AIMS

EPSRC Centre for Doctoral Training in
Autonomous Intelligent Machines & Systems

Annual Review 2018/19
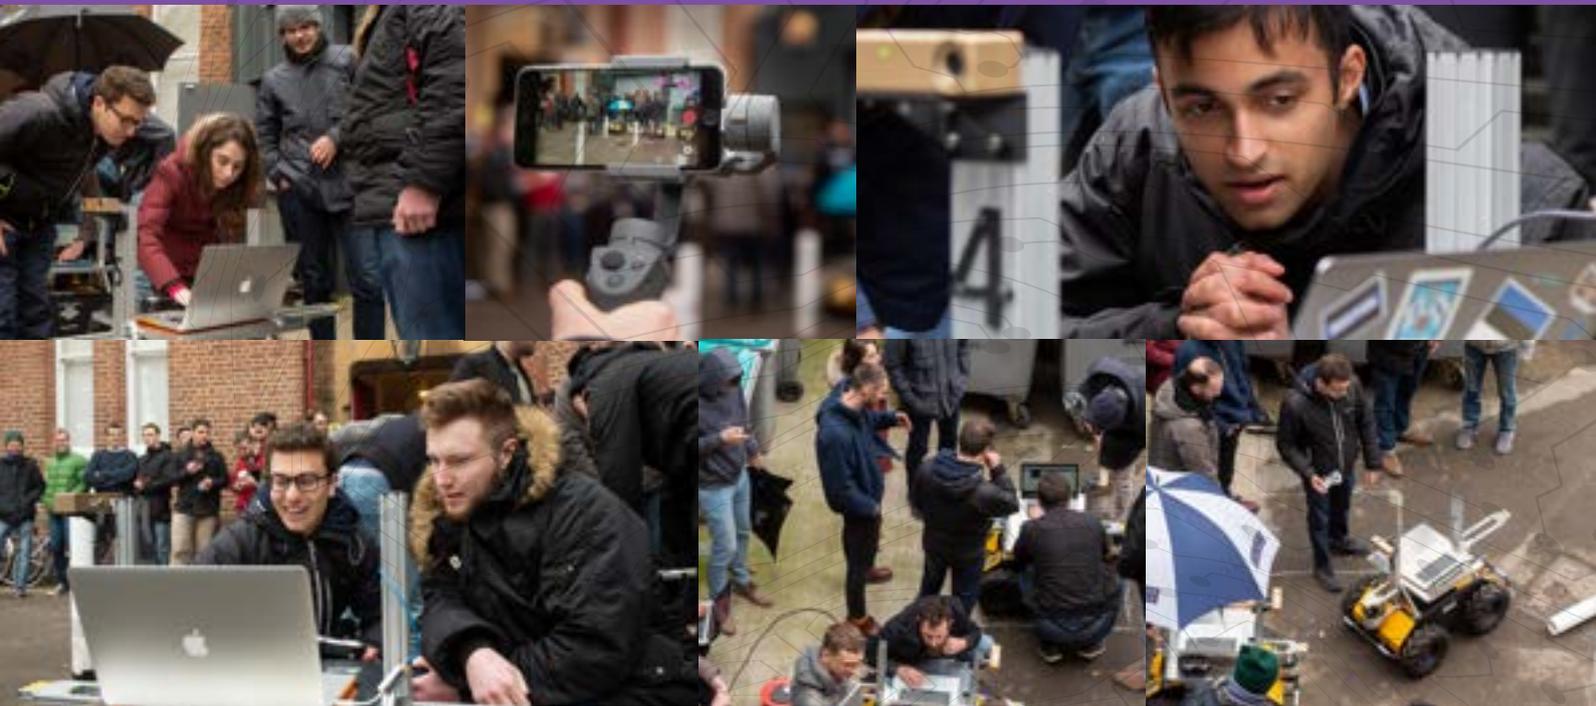
Siddhant Gangapurwala – Bhagyalaxmi Sarla Guru Sahani Memorial School, Aurangabad

Foreword

Welcome to the fifth annual review highlighting key aspects and activities of staff and students in AIMS during 2018/2019. This has been the fifth full year of the EPSRC Centre for Do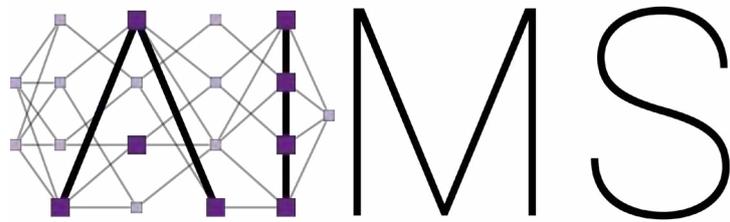ctoral Training in Autonomous Intelligent Machines & Systems. Following intakes in 2014, 2015, 2016, 2017 and now 2018, the Centre now has 62 full time Ph.D. students engaged in the four year Ph.D. programme. It's been another busy year, but again an enormously rewarding one. The CDT is thriving with applications escalating at a pace, along with more Industry partners joining the CDT.

We would like to thank everyone with their support with the renewal bid, this was successful, and the CDT has now been refunded for a further 8.5 years.

In 2018 we received over 190 applications for the CDT programme. We have recruited 17 students, eleven of whom are fully or half funded by EPSRC, six who are either fully or partially funded through Industry or University scholarships.

This year has seen our fourth cohort transfer from PRS (Probationary Research Student) to full DPhil Status and move in to their third year, with several of them producing some outstanding papers submitted to conferences and journals. Our third years confirming their status, and our first cohort submitting their thesis, and have now become our first AIMS alumni. These students have gone to work at Google DeepMind, a start-up company in London, Shell, and also academia. Further details can be found at: www.aims.robots.ox.ac.uk/alumni

The current first year students have now completed their training year and now move in to the substantive research phase of their PhDs.

All students have acted as ambassadors for the CDT at events joint with other CDTs and outreach events and have set a bar of excellence. It is truly an exciting time to be involved in autonomous and intelligent systems in the UK.

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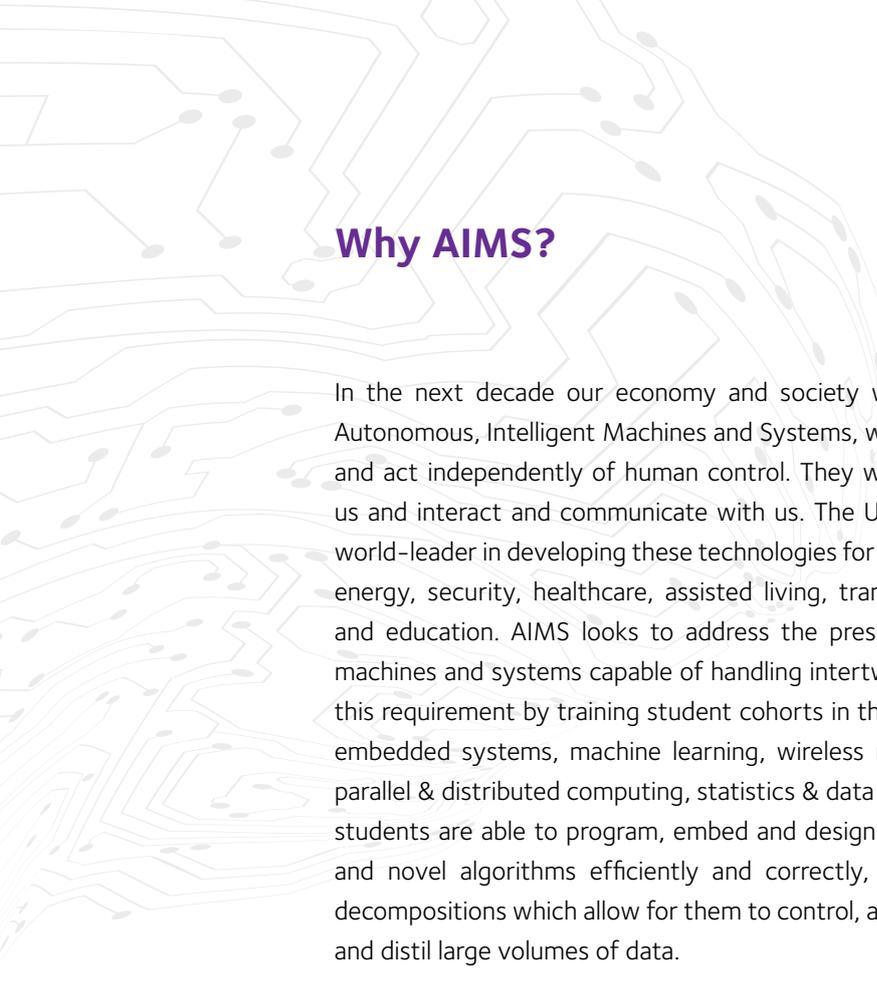
Mike Osborne
Acting Director

Alex Rogers
Acting Co-Director

Wendy Poole
Centre Administrator

Why AIMS?

