Adam Foster

a.evan.foster@gmail.com ae-foster.github.io
uk.linkedin.com/in/adamefoster scholar.google.com/citations?user=1MsXZJ0AAAAJ

RESEARCH STATEMENT

I am interested in devising machine learning and AI algorithms that will have an impact in natural science and scientific industry, as I believe this is where the important and interesting problems are. My research experience to date has focused on the following: 1) Bayesian optimal experimental design and active learning, particularly for designing science experiments, 2) deep learning methods for electronic structure, focusing on Quantum Monte Carlo wavefunction models, 3) generative models for proteins structural ensembles, 4) probabilistic programming as an emerging abstraction for probabilistic AI. My work has been published in top journals and machine learning conferences.

EXPERIENCE

2021 – PRESENT | MICROSOFT RESEARCH

Senior Researcher

I currently work on the Biomolecular Emulator team within Microsoft Research AI for Science, reporting to Frank Noé. We are building generative models for protein structure ensembles and recently published in Science. I focus on bringing new data modalities to bear on this project and play a major role in our collaboration with the Crick Institute. Previously, I was the lead researcher on a team that developed a Quantum Monte Carlo foundation model of wavefunctions, with a focus on multireferential quantum states that occur in bond-breaking: a phenomenon that is extremely difficult to accurately capture with existing computational chemistry tools.

2017 – 2021 | UNIVERSITY OF OXFORD, UNIVERSITY COLLEGE

DPhil Statistical Machine Learning

I completed my DPhil in the Oxford Computational Statistics and Machine Learning group, under the supervision of Yee Whye Teh and Tom Rainforth. A large part of my work in Oxford was on Bayesian optimal experimental design: how do we design experiments that will be most informative about the process being investigated? I worked on efficient estimators of expected information gain, a key objective function in experimental design, as well as sequential adaptive experimental design using policies. I also studied contrastive representation learning through the lenses of mutual information and invariance. I published numerous papers in top machine learning venues and delivered oral presentations at ICML 2021 and 2022. I continue my collaboration with Oxford by supervising a current PhD student.

SUMMER 2020 | BENEVOLENTAI

AI Science intern

I investigated deep representation learning methods for single-cell RNA sequence data in genomics. Inspired by several 'grand challenges' in computational biology, we developed an approach for data integration, and counterfactual inference to predict the effects of drugs and gene knock-outs. Our work received an ICML oral.

SUMMER 2018 | UBER AI LABS

Research intern

I interned with the Pyro team under the supervision of Noah Goodman. My contributions to the probabilistic programming language Pyro were part of a project to automate optimal experimental design for adaptive experimentation in science. We investigated estimators for mutual information in Bayesian experimental design, and implemented them in Pyro. Our work was published at NeurIPS 2019.

2016 – 2017 | ROAM ANALYTICS

Machine learning engineer

Through the Silicon Valley Internship Program, I spent a year in San Francisco working for a startup as a machine learning engineer. I helped build a knowledge graph of medical and pharmaceutical concepts and data.

2015 – 2016 | UNIVERSITY OF CAMBRIDGE, QUEENS' COLLEGE

MMath Mathematics

Grade Distinction | **Rank** 6th | **Awards** 2016 Wishart Prize, 2016 Foundation Scholarship

In Part III, "the oldest and most famous mathematics examination in the world", I chose to focus on statistics and probability. I was ranked 6th in a cohort of ∼250. I took courses in Advanced Probability, Stochastic Calculus, Modern Statistical Methods, Advanced Financial Models and Applied Stats.

2012 – 2015 | UNIVERSITY OF CAMBRIDGE, QUEENS' COLLEGE

BA Mathematics

Grade First | **Rank** 14th | **Awards** 2014, 2015 Foundation Scholarship; 2015 Colton Prize, 2013 Braithwaite Prize

During my undergraduate years at Cambridge I studied a broad range of pure and applied mathematics. My focus tended towards applied maths and theoretical physics. My Part II courses included Principles of Quantum Mechanics, Dynamical Systems, Classical Dynamics, Mathematical Biology and Applied Probability.

SUMMER 2015 | UCLA INSTITUTE FOR PURE AND APPLIED MATHEMATICS

Research in industrial projects for students

I led a team of four students on a project sponsored by the USC Shoah Foundation. We used Latent Dirichlet Allocation for topic modelling, and developed an extension of the model suited to our particular data modalities.

SELECTED PUBLICATIONS

Sarah Lewis, et al. Scalable emulation of protein equilibrium ensembles with generative deep learning. **Science, 2025.**

Adam Foster, Zeno Schätzle, P Bernát Szabó, Lixue Cheng, Jonas Köhler, Gino Cassella, Nicholas Gao, Jiawei Li, Frank Noé, Jan Hermann. An ab initio foundation model of wavefunctions that accurately describes chemical bond breaking. **Preprint, 2025.**

Lixue Cheng, P Bernát Szabó, Zeno Schätzle, Derk P Kooi, Jonas Köhler, Klaas JH Giesbertz, Frank Noé, Jan Hermann, Paola Gori-Giorgi, Adam Foster. Highly accurate real-space electron densities with neural networks. **JCP, 2025.**

Thomas Rainforth, Adam Foster, Desi R Ivanova, Freddie Bickford-Smith. Modern bayesian experimental design. **Statistical Science, 2023.**

Adam Foster, Árpi Vezér, Craig A Glastonbury, Páidí Creed, Sam Abujudeh, Aaron Sim. Contrastive Mixture of Posteriors for Counterfactual Inference, Data Integration and Fairness. **ICML 2022 (oral)** .

Emile Mathieu, Adam Foster, Yee Whye Teh. On Contrastive Representations of Stochastic Processes. **NeurIPS 2021.**

Adam Foster, Desi R Ivanova, Ilyas Malik, Tom Rainforth. Deep Adaptive Design: Amortizing Sequential Bayesian Experimental Design. **ICML 2021 (long presentation).**

Adam Foster, Martin Jankowiak, Matthew O'Meara, Yee Whye Teh, Tom Rainforth. A Unified Stochastic Gradient Approach to Designing Bayesian-Optimal Experiments. **AISTATS 2020.**

Adam Foster, Martin Jankowiak, Eli Bingham, Paul Horsfall, Yee Whye Teh, Tom Rainforth and Noah D Goodman. Variational Bayesian Optimal Experiment Design. **NeurIPS 2019 (spotlight).**

CODE

OneQMC. I maintain the open-source release of the Orbformer wavefunction foundation model.

Folx. We released this open-source tool for efficiently computed the Laplacian of neural networks.

Pyro OED. I am the main author of the `pyro.contrib.oed` subpackage, which supports automated experimental design for models written in Pyro. Our Deep Adaptive Design implementation that uses PyTorch and Pyro is also open source.

PyTorch SimCLR and InvCLR. I wrote an open-source PyTorch implementations of SimCLR. The PyTorch implementation of our invariant constrastive learning algorithm, InvCLR, extends the SimCLR implementation.

Datasketch. I contributed Redis support to the Locality Sensitive Hashing library Datasketch.