AIMS

EPSRC Centre for Doctoral Training in
Autonomous Intelligent Machines & Systems

Annual Review 2021/22

Engineering and Physical Sciences Research Council

UNIVERSITY OF OXFORD


IAC
PARIS
2022
73rd INTERNATIONAL
ASTRONAUTICAL
CONGRESS

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 seventh annual review highlighting key aspects and activities of staff and students in AIMS during 2021/2022. This has been the eighth 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 Hyundai 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, Deimos Space, GoogleBrain, DeepMind, Toshiba, Amazon Web Services, Mission Control Space Services as well as universities across the world.

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 2021/2022 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, Security in Wireless and Mobile Robotics, Computational Game Theory, Reinforcement Learning, Internet of Things, Autonomous Robotics 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: AAMAS (Autonomous Agents and Multiagent Systems, MICCAI (Medical Image Computing & Computer Assisted Intervention), ICCRA (International Conference on Robotics & Automation), ICLR (International Conference on Learning Representations), NeurIPS (Neural Information Processing Systems), Royal Society Open Science to name but a few.

Publications

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

- **A. Gautier, A. Stephens,** B. Lacerda, N Hawes and M Woolridge. *Negotiated Path Planning for Non-Cooperative Multi-Robot Systems.* AAMAS 2022.
- **A. Gautier** & M. Woolridge. *Understanding Mechanism Design.* IEEE Intelligent Systems.
- H. Berg, M. Hall, **Y. Bhalgat**, W. Yang, H. R. Kirk, **A. Shtedritski** and M. Bain. *A Prompt Array Keeps the Bias Away: Debiasing Vision-Language Models with Adversarial Learning.* AACL-IJCNLP 22.
- **A. De Palma**, R. Bunel, K. Dvijotham, M. P. Kumar, R. Stanforth. *IBP Regularization for Verified Adversarial Robustness via Branch-and-Bound.* ICML 2022 (best paper award).
- **V. Kurin*, A. De Palma*,** I. Kostrikov, S. Whiteson, M. P. Kumar. *In Defense of the Unitary Scalarization for Deep Multi-Task Learning.* (NeurIPS) 2022.
- **J. Brauner, M. Sharma** et al. *Prioritized Training on Points that are Learnable, Worth Learning, and not yet Learnt*. Proceedings of Machine Learning Research.
- **J. Brauner, M. Sharma** et al. *Mask Wearing in Community Settings Reduces SARS-CoV-2 Transmission.* PNAS vol. 119 / No. 23.
- **J. Brauner, M. Sharma** et al. *A Dataset of Non-Pharmaceutical Interventions on SARS-CoV-2 in Europe*. Scientific data.
- **T. Rudner, F. Bickford Smith**, Qixuan Feng, Yee Whye Teh, Yarin Gal (2022). Continual learning via sequential function-space variational inference. International Conference on Machine Learning.
- **T. Reichelt, A. Goliński,** Luke Ong, Tom Rainforth. *Expectation programming: Adapting probabilistic programming systems to estimate expectations efficiently.* Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence, 2022.
- C. Briand**, K. Doerksen** & F. Deleflie. *Space Weather.* Advancing Earth & Space Science.
- **K. Doerksen**. *A Multi-Lander New Frontiers Mission Concept Study for Enceladus: SILENUS* accepted for publication in Frontiers in Astronomy and Space Sciences, section Planetary Science.
- **Panagiotis Tigas*,** Yashas Annadani*, Andrew Jesson, Bernhard Schölkopf, Yarin Gal, Stefan Bauer. *Interventions, Where and How? Bayesian Active Causal Discovery at Scale.* Neurips 2022.
- **M. Newton** & A. Papachristodoulou. *Stability of Non-linear Neural Feedback Loops using Sum of Squares.* 61st IEEE Conference on Decision & Control 2022.

Internships

- DeepMind, CPP Investments, NASA Jet Propulsion Lab, Microsoft Research Cambridge, Amazon and Waymo.

Outreach/Invited Speaker

- AI4Space: Perspective from the Next Generation Plenary Panellist at the International Astronautical Congress (Sept 2022) https://www.youtube.com/watch?v=Hmi1t1tLMtw&ab_channel=InternationalAstronauticalFederation
- Panellist for the Zenith Canada Pathways Foundation Summit – Opportunities in Space
- Organizer and Panellist for the Oxford Women in Computer Science Society *Highlighting Wom*n Researcher in ML* ICML panel: https://icml.cc/virtual/2022/social/20899
- Invited Speaker for Carleton University SOAR Student Leadership Conference: https://carleton.ca/seo/soar/#sect1
- Mentor for the Zenith Canada Pathways Foundation: www.zenithpathways.ca
- Interview by CGTN Europe about AI for space: https://twitter.com/CGMeifangZhang/status/1573063298295402496?s=20&t=49gvNgUSw5GblNIidIm9Xw
- 1st years ran a masterclass on How do Machines Learn: An Interactive Introduction to Machine Learning, session as part of the Oxford Computer Science Royal Institution Masterclasses, with a section and discussion on AI ethics/biases/safety.
- Attended the Didcot/Vauxhall Barracks play & Activity Day. The activity was a "Crash course in Reinforcement Learning"
- Outreach at St Clare's Oxford. "How do Machines Learn: An Introduction to Machine Learning.
- Attended a Computer Science Summer School at Cherwell College. "AI in Production: Good, Bad and the Ugly.