In the next decade our economy and society will be revolutionised by ubiquitous Autonomous, Intelligent Machines and Systems, which can learn, adapt, take decisions and act independently of human control. They will work for us and beside us, assist us and interact and communicate with us. The UK has the opportunity to become a world-leader in developing these technologies for sectors as diverse as manufacturing, energy, security, healthcare, assisted living, transport, environment, entertainment and education. AIMS looks to address the present need for smarter, more useful, machines and systems capable of handling intertwined heterogeneous data. We meet this requirement by training student cohorts in the underpinning sciences of robotics, embedded systems, machine learning, wireless networks, control, computer vision, parallel & distributed computing, statistics & data analysis, design and verification. Our students are able to program, embed and design software, to implement established and novel algorithms efficiently and correctly, to develop and apply models and decompositions which allow for them to control, access, leverage, learn from, interpret and distil large volumes of data.

Our research themes

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

Robotics, Vision and Perception: The first key skills area is in enabling autonomous systems to identify and interpret complex scenes, from moving vehicles to human activity and form robust situation assessments to enable appropriate action and decision making. For example, robotic systems require such capabilities so that they can navigate in unknown environments; augmented reality systems require methods for scene perception and object identification. Our vision is to train a new generation of researchers that will be able to understand and embed such intelligent machines across sectors, from home health care to driver-less cars. Such applications are particularly challenging because they require autonomous systems to operate in environments that are inherently unpredictable, continually changing, and impossible to directly model. We infuse expertise in Robotics, Vision and Perception in a unique educational curriculum that cuts across theoretical developments in vision and robotics, scene understanding and perception and state-of-the-art systems research in mobile robot autonomy, navigation and mapping.

Machine Intelligence & Multi-Agent Systems: The second key skills area is in making machine autonomy and intelligence ubiquitous; allowing machines to discreetly pervade the world around us and assist us. Our students are equipped to answer questions like "how can we make machines part of our daily lives without having to continually give them instructions, maintain, repair and look after them?" and "how can machines increasingly learn our objectives, sense our frustration, and help us achieve our goals with minimum interference?" With strong multi-disciplinary expertise in the areas of artificial intelligence, machine learning, crowd-sourcing, participatory systems, language understanding, scalable inference, decentralised information systems, agent-based computing and game theory, the CDT promotes a training foundation for students to inject machine intelligence into real-world applications, such as the critical domains of healthcare, smart grids and energy resources, big data analytics, disaster response, citizen science, human-in-the-loop systems and the environment.

Control & Verification: Our third skills areas lies in developing effective techniques to monitor and control intelligent machines, such as those used in manufacturing, transportation and biosensing/healthcare systems, and to ensure their safety and dependability. For example, how do we ensure that the embedded software controller of the self-driving car does not crash, or that the implantable blood glucose monitor correctly identifies an abnormal range and raises an alarm? Verification via model checking provides automated methods to establish that given requirements are satisfied, but is challenged by the need to consider the complex interplay of discrete, continuous and probabilistic dynamics. This problem is exacerbated in the context of multi-agent systems interacting in uncertain environments. Although there are many new results in the emerging area of hybrid and probabilistic systems, there is a clear gap in developing computational tools that make use of solid theoretical foundations to solve practical problems. Our CDT combines robust control methods with approximate computation methods in stochastic hybrid systems and symbolic model checking & synthesis of embedded software.

Machine-to-Machine (M2M), Secure Sensing & Actuation: The fourth skills area underpins the vision of connecting intelligent devices seamlessly, allowing them to share their sensing, monitoring and actuating capabilities. This is often referred to as "M2M" or the "Internet of Things". Although this vision is not new, there are key technical barriers in the widespread adoption of "intelligent networked" devices. First, machine interaction typically relies on context-awareness (e.g. location) which is problematic in indoor environments. Second, sensors and actuators are inherently unreliable, often lacking calibration, quality estimation, energy management and fault detection capabilities. This compromises their practical use. Third, most M2M solutions have been designed to meet functional requirements, ignoring security and privacy
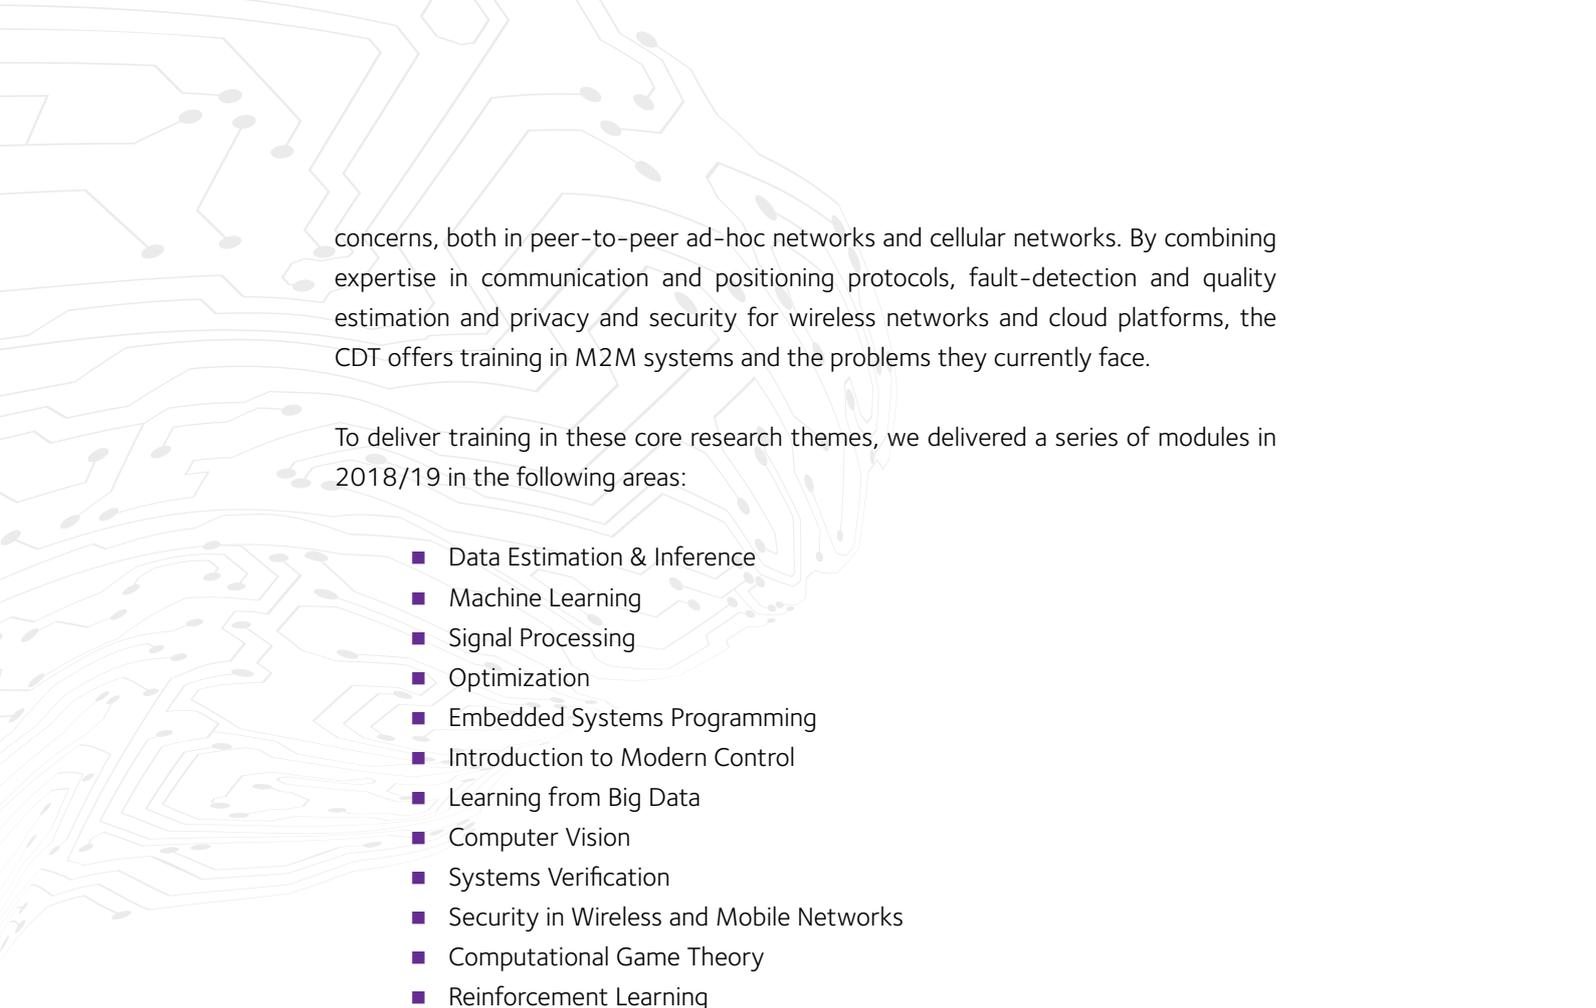
concerns, both in peer-to-peer ad-hoc networks and cellular networks. By combining expertise in communication and positioning protocols, fault-detection and quality estimation and privacy and security for wireless networks and cloud platforms, the CDT offers training in M2M systems and the problems they currently face.

