InspiredResearch

NEWS FROM THE DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF OXFORD

THE FUTURE OF ARTIFICIAL INTELLIGENCE

Google DeepMind's
Demis Hassabis gives
Strachey Lecture – p26

IN SEARCH OF THE NEW FOREST CICADA:
Developing new tools for acoustic ecology – p6

PREVENTING LUNCHTIME ATTACKS:
Fighting insider threats with eye movement biometrics – p12

ROBOTS AND THE MORAL MAZE:
Towards a code of ethics for artificial intelligence – p23

UNIVERSITY OF OXFORD

DEPARTMENT OF
COMPUTER
SCIENCE

🐦 CompSciOxford   f DepartmentOfComputerScienceUniversityOfOxford

Letter from the Head of Department

One of the many rewards of academic life at Oxford is the constant stream of incredible academic speakers that pass through our doors. A quick look at our forthcoming seminar series tells me that in the coming week, I can hear about 'The expressiveness of decidable fixpoint logics', 'Fighting modern cyber criminals', and 'Semantic alignment for agent interactions'. This is not a phenomenon restricted to the Department of Computer Science, of course: Oxford is one of the world's great intellectual meeting places, and a visit to Oxford is very like a pilgrimage for most academics.

Of all the remarkable speakers I have had the pleasure to see at Oxford, none has been more remarkable than Demis Hassabis, who gave the Hilary Term Strachey Lecture in February 2016 (see page 26). Demis is a founder and CEO of DeepMind, the London-based artificial intelligence company that was acquired by Google in 2014 for $400m. We invited Demis to give this lecture a year ago, and had a strong hunch that it was going to be a popular event. In anticipation of this, and thanks to generous financial support from Oxford Asset Management, we were able to book Christopher Wren's Sheldonian Theatre for the lecture, with a reception afterwards in the Divinity School. The stage was set for a special occasion. However, nothing prepared us for the deluge of interest that Demis's lecture generated. All 800 tickets were taken within 48 hours of release, and we were so overwhelmed by pleas for additional tickets that we booked another 300-seat lecture theatre and had the event live streamed there.

On the day, Demis arrived complete with a camera crew making a film of his life, and I had the very odd experience of walking through Oxford with Demis making smalltalk, while a camera crew filmed us from across the road. (A PhD in Computer Science doesn't prepare you for things like this…) The lecture itself was a joy – a model of what a Strachey Lecture should be, pitched perfectly, with plenty of good humour and much food for thought. So, you are probably asking yourself, why all the fuss? Well, DeepMind's core technology is deep learning – probably the most active research area in Computer Science at present. Deep learning has made huge progress over the past decade, and no clearer evidence of this is provided by the fact that, just a week after Demis's Oxford lecture, DeepMind triumphed in a series of matches in which a self-taught AI programme comprehensively defeated one of the world's best Go players. The series of matches made headlines across the globe. With Demis giving his Oxford lecture just a week before the Go tournament, everyone present at the lecture felt they were witnessing history in the making. And indeed, they were.

In other exciting news, I'm delighted to be able to say that Sir Jonathan Ive, head of design at Apple, Inc., will receive an honorary degree at the annual encaenia event in June (page 3). Sir Jonathan is probably the most influential industrial designer of the past 50 years: amongst other things, he designed the iPad and iPhone. He designed the iPod!

Professor Michael Wooldridge
May 2016

Apple's Chief Design Officer Sir Jonathan Ive to receive honorary degree

We are delighted to announce that Sir Jonathan Ive, the Chief Design Officer at Apple, Inc., will receive an honorary degree from the University of Oxford at the University's annual Encaenia ceremony, this year being held on 22 June 2016.

Sir Jonathan's nomination from the Department of Computer Science recognises his iconic, hugely influential, and incredibly successful contributions to industrial design. Sir Jonathan was the chief designer for a range of Apple products, including the ground-breaking original iMac in the late 1990s, the iPod, and the iPhone. Apart from being commercially successful on a truly global scale, his trailblazing designs have dramatically influenced not just the look and feel of these devices, but the way we think about and interact with technology itself. Perhaps the best way we can summarise Sir Jonathan's influence would be to point out that most people reading this announcement online will do so on a device either designed or directly influenced by him.

We very much look forward to welcoming Sir Jonathan to Oxford in June.