participants playing crash course in Reinforcement Learning

Other achievements

- Yash - open-sourced a pure-Pytorch implementation of NVIDIA's Instant-NGP algorithm, which got quite popular achieving 500 stars on GitHub recently! GitHub link: https://github.com/yashbhalgat/HashNeRF-pytorch. This was a result of my first mini project.
- Mrinank and Jan receive Impact Award in the Social Impact category. 2022 MPLS Impact Award - Awarded for social impact of research on ``Understanding effectiveness of interventions against Covid-19 using Bayesian models''. Winner among researchers in any career stage across all MPLS departments (mathematics, physics, engineering, computer science, life sciences) at the University of Oxford.
- Gunshi Gupta is the MPLS Equality, Diversity, and Inclusion Fellow.
- Pan Tigkas - won the first place at GSKs Gene Disco challenge https://www.gsk.ai/genedisco-challenge/
- AIMS students win award for their public engagement with research activities. The Director's Choice award - "Crash Course in Robot Learning" – AIMS Kelsey Doerksen, Patrick Benjamin, Benjamin Gutteridge, Luke Richard, Sebastian Towers.
- This award was chosen by Prof Nick Hawes, Director of the Oxford Robotics Institute and MPLS Academic Champion for Public Engagement with Research).
- Oxford won the Best Presentation (Charig Yang – It's about time: Analog Clock Reading in the Wild" and Best Poster (Matt Jackson, Shreshth Malik – Improving Few-Shot Learning Task-Informed Meta-Initialisation) at the Joint CDT conference with Bristol, Lincoln, and Edinburgh. This was held in June 2002 in Bristol.
- Kelsey Doerksen is one of the lead organizers for the Oxford-Cambridge Women in Computer Science Conference in May 2022, and Oxford Women in Computer Science Society Treasurer.

Best Poster and Presentation winners at the Joint CDT conference with Bristol, Edinburgh, Lincoln, and Oxford

Case studies

Zheng Xiong - Towards Efficient and Generalizable Dexterous Manipulation with Reinforcement Learning using AWS

Despite that robot driven by hard-coded instructions have achieved great success in the past decades, their applications are still quite limited in highly controlled scenarios due to the lack of adaptability to flexible environments with unforeseen variations. Moreover, robotic programming requires extensive human expertise and trial-and-errors, and the programming complexity of some dexterous manipulation tasks is even beyond the scope of human experts.

Reinforcement learning (RL) has the potential to tackle these challenges by automatically learning the optimal control policy. However, existing RL algorithms are mainly limited in the following aspects:

Sample inefficiency. It usually requires millions of steps to train an RL policy, which is very time consuming even in simulation, let alone on real-world robots.

Sparse reward. The intrinsic rewards of robotic manipulation are usually sparse, such as a binary signal of whether a destination is reached or not, which makes learning much harder. Although reward shaping can provide more feedback with hand-crafted dense rewards, it is very laborious and may introduce human bias which leads to sub-optimal or even unexpected learning outcomes.

Poor generalization. Existing RL algorithms are usually trained specifically for a certain task or a small set of tasks, thus are very prone to task variations in dynamic environments.

This project aims to tackle these challenges for a specific type of dexterous manipulation task of moving a block to arbitrary destinations with a tri-finger robot, which presents a useful yet challenging skill to learn. This figure illustrates the Markov Decision Process (MDP) framework for solving the tasks considered in this project. We decompose the environment for the RL controller into three parts as *robot*, *physical world*, and *task*. The *robot* consists of the robot state which can be directly manipulated by the controller. The *physical world* consists of object state (such as block pose) which can be indirectly manipulated by the controller through interaction, task-relevant context (such as block mass) which can influence the state transition dynamics, and task-irrelevant context (such as block colour) which has no effect on transition dynamics. The *task* consists of goal context (such as goal pose) which, together with the object state, determines the sparse reward in each step. Based on this framework, we highlight four desired properties of a good RL algorithm for dexterous manipulation:

Generalization to different goals. We want to learn a versatile multi-goal RL policy which can adaptively achieve different goals in the same task family, such as pushing a box to different positions in a room.

Learning with little reward engineering. Little human intervention should be involved in reward shaping to reduce human bias and workload.

Adaptation to task-relevant context. The controller should be robust to variation of task-relevant context. For example, a picking policy should be adaptive enough to lift blocks with different mass.

Invariance to task-irrelevant context. The control's behaviour should not be influenced by task-irrelevant context for robust control. For example, the colour of walls in a room should not influence a robot's behaviour to manipulate objects.

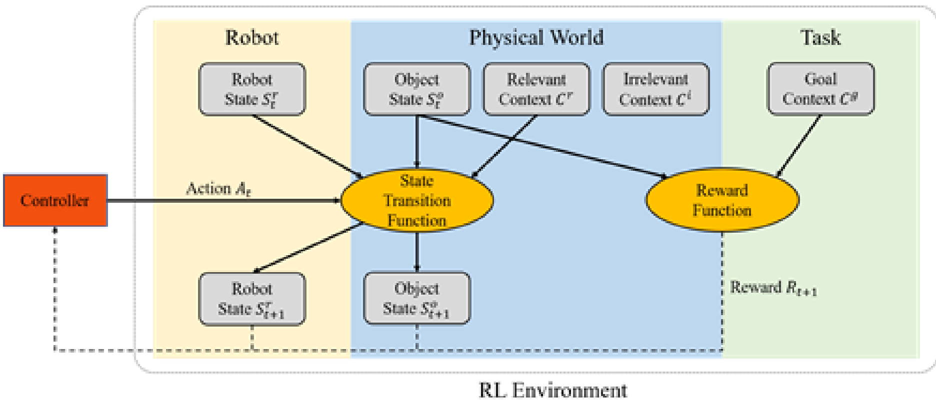

A framework of RL for robotic manipulation. The red box represents the RL controller. The grey boxes represent the environment variables, where *state* represents manipulable variables, and *context* represents variables which cannot be changed by the agent but may differ across environments. The transition dynamics of the MDP is decomposed into the state transition function and the reward function, represented by the two orange ellipses. The state transition function takes robot state, object state, task-relevant context and action as input, and outputs the updated robot state and object state. The reward function takes goal context and object state as input and outputs the step reward.