To deliver training in these core research themes, we delivered a series of modules in 2018/19 in the following areas:

- Data Estimation & Inference
- Machine Learning
- Signal Processing
- Optimization
- Embedded Systems Programming
- Introduction to Modern Control
- Learning from Big Data
- Computer Vision
- Systems Verification
- Security in Wireless and Mobile Networks
- Computational Game Theory
- Reinforcement Learning
- Dynamic Robot Locomotion and Motion
- Sensor & Actuator Networks
- Mobile Robotics

Events, highlights & outreach

AIMS students have taken part in a wide range of research and outreach events this year; these include internships, summer schools, outreach and a large number of publications to top machine learning conferences. These these include: AAAI, ICML, NeurIps, ICLR, Interspeech, CVPR, AAMAS, ICRA and AISTATS.

- A number of AIMS students attended the annual joint CDT conference in Edinburgh with CDT's in Future Autonomous and Robotics Systems (FARSCOPE) and Robotics and Autonomous Systems (RAS). There were two keynote speakers; Dr Ali Eslami (Google DeepMind) and Dr Margarite Chli (ETZ Zurich), with a number of interesting talks from all three CDT's.

- Bryn Elesedy – In July I spent two great weeks at the Machine Learning Summer School (MLSS) in London, organised by Marc Deisenroth and Arthur Gretton. The summer school consisted of intensive courses that covered a broad range of topics in machine learning and related fields, such as optimisation and submodularity.

MLSS was not only a great opportunity to learn from the best through the lectures and tutorial sessions, but also a chance to meet students from all over the world working on the whole spectrum of problems in machine learning and its applications. It was an enriching experience and I would recommend anyone interested in MLSS to apply for the next course.

■ Thomas Steeples – During the summer of 2019, I was able to attend the 31st occasion of the European Summer School in Logic, Language and Information (ESSLLI).

This takes place in a different European country each year, and was at the University of Latvia in Riga this year. It was an intensive, two week summer school, with lectures from 09:00 to 18:30 each day, with special evening lectures afterwards. The lectures spanned across three tracks: Language and Logic, Language and Computation and Logic and Computation. Given my research, I mainly attended lectures from the Logic and Computation track but in the interest of breadth, I did dip into the other tracks.


The Logic and Computation lectures explored exciting new areas of modern logic research that I hadn't been exposed to before. Highlights for me included a course on reasoning about knowledge and resources in games, and another course on responsibility, that is, can we attribute certain outcomes to particular agents? – An important question in this AI age.

The summer school was worthwhile just by the virtue of it introducing me to things I was previously unaware of.

On a softer note, Riga was a beautiful city to explore, with wonderful food and drink. There were also plenty of opportunities to network and get to know fellow logic PhDs from across a range of disciplines from around the world. I would highly recommend doing a summer school to any DPhil student considering it!"
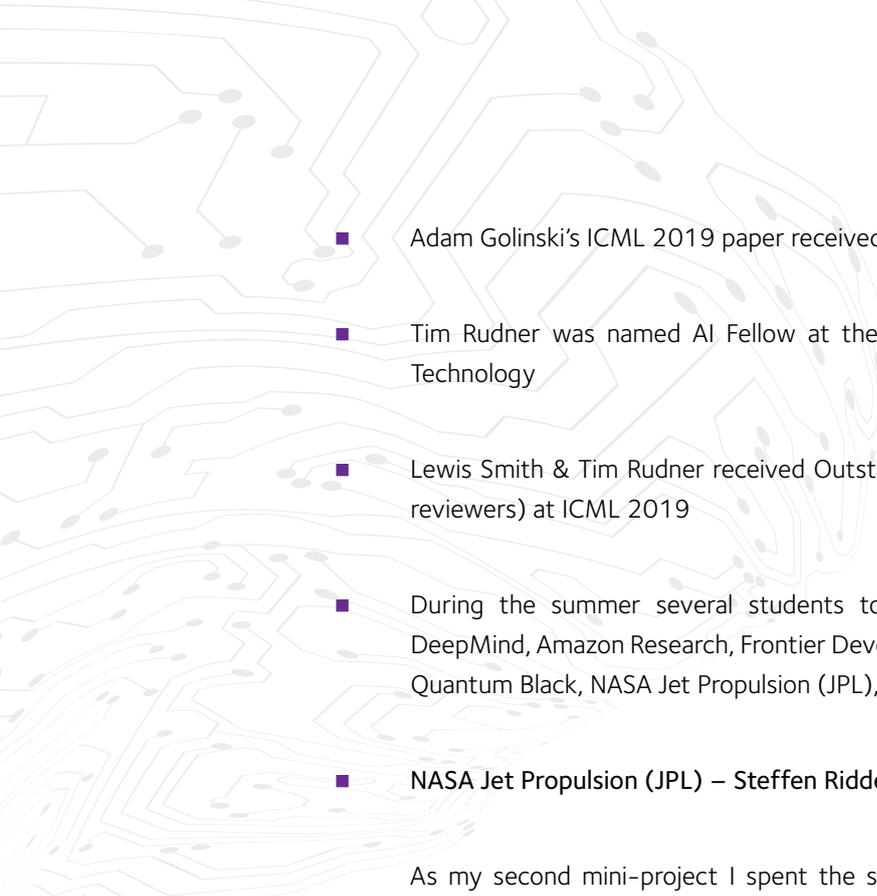
- Adam Golinski's ICML 2019 paper received a Best Paper Honourable Mention.

- Tim Rudner was named AI Fellow at the Centre for Security and Emerging Technology

- Lewis Smith & Tim Rudner received Outstanding Reviewer Awards (top 5% of reviewers) at ICML 2019

- During the summer several students took up internships, these were at DeepMind, Amazon Research, Frontier Development Lab, Facebook AI Research, Quantum Black, NASA Jet Propulsion (JPL), Nvidia and Uber.

- NASA Jet Propulsion (JPL) – Steffen Ridderbusch

As my second mini-project I spent the summer at the NASA Jet Propulsion Lab (JPL) in Pasadena. My task was to simplify the limb design for a wheeled robot with active suspension with the goal of developing a developing a mobility platform for Ocean Worlds. The moons Europa and Enceladus are thought to contain vast amounts of water with a chance to find live, but are expected to be much harder to traverse than Mars or the Moon. My mentors at JPL used the highly capable RoboSimian platform in terrain in Alaska and Death Valley thought to be analogous to Ocean Worlds to identify the necessary capabilities and mobility modes. Based on this I have developed a framework to optimize different limb prototype based aspects like energy efficiency or manipulability while achieving the low complexity. From there, I have proposed a number of candidate designs, which are planned to be prototyped and tested to form a proposal for a next-generation extra-terrestrial mobility platform.