DPhil student Elizabeth Phillips has been selected as the only student alongside nine UK government officials to participate in a US government-sponsored exchange programme to identify the top ten leaders in UK Cyber Security. Participants were selected by the US Department of State and the US Embassy in London. The programme enables participants to learn about US efforts to protect digital infrastructure, ensure free flow of information, security and privacy of data, and the integrity of networks. Elizabeth met President Obama at the 'Town Hall' he hosted in London in April.

Oxford involved in new Internet of Things (IoT) Research Hub

Department of Computer Science academics, Professors Andrew Martin, Sadie Creese and Kasper Rasmussen are involved in the newly announced £23 million Internet of Things (IoT) Research Hub. They have been working with fellow Oxford Professors, Luciano Floridi of the Oxford Internet Institute and David De Roure of the Oxford e-Research Centre.

Oxford is one of a consortium of nine UK universities which will work over a three-year period to explore issues in privacy, ethics, trust, reliability, acceptability, and security (known as PETRAS).

The project is part of IoTUK, an integrated £40 million initiative between government, industry, the research community and the public sector to help advance UK global leadership in the IoT and increase adoption of high quality IoT technologies and services by businesses and the public sector in the UK, benefiting citizens. The hub will draw in substantial support from over 47 partners from industry and the public sector.

Professor Philip Nelson, the Engineering Physical and Live Sciences Research Council's (EPSRC) Chief Executive, said, 'In the not too distant future almost all of our daily lives will be connected, in one way or another, to the digital world. Physical objects and devices will be able to interact with each other, ourselves, and the wider virtual world. But, before this can happen, there must be trust and confidence in how the Internet of Things works, its security and its resilience. By harnessing our world-leading research excellence this PETRAS research hub will accelerate IoT technology innovation and bring benefit[s] to society and business.'

NEWS

News in brief

Oxford climbed to top spot in The Sunday Times University Guide 2016 league table for Computer Science. It also took third place (best in the UK) in the QS World University Rankings 2016 for Computer Science and Information Systems. In the Times Higher Education World University Rankings Oxford has consolidated its position as one of the world's top universities. Overall Oxford was the top European University in the rankings (2nd in the world), and in subject-specific tables for Engineering and Technology the University was rated as 2nd in Europe and 6th in the world.

....

The Alan Turing Institute (ATI) held a symposium on reproducibility for data-intensive research. Reproducibility is both a technical and a socio-cultural issue, requiring new metadata systems, technical architectures, workflows and research working practices in an increasingly open and transparent environment. The ATI convened an inter-disciplinary group of researchers and data scientists to discuss these challenges in sectors in which the ATI intends to have greatest impact, for example in health, medicine, bioscience, urban environments, finance, transport, social sciences and digital media.

....

University of Oxford Women in Science is a new project which has interviewed 39 women scientists, all working at Oxford. Blanca Rodriguez, Professor of Computational Medicine, was one of the women interviewed. The aim of the project is to share a broad range of successful women's experience, as a way to help and encourage other women's career decisions and choices. You can read the pages on Blanca and the other inspiring women who took part at the project website: http://goo.gl/xh2zCD

....

Professor Bill Roscoe wins MPLS Lifetime Award

Professor Bill Roscoe has been given a lifetime award by the Oxford University Mathematical, Physical and Life Sciences Division for external engagement and promoting impact.

The former Head of Department has been recognised for his pioneering work applying formal verification tools in industry since the late 1980s, including the invention of the FDR Software Verification Tool.

He says of his award, 'It has been wonderful to see so many real-world users of what originally seemed like theoretical research. Furthermore, interaction with these users has been a regular source of inspiration for my research: a real virtuous circle!'

Degree application numbers climb

Application figures for the Oxford undergraduate Computer Science degree programmes have risen for the 9th consecutive year. The department received 683 applications across the 3 degrees: an increase of around 18% on last year's figures. The biggest increase was seen in applications to the Computer Science and Philosophy course, with applications up by approximately 70%.

Computer Science now has the second highest ratio of applications to places of any subject at Oxford – more competitive even than Law. Just over 80 students were offered places.

There has also been a significant increase in applications to the full-time MSc Computer Science degree for entry in 2016. The course saw a total number of applications of 433; more than 28% higher than the figure for entry in 2015.

