuals)
- Research in causal discovery with multiple datasets lead to publication at AAAI 2020. Multiple studies with overlapping variables are common in medicine. Presented a poster and a gave a talk at AAAI 2020. Work was featured in MIT technology review: [Link](#)
- Used research outcomes to analyse effect of interventions on real user interaction data.

EDUCATION

Imperial College London

London

*PhD in Machine Learning**Oct. 2021 - Present*

- Supervised by Dr. Mark van der Wilk
- Visiting at the University of Oxford since October 2023.
- First author publications in top conferences such as ICML, ICLR, Neurips.

University College London

London

*Computational Statistics and Machine Learning MSc; Distinction (Dean's List)**Sep. 2017 – Sep. 2018*

- **Thesis:** Created and tested a novel algorithm for conditionally executed learning in modular networks to ensure the most efficient parameter usage. Advised by David Barber.

Imperial College London*Physics MSc; Distinction*

London

*Sep. 2016 – Sep. 2017***Imperial College London***Physics with Theoretical Physics BSc; 2:1*

London

*Sep. 2013 – July 2016***AWARDS**

- Dean's List UCL 2017/18
- G-Research Early Career Researcher Grant.

PROGRAMMING SKILLS

Languages: Matlab, Python, limited experience in C++; Frameworks: TensorFlow, Jax, PyTorch, Numpy; Experienced with HPC, AWS, distributed training of large models.

OTHER EXPERIENCE

Teaching Assistant at Imperial College

London

*TA**Sep. 2021 – Current*

- Conducted seminars and tutorials for groups of students in courses ranging from undergraduate to masters level.

Private Tutor

London

*Tutor**Sep. 2013 – Current*

- Tutored students from secondary school to university level on subjects such as Maths, Physics, Computer Science