Frontier Development Lab – Lewis Smith

I took part in the 2019 FDL Europe program as part of the disaster prevention and response team, working on flood detection and response. Satellite imagery is extremely important and useful in flood disaster response, but there are limitations to existing infrastructure. One of the most fundamental is the revisit time - a satellite in a polar orbit may only be able to image a particular area once every few days, depending on its latitude. For example, ESA's sentinel mission has a worst case revisit time of 5 days. The most obvious way to reduce the revisit time is to have a larger constellation of satellites, but this has obvious implications for the cost of a mission. Again to take sentinel as an example, the pair of satellites cost €350 million. One solution to this problem is to look into the potential of 'cubesats', smaller satellites made out of cheaper, commercial components and assembled out of fixed size 'units' of 10x10x10cm. While originating in academia as an inexpensive way to demonstrate new technology, they have recently been taken more seriously as a commercial or scientific proposition.

You could feasibly launch a fairly large constellation of cubesats for the cost of a single sentinel, massively improving the revisit time. However, cubesats have their own challenges. There is an obvious trade of in the quality of imaging that can be produced relative to a more expensive satellite like the sentinels, but there are also issues with communication and battery life. Scheduling downlink communications with a lot of satellites simultaneously can be a challenge, and the small size of cubesats means they cannot transmit continuously for long duoto their battery life. As a result, reducing the amount of data that a cubesat constellation needs to send to the ground would be an important step in making this proposal practical.

The solution to this that our project explored is carrying processing out on board the satellite itself. The typical end product that a user in the field (say a government or international aid agency, like UNICEF) is after is a map of an area with annotated flood extent. This can be a segmented image, say a single channel image with a 3 bit depth (water, land, cloud). The image from the original camera can be much larger- on the satellite that this project targeted, the camera is a hyperspectral one with 50 input bands, each with 16 bit depth, so the processed image is almost 1000 times smaller than the raw data product, meaning that there is a significant saving in data downloaded if it can be processed on board and only the end product downloaded to the ground. While the on-board processing does in theory increase battery usage, this can be mitigated by taking advantage of dedicated hardware, such as the intel myriad chip, which can efficiently process images with neural networks. These are cheap, lightweight and commercially available, and so can easily be included in future
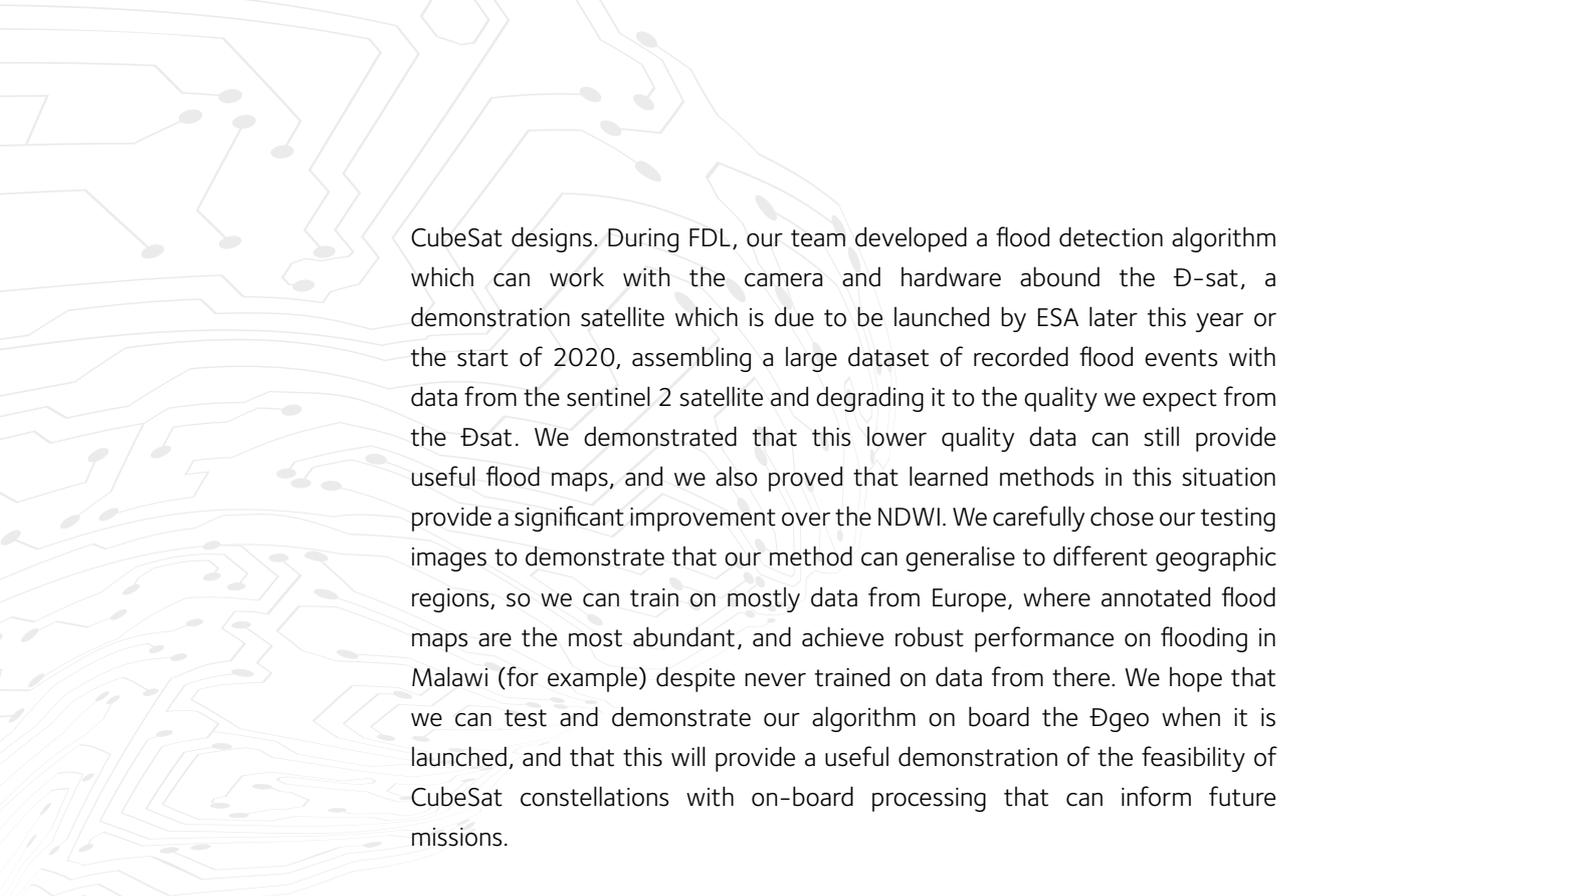
CubeSat designs. During FDL, our team developed a flood detection algorithm which can work with the camera and hardware abound the Đ-sat, a demonstration satellite which is due to be launched by ESA later this year or the start of 2020, assembling a large dataset of recorded flood events with data from the sentinel 2 satellite and degrading it to the quality we expect from the Đsat. We demonstrated that this lower quality data can still provide useful flood maps, and we also proved that learned methods in this situation provide a significant improvement over the NDWI. We carefully chose our testing images to demonstrate that our method can generalise to different geographic regions, so we can train on mostly data from Europe, where annotated flood maps are the most abundant, and achieve robust performance on flooding in Malawi (for example) despite never trained on data from there. We hope that we can test and demonstrate our algorithm on board the Đgeo when it is launched, and that this will provide a useful demonstration of the feasibility of CubeSat constellations with on-board processing that can inform future missions.

- Filip Svoboda - was the course lead for the Introduction to Machine Learning masterclass: http://www.cs.ox.ac.uk/conferences/RIMC/index.html

This is a program for high school students interested in studying at Oxford CS. My module aimed to introduce them to the foundation concepts in machine learning, with a particular focus on deep models. After a short intro lecture, they worked with deep dream, style transfer, and with a keyword recognition demo. That is they saw what different parts of a model are sensitive to in the input, how this can be used to transfer style and content between images, and finally they saw how different input modalities can be accounted for.