© Paul Tait

Michael Wooldridge made an ACM Fellow

We are delighted to announce that our Head of Department, Professor Michael Wooldridge has been made a Fellow of ACM, the Association for Computing Machinery. Mike is the only new Fellow from the UK.

Mike has been specifically recognised for his contributions to multi-agent systems and the formalization of rational action in multi-agent environments.

© Paul Tait

Three Computer Science academics nominated for teaching awards

Three of the department's academics have been shortlisted in the 2016 Oxford University Student Union (OUSU) Teaching Awards. Software Engineering lecturers Professor Andrew Markham, Dr Jeffrey Ritter and Dr Stephan Murer have all been nominated in the 'Most Acclaimed Lecturer' category. The OUSU student-led Teaching Awards provide students with direct opportunities to recognise excellence in teaching. Each year award winners are selected in six categories by a panel of students, and the winners are then referred to as examples to promote excellent teaching around Oxford.

Yelp Dataset Challenge round five win!

Professor Nando De Freitas and DPhil student Misha Denil were part of the winning team in the fifth round of the Yelp Dataset Challenge for their paper 'From Group to Individual Labels Using Deep Features'.

The Yelp dataset includes information about local businesses in ten cities across four countries. The challenge awards $5,000 to work that uses this data in an innovative way and breaks new ground in research.

This paper proposes a novel approach to using group-level labels (eg the category of an entire review) to learn instance-level classification (eg the category of specific sentences inside this review). The authors designed a new objective (cost) function for training a model which uses features from a deep learning convolutional neural network. This trained neural network can, in turn, be used as a classifier predicting which category a specific instance belongs to. Their innovative research has broad implications for a variety of fields, and not just text classification.

A PDF of the winning paper can be found at: http://goo.gl/n2Mzxz

News in brief

The winners of the 2015 Valuable Artifacts Prize were Departmental Researchers Benoit Barbot, Alexandru Mereacre, and Nicola Paoletti, along with Andrea Patanè (University of Catania). The winning project was HeartVerify, which is a plug-and-play framework for the analysis and verification of pacemaker software and personalised heart models, compatible with MATLAB. The winners collected a cheque for £500. The prize is sponsored alumnus Dr Steve Moyle.

The Vienna Center for Logic and Algorithms (VCLA) has awarded DPhil student Maximilian Schleich, the Honourable Mention for Outstanding International Master Thesis. Maximilian's thesis entitled 'Learning Regression Models over Factorised Joins' bridges databases and machine learning. His work led to a publication in ACM SIGMOD 2016, and has been previously awarded the 2015 Hoare Prize for best MSc project in Computer Science at Oxford.

Every two years, IEEE Intelligent Systems acknowledges and celebrates 10 young stars in the field of AI as 'AI's 10 to Watch.' This year Gerardo I. Simari, a former Researcher in the department, and Haris Aziz, who did an MSc in Mathematics and Foundations of Computer Science at the University of Oxford were both featured. Gerardo is now an Assistant Professor at Universidad Nacional del Sur in Bahía Blanca and a researcher at CONICET, Argentina. Haris has gone on to be a Senior Researcher at Data61 and a Conjoint Senior Lecturer at the University of New South Wales, Australia.

Developing new tools for acoustic ecology:
in the search for the New Forest cicada

Acoustic ecology is a rapidly growing area that seeks to understand an environment, and the dynamics of animal populations living within it, through ambient sound and the sounds that individual animals make. Oxford's Professor Alex Rogers discusses his research.