In this project, we aim to learn RL policies which satisfy the above four criteria for dexterous manipulation. The first two goals are closely related to the reward function, while the last two goals are involved in the state transition function. Thus, we design a two-stage strategy to accomplish these goals, while different techniques are utilized to tackle the key challenge in each stage.

In the first stage, we try to achieve the first two goals by learning a goal-conditioned policy with sparse reward. The context variables are fixed to some default values, so we do not need to consider context adaptation in this stage yet. However, even without context variations, learning goal-conditioned policies with sparse reward is still challenging. To tackle this challenge, we utilize curriculum learning to improve learning efficiency, which automatically proposes new tasks with increasing difficulties to guide the controller's learning. Specifically, as all the manipulation tasks considered in this project should be solved by moving an object to a goal position, we use a path planner to generate pseudo goals with adequate difficulty as the learning curriculum.

In the second stage, we continue to fine-tune the policy learned from the first stage in different environments with context variations to improve its generalization. In this stage, we first do an intervention-based simulation test to identify task-relevant and task-irrelevant context. Then we train the policy in a compact search space formed by the identified task-relevant context. Furthermore, instead of randomly sampling in the search space, we propose a particle filter-based curriculum to focus more on the harder context variations to achieve better generalization.

We conduct simulated experiments on three single-object manipulation tasks from the CausalWorld benchmark. Experimental results show that in the first stage, our path planner-based curriculum improves the learning efficiency on simple tasks, and better solves the harder tasks where vanilla RL algorithms fail; In the second stage, our particle filter-based curriculum helps improve the policy›s generalization to context variations. This video shows some example trials of solving a very hard pick and place task with the tri-finger robot.

https://s3-us-west-2.amazonaws.com/secure.notion-static.com/9592c379-ba11-4ae8-9ff1-c45148813bd9/path_planningRL(with_watermark).mp4

The computational resources for the experiments in this project are kindly supported by Amazon Web Services (AWS) with £500 of cloud credits. Reinforcement learning experiments are known to be computationally expensive, and the results of this project cannot be achieved without the computational resources supported by AWS.

I mainly used the EC2 service in AWS to run the reinforcement learning experiments. The web interface for instance management is user-friendly, which helps me conveniently manage multiple instances in a very clear way. The machine learning instance templates provided in EC2 are also quite helpful, which make it very easy for me to setup the experimental environments.

Dominik Kloepfer – Learning How to Learn Where you are: Meta-Learning for Few-Shot Camera Localization using AWS.

In this project, we looked at the task of camera localization: given several images of a scene (a room, a street, ...) and the positions and viewing directions from which they were taken (together, these are the **camera pose**), can we determine the camera pose from which a new, unseen image was taken?

A method that could do so reliably and quickly would find use in many different fields, from robotics and autonomous driving (to determine where to go, the self-driving car needs to know where it is!) to augmented reality applications (to determine what virtual objects to overlay at what point in the scene, the AR-device needs to know where it is itself).

Given how widely applicable camera localization is, it is no surprise that many different methods have been proposed to solve this problem. Some aim to use correspondences between distinctive features in the image and these features in some internal 3D-representation of the scene to then triangulate the camera pose, while others train neural networks to directly predict the camera pose from an image directly.

These methods however have some significant drawbacks: The former class of methods usually requires a detailed 3D-scan of the scene to establish these correspondences between image and the 3D-structure of the real world, and these scans are time-consuming and expensive to collect. The neural-network-based methods on the other hand usually require many (several thousand) images to learn a single scene and need to be re-trained from scratch for each new scene. Neural-network-based systems also usually do not perform as well as the 3d-structure-based methods achieve on most camera-localization benchmarks.

Method

DSAC

To address these challenges, we built our method based on a previous method, DSAC [1], which achieved state-of-the-art results on several camera localization datasets.

DSAC in a way is a hybrid of the 3D-structure-based and the neural-network-based methods. It uses a neural network to predict the 3D-world-coordinates of the pixels in the image to establish the correspondence between the image and the real world, and then uses these correspondences to triangulate the camera pose. Establishing these image-to-world correspondences using a neural network means that DSAC does not require a detailed point cloud, while its use of triangulation still allows it to perform very well.

Intuitively, the neural network picks out some distinctive landmarks (e.g., a chair) in the scene and memorises their 3D real-world coordinates. If it then sees some of these landmarks in a test image, it can solve the question "if this chair appears in the left part of the image, and I know that it is at this position in the scene, where must I (the camera) be for the chair to appear where it does?" to determine the camera pose. This is visualised in Fig. 1.

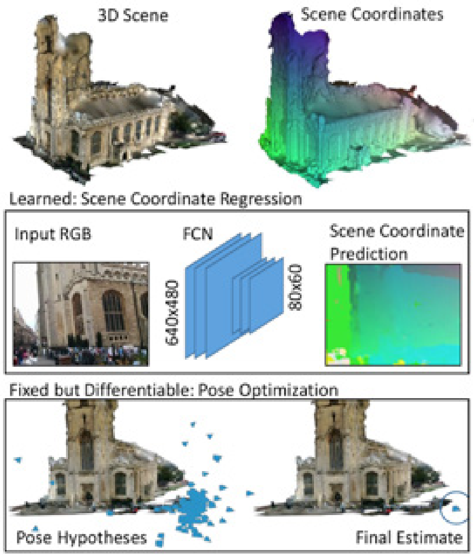

Fig. 1: Visualisation of the method of DSAC. Image taken from [1].