I also took part in the Bodleian library event. This is a summer school which runs at the Visiting Scholars Centre of the Bodleian. Across the summer school, students undertook various creative writing activities with the Bodleian's collections, looking at material that shows how the different ways that people in the past have envisioned 'the future'

- Adam Cobb – In June 2019, I gave a talk at NASA Goddard entitled "Using machine learning in the astrophysical sciences". There was a fair bit of interest with around 150 people attending. I based the talk on my recent paper, which used machine learning techniques to retrieve atmospheric parameters for exoplanets. Through using this example, I introduced the audience to the process that is often required when applying machine learning techniques in the astrophysical sciences. My motivation was to encourage the scientists in the audience to see parallels in their own work and possibly use similar techniques in the future. The talk was well-received, which led to many interesting questions and conversations afterwards.

I gave a similar talk at the University of Maryland, where the focus was more directed at how members of the Planetary Group could use my code directly in their research. In fact, there were already a few PhD students who were building on our work and I met with them to see if I could help with their research.

- **Robotics team showcase robot's "humansupport" capabilities at Sydney's RoboCup contest**

This July, a team from the Oxford Robotics Institute (ORI) competed in their first RoboCup, an annual international robotics competition held in Sydney, Australia. With 13 different leagues attracting teams from around the globe, it is one of the world's largest robotics events.

Comprising Engineering undergraduates, DPhil candidates and postdocs, Team ORIon previously competed at last year's World Robot Summit in Tokyo, Japan. They spent the intervening months developing and refining their Toyota Human Support Robot (HSR) to give themselves the best possible chance of success at RoboCup.

AIMS student Mark Finean, who led the nine-person team, explains: "We competed in the '@Home' league, which requires the robots to assist humans in household environments such as a living room or a kitchen. The tasks included storing groceries, carrying luggage and taking out the garbage.

Team ORIon received a round of applause from the audience when their HSR succeeded in autonomously taking out the garbage. "This was a highlight of the competition for us," Mark reflects. "Many teams stumbled at this final hurdle as the garbage bag would get caught on the robot's gripper. This was no problem for us as we implemented functionality to detect whether the bag could still be felt after letting go. If the bag could be detected then the robot would shake its gripper appropriately to remove it."

After an intense week of programming, testing, and integrating, Team ORIon finished in 6th place out of 10 with a total score of 600, missing out on the semi-finals by just one place. The team earned points for successfully taking out the garbage and for an impressive poster and presentation session delivered by their undergraduate members. Mark says: "This is just the beginning of our journey and we will definitely be back for more next year!"

"To excel required a combination of all of the robot's capabilities, including navigation, task-level planning, manipulation, voice recognition, and vision."

- Bradley Gramm-Hansen's Machine Learning Skills Help NASA to map settlements

  Currently, nearly one quarter of the world's urban populationlives in informal settlements or encampments, most in developing countries but increasingly also in the most affluent ones

  Largely ignored by governments, featuring little to no infrastructure, and home to startling levels of inequality, informal settlements are home to some of the poorest people on the planet.

  Efforts to improve living standards within these settlements are hampered by a lack of information on them. Their sizes and locations are often unknown, and with no security of residence, housing is unlikely to comply with building, planning and safety regulations.

  For this reason, 2018 saw the Frontier Development Lab (FDL) announce their goal of mapping informal settlements. The FDL is a collaboration between NASA, the European Space Agency and private enterprise, and they unveiled this project as part of their aim to create a 'Mission Control for Planet Earth', using satellites to make a difference to people on the ground.

He explains: "Mapping informal settlements is vital for providing local government and non-governmental organisations with geographic locations of informal settlements, allowing the delivery of targeted economic and social aid.

"However, the process of generating these maps is currently very costly; for example it costs approximately $10,000 to obtain a high-resolution image the size of the urban area of Nairobi."

The team knew that this problem could be surmounted with the help of machine learning, making it possible even for cash-strapped NGOs. "For the first time," says Bradley, "we can make use of freely available global low-resolution satellite imagery, applying machine learning to predict and generate maps providing the locations of these settlements."

Bradley worked on the project alongside two other Oxford researchers, Yarin Gal and Piotr Bilinski (now based at the University of Warsaw). "In terms of my contribution, I helped to determine and implement the methodologies that could be utilised to solve this problem. In regards to the classification algorithms used, Canonical Correlation Forests and using the spectral information contained within the ESA satellite images to extract information pertaining to the types of materials in each pixel."

The project involved researchers from institutions around the world. They spent a week at the ESA Earth Observation headquarters in Esrin, Italy, where they were introduced to Do-Hyung Kim of UNICEF. "He really opened our eyes to the problems that they are trying to solve," says Bradley. "In addition to this, Satellite Applications Catapult were able to facilitate our access to very-high resolution data via Digital Globe and Nvidia provided FDL Europe with $1 million of computing, which was shared between the two teams that worked on separate projects over the 8 weeks. Google also provided additional cloud services."
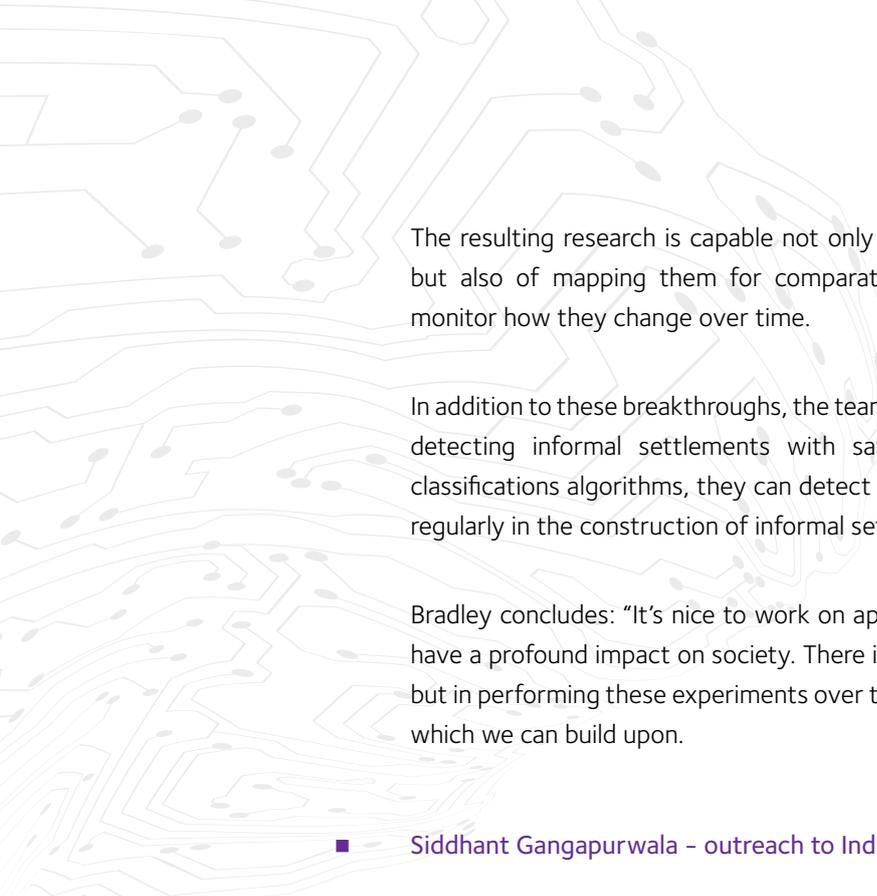
The resulting research is capable not only of identifying informal settlements, but also of mapping them for comparative purposes, allowing agencies to monitor how they change over time.

In addition to these breakthroughs, the team demonstrated a second method for detecting informal settlements with satellite images. By applying simple classifications algorithms, they can detect materials that are known to be used regularly in the construction of informal settlements.

Bradley concludes: "It's nice to work on applications that have the potential to have a profound impact on society. There is still an awful lot of work left to do, but in performing these experiments over the summer, we have provided a base which we can build upon.

■ Siddhant Gangapurwala – outreach to India

During a visit to India, I took part in some outreach events where I attended a number of deprived schools. I gave talks on various research topics related to my studies.