A typical ecoacoustics survey involves deploying a small number of battery-powered recorders that make continuous sound recordings over a period of weeks to months. These recordings are then analysed by ecologists to identify processes operating in the environment (both natural and anthropogenic), the animal species present, and their behaviour. To date, these surveys have been expensive and limited in scale since each recording device costs close to £1,000 and the manual analysis of the resulting recordings is particularly labour intensive.
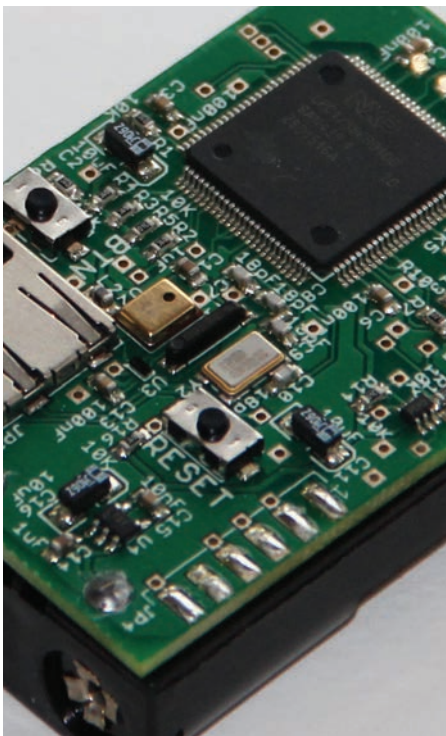
The high cost of commercial acoustic recorders is not particularly surprising given that these are custom devices serving a relatively small market. However, in many areas of science, particularly in laboratory work, high-end one-size-fits-all commercial devices are rapidly being superseded by locally developed solutions tailored to address specific settings. This is driven by the emergence of low-cost development boards, such as Arduino and Raspberry Pi, the online community sharing designs, and online manufacturers offering PCB assembly, 3D printing and laser-cutting services, which have made low-volume manufacturing of custom devices more accessible than ever before.

In the Soundtrap project (www.soundtrap.io) we are exploiting these trends to develop an open-source low-cost acoustic recorder for environmental monitoring applications. Our prototype logger combines a capable low-power 32-bit microcontroller with a sensitive MEMS microphone from a smartphone and an SD-card to store audio recordings. The total cost is less than £20 and the device is capable of recording 18 hours of high quality audio, sampled at 48kHz, using just three AAA batteries for power. To reduce the time-consuming task of analysing the resulting recordings, and to maximise the battery life of the device itself, we use acoustic recognition algorithms on the device to decide in real time whether a particular sound is interesting and should be recorded.

To evaluate this prototype, we are deploying it in May in Hampshire to search for the New Forest cicada; the UK's only native cicada species. The New Forest cicada has been known in the UK since 1812 but has not been officially sighted for over 20 years. Unlike more common species of cicada encountered in southern Europe, whose song is very familiar, the New Forest cicada sings continuously for up to thirty seconds at a frequency around 15kHz. This is close to the limit of most adults' hearing range, although children can typically hear it clearly from up to 100 meters away. The reason for its disappearance from a small number of well-studied breeding sites within the forest is unknown; possibly due to climate or land use changes. However, there remains a good chance that colonies still survive in as-yet-unknown sites within the forest, and we hope to discover these.

For the previous three years, in conjunction with the New Forest National Park Authority, the Forestry Commission and BugLife (the Invertebrate Conservation Trust),

we have been searching for the New Forest cicada using a citizen science smartphone app (www.newforestcicada.info) which is capable of recognising and recording the song of the cicada. However, after 6,000 survey reports from the forest, there is still no sign of the cicada. This might be because the cicada is very particular about when it sings, favouring warm sunny days with little wind between May and June, and our citizen scientists just haven't been there when they are singing. We hope that deploying 50-100 fixed recorders at the most promising sites, particularly less accessible, sunny, south facing clearings in the forest, may be more successful.

The recorders to be deployed this spring use the Goertzel filter (an efficient form of discrete Fourier transform) to detect the main frequency components of the cicada's song, and a simplified hidden Markov model to recognise the song's characteristic time-domain structure. The recorders, by waking up every five seconds and listening for 200ms, can operate for over eight weeks, continuously monitoring an area and making recordings of any cicadas heard. The resulting recordings will be retrieved after the deployment period and uploaded to a website, providing the ecologist and entomologists at BugLife with tools to view, sort and assess them. Depending on the false positive rate of the cicada detection

algorithm on the device itself, we may have over 50,000 recordings, each 30 seconds long, to analyse. However, the cloud-based platform allows us to run more sophisticated machine learning algorithms, recognising and identifying bird and insect recordings, as well as confirming any New Forest cicada recordings.

In addition, to the cicada deployment, we are working with a number of other groups to understand the use cases and resulting operational requirements of future devices. For example, we are working with Professor Patrick Doncaster and Dr Jake Snaddon at the University of Southampton to develop monitoring tools for areas of protected tropical forests in Belize. These areas serve as corridors for native jaguar populations but are threatened by the illegal logging of timber and the the illegal hunting of wild boar; both of which disrupt the jaguar's environment and its food chain. To this end, we are developing acoustic recognition algorithms that can reliably detect the sound of chainsaws and gun shots in the forest. While chainsaws are somewhat similar to the cicada detection problem in that they are long duration sounds which offer many opportunities for detection, gunshots present a more challenging problem requiring ultra-low power devices that can continuously listen and analyse sound recordings, looking for the very short signature of a gunshot.

