

LEWIS HAMMOND

Hertford College • Catte Street • Oxford • OX1 3BW

lewis@lewishammond.com • +44 7530 981741

EDUCATION

University of Oxford

2019 - 2023

DPhil Computer Science

- Focusing on the development of formal methods for multi-agent and machine learning systems, supervised by Michael Wooldridge, Julian Gutierrez, and Alessandro Abate
- Partially based at the Future of Humanity Institute as a DPhil Affiliate working on AI safety

University of Edinburgh

2017 - 2018

MSc Artificial Intelligence

Distinction

- Grades: 85% in both taught component and research component (85% overall)
- Courses in machine learning, probabilistic modelling, algorithmic game theory, and formal methods
- Dissertation, *Attributing Blame To Decisions Using Tractable Probabilistic Models*, embedded and implemented a theoretical framework for computing blame within a particular class of tractable probabilistic model, before experimentally evaluating the system on moral decision-making scenarios
- Regularly attended the Philosophy Society lectures and the CISA, LFCs, and PVSS seminar series

University of Warwick

2013 - 2017

BSc (Hons) Mathematics and Philosophy (with Intercalated Year)

First Class

- Grades: 71%, 77%, 80% in first, second, and fourth year respectively (77% overall)
- Courses in logic, algebra, epistemology, algorithms, and philosophy of mathematics and language
- Spent third year abroad as an exchange student at Uppsala Universitet in Sweden
- Regularly attended the Non-classical Logic Reading Group and the MathPhil seminar series

EMPLOYMENT

Imandra

July 2019 - September 2019

Research & Engineering Intern

- Implemented a probabilistic programming library and proof calculus for *Imandra* that allows users to easily define statistical models for probabilistic symbolic execution and probabilistic verification
- Produced several documentation notebooks and blog posts demonstrating the use of *Imandra* for analysing models learnt from data, and for verifying fairness properties of algorithms

University of Edinburgh - CISA

November 2018 - June 2019

Research Assistant

- Worked on self-proposed project, *Safe Deep Symbolic Reinforcement Learning*, that used symbolic reasoning and causal models within reinforcement learning to provide guarantees of safe behaviour

TEACHING

University of Oxford

**Private*

- Tutor & Marker: Artificial Intelligence (2020)
- Tutor & Marker: Computational Game Theory (2020)

University of Edinburgh

- Tutor: Probability* (2018)
- Tutor: Logic & Computation (2017)
- Tutor & Marker: Discrete Mathematics & Mathematical Reasoning (2017)

SERVICE

Academic

- Reviewer: IJCAI-PRICAI (2020)
- Interviewer: BSc Computer Science at St. Hugh's College, University of Oxford (2019)
- Co-organiser: Oxford AI Safety Reading Group (2019)
- Interviewer: AI Safety Research Programme (2019)
- Content developer: Road To AI Safety Excellence (2018)
- Volunteer: IJCAI-ECAI (2018)
- Volunteer: Designed Mind Symposium (2017)

Other

- Volunteer: Edinburgh Nightline (2018)
- Volunteer: Leamington Winter Support Homeless Shelter (2017)
- Member: Amnesty International UK Youth Advisory Group (2013)
- Governor: Shrewsbury Sixth Form College (2012)

AWARDS

Full Studentship (EPSRC Doctoral Training Partnership) <i>For my DPhil at the University of Oxford</i>	2019
Go Abroad Fund (University of Edinburgh) <i>For attending IJCAI-ECAI 2018 in Stockholm</i>	2018
Informatics EU/UK Master's Scholarship (University of Edinburgh) <i>For academic merit, one of only ten such awards</i>	2017
Philosophy Examiners' Prize (University of Warwick) <i>For the highest final year mark in the department</i>	2017

PUBLICATIONS

- Lewis Hammond** and Vaishak Belle (2019). *Learning Tractable Probabilistic Models for Moral Responsibility and Blame*. Currently under peer review, available upon request.
- Lewis Hammond** (2019). *Safe Deep Symbolic Reinforcement Learning*. Working paper/in progress.
- Nikolas Bernaoula, **Lewis Hammond**, and Saasha Nair (2019). *Learning Models of Mistakes*. Working paper/in progress.

PRESENTATIONS

- Lewis Hammond** and Vaishak Belle (2019). *Tractable Probabilistic Models for Moral Responsibility*. Presented at the Knowledge Representation Reasoning Meets Machine Learning (KR2ML) Workshop at NeurIPS 2019.
- Lewis Hammond** (2019). *Machine Ethics & Formal Ethical Guidelines*. Presented at Fair machines: Student perspective on Data Justice and Ethics as part of the University of Edinburgh's Data Justice Week 2019.
- Lewis Hammond** and Vaishak Belle (2018). *Deep Tractable Probabilistic Models for Moral Responsibility*. Presented at the Human-Like Computing Third Wave of AI Workshop (3AI-HLC 2019).

LANGUAGES

Natural	English (native), Swedish (intermediate)
Programming	Python, Matlab, OCaml, Isabelle/HOL, Prolog, WebPPL