Holy Crescent English High School, Aurangabad – Lectured students of the 7th, 8th, 9th and 10th grade, with about 60 students in each grade, on a general topic of Artificial Intelligence (AI) and Robotics. I briefly described AI, focussing on the mathematical relevance of the field, and described its use in some applications that the students were familiar with such as Google Translate, FaceApp, Siri etc. I then presented research from the Reinforcement Learning (RL) community in the form of videos of RL agents trained for various tasks such as locomotion. Following this, I talked about Robotics and briefly explained my research. The Q&A session was quite long and lasted for about 45 minutes with about 30 questions answered. Surprisingly, the questions asked by students were quite technical even for 10th graders.

Maharashtra Public School, Aurangabad – Lectured students of the 8th, 9th and 10th grade, with about 50 students in each grade. The lecture was similar to the one for Holy Crescent English High School, with slight difference in the form of presentation. Maharashtra Public School provided a projector for computer presentation which made it easier to show the students multimedia content, making the lecture quite interactive.

Bhagyalaxmi Sarla Guru Sahani Memorial School, Aurangabad – Lectured students of the 9th and 10th grade, with about 50 students in each grade. The lecture was similar to the one for Maharashtra Public School.

Dwarkadas J. Sanghvi College of Engineering, Mumbai - Discussed with about 70 undergraduate students of years 1 to 4 from the Computer Science, Electronics and Biotechnology departments, the theory of Machine Learning and Robotics, detailing upon the AIMS research themes. The session was highly interactive and technical.

- **Can we track elephants using the vibrations of their footsteps?**

Ben Moseley, Alexandre Szenicer (Earth Sciences, Dphil), Paula Koelemeijer (Earth Sciences, research fellow (Oxford & UCL), Michael Reinwald (Zoology, post-doc), Tarje Nissen-Meyer (Earth Sciences, Associate Professor), Beth Mortimer (Zoology, Research Fellow)

We went to Kenya to see if it is possible to track elephants using the vibrations they make through the ground. The outputs of this unique project could provide better tools for conservation and new insights into their behavior.

Elephants are fantastically intelligent and social animals, yet many are endangered and living with threats such as poaching, habitat destruction and conflicts with local people. Conservation is crucial but monitoring these animals in real-time and over large areas is expensive and technologically challenging.

Furthermore, there is still much to discover about how elephants interact with others and their environment. They are socially complex animals and they communicate in a variety of different ways. Our recent Ted Talk showed how they use vibrations through the ground to communicate: https://www.ted.com/talks/beth_mortimer_and_tarje_nissen_meyer tedwomen_2018

There are many different types of elephant behavior which generate ground vibrations. These include running, playing, bathing and walking. Elephants also make vocalizations, such as trumpets and "rumbles", which are known to propagate vibrations into the ground.
In this project we have been researching whether it is possible to track elephants by detecting these vibrations. A key benefit of using ground vibrations is that we think they can travel long distances, potentially allowing us to track the elephants at low cost, non-invasively and over large areas.

To carry out our research we traveled to the Mpala research center in Nanyuki, Kenya in February 2019. This center is home to nearly 7,000 elephants and stretches over 48,000 acres of savanna, bushlands and rivers, providing an excellent place to study elephants in their natural habitat.

To detect vibrations, we used seismometers. These are low-powered, buried instruments which can detect very small movements of the ground. We deployed 20 of them, covering an area of nearly 50 square kilometers. We also deployed camera traps to help us identify elephant behaviors. We mounted cameras on poles and trees which took photographs of any wildlife that moved across their field of view.

We used seismometers buried in the ground to detect the vibrations made by elephants. This image shows us installing a seismometer.

Over three weeks of field work we recorded over 10,000 hours of seismic recordings and 50,000 images and videos. We observed a variety of different elephant behaviors, including running, playing, bathing, "rumbles" and fighting

We are currently in the data analysis stage, where we are testing many different detection and localization methods on this dataset. With the camera images as a guide, our goal is to use AI to detect elephants from the recorded seismic signals.

Once we are able to detect elephants, we will use triangulation methods to locate and track them.

Alongside tracking our research could provide new insights into elephant behavior, as well as their long-distance communication.

This work is funded by National Geographic and the John Fell Fund. It is a cross disciplinary effort between the Earth Science, Computer Science and Zoology departments.

- **Vitaly Kurin – Improving SAT Solver Heuristics with Graph Networks and Reinforcement Learning**

As part of a second mini-project I spent the summer at Nvidia working on Boolean satisfiability (SAT).

Boolean satisfiability (SAT) is an important problem both for the industry and academia. It impacts computer security, circuit design, theorem proving, formal verification, and combinatorial optimisation, to name a few. As a result, modern SAT solvers are highly sophisticated, reliable and well-tuned pieces of software which can solve extremely large problems.

To solve a SAT problem means to find such assignments of Boolean variables which make a given Boolean formula true or to prove that such assignment does not exist.

SAT solvers rely on multiple heuristics designed with a lot of scrupulous work, often after a long trial-and-error process. Can machine learning help humans design better heuristics? To answer this question, Vitaly Kurin, a DPhil student (part of the Department's Centre for Doctoral Training in Autonomous Intelligent Machines and Systems, which is training a world-class cohort of researchers in the theory and practice of a new generation of autonomous intelligent machines and systems), spent summer 2019 in Santa Clara, California, interning at NVIDIA, an AIMS CDT industrial partner.

The results he and his colleagues managed to achieve are promising. They replaced the branching heuristic in a SAT solver thus keeping it complete, i.e. always solving a SAT problem correctly. The improved heuristics called Graph Q-SAT (GQSAT) is able to reduce the number of iterations of the algorithm more than by a half. Moreover, it is able to generalise from satisfiable instances to unsatisfiable instances, a problem, with what several previous machine learning approaches struggled. Apart from that, GQSAT is able to generalise to problems up to five times larger than those it was trained on.
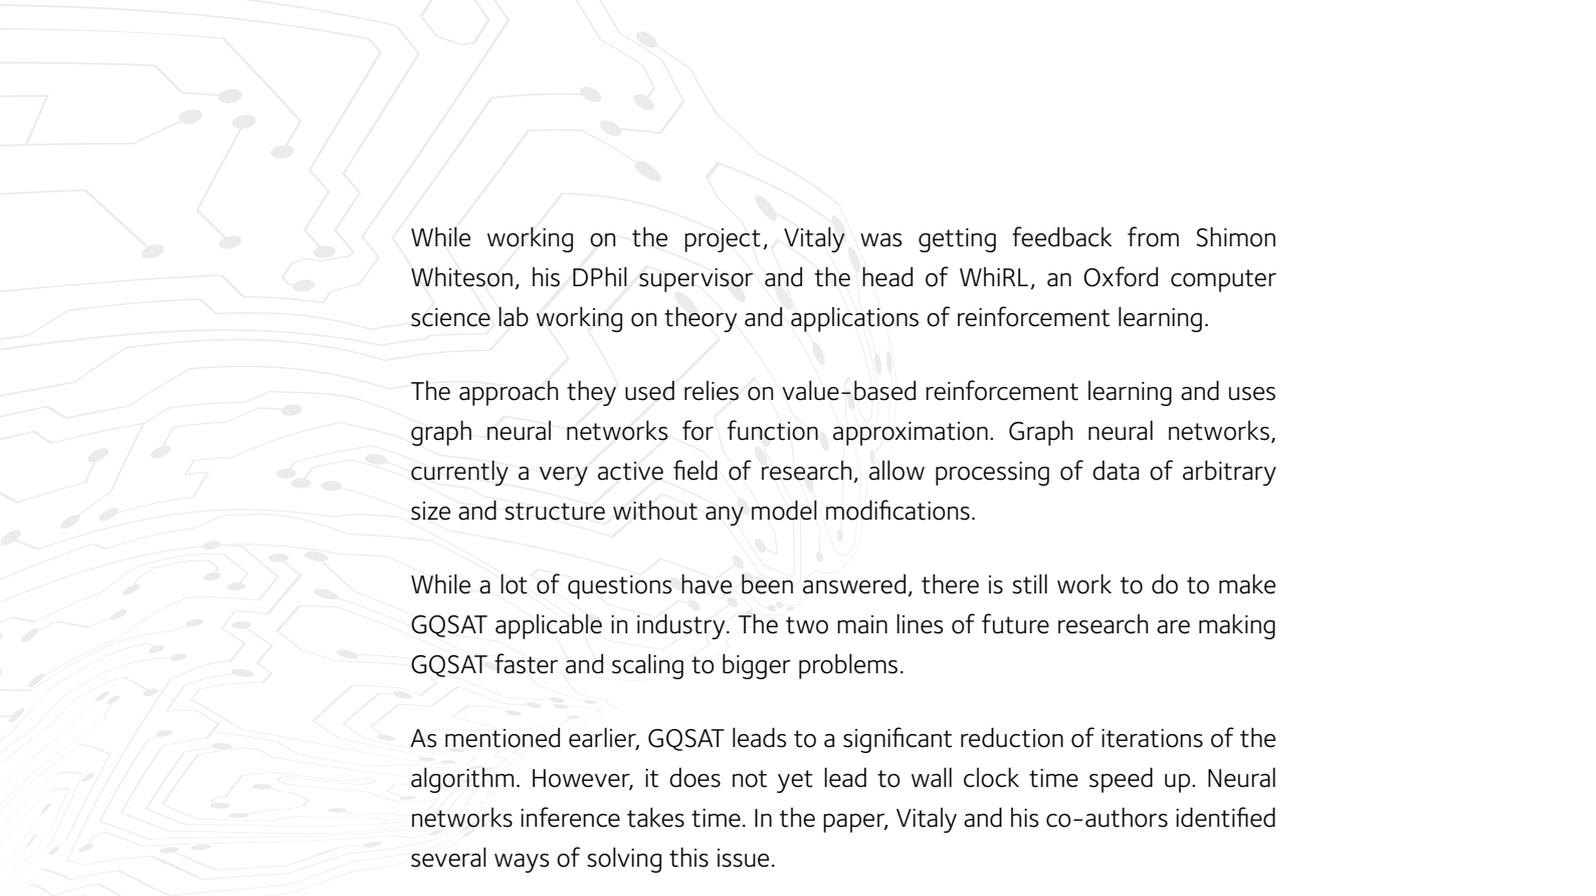
While working on the project, Vitaly was getting feedback from Shimon Whiteson, his DPhil supervisor and the head of WhiRL, an Oxford computer science lab working on theory and applications of reinforcement learning.