In addition, while a recorder is a useful first tool for the rangers currently working in the area, a more developed solution will also need a low-power long-range radio capable of providing alerts to the rangers in real time.

In addition, we are working with Dr Davide Zilli and Professor Steve Roberts of the University of Oxford, and Professor Kathy Willis of the Royal Botanic Gardens, Kew to explore low-cost devices for mosquito species recognition in a project funded by a Google Impact Award. The project aims to recognise species of mosquito that act as disease carriers by their wing beat frequency, something that has proved to be possible using direct laser measurements of a mosquito wing, but has yet to be deployed at scale using acoustic detectors.

Our aim in all these projects is to develop an open platform consisting of both hardware and software components, that can rapidly be re-purposed to any new application, providing a cheaper and more capable solution than current commercial devices.
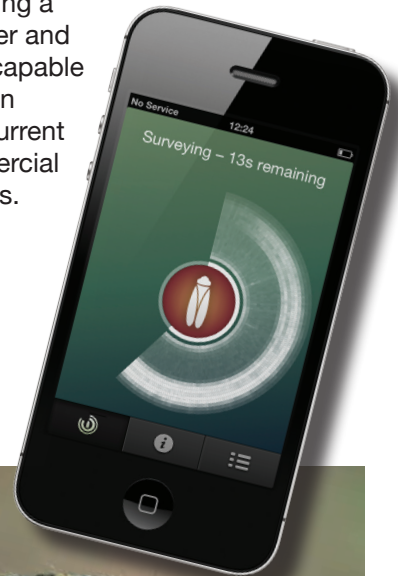
Beyond systems biology

Can we model entire social–ecological systems?

The development of a systems approach to biological and biomedical research that is both quantitative and predictive is one of the key research goals for the Computational Biology Group (CBG). Whilst systems biology aims to revolutionise biomedicine and thus enhance human wellbeing, human wellbeing is also inextricably linked to the health of societies as a whole and the natural and built environments within which those societies live. How far an integrative systems modelling approach can be applied to such social-ecological systems is the focus of the work of the IDEA (Island Digital Ecosystem Avatars) Consortium, with the Oxford Computer Science input into the consortium being provided by CBG researchers Dr Greg McInerney (recently appointed to a lectureship at Warwick) and Professor David Gavaghan.

Simulations of (potentially entire) social-ecological systems are a daunting task, and so we are starting our work by looking at the most scientifically tractable locations for these studies – islands – commencing with one of the most intensively studied and researched: Moorea. Moorea is a small, volcanic, oceanic island about 15 km northwest of Tahiti in French Polynesia. It has significant scientific capacity through its two international research stations (CNRS-EPHE CRIOBE since 1971 and UC Berkeley Gump Station since 1985). In terms of its biodiversity, Moorea is probably the best-known tropical ecosystem in the world, hosting France's Center of Excellence for Coral Reef Research and the US National Science Foundation's only coral reef long-term ecological research site. Moorea is therefore a natural laboratory spanning marine and terrestrial environments (to 1,207m) that is constrained enough to be tractable, but sufficiently large (132 sq. km) to contain all the elements of a complex socio-ecosystem, including a sizable human population (approx 17,000).

From an ecological perspective, coral reefs are the rainforests of the sea and millions of people depend on them for food security and other ecosystem services, such as storm protection and sources of medicines. Tropical island reefs, however, are under threat and represent 'canaries in the mine' due to their sensitivity to global change. Understanding interactions between island social systems and environments, and of differential cultural responses to ecosystem change, informs policies for sustainability and resilience.

The presence of long-term research centres on Moorea means that the consortium has available to it what is probably the world's most complete single-site ecological data set. This is further enhanced by the data collected by the Moorea Biocode Project, which has characterised every species on the whole island of dimension greater than 1mm (this includes genetic sequences, museum specimens, and digital photographs). Although further data collection on the impact of human society and at the microbial level is still needed, the existing physical and biotic databases, combined with a wealth of data on the resilience of Moorea's ecosystems, means that Moorea provides an ideal starting point to develop a systems approach