

# AMMS

EPSRC Centre for Doctoral Training in  
Autonomous Intelligent Machines & Systems

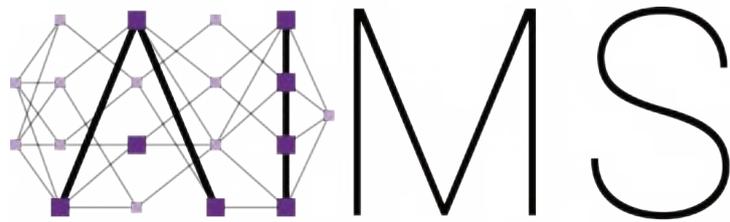
Annual Review 2022/23



Engineering and  
Physical Sciences  
Research Council







EPSRC Centre for Doctoral Training in  
Autonomous Intelligent Machines & Systems

## Foreword

---

In this document we highlight what has happened in the past year of the AIMS CDT.

Welcome to the eighth annual review highlighting key aspects and activities of staff and students in AIMS during 2022/2023. This has been the ninth full year of the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines & Systems. The CDT is thriving with applications in excess >250 in the past year.

We welcome SNAP to our list of Industry Partners. Thanks to their generosity, we have been able to fully-fund more students this past year, and for future years.

The last year has been another successful year for publications to top conferences, students submitting their dissertations, and going to work for several companies, as well as in academia.

We held a very successful AIMS seminar series with speakers from the AIMS supervisory pool, industry representatives, including Mind Foundry, Nvidia, FAIR, Microsoft Research NYC, UCL Dark & Meta AI as well as universities across the world and AIMS faculty.

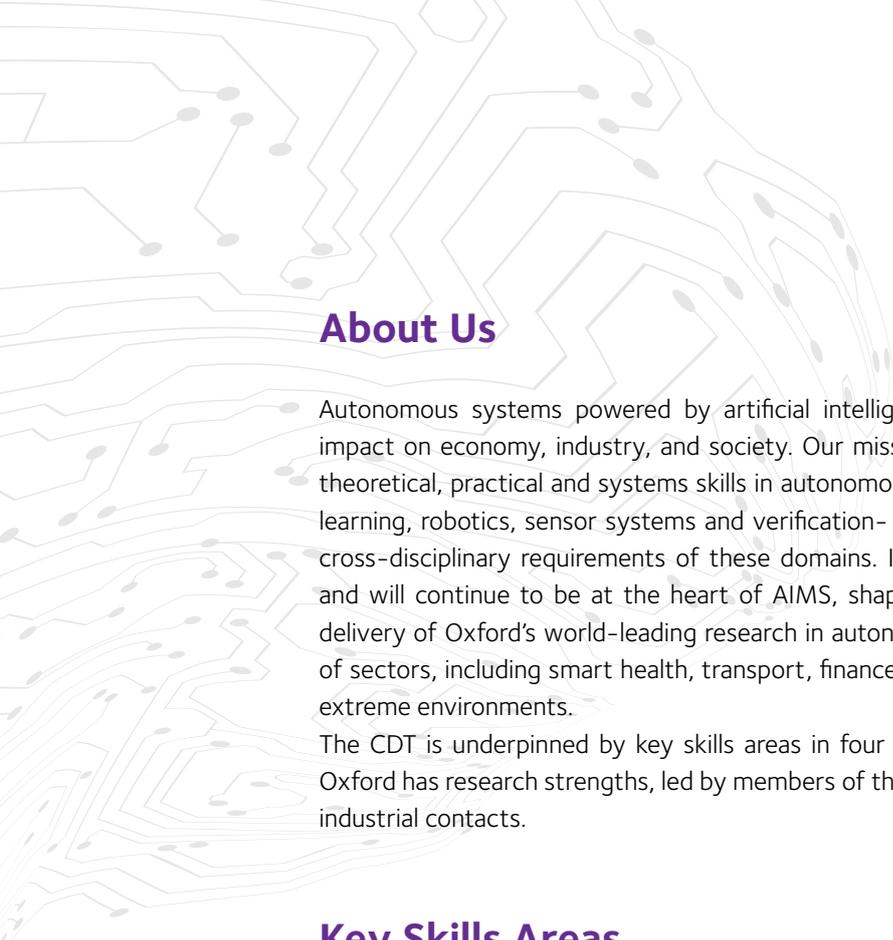
We would like to warmly acknowledge EPSRC and our industry partners for their continued support of studentships and internships.

**Mike Osborne**  
Director

**Alex Rogers**  
Co-Director

**Wendy Poole**  
Centre Administrator





## About Us

Autonomous systems powered by artificial intelligence will have a transformative impact on economy, industry, and society. Our mission is to train cohorts with both theoretical, practical and systems skills in autonomous systems - comprising machine learning, robotics, sensor systems and verification- and a deep understanding of the cross-disciplinary requirements of these domains. Industrial Partnerships have been and will continue to be at the heart of AIMS, shaping its training and ensuring the delivery of Oxford's world-leading research in autonomous systems to a wide variety of sectors, including smart health, transport, finance, tracking of animals, energy, and extreme environments.

The CDT is underpinned by key skills areas in four interconnected themes, in which Oxford has research strengths, led by members of the CDT team, and strengthened by industrial contacts.

## Key Skills Areas

What's holding up the real-world impact of Artificial Intelligence? Today, too often, innovation is overly focussed on new component algorithms, particularly those from Machine Learning. To realise impact on the world, however, such algorithms must be integrated with complete autonomous systems – in which there are far-too-few trained experts. AIMS imparts unified training in four important and intimately connected components of such systems:

- Machine Learning, as a unifying core.
- Robotics & Vision.
- Cyber-Physical Systems (e.g., sensor networks); and
- Control & Verification.

As examples of autonomous systems, AIMS aim is at building systems to impact upon

- sustainable urban development (transport, financial services, and smart infrastructure),
- extreme and challenging environments (space robots and satellite data) and
- smart health (cancer diagnosis).

To deliver training in these core research themes, we delivered a series of modules in 2022/2023 in the following areas: Data Estimation & Inference, Machine Learning, Signal Processing, Optimization, Embedded Systems Programming, Introduction to Modern Control, Discriminative & Deep Learning for Big Data , Computer Vision, Autonomous Systems Safety & Governance, Systems Verification, Privacy & security, Computational Game Theory, Reinforcement Learning, Internet of Things, Autonomous Robotics, Cooperative AI: Foundations & Frontiers and Deep Learning in Distributed and Constrained Systems.



## Events, highlights, outreach, and publications

AIMS students have taken part in a wide range of research and outreach this last year. They have also published many papers at top conferences. These include: ICCV (International Conference on Computer Vision), BMVC (British Machine Vision Conference), TARK (Theoretical Aspects of Rationality and Knowledge), ICML (International Conference on Machine Learning), CVPR (Conference on Computer Vision and Pattern Recognition), NeurIPS (Neural Information Processing Systems).

## Publications

(A full list of publications can be found at: <https://aims.robots.ox.ac.uk/publications/>)

- **Dominik Kloefer**, Dylan Campbell and Joao Henriques. “LoCUS: Learning Multiscale 3D-consistent Features from Posed Images”. ICCV 2023.
- **Benjamin Ellis, Jonathan Cook**, Skander Moalla, Mikayel Samvelyan, Mingfei Sun, Anuj Mahajan, Jakob Foerster, Shimon Whiteson. “SMACv2: An Improved Benchmark for Cooperative Multi-Agent Reinforcement Learning”. NeurIPS 2023
- **Aleksandar Shtedritski**, Christian Rupprecht, Andrea Vedaldi. “What does CLIP know about a red circle? Visual prompt engineering for VLMs”. ICCV 2023
- **Aleksandar Shtedritski**, Andrea Vedaldi, Christian Rupprecht. “Learning Universal Semantic Correspondences with No Supervision and Automatic Data Curation”. ICCV 2023
- Siobhan Mackenzie Hall, Fernanda Gonçalves Abrantes, Hanwen Zhu, Grace Sodunke, **Aleksandar Shtedritski**, Hannah Rose Kirk. “VisoGender: A dataset for benchmarking gender bias in image-text pronoun resolution”. Neurips Datasets and Benchmarks 2023
- **Aleksandar Petrov**, Emanuele La Malfa, Philip H.S. Torr, Adel Bibi. “Language Model Tokenizers Introduce Unfairness Between Languages”. NeurIPS 2023
- **Yash Bhalgat**, Iro Laina, João F. Henriques, Andrew Zisserman, Andrea Vedaldi. “Contrastive Lift: 3D Object Instance Segmentation by Slow-Fast Contrastive Fusion”. NeurIPS 2023.
- **Cong Lu**, Philip J. Ball, Yee Whye Teh, Jack Parker-Holder. “Synthetic Experience Replay”. NeurIPS 2023.
- **Lars Holdijk**, Yuanqi Du, Ferry Hooft, Priyank Jaini, Bernd Ensing and Max Welling. *Stochastic Optimal Control for Collective Variable Free Sampling of Molecular Transition Paths*. NeurIPS 2023.
- **Niki Amini-Naieni**, Kiana Amini-Naieni, Tengda Han and Andrew Zisserman. “CounTX: Open-world Text-specified Object Counting”. BMVC 2023
- A. Mitchell *et al.*, “VAE-LoCo: Versatile Quadruped Locomotion by Learning a Disentangled Gait Representation,” in *IEEE Transactions on Robotics*, doi: 10.1109/TRO.2023.3297015.
- Lewis Hammon, **James Fox**, Tom Everitt, Ryan Carey, Alessandro Abate & Michael Wooldridge. “Reasoning about Causality in Games”. The AI Journal.



- Alessandro Abate, Yousif, **James Fox**, David Hyland and Michael Wooldridge. "Learning Task Automata for Reinforcement Learning using Hidden Markov Models". ECAI 2023.
- **James Fox**, Matt MacDermott, Lewis Hammond, Paul Harrenstein, Alessandro Abate & Michael Wooldridge. "On Imperfect Recall in Multi-Agent Influence Diagrams". TARK 2023 (Best Paper Award)
- **K.Doerksen**, Y.Gal, F. Kalaitzis, C. Rossi, D. Petit, S. Li, S. Dadson, Precipitation-triggered landslide prediction in Nepal using Machine Learning and Deep Learning, *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, Pasadena, California, 2023.
- **Benjamin Ramtoula**, Matthew Gadd, Paul Newman, Daniele De Martini. "Visual DNA: Representing and Comparing Images using Distributions of Neuron Activations". CVPR 2023.
- **Aleksandar Petrov**, Emanuele La Malfa, Philip H.S. Torr, Adel Bibi. *Language Model Tokenizers Introduce Unfairness Between Languages*. Challenges of Deploying Generative AI workshop at ICML 2023.
- **Aleksandar Petrov, Francisco Eiras**, Amartya Sanyal, Philip H.S. Torr, Adel Bibi. *Certifying Ensembles: A General Certification Theory with S-Lipschitzness*. ICML 2023.
- **Zheng Xiong**, Jacob Beck, Shimon Whiteson. *Universal Morphology Control via Contextual Modulation*. ICML 2023.
- **Luke Rickard**, Thom Badings, Licio Romao, Alessandro Abate. *Formal Controller Synthesis for Markov Jump Linear Systems with Uncertain Dynamics*. QUEST 2023.
- **Shreshth A Malik**, Salem Lahlou, Andrew Jesson, Moksh Jain, Nikolay Malkin, Tristan Deleu, Yoshua Bengio, Yarin Gal. *BatchGFN: Generative Flow Networks for Batch Active Learning*. ICML 2023: Structured Probabilistic Inference & Generative Modelling Workshop.
- **Bajgar, O.**, & Horenovsky, J. (2023). Negative Human Rights as a Basis for Long-term AI Safety and Regulation. *Journal of Artificial Intelligence Research*, 76, 1043-1075.
- **Jonathan Carter** João Jorge, Bindia Venugopal, Oliver Gibson, Lionel Tarassenko. "Deep Learning-Enabled Sleep Staging From Vital Signs and Activity Measured Using a Near-Infrared Video Camera", was recently published at the 6<sup>th</sup> International Workshop on Computer Vision for Physiological Measurement (CVPM) at CVPR 2023 in Vancouver. It received an Honourable Mention for the Best Paper award.
- **Benjamin Gutteridge**, Xiaowen Dong, Michael Bronstein and Francesco Di Giovanni. "DRew: Dynamically Rewired Message Passing with Dela". ICML 2023.
- **J. Beuchert**, M. Camurri and M. Fallon, "Factor Graph Fusion of Raw GNSS Sensing with IMU and Lidar for Precise Robot Localization without a Base Station," 2023 IEEE International Conference on Robotics and Automation (ICRA), London, United Kingdom, 2023, pp. 8415-8421, doi: 10.1109/ICRA48891.2023.10161522.

- Yamada Jun\*, **Chia-Man Hung\***, Jack Collins, Ioannis Havoutis, Ingmar Posner. "Leveraging Scene Embeddings for Gradient-Based Motion Planning in Latent Space." IEEE International Conference on Robotics and Automation (ICRA), 2023. \*Equal contribution.
- **Yash Bhalgat**, Joao Henriques and Andrew Zisserman. "A Light Approach to Teaching Transformers Multi-view Geometry".(CVPR 2023).
- **Anna Gautier**, Marc Rigter, Bruno Lacerda, Nick Hawes, and Michael Wooldridge (2023). "Risk-Constrained Planning for Multi-Agent Systems with Shared Resources". In Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023).
- **Anna Gautier**, Bruno Lacerda, Nick Hawes, and Michael Wooldridge (2023). "Multi-Unit Auctions for Allocating Chance-Constrained Resources". In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2023).
- **Jonas Beuchert**, **Amanda Matthes** and Alex Rogers. "SnapperGPS: Open Hardware for Energy-Efficient, Low-Cost Wildlife Location Tracking with Snapshot GNSS". Journal of Open Hardware, 7(1): 2, pp. 1–13. 2023. DOI: <https://doi.org/10.5334/joh.48>
- **Freddie Bickford Smith\***, **Andreas Kirsch\***, Sebastian Farquhar, Yarin Gal, Adam Foster, Tom Rainforth. "Prediction-oriented Bayesian active learning". International Conference on Artificial Intelligence and Statistics (AISTATS 2023).
- **Gunshi Gupta**, **Tim G.J. Rudner**, Rowan Thomas McAllister, Adrien Gaidon, Yarin Gal. "Can Active Sampling Reduce Casual Confusion in Offline Reinforcement Learning".(Clear 2023).
- Hugo Berg, Siobhan, **Yash Bhalgat**, Hannah Kirk, **Aleksandar Shtedritski**, Max Bain. "A Prompt Array Keeps the Bias Away: Debiasing Vision-Language Models with Adversarial Learning". (AAACL-IJCNLP 2022).
- Taras Rumezhak, **Francisco Eiras**, Philip HS Torr, Adel Bibi. "RANCER: Non-Axis Aligned Anisotropic Certification With Randomized Smoothing." IEEE Winter Conference on Applications of Computer Vision (2023).
- **Francisco Eiras**, Motasem Alfarra, Philip HS Torr, M. Pawan Kumar, Puneet Dokania, Bernard Ghanem, Adel Bibi. "ANCER: Anisotropic Certification via Sample-wise Volume Maximization." Transactions of Machine Learning Research (2022).



## Internships

- AstraZeneca, FDL, Oxa, Huawei and NASA

## Outreach/Invited Speaker

- Oishi Deb too part in a panel discussion on “Inspiring People and Tech” at the Inspiring Entrepreneurs: She’s in CTRL by Dr Anne-Marie Imafidon MBE. (She’s in CTRL - YouTube).
- Kelsey Doerksen was an invited speaker on the AI panel for the Canadian Space Conference.
- AIMS/StatML students win the people choice award for “Explaining Computer Vision and Face Detection with Where’s Wally” which they delivered at the Maths Institute as part of the Maths Festival (*Uljad Berdica, Anjun Hu, Marcos Tapia Costa, Linying Yang, Alex Goldie, Mark Eid*).
- The Directors choice was awarded to “Probability, Intuition, and the Monty Hall Problem” which was delivered at the Begbroke Science Wonder Garden Party. (*Mathias Jackermeier, Nicolas Petit, Deepak Badarinath, Alexander Forster*).
- This award was chosen by Prof Nick Hawes, Director of the Oxford Robotics Institute and MPLS Academic Champion for Public Engagement with Research).



*Begbroke Science Wonder Garden Party*



## Other achievements

- Kelsey Doerksen is the MPLS Equality, Diversity, and Inclusion Fellow.
- Wendy Poole (AIMS Administrator) winner of the Outstanding ED&I champion. An award made to an academic, researcher or professional services member of staff who has shown outstanding dedication, sustained commitment, and effective leadership or impact in advancing equality, diversity and inclusion.
- Honourable mention to **Jonathan Carter** who received the best paper award (CVPM at CVPR 2023) on *Deep Learning-Enabled Sleep Staging From Vital Signs and Activity Measured Using a Near-Infrared Video Camera*
- **James Fox**, Matt MacDermott, Lewis Hammond, Paul Harrenstein, Alessandro Abate & Michael Wooldridge. *"On Imperfect Recall in Multi-Agent Influence Diagrams"*. TARK 2023 (Best Paper Award)
- **Oliver Sourbut** took part in University Challenge on behalf of Hertford College. (<https://www.oliversourbut.net/p/hertford-sourbut>)



Oliver Sourbut - University Challenge



## Oxford DPhil Siddhant Gangapurwala Honoured with Best UK Robotics PhD Award

The realm of robotics is vast and continuously evolving. Among the many dedicated researchers in this field, Siddhant Gangapurwala from the University of Oxford's Department of Engineering Science has recently distinguished himself by earning the prestigious Queen Mary UK Best PhD in Robotics Award at the TAROS 2023 conference. Siddhant's foray into robotics began with his participation in the AIMS CDT program at the University of Oxford. Guided by esteemed researchers, Prof. Ioannis Havoutis and Prof. Ingmar Posner from the Oxford Robotics Institute, he delved deep into the intricacies of reinforcement learning (RL) - a method where machines learn to make decisions by trying out actions and seeing which ones yield the best results over time. With a clear objective in mind, Siddhant pioneered new approaches, aiming to craft adaptive RL policies suitable for industrial-grade robots. The goal was straightforward: to equip robots with the capability to seamlessly adjust to fluctuating system dynamics. This would enable them to execute intricate tasks in challenging environments, such as loading/unloading and navigating unpredictable terrains.



During his research journey, Siddhant introduced "Guided Constrained Policy Optimization for Dynamic Quadrupedal Robot Locomotion," abbreviated as GCPO. This pivotal work highlighted a novel approach to ensuring that robotic locomotion operations, once learned, remained safe, especially when confronted with unforeseen situations. By integrating safety-critical boundaries into the training methodology, Siddhant's solution showcased the potential to adapt dynamically while ensuring safety, paving the way for enhanced robotic applications in practical scenarios.

Siddhant's initial research laid the foundation for profound advancements in robotics. One of his standout projects was "RLOC: Terrain-Aware Quadrupedal Locomotion



using Reinforcement Learning and Optimal Control.” This work not only showcased state-of-the-art results but also attracted the attention of prominent figures in the industry. Specifically, Benjamin Swilling, a lead on the Spot project at Boston Dynamics, recognized the implications of Siddhant’s RL framework during the International Foundation of Robotics Research (IFRR) colloquium on quadruped robotics in February 2021. Beyond academic circles, Siddhant’s innovations found tangible real-world applications, evident in the deployment of RLOC at Chevron’s industrial facility in Belgium. Additionally, his contributions were highlighted in the esteemed “State of AI” report and were featured on the Computerphile YouTube Channel and IEEE Video Friday. One pivotal takeaway from Siddhant’s research was the introduction of a control framework capable of seamlessly transferring an RL policy from one robot to another. This significant capability promotes scalability and adaptability, addressing a major challenge in the robotics domain. The benefits are evident: accelerated development cycles and more resource-efficient robotic control systems.

Siddhant’s work extends beyond just papers and projects. He’s been a global collaborator, establishing connections from the snowy peaks of Switzerland at ETH Zurich to the bustling tech-hubs of South Korea at KAIST. In addition to his academic pursuits, Siddhant’s endeavours have generated commercial interests, indicating the market potential of his innovations. His research associations with luminaries in Robotics and AI further underscore his influential position in this domain.

Siddhant Gangapurwala’s doctoral research journey showcases a blend of academic rigor and practical application. His dedication to bridging theoretical frameworks with hands-on solutions has garnered respect in academia and industry. His notable recognition at TAROS 2023 is a testament to his substantial impact in the field of robotics.



## Case Study

### Helping to protect endangered sea turtles – Jonas Beuchert

Student Jonas Beuchert is in the final year of a DPhil in Autonomous Intelligent Machines and Systems. He recently tweeted a photo of a sea turtle tracking tag under test in the Kellogg College gardens. Here he talks about his involvement in the development of the tracking device.

#### What is a sea turtle tracking tag?

Simply put, it's a small piece of electronics with an antenna and a battery that goes into a waterproof enclosure and is glued on the carapace of a sea turtle. Whenever the turtle surfaces to breathe, the tag wakes up and captures signals from navigation satellites that orbit the Earth (think of GPS). Once we get the tag back, we do some maths to discover where the turtle spent the past weeks.

#### Why do people want to track sea turtles?

Nearly all sea turtle species are endangered. Marine pollution, fishing, poaching, and climate change have all led a reduction in their numbers. If we know where they go when they are not at the beach, we can better protect these areas and help populations to survive. (Arguably, protecting them from climate change is a bit trickier...)

#### What is your role?

I develop the hardware and the software for the tags, but it is not just me: I work with Professor Alex Rogers and in the past I worked with Amanda Matthes, another DPhil student.



## Why did you choose this subject for your DPhil?

Existing turtle tags can cost several thousand US Dollars each, making them inaccessible to many wildlife conservationists, this is partially due to expensive electronics. In contrast, we have developed a tag design that is low-cost, highly energy-efficient and will be fully open sourced. In addition, the tag needs only twelve milliseconds of satellite data to resolve a position fix, allowing us to locate a turtle that surfaces only very briefly.

## How is the project going?

In summer 2021, working with a local conservation organisation, we had the chance to deploy the first version of our tags on sea turtles in Cabo Verde. Since then, we have open-sourced a general-purpose version of our tracking tags, distributing more than 100 to interested conservationists and researchers for different kinds of animal tracking projects around the world, ranging from sea birds to venomous snakes.

## What comes next?

We are planning to work with several collaborators on more joint deployments this summer, including updated versions of our turtle tag. On a personal level, I also need to find a position after finishing my DPhil in autumn, ideally working at the intersection of software and hardware and maybe even on wildlife conservation technology and open hardware and software.

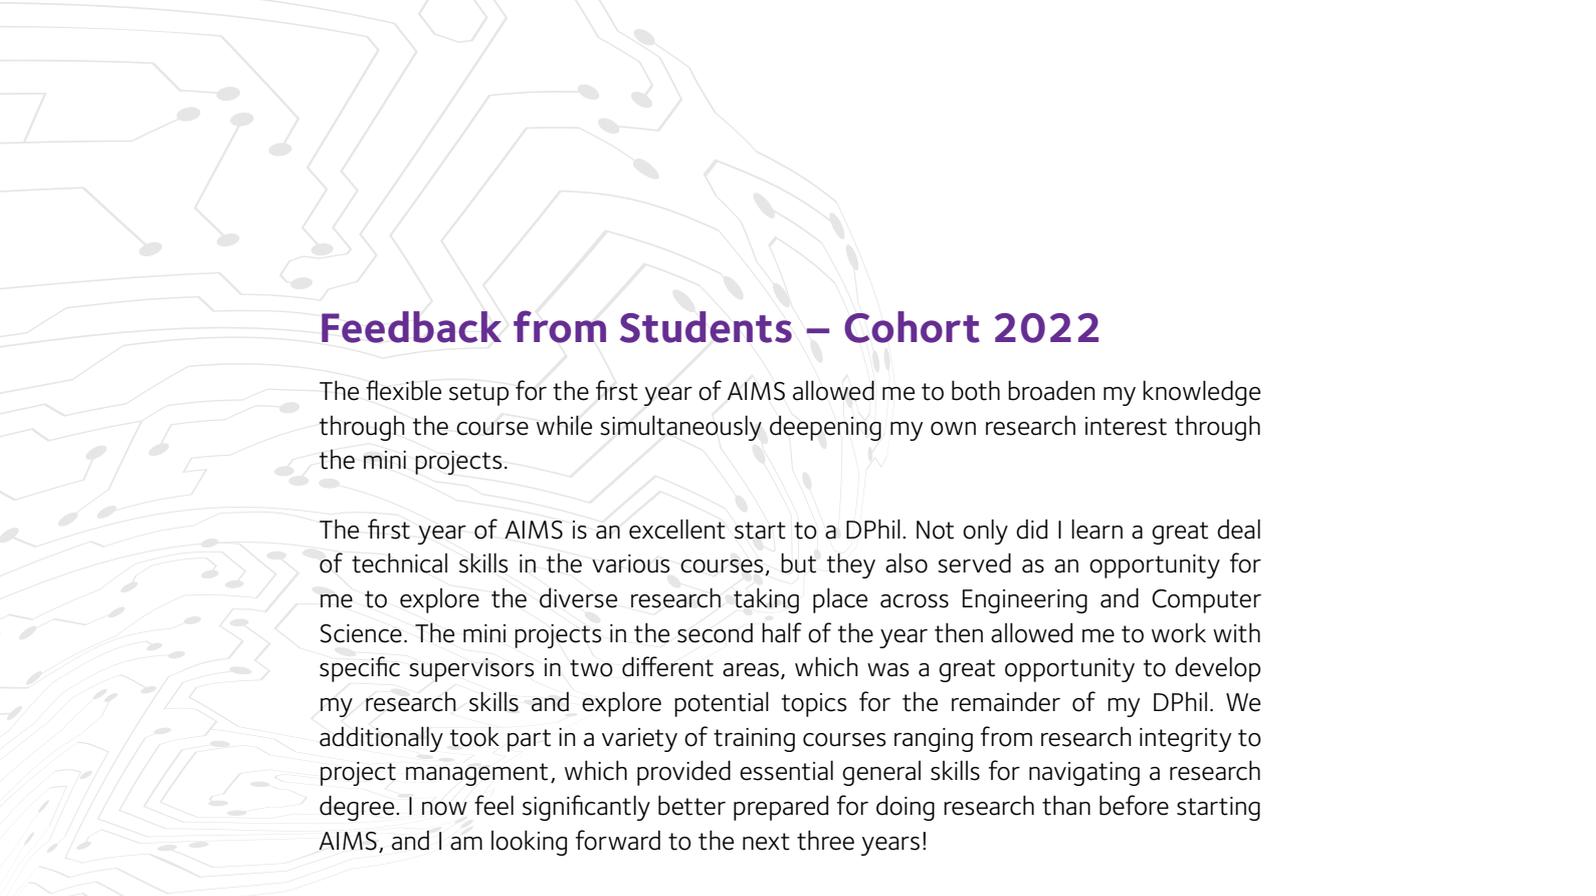
## What support have you received?

Recently, I received support from a Kellogg Travel Grant, and I also hope to benefit from a Kellogg Research Grant later this year. Also, Kellogg's tables in the quad are great for assembling and testing sea turtle tags! I am glad to be part of a college where you are allowed to step on the grass.

## How can I learn more about your work?

Follow me on Twitter @JonasBchrt or check out our project website <https://snappergps.info>. Reproduced by kind permission of Kellogg College





## Feedback from Students – Cohort 2022

The flexible setup for the first year of AIMS allowed me to both broaden my knowledge through the course while simultaneously deepening my own research interest through the mini projects.

The first year of AIMS is an excellent start to a DPhil. Not only did I learn a great deal of technical skills in the various courses, but they also served as an opportunity for me to explore the diverse research taking place across Engineering and Computer Science. The mini projects in the second half of the year then allowed me to work with specific supervisors in two different areas, which was a great opportunity to develop my research skills and explore potential topics for the remainder of my DPhil. We additionally took part in a variety of training courses ranging from research integrity to project management, which provided essential general skills for navigating a research degree. I now feel significantly better prepared for doing research than before starting AIMS, and I am looking forward to the next three years!

My first year in AIMS exceeded my expectations in a few ways. I was particularly excited to get to know the cohort and the previous years' cohorts, as well as the various academics and administrators who form the community, and this has been a really welcoming and fruitful experience. The classes in the first two terms were interesting for the most part, though by term two I was eager to get more into research! The miniprojects period in the latter part of the year provided some good time to brainstorm and also resulted in some work I'm quite proud of. Besides AIMS, getting to know Oxford and my college has also been very welcome and rewarding, with quiz groups, reading clubs, and interdisciplinary discussions all definitely enriching my time here.



## Feedback from Courses – Cohort 2022

**Autonomous Robotics** - Excellent course. Very well structured and offered great exposure to the ORI. The lab demonstrators were also all brilliant. We all got very into it, and were getting very competitive at the end! Loved how every lab was related to that morning's lecture too.

**Computer Vision** - The lectures were very helpful and provided a broad understanding of computer vision from older methods to deep learning.

**Cooperative AI: Foundations & Frontiers** - Very diverse, intriguing (and condensed!) set of problems that I haven't been exposed to during my previous studies and research. Really enjoyed this more discussion-based course. The content was quite generalist, and didn't focus as much on AI as I expected - but I really enjoyed the interdisciplinary conversations that this approach brought. Christian was a great lecturer and I found myself highly engaged in his slides.

Introduction to Modern Control - I enjoyed Alessandro's seminar on his current research on one of the days, which made it clear how some of the ideas in control are now being applied for various different problems.

Lecturers were all excellent. The problem sheets were engaging but quite difficult - not a bad thing!

**Date Estimation & Inference** - I thought the lecturers were very knowledgeable and encouraging, and open to any questions. The coursework was difficult, but enjoyable. It has been interesting to learn about and gain experience with Gaussian Processes, and also to learn more about Bayesian Deep Neural Networks.

**Machine Learning** - Lab tasks were excellent and varied, and I appreciated the scalability of the assessed task (I chose to challenge myself a bit but I was able to keep it within the time scope available). Exposure to a varied set of techniques, especially the less mainstream (e.g. inverse models) was really insightful and appreciated. The labs with solutions were great. I loved learning about Pyro. The instructors were also incredibly nice. I wish we had more time to explain Pyro concepts and the training loops in Pyro in more detail.

**Optimization** - Paul was an excellent lecturer and this was a really well set up course. I genuinely have no complaints at all.

### Challenging coursework

The labs were a great mix of interesting mathematical challenges and coding implementation tasks. The coding demos were satisfying.



## AIMS CDT Cohort 2023



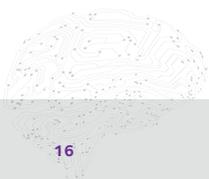
**JACQUES CLOETE** - I grew up on Guernsey in the Channel Islands, and graduated from the University of Oxford with a MEng in Engineering Science. I have a great passion for robotics; I undertook a summer internship at the Oxford Robotics Institute, learning to use ROS and acquiring a great deal of practical experience in robotics and cyber-physical systems. I also ran a robotic arm project in the OxRAM student society, wherein we 3D-printed, assembled and programmed a robotic arm for completing

simple manipulation tasks. For my final year project, I worked on loco-manipulation skills for quadruped robots with arms, creating motion- and task-planners for tasks such as autonomously approaching then tightening a valve. This work evolved into researching adaptive robot behaviour that allows a robot to intelligently react to observations it makes during task execution, such as adjusting its manipulation strategy to apply more torque to a stiff valve. Machine learning has plenty to offer for this work, and I am thrilled to be able to continue my research through the AIMS CDT. Outside of research, I love skiing, model-building and playing the saxophone.



**THOMAS FOSTER** - I graduated from Oxford in 2020, coming 2nd in my cohort and winning the Microsoft Prize for best undergraduate dissertation. After graduating I tried my hand at entrepreneurship, founding 2 successful companies. At Halo we developed a saliva-based PCR test that we scaled rapidly in response to Britain's COVID-19 testing shortages. After this I founded Genei, where we used LLMs and RLHF to help speed up research. We went through YC combinator in 2021 and raised funding from some great

people. In late 2022 I left Genei to focus on research. I'm particularly interested in how we can automate the evaluation and training of deep neural nets. This means I'm interested in how we evaluate models on complex criteria (such as with learned reward functions), how we can systematically generate examples the model finds challenging (as in curriculum/active/adversarial learning) and how we can turn this into a useful training paradigm (such as RLHF or open ended RL). I spent the summer of 2023 exploring these themes at V7, developing in-context learning models for video segmentation.





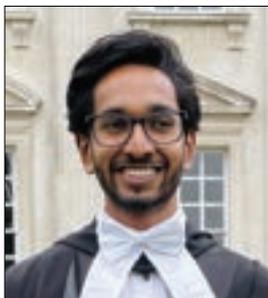
**YOAV GELBERG** - I graduated with a bachelor's degree in mathematics (summa cum laude) from Technion – Israel Institute of Technology in 2018, where I was a part of the Technion excellence program. As an undergraduate, I took part in two research projects: one applying metric learning to image processing at Prof. Ron Kimmel's Technion-based Geometric Image Processing lab, and the other studying robustness of expansion properties of graphs, with Prof. Irit Dinur at the Weizmann Institute funded by the Weizmann

Young scholars program.



**YONATAN GIDEONI** - Before starting AIMS I completed an MPhil in Computer Science at the University of Cambridge and a BSc in Physics from the Hebrew University of Jerusalem. There I also completed a minor in computer science and took a diverse assortment of courses, ranging from astrobiology to algebraic structures. In parallel I also worked at Mobileye on maps for autonomous vehicles, taught at the Israeli Arts and Sciences Academy, and had a machine learning for quantum computing research

internship. Outside of my studies, I enjoy reading literature and learning about different cultures.



**DULHAN JAYALATH** - I care about understanding the nature of learning and consciousness in intelligence of any form, whether human-like or otherwise. My interest in AI research began as an undergraduate at the University of Southampton, where I investigated how we can navigate autonomously in natural environments using only the visible spectrum as humans do. Then, as a master's student at the University of Cambridge, I delved into multi-agent learning.

My work aimed to understand how we can encode information about environments so that agents can leverage their prior knowledge to learn faster and with less data, much like how humans map our existing knowledge to new scenarios to learn quickly. Now, on the AIMS CDT at Oxford, I will be working on neural decoding, researching how we can map brain signals to representations like text and images to help those with limited ability to communicate. I hope that this work will be a step towards another paradigm of AI---perhaps one controlled by human thought. Beyond university, I spent one summer as a machine learning research intern at Speechmatics, working on representation learning for speech recognition technologies. Before that, I spent two summers as a kernel developer at Arm, developing algorithms for machine learning kernels that utilise new hardware efficiently. In my free time, I enjoy sports (mostly tennis), films, and learning about anything that captures my curiosity.





**BEN KAYE** - I have received an MEng in Engineering Science from Oxford University, and I am super excited to be working on robotics problems. My interests lie in optimisation and control, and on the side I enjoy climbing, Taekwondo, and travelling. Prior to joining AIMS, I worked as a control engineer in the medical device industry.



**LUISA KURTH** - My motivation is to help improving people's lives through AI. Currently, I am mostly interested in the challenges of advancing machine learning for medical image analysis. The CDT in AIMS offers the perfect platform for this journey and I'm excited to connect with anyone sharing my interest. I hold a Bachelor's degree in Psychology from the University of Mannheim and a Master's degree from the University of Oxford's Internet Institute. During my time at Oxford, I participated in cutting-edge research on the societal and ethical aspects of AI. This experience fueled my fascination for machine learning, leading me to pursue a second Master's degree at the University of Tübingen, where I focused on the foundations of deep learning, large language models and explainable AI. Along the way, I've conducted brain research at the Max-Planck-Institute and worked as a policy researcher at the OECD. Outside of research, I enjoy reading, socializing with friends, and visiting art galleries.



**BERNARDO LUSTRINI** - I was born in Naples, Italy, but grew up in South Holland in the Netherlands. I moved to the UK for University and graduated with an MEng in Engineering from the University of Cambridge (Queens' College) in 2021, specialising in Information Engineering. My master's research was in the Cambridge Computer Vision & Robotics group, where I looked at developing techniques to predict the fit-quality of rigid face masks from 3D scans of human faces. After graduating I started working as a Computer Vision research engineer at Oxehealth: an Oxford university spin-out that specialises in non-contact vision-based patient monitoring systems. I've worked on many projects involving all things Computer Vision, but I have particularly enjoyed my recent work on multi-person tracking. Oxehealth is also sponsoring my DPhil research, which will focus on pushing the state-of-the-art in non-contact health monitoring, with a particular focus on inpatient mental health settings. I'm happiest on a football pitch or with a guitar in hand, preferably both. I've done a bit of climbing and would like to pick this up again in my coming years, as well as hiking as I'm always looking for excuses to go out into nature.



**HARRY MEAD** - I grew up in Australia and graduated in 2023 from the University of Oxford with an MEng in Engineering. Throughout my time at Oxford, I undertook a number of internships including at the Oxford Robotics Institute as well as Archangel Imaging, the latter of which focused on researching few-shot learning applications in computer vision. For my final year project, I explored the effectiveness and limitations of imitation learning for applications where labelled demonstration data is unavailable, with comparisons to reinforcement learning in terms of agent performance and similarity to human behaviours. Outside of study, I enjoy playing tennis, as well as watching any sport that is on.



**DARIUS MUGLICH** - I graduated with an Honours BSc in Mathematics and Computer Science from the University of British Columbia. I then completed a MSc in Computer Science from the University of Oxford. My MSc research focused on improving coordination amongst agents in cooperative multi-agent systems. I am broadly interested in improving robustness and generalization capabilities of ML systems, particularly in multi-agent contexts, as well as in human-AI coordination. I am generously funded by Toshiba Research and the EPSRC for my doctorate, so to continue developing methods for agents to better generalize and adapt to novel tasks and environments.



**ALEX PONDAVEN** - I am interested in working with generative models that can understand the structure of data in a compositional and generalisable manner. After working with object recognition models and video data pipelines for road users at Humanising Autonomy, a computer vision startup, I realised how current state of the art models are very dependent on large amounts of data and lack robustness to new problems. I have had the opportunity to apply meta learning techniques for satellite image inpainting as a research project while studying at Imperial College London. As part of my Master's Project at Imperial supervised by Yingzhen Li, I worked with text-to-image diffusion models and investigated how the diversity of generated images can be improved efficiently. I just graduated from my MEng in computer engineering in 2023 and plan on continuing to work on these generative models for vision in the AIMS CDT. Outside of machine learning, I love running, climbing, working out, gaming and playing piano.





**ELEANOR TROLLOPE** - Prior to pursuing my PhD, I completed my master's in Mathematics and Foundations of Computer Science at the University of Oxford in 2023. My thesis explored solving differential equations using Physics-Informed Neural Networks. During my master's, I also worked with Quantinuum on a paper involving quantum circuits for natural language processing. Prior to this, I did an internship funded by UniQ+ DeepMind in the Torr Vision Group in Oxford where I worked on equivariance and self-supervised learning. This internship was my step back into academia after working as a quantitative analyst in finance for 2.5 years. Before this, I graduated with a degree in Natural Sciences from the University of Exeter, with a focus on mathematics and computer science. Outside of my research, I enjoy travelling, dancing, and reading about science.



## AIMS Contacts

The AIMS administration team comprises the Director, the co-Director and the Centre Administrator.



### MICHAEL OSBORNE

Michael A Osborne is an expert in the development of intelligent algorithms capable of making sense of complex big data. His work in Machine Learning and non-parametric data analytics has been successfully applied in diverse and challenging contexts. For example, in astrostatistics, Michael's probabilistic algorithms have aided the detection of planets in distant solar systems, and in autonomous robotics, his work has enabled self-driving cars to determine when their maps may have changed due to roadworks. More recently, he has addressed key societal challenges, analysing how intelligent algorithms might soon substitute for human workers, and predicting the resulting impact on employment. Michael is an Associate Professor in Machine Learning, an Official Fellow of Exeter College, and a Faculty Member of the Oxford-Man Institute for Quantitative Finance, all at the University of Oxford.



### ALEX ROGERS

I originally studied Physics at Durham University before joining Schlumberger as a wireline logging engineer. After five years working in various oilfields around the world, I took suspended employment to study for a PhD applying statistical physics to models of evolving populations. Upon completing my PhD, I worked for a spin-out from the Santa Fe Institute applying complexity science to business problem before returning to academia, initially at the University of Southampton, and now at the University of Oxford.



### WENDY POOLE

I have been working in the University for 28 years now. I accepted the position as CDT Centre Administrator, after working in the Department of Computer Science as the MSc Course Administrator for 20 years.

## Academic Supervisors

A full list of academic supervisors can be found at:

<http://aims.robots.ox.ac.uk/academics-and-staff/>







# THE MONTY HALL PROBLEM

Play this game to win a prize!

You are a contestant on a game show. In front of you are **three** doors. Behind one of the doors, there is a car! Behind the other two doors, there are goats. Of course, you would prefer to win a car over a goat. The game show host, Monty, then says...

Pick a door!

Suppose you picked door number **1**. Monty, who knows where the car is located, is being mischievous. Before opening door number **1** and revealing whether or not you win the car, Monty opens door number **2** to reveal... a goat!

Monty then gives you an offer: you can stick with your original choice of door **1**, or you can switch to door number **2**.

Stick with door number **1** or switch to door number **2**?

What do you do? Do you stick or do you switch? Which option gives you the best chance of winning the car? Or do you have the same chance of winning the car regardless of your decision?

# A GAME OF CHANCE

Probability is the chance that an event will occur, or how likely it is something will happen. It helps us reason about situations where we don't know what will happen. Here are a few scenarios where probability comes to the rescue.

Probability and statistics have applications in cool areas such as:

### Weather prediction

The weather forecast uses a combination of past weather data, temperature, air pressure, etc to predict the weather for tomorrow.



### Healthcare

Will this treatment work? Is this headache caused by the pill? The test results seem back negative. How sure are we that they are correct?



### Robotics

Statistics is used by a robot to find its position. Statistics is also used to predict which direction to move to reach the target.



### Games

We can analyze all sorts of word and dice games with probability and statistics. For example, we can estimate the odds of winning or even build computer programs that beat human players!



# HISTORY

The Monty Hall Problem first appeared in 1975. It was based on the American TV gameshow *Let's Make a Deal*, and the problem is named after *Monty Hall*, the presenter of this gameshow.



The problem became famous when it appeared in the American magazine *Parade*. In Marilyn vos Savant's "Ask Marilyn" column.

In this column, Savant would answer questions sent in by readers. In 1990, one reader asked Savant to solve the Monty Hall problem. Savant was famous for having the Highest IQ in the Guinness Book of World Records, with an IQ of 228.

Savant found the correct answer to the problem - can you?





EPSRC Centre for  
Doctoral Training in  
Autonomous Intelligent  
Machines and Systems (AIMS)  
Department of Engineering Science  
University of Oxford  
Parks Road  
Oxford OX1 3PJ

**Director**

Mike Osborne  
Email: [aims-cdt@robots.ox.ac.uk](mailto:aims-cdt@robots.ox.ac.uk)  
Tel: 01865 616622

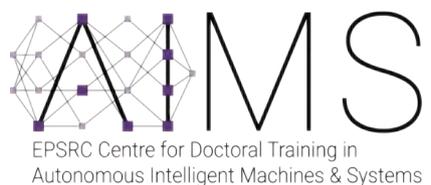
**CDT Co-Director**

Alex Rogers  
Email: [aims-cdt@robots.ox.ac.uk](mailto:aims-cdt@robots.ox.ac.uk)

**CDT Centre Administrator**

Wendy Poole  
Email: [aims-cdt@robots.ox.ac.uk](mailto:aims-cdt@robots.ox.ac.uk)  
Tel: 01865 616632

[www.aims.robots.ox.ac.uk](http://www.aims.robots.ox.ac.uk)



EPSRC Centre for Doctoral Training in  
Autonomous Intelligent Machines & Systems