The approach they used relies on value-based reinforcement learning and uses graph neural networks for function approximation. Graph neural networks, currently a very active field of research, allow processing of data of arbitrary size and structure without any model modifications.

While a lot of questions have been answered, there is still work to do to make GQSAT applicable in industry. The two main lines of future research are making GQSAT faster and scaling to bigger problems.

As mentioned earlier, GQSAT leads to a significant reduction of iterations of the algorithm. However, it does not yet lead to wall clock time speed up. Neural networks inference takes time. In the paper, Vitaly and his co-authors identified several ways of solving this issue.

The second problem is the most challenging one. Scaling to larger problems will require much more research tackling both fundamental machine learning problems and implementation challenges.

The results achieved within such a limited time budget are very encouraging, providing a base for further improvements. Vitaly and his colleagues hope that GQSAT is a stepping stone providing grounds for the next generation of SAT solvers leveraging data to design new heuristics automatically.

The paper "Improving SAT solver heuristics with Graph Networks and Reinforcement Learning" is currently under review at ICLR 2020. The preprint is available on Arxiv (https://arxiv.org/abs/1909.11830). The implementation will follow shortly.

Publications & Output

AIMS students have had another successful year so far in both submitting and having papers accepted for this year's big machine learning, computer vision and programming conferences, such as: NeurIPS, ICML, ICLR, ECML, KDD, POPL, AISTATS, AAAI and ICCV.

An updated list can be found on our web pages at: http://aims.robots.ox.ac.uk/publications

DPhil Proposals
Cohort 2018 will now work on their PhDs. A description of their research proposals can be found at: http://aims.robots.ox.ac.uk/research/dphil-projects/

Feedback from Students – Cohort 2018

The first year gave a wide exposure to the state of the art across the most pressing topics in AI and ML. It is hard to understand how advantageous it is to spend a year studying the latest developments across the field with the lecturers that are actively contributing to them before making the all-important choice of supervisor and topic specification. Beyond this, it also formed a theoretical and empirical solid foundation on which I started building my core DPhil project. In short it lets you to make a well informed choice about your research topic and to hit the ground running once you transition to the second year.

The first year was great at providing exposure across a number of fields and the classes were strongly supplemented by rigorous math which I found invaluable. In addition, the cohort felt like a family, and we learned as much from each other as we did from the classes.

The first year of the AIMS CDT was very positive for me overall. I very much enjoyed the breadth of courses and what I initially thought would not be within my interests at all, I found really quite interesting. I found all of the courses engaging and I felt there was a very nice blend between the computer science and engineering flavours of AI. I think the real strength of the CDT is the range of the interests of the cohorts and by doing all the different courses, it allows us to hold sane conversations with each other about what we're working on, which I think is incredibly valuable. I also think that the two mini projects are great for helping us understand what we might want to work on in the future, but they could debatably each be a couple of weeks longer.

In terms of things that could be improved, I feel that the 'level' of courses was somewhat inconsistent. Some gave a broad, 'bird's eye view' whilst others quickly went into very deep detail. In general, I believe that it is better to give a broader, higher-level, foundational view of a subject rather than delve deep into particulars. This gives people with no prior knowledge a good introduction to the area, and lets them explore it further if they want to, rather than having to digest a great deal of new information too quickly. I also feel that the AIMS programme has a blind spot in that there is almost no coverage of classical AI techniques that are still very relevant today.

My first year with AIMS was amazing. I had a wonderful time with the people in my cohort. There was a lot to learn from each other. I particularly enjoyed the small class size and close interaction with lecturers during the course phase and the shift into more research-centered work with the mini-projects. In retrospect, I wouldn't have wanted to miss this first year and would always choose it over a regular 3-year DPhil. Overall, it was a very wholesome and supportive environment to take your first steps in as a researcher.

The AIMS CDT was by far my first choice in for a DPhil/PhD and the first year has delivered on my expectations. As I had picked a research topic already, I was sometimes unsure about spending another year broadening my knowledge, but I can conclude that it was time well spent. The courses have exposed me to topics that I knew little about and that I never would have learned about otherwise. Likewise, the mini-projects gave me an opportunity to explore topics that I was very curious about but that were only partly connected to my chosen research. I am sure that my upcoming research will be greatly benefited from these experiences.

The excellent location of the student office and the efficient organization were another cherry on top.

The first year is designed to provide a broad perspective in the core areas of the CDT: machine learning, robotics and control. The taught modules provided a valuable point of contact with the researchers in the Engineering Science and Computer Science departments and, through teamwork, a way to build friendships with our colleagues. The mini-projects, in the second half of the year, are a good opportunity to get research experience in the topics of one's interest, and to test whether one's expectations correspond to reality. Overall, I would say that the first year is definitely a good "warm-up" for the remaining three years of DPhil.

The first year of the AIMS CDT has been a great academic environment. For someone who was not well versed with the themes of the programme, I found the balance of taught courses and research projects a great way to build foundational knowledge while pursing original research. The clear definition of class cohorts, combined with my own two lab rotations, exposed me to many different research areas, and allowed me to make an informed decision for a substantive DPhil area.

I enjoyed my first AIMS CDT year. On the one hand, some of the topics were completely new to me and it was useful to learn something about them. On the other hand, I already had an idea of what I want to do for my research, and I was able to approach every course from the perspective of it. I am very supportive of the idea of cohorts where you do not only have work-related discussions, but you can find new friends and avoid being lonely for the whole duration of the DPhil program. The organisation in general and trips to Edinburgh or the Ocado facilities were amazing.

Feedback from Courses – Cohort 2018

Feedback we have received from the students withb regards the lecture courses.