Meta-Learning and Our Method

Despite its advantages, DSAC still requires a large number (at least 1000) images to learn to localise a camera in even small scenes. Inspired by previous research on few-shot-learning (that is, machine learning based on only very small datasets), we decided to extend DSAC with **meta-learning** techniques. The field of meta-learning generally aims to use a network's experience on many related, but different, tasks to enable it to learn how to perform well on a new task more quickly. This is often described with the phrase "learning how to learn".

Following the technique described in [2], we split the network that predicts the world-coordinates of pixels in two: One part, the **feature extractor** stays fixed for a single scene, while the second part, a simple linear regression, is adjusted for each new scene (this can be done efficiently simply by computing the closed-form solution). The weights of the feature extractor are then adjusted after each scene so that the linear regression will do better at predicting the 3D world-coordinates of pixels.
The idea behind this is that the feature extractor will learn to recognise parts of scenes that would be god landmarks, making it easy for the linear regression to remember the 3D-world coordinates of these landmarks.

Because for each scene only a small number of weights is adjusted, fewer datapoints (i.e., fewer images of the scene) are required to memorise a scene. Due to the simplicity of the linear layer, overfitting of the network is also unlikely to be a problem.
Fig. 2 shows a diagram of our method.

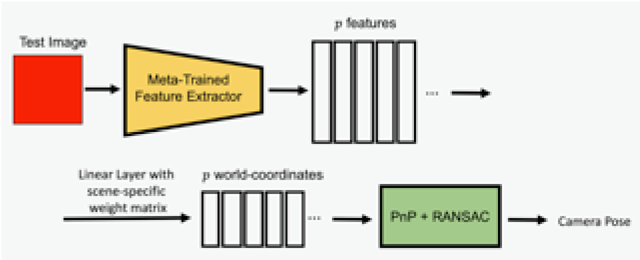

Fig. 2: Diagram of our method, which extends DSAC [1] using the meta-learning technique from [2]. The meta-trained feature extractor embeds a test image to make it easier for the linear layer to predict the scene coordinates. Using a Perspective-n-Pose algorithm in a RANSAC loop, the camera pose is then triangulated.

Training using AWS

Our method has one drawback: the training of the linear layer requires all the images from one scene to be loaded into the computer's memory at the same time. Even though we reduce the number of images required for each scene by a factor of more than 10, the memory required still exceeds that available to most GPUs. Even when distributing the training across multiple GPUs, the memory requirements remained substantial.

Here, access to cloud services from Amazon Web Services (AWS) was very valuable. The Elastic Compute Cloud (EC2) service allowed us to select the instance type with the necessary GPU capacity, which in this case was the g4dn.12xlarge instance. To safeguard against possible crashes and to be able to straightforwardly continue training after terminating the instance for a short period, we added 64GB of Elastic Block Storage (EBS) to that instance.

The g4dn.12xlarge instance gave us access to 4 high-end GPUs with 16GB of memory each for 64GB of GPU memory in total. Even when ignoring the benefit of increased GPU memory, using higher-end and more powerful GPUs on its own increased training speed, even when increasing the batch size. On the departmental server, jobs would usually run for six days to a week, while we achieved a similar number of training steps in roughly four-and-a-half days on the AWS instance, and that while using more GPU memory.

Citations

[1] Eric Brachmann and Carsten Rother. Learning less is more - 6d camera localization via 3d surface regression. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4654– 4662, 2018

[2] Luca Bertinetto, João F. Henriques, Philip Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. In International Conference on Learning Representations, 2019.

Ben Ellis – Using AWS to Analyse Elephant Rumbles

As part of my second mini project in the AIMS CDT, I investigated the communication of African Elephants. This work was supervised by Andrew Markham and Beth Mortimer.

Elephants generate a wide range of vocalisations to communicate with one another, such as snorts, trumpets, and infrasonic rumbles. Since these rumbles are so low frequency, they cannot be heard by the human ear, but can be detected either by microphones or seismometers. Some previous field work [1] had collected a large array of both seismic and microphone data over a 3-week deployment. The purpose of my project was to build a pipeline to cluster the data and try to derive some insight from that about elephant communication.

I decided to analyse the seismic data for this purpose because rumbles are very clearly visible there. However, even loading the data is quite a problem because there are 135 GB of these rumbles. Computing spectrograms on this much data is quite computationally intensive. Additionally, GPUs are required for running the neural networks in the detector, which detects elephant rumbles, and the clusterer, which clusters the resulting rumbles. I used the GPUs on AWS EC2 g4dn.4xlarge instances to train these networks. I also performed much of the code development on AWS because it was infeasible to run the loader, detector or clusterer locally. By caching the results of subsequent spectrogram computations, and using multiple loading threads, I was able to load this data relatively efficiently, reducing the time per epoch from over 100 hours to 22 minutes. I used contrastive predictive coding [2] to cluster the rumbles.

Some example clusters for the elephant data are shown below. In the data used here, the rumbles were detected by hand. There are two clear clusters here which correspond to two different seismic stations. This suggests that CPC is separating the two based on some properties of the background noise, or different set-up of the seismometer, rather than some aspect of elephant communication.

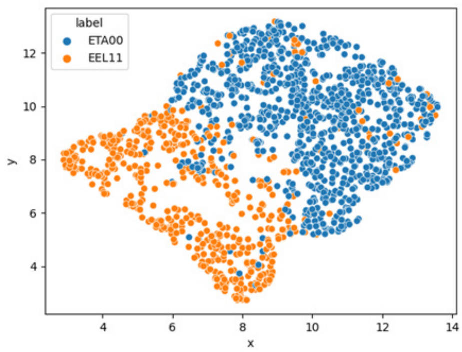

This could possibly be addressed by looking at the rumbles that were detected on multiple microphones and trying to predict the same rumble on another microphone. Tighter cropping of the rumbles would also possibly help.

We also tested clustering on some urban noises. The below image shows the clusters generated for the UrbanSound8K dataset.

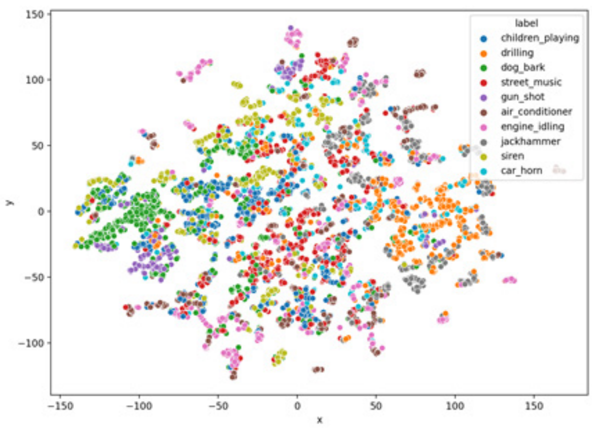

You can see that while some of the classes, such as dog barks and gun shots, are well separated, others, such as air conditioners, are not convincingly clustered at all. In general, the method seems better at clustering shorter sounds. This might be because it uses a fixed look-ahead parameter, which could be adjusted based on the input length.

Although the rumbles did not cluster in the way that was expected, the pipeline to load the data will be reusable in the future, and the clustering method demonstrated that it could easily separate the data. These will both be useful in future work to gain more insight into elephant communication.

[1] Michael Reinwald, Ben Moseley, Alexandre Szenicer, Tarje Nissen-Meyer, Sandy Oduor, Fritz Vollrath, Andrew Markham, and Beth Mortimer. 2021 Seismic localization of elephant rumbles as a monitoring approach. *J. R. Soc. Interface.* **18**. 20210264. http://doi.org/10.1098/rsif.2021.0264

[2] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2019.

Amanda Matthes – Tracking Manx Shearwaters with SnapperGPS

Last summer we wrote about how SnapperGPS was used to track nesting sea turtles in Cape Verde. SnapperGPS is an open source data logging system for wildlife tracking developed by Jonas Beuchert, Amanda Matthes and Alex Rogers. It uses a snapshot GNSS approach which has very low power consumption needs. The receiver can therefore run on very small and light batteries. This makes SnapperGPS ideal for wildlife tracking applications where low size and weight are crucial.

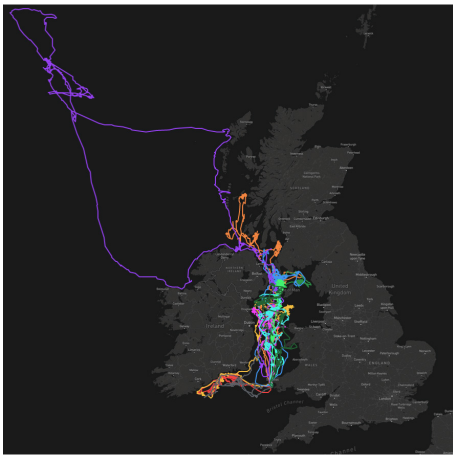

This year, I developed a modified version of the SnapperGPS hardware to support a Manx shearwater tracking project with the Oxford Navigation Group (OxNav). The Manx shearwater is a medium-sized seabird that migrates to South America but returns to the northern hemisphere for the breeding season. The UK is home to multiple large breeding colonies on islands off the coasts of Wales, Scotland and Northern Ireland. For this project, we mostly tracked nesting birds on the island of Skomer off the coast of Pembrokeshire in Wales.

This SnapperGPS version has a slightly slimmer profile to fit more easily between the wings. It also features a larger data storage as well as a charging circuit so the battery can be charged via the USB-C port. The entire tag weighs under 9 g and can take up to 20 thousand location fixes.

We tracked individuals during two important periods, egg incubation and chick rearing. In the incubation period, parents take turns to go on feeding trips for up to two weeks while the other incubates the egg. Once the chick hatches, the parents continue to alternate between foraging and staying with the chick. However, instead of taking one long feeding trip, they take shorter trips and return to their nest more frequently to feed their offspring.

An outbreak of avian influenza stopped this field work early, but before that, 19 tracks were recovered. Most foraging trips were focused on the Irish Sea but the tracks show a lot of diversity. One individual visited the northern coast of France and another one flew half-way to Iceland before returning. This highlights the importance of accessible technology for wildlife research as high sample sizes are necessary to make inferences about species behaviour.

Thank you to the Oxford Navigation Group for making this collaboration possible! Stay tuned for publication of the results.