Game Theory
- the appropriate level of difficulty and pace
- good introduction to Game Theory

Data Estimation & Inference
- Lecturer was fun & engaging, always prepared to answer questions
- The lab was very useful for learning about the practical implementation of
- GPs and complemented the theory aspect of the lectures well

Dynamic Robot Locomotion & Motion
- Really fun way to structure the course
- The course was very motivating and interesting
- The lab demonstrators were helpful and encouraging

Learning from Big Data
- The practical was extremely good and well thought out
- The lecture course was very interesting especially the deep neural
- networks and their analysis

Optimization
- Clear overview of the concepts was delivered, great lecture style, very engaging
- The materials of the course was all very interesting and valuable Sensor & Actuator Networks
- Lectures were extremely well taught. The lecturer consistently gave helpful and interesting examples
- Second lab was awesome

Mobile Robotics
- It was fun and the environment as fairly relaxed as it should be.
- Nice to work on a project, rather than have lectured and problem sheets

We placed multiple seismometers around a dam which elephants often visited. We observed a lot of activities here; elephants running, playing and bathing. We also recorded "rumbles" – vocalizations which propagate vibrations into the ground.

Holy Crescent English High School, Aurangabad

Student Biographies – Cohort 2019

Jonas Beuchert

At TU Berlin, I studied Electrical Engineering with a focus on automation and control towards my Bachelor's and Master's degrees. During this valuable time, I gained deep knowledge in many areas of control and measurement data processing, including classic and modern control concepts, modelling, simulation, optimisation, and machine learning. Additionally, I had the chance to apply my skills throughout multiple challenging employments in Germany and the United States, covering test engineering for driver-assistance systems, signal processing for medical applications, and machine learning for production optimisation in the aerospace industry. My previous research publications deal with iterative learning control and event-triggered learning. To conclude, I am at the crossroads between measurement data processing, control engineering, optimisation, and learning algorithms. Since autonomous intelligent systems in particular often require unified solutions with contributions from all those fields, the AIMS CDT at the University of Oxford is an excellent place to shed some light on dark corners of adjacent streets.

Jan Brauner

I studied medicine at the University of Erlangen-Nuremberg and the University of Wuerzburg, Germany. Parallel to studying, I worked as a research assistant in neuroscience, immunology and global health. I developed algorithms and mathematical models for biomedical problems but also got my fair share of experience working in wet-lab research and designing and conducting clinical trials. During that time, I co-authored 9 peer-reviewed publications which were cited more than 180 times. In 2018, I started a one-year master's degree in Operational Research with Data Science at the University of Edinburgh. In my MSc project, I worked on innovative deep learning-based approaches to medical image analysis with Prof. Amos Storkey. My main research interests are AI safety, and applications of AI in medicine and biomedical research. I care a lot about making a positive contribution to society. In practise, this means that I tend to gravitate towards questions that I perceive to be under-researched relative to their potential relevance, and always try to keep an eye open for possible applications of my research. In my spare time, I co-supervise high-impact clinical trials (currently one on cognitive enhancement and one on depression treatment), make music (badly), like to nerd out on 'Effective Altruism', and enjoy various ways of moving my body (running, weightlifting, salsa dancing,). Google scholar profile: https://scholar.google.de/citations?user=tNZUnjcAAAAJ&hl=de&oi=sra

Alec Edwards

I graduated with a MEng in Engineering Science from the University of Oxford in 2018, specialising in Control and Information Engineering, after completing a final year project on distributed optimisation for energy management in buildings. Over the last year I have been working at Energy Futures Lab at Imperial College London conducting a scoping study on electric mobility in urban areas across Sub-Saharan Africa, investigating the key challenges and opportunities. My research interests include control strategies in multi-agent systems, with particular regard to privacy and security concerns.

James Fox

I grew up in Dorking, Surrey before studying Natural Sciences (specialising in Physics) at Christ's College, Cambridge. I spent a summer as a research intern for ATLAS, CERN searching for signal of a specific Top quark decay channel and another summer at Oxford writing scripts to scan for possible dark matter particle candidates from supersymmetry. I graduated from Cambridge with an MSci in 2018 and, since then, I have worked for the Centre for the Study of Existential Risk, a think tank dedicated to the study and mitigation of risks threatening human extinction.

In my spare time, I'm often running, cycling or debating politics and philosophy. I am excited by the breadth of topics covered in the AIMS CDT programme and I look forward to joining the mission of building safe, scalable and beneficial AI systems.

Shu Ishida

I grew up in Japan, and completed MEng in Engineering Science at the University of Oxford. I worked on robot path planning for multiple target regions for my master's thesis at the Oxford Robotics Institute. I am interested in applications of machine learning to education, healthcare, environmental issues, and other social issues. I developed a prototype of a chatbot to help students manage stress and anxiety for Microsoft Imagine Cup UK 2018. For RoboCup 2019, I was in charge of speech recognition as part of the Oxford team. I am particularly interested in explainable deep learning models, reinforcement learning and active learning, and wish to better understand the mechanism of logic and reasoning. I enjoy musicals, composing music, web design and travelling in my free time.

Cong Lu

I grew up in Aberdeen in the north of Scotland and graduated with a Master's in Mathematics and Computer Science from Balliol College, Oxford in 2019. My Master's project was in reinforcement learning with corrupt or noisy reward and involved the study of randomised agents to overcome reward corruption. My areas of interest are in reinforcement learning, deep learning and multi-agent systems, and I'm excited to continue exploring these areas in the AIMS CDT. During my undergraduate degree, I had the opportunity to undertake internships at Metaswitch, Barclays and Jump Trading where I worked in software engineering and systematic trading. Outside of academia, I'm a keen runner, skier and foodie, and you can catch me jogging down Port Meadow and the Thames!

Amanda Matthes

Before joining the AIMS CDT, I was a physics student at the University of Heidelberg in Germany, where I actually ended up doing a lot of computer architecture related work. I even got to be a research assistant and tutor for technical computer science. Today, I am particularly interested in alternative approaches to computing, like quantum and neuromorphic architectures and their connection to AI research. Talk to me about food and video games! You can find a more detailed CV at amanda-matthes.github.io.

Akam Rahimi

After growing up in a beautiful mountainous area of Kurdistan I moved to the UK and studied Computer Science and Artificial Intelligence at the University of Sussex. My final year project focussed on using deep learning to infer connections between neurons in biological neural networks. It introduced a novel neural network architecture which set the benchmark and outperformed existing computational methods for this problem. I was proud to be awarded the Brandwatch 'Best Artificial Intelligence Project' prize for my achievement. In a professional capacity I have developed a bespoke recruitment and management system used by Further Education colleges across the UK. My main area of interest is the development of machine learning algorithms across the interdisciplinary fields of health and robotics.

Tim Reichelt

I grew up in Germany and graduated with a BSc Hons in Computer Science from the University of Edinburgh in 2019. In my honours project I worked with the Institute for Genetics and Molecular Medicine in Edinburgh on analysing images of cancer tissue. During my studies I did internships in the Statistical Machine Translation group in Edinburgh and at the Bosch Center for Artificial Intelligence in Stuttgart. Outside of studying I enjoy running, playing football and exploring all kinds of sports.

Mrinank Sharma

I'm Mrinank. I grew up in the Wirral, right by Liverpool before starting my undergraduate studies at the Engineering Department at the University of Cambridge, during which I was at Pembroke College. I graduated in Information and Computer Engineering, primarily specialising in Machine Learning with additional modules in Bioengineering. I worked with Dr Richard Turner on my masters' project: Differential Privacy & Approximate Bayesian Inference. I am particularly interested in probabilistic machine learning and hope to not only contribute to the field but also use advanced AI techniques to make a positive social impact. I am also interested in AI ethics, safety and policy. In my free time, I especially enjoy playing and watching football (I am a massive Liverpool fan) as well as reading.

Student Biographies from other cohorts can be found at: www.aims.robots.ox.ac.uk/students

Our Alumni page can be found at: www.aims.robots.ox.ac.uk/alumni

AIMS Contacts

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

Michael Osborne
Director

Alex Rogers
Co-Director

Wendy Poole
Centre Administrator

Academic Supervisors

A full list of academic supervisors can be found at:
http://aims.robots.ox.ac.uk/academics-and-staff/

Kenya – tracking elephants

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 283155

http://aims.robots.ox.ac.uk/

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

AIMS

EPSRC Centre for Doctoral Training in
Autonomous Intelligent Machines & Systems