Learn more about SnapperGPS on the project website: https://snappergps.info/
Follow the project on Twitter: https://twitter.com/SnapperGPS
Read about the hardware in this paper preprint: https://doi.org/10.48550/arXiv.2207.06310

Feedback from Students – Cohort 2021

I had an exciting 1st year in AIMS CDT. Having a chill first year allowed me to work with different labs and collaborators and get a solid understanding of what I would like to achieve during my PhD. Also, I was able to develop an organic relationship with my potential PhD advisors without being under any pressure to secure funding (since AIMS guaranteed it). Last but not the least, Wendy always had my back -- be it in figuring out the formalities or even when I didn't feel confident about my research or was struggling to find suitable PhD advisors. :)

I really enjoyed the year. The courses were well thought out and enabled us to learn a wide range of subjects I wouldn't have had exposure to otherwise. This was complemented by useful and important wider skills and professional development sessions. The small cohort-based learning this year and the CDT room was exactly what I'd hoped from the programme in how it enables us to learn and work with each other. The course is extremely well resourced and organised and helps/gives you what you need to do well in the PhD.

Highlights:

- Small group teaching and diversity of courses.
- Range of interesting mini projects on offer.
- PER was fun and I think it is important, particularly outreach activities.
- Very well structured and supported course and admin – great resources, flexible when needs to be, has your back.

In my opinion the first year of the AIMS CDT offers the best possible route into ML/AI research for early career researchers. I really enjoyed the first-year lectures and labs, and they provided a solid grounding into the different areas of the CDT. By having the opportunity to explore different areas to those I had seen at undergrad or through internships, I changed the direction of my research – this wouldn't have been possible for a regular DPhil. The mini-projects give a lot of much-needed structure to the turbulent first year of a PhD, allow you to try out different areas, work in different groups and with different supervisors. I would recommend AIMS to anyone interested in research in AI/ML, especially those applying to PhDs directly after 4-year undergrad or Masters.

The cohort, courses, and breadth of the first year of AIMS make it a perfect introduction to a DPhil. One of the greatest difficulties of traditional DPhil programs is selecting a supervisor and research topic, without having significant experience with either. From the first week of AIMS, you meet members of the program in every field, giving you an immediate perspective of the different areas of research in Oxford. Completing the courses and mini projects then provides an opportunity to meet and learn from many of the professors you may later work with, as well as real research experience. After this, I am significantly more confident in my research direction than I was before starting AIMS and feel prepared for a productive DPhil!

I enjoyed the taught courses component, in particular the modules "Data, Estimation and Inference", "Optimization" and "Robotics".

I feel a good variety of courses were provided, allowing everyone in the cohort to find their own research area. We were lucky of having many of the courses in person, or at least the labs sessions. The lab room contributed creating a strong bond within the cohort, as we had a shared space where we could go to work and help each other out. The Joint CDT Conference in June was a great moment to interact more with people from the previous cohorts and to see what kind of work the other CDTs focus on.

In AIMS I feel there are always opportunities to develop relevant research skills: for example, we present our mini projects or personal research at the CDT annual events and conferences, which are a good way to work on our presentation skills.

I am glad of being part of the AIMS program, as I can balance independent work and networking with other researchers and have professional and personal interaction with them. PhD research can be tough, so it is nice to have a cohort of people who understand what you are going through and with whom you can relax and chat when you need to.

After completing the first year of the AIMS CDT, I can confidently say that it is the perfect program for my career aspirations. During the courses, we were able to explore a variety of Machine Learning, Cybersecurity, Control, and Robotics topics, which has prepared me well for the next three years of research. The support from the AIMS CDT to develop your technical as well as professional skills is extremely valuable, from the engaging weekly seminars, scientific writing workshops, to the interdisciplinary lab and course work. The community that AIMS has built is a welcoming and inclusive space to learn with and from your peers, and I am deeply thankful for the connections I have made that will grow throughout the next years of my DPhil and beyond. I would strongly recommend AIMS to anyone interested in the program!

I'm currently in my second year of the CDT, and I really liked my first year in the programme.

I like the concept of mini-projects and I do think these helped people to narrow down their advisors in a more informed way. Most of all, I think having a cohort was nice as a D. Phil can get very isolating if we only ever meet people from the same lab, and I appreciate the different viewpoints I got exposed to because of it.

As someone who doesn't particularly enjoy taught courses, I found the weekly change of subject a bit hard to keep up with, especially alongside so many other things going on like the PER and talks and events. However, the structure does make sense given that the purpose is to allow people of different backgrounds to pick and choose what interests them (and sometimes learn something new), while maybe having the flexibility to not take every subject equally seriously.

Feedback from Courses – Cohort 2021

Signal Processing
Very interesting, found the practicals very interesting and visual, making them enjoyable Xiaowen' s lectures were excellent, and I especially enjoyed the graph sig proc stuff - I'd have liked to have learned more about this.

Data Estimation & Inference
The lab is a really good GP primer and has helped me understand a lot of the core concepts. I think it would be good to have some exposition to more practically used GP tools.

Very interesting extra content was very valuable, it's very interesting to hear Michael talk about the things he's passionate about.

Optimization
Pawan is perhaps the most considerate lecturer I have had the pleasure to be taught by. He not only went slow enough for us to follow, but constantly was checking in if we are indeed understanding and following and not just sitting quietly not understanding things. Thanks to him I finally understand several concepts which I've seen before but never made sense.

Good coverage on projected subgradient descent, labs were useful to understand implementation of optimisation techniques good coverage on projected subgradient descent, labs were useful to understand implementation of optimisation techniques.

Autonomous Robotics
This course was so much fun! I really enjoyed getting to work on a team to code an Autonomous rover. The course pushed you academically to rise to the challenge and I really enjoyed it.

The TAs were knowledgeable and eager to help.

Computer Vision
It is an interesting topic and I like they did not cover just the ML approach to it, but also more traditional CS techniques.

The lab sessions were very well done and a fantastic resource for learning about the concepts of computer vision.

The practical sessions will be a great resource for my future PhD research, I really appreciate the time that went into developing the code notebooks.

Game Theory

I think the lecturer did a great job explaining the concepts of Game Theory in an engaging way and encouraged input from the students during lectures which I appreciated.

Great lecturer/lectures, all the lab demonstrators were helpful and engaged. The problem sheets were good.

AIMS CDT Cohort 2022

NIKI AMINI-NAIENI

I received my bachelor's degree in computer science with summa cum laude distinction from Cornell University's College of Engineering in 2022. As an undergraduate, I conducted research and development at Project Kuiper, Amazon's initiative to launch a constellation of satellites into low Earth orbit (LEO) to provide internet to unserved and underserved communities around the world. I have also interned at SpaceX's Starlink constellation prior to joining Amazon. At Amazon, I contributed to two main projects. For my first project, as the first intern at Kuiper, I developed a software engine using novel mathematical techniques I invented for the rapid analysis and synthesis of phased array antenna beamforming. This novel engine is two to three orders of magnitude faster than off-the-shelf versions. For my second project, I created an intelligent system for phased array antenna calibration and measurements. For my last year as an undergraduate at Cornell, I studied abroad at Oxford in maths and computer science. At Oxford, I conducted a research project on domain generalization with the Torr Vision Group. I look forward to exploring topics in space systems, computer vision, and machine learning during my first year of AIMS. Outside of research, I enjoy kickboxing, yoga, running, healthy cooking, and spoken word poetry.

ULJAD BERDICA

I received a Bachelor of Science in Electrical Engineering from NYU in May 2022. My research experience spans across many disciplines including Computational Physics, Robotics, Electronics and Hardware Cybersecurity. My capstone project was on the implementation of the Advanced Encryption Standard on hardware and effective statistical attacks on its power side channels. I have worked on physical simulations of multi-robot mobile 3D printing, computer vision methods to streamline biomedical image annotations and used various controls and automations methods to test analog circuit designs for Electrochemical Impedance Spectroscopy. I look forward to furthering my knowledge in advanced statistics, Machine Learning, and sensor fusion to solve odometry problems in autonomous robotics. I am an avid long-distance runner and stand-up comedian with multiple performance credits in English and Albanian. I love cooking, reading poetry and all genres of heavy metal music.

JONATHAN COOK

I grew up in Leicestershire before heading down to University College London (UCL) where I graduated with a BSc in Theoretical Physics. I was then funded and mentored by the DeepMind Scholarship to study for an MSc in Computational Statistics and Machine Learning, also at UCL. My Master's research focused on the generalisation capabilities of multi-agent reinforcement learning models, but I'm also deeply interested in the collaborative AI research agenda; the problem of human-AI coordination in real-world systems; and how artificial agents can learn grounded communication strategies. I've previously had the pleasure of working on several AI for social good projects, publishing research in deep learning for online misinformation detection, natural language processing for time-varying offensive language detection and working with the United Nations to predictively map the world's offline population. Rowing, rugby, and films are a few of the ways I like to spend my free time.

SAMUEL COWARD

Before joining the AIMS CDT, I worked as a machine learning engineer investigating: GNNs, attention-based models, and reinforcement learning towards solving problems related to electronic design automation. It was during this time that I became interested in applying reinforcement learning to large action space problems. Prior to this, I received my MEng in computer science from Bristol University, with extra interest in computational neuroscience, followed by a brief stint in the games industry. In my spare time I enjoy swimming, piano, and cooking.

OISHI DEB

Prior to starting her PhD at the University of Oxford, Oishi gained industrial experience at Rolls Royce working in Software Engineering, Data Science and Machine Learning. Oishi was named 2017 Student of the Year by the President and Vice-Chancellor of the University of Leicester based on her academic performance on her undergraduate degree in Software and Electronics Engineering with Industry. Oishi was selected for the DeepMind Scholar programme, as part of which DeepMind funded her Master's in Artificial Intelligence. The DeepMind scholars were chosen based on their academic records and industrial experience. Prior to starting her PhD Oishi has presented at the NeurIPS WiML workshop and has also been involved as an ICML workshop reviewer.

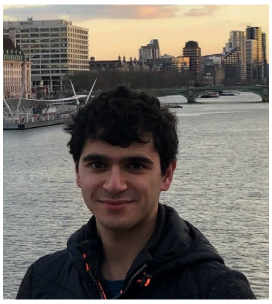

MARK EID

I grew up in Hertfordshire, and recently graduated from Imperial College London with an MEng degree in Mechanical Engineering. The robotics and computing elements have been the highlights of my course, together with the designing of an amphibious waterway cleaner robot for my 3rd year group project. As such, I am really excited to further expand on this and delve deeply into the details of many new cutting-edge fields as part of the AIMS CDT. I enjoy playing hockey and squash on a regular basis, and whenever possible I love to ski, kite/windsurf, and paraglide.

ALEXANDER GOLDIE

Having grown up in Cheshire, I graduated from the University of Oxford with an MEng in Engineering Science in 2022. Throughout my time in Oxford, I explored a range of different engineering areas; I discovered my passion for machine learning on a summer research internship at the Oxford Robotics Institute, where I considered an application of computer vision for semantic segmentation of off-road driving terrains. The skills I developed on this internship were extended in my final year project, where I researched natural language processing techniques and graph neural networks to understand the similarities of different scientific fields. I can't wait to explore the wide-ranging subject areas of AIMS, develop new passions as part of the course and develop impactful research on my PhD. Besides my background as an engineer, I love playing hockey and immersing myself in a good fantasy book.

LARS HOLDIJK

Originally from the Netherlands, I obtained my bachelor's degree in Computer Science from the University of Groningen and my master's degree in Artificial Intelligence from the University of Amsterdam. During my master's degree, I collaborated with professor Max Welling and Professor Yarin Gal as part of the ELLIS honors program. Consequently, I spend the last year of my degree at the University of Oxford. Between my bachelor's and master's degrees, my interest in AI brought me to work as a researcher in autonomous driving in the Innovation Campus at Porsche AG in Germany.

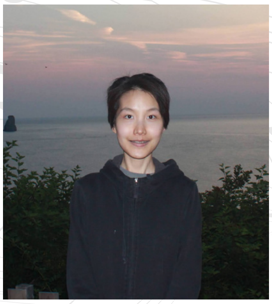

ANJUN HU

I graduated from McGill University in 2022 with a bachelor's degree in Honours Electrical Engineering. Under the supervision of Professor Tal Arbel, I have worked on several projects involving machine learning for medical image analysis. I have also interned at Ericsson Canada as an IoT developer. I'm very excited to start my graduate studies at AIMS CDT and continue exploring various probabilistic inference methods and their role in healthcare. In my spare time, I enjoy baking, gardening, and road trips

MATHIAS JACKERMEIER

I am originally from Germany and hold a bachelor's degree in Informatics from the Technical University of Munich. In my final year I worked on a project combining machine learning and verification for efficient strategy representation in controller synthesis. After that, I pursued an MSc in Computer Science at the University of Oxford with a focus on machine learning. My Master's project involved the intersection of machine learning and symbolic reasoning techniques to develop efficient methods for reasoning and prediction in large-scale knowledge graphs. During the AIMS CDT, I wish to continue to work on the intersection of machine learning with other areas of computer science, such as verification and knowledge representation. When I'm not working, you might find me playing the piano, reading a good book, or working out in the gym.

SCOTT LE ROUX

Growing up in Dublin, Ireland, I pursued my studies at Trinity College Dublin obtaining a BA in Mathematics. I also spent a year abroad at the University of Pennsylvania. During my studies, I undertook a few research projects in the areas of Computational Geometry and Topology. After this, I turned my interest to Machine Learning undertaking a summer internship with Partnership for Advanced Computing Europe and STU Bratislava. During this I researched the use of Neural Networks for Quantum Chemistry problems. I look forward to continuing my studies within the AIMS CDT. I am interested in researching topics in the growing field of Geometric Machine Learning but am also excited to discover the range of topics offered in the program. Outside of academics I enjoy reading, running, hiking, and playing or watching most sports, especially an Arsenal game.

JAKE LEVI

I'm passionate about trying to overcome the shortcomings of machine learning relative to human intelligence, particularly by applying insights from neuroscience to current machine learning methods. As a student on the AIMS CDT, I intend to develop machine learning models which learn useful representations from video using unsupervised learning, and ultimately, I'd like to extend and apply such models to challenging problems in reinforcement learning, with an emphasis on improving sample-efficiency. I have worked from 2019-2022 writing software for wireless communications at Cambridge Consultants, and previously studied at Jesus College Cambridge from 2015-2019, graduating in "Information and Computer Engineering", including a 4th year project on meta-learning with Professor Richard Turner. In my free time I enjoy climbing (trad, sport, and bouldering), hiking, watching movies, and playing classical guitar.

OLIVER SOURBUT

I grew up in Bath, England. Mathematics and computer science have always been major interests, along with philosophy, music, and biology. For several years I've worked in software R&D, most recently as a Senior Data Scientist at The Trade Desk, designing and implementing optimisation, control, and forecasting algorithms in real-time bidding. I also spend some intermittent time with the Stanford Existential Risks Initiative, researching methods for safety and beneficence of highly capable artificial systems, a difficult and important task which the research community must consider carefully. To date I have focused on the (dis)analogies between biological and artificial systems, and the nature of intelligence, goals, and learning. I am especially interested in reinforcement learning, exploration, and curiosity, and the properties of training trajectories. I also look forward to learning about a breadth of interrelated topics with the CDT! Outside research, I enjoy reading, music, exploring nature, and gaming (especially cooperative).

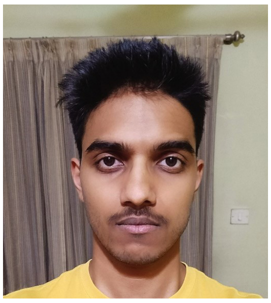

KALYAN RAMAKRISHNAN

Starting out as an aerospace engineering undergrad at IIT Madras, my varied experiences and independent study have led to unexpected research interests. Today, I want to pursue research in machine learning, computer vision, and ways to take inspiration from human intelligence to improve artificially intelligent systems. I am also keen on exploring the theoretical foundations of machine learning and devising algorithms that humans can easily interpret. Given my wide-ranging interests, the AIMS CDT is ideal since its preparatory phase would allow the time and freedom to formulate specific research questions for my dissertation. I look forward to learning from the world-renowned faculty members at Oxford and making an impactful research contribution.

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.

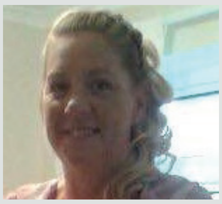

WENDY POOLE

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/

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
Tel: 01865 616622

CDT Co-Director
Alex Rogers
Email: aims-cdt@robots.ox.ac.uk

CDT Centre Administrator
Wendy Poole
Email: aims-cdt@robots.ox.ac.uk
Tel: 01865 616632

www.aims.robots.ox.ac.uk

AIMS
EPSRC Centre for Doctoral Training in
Autonomous Intelligent Machines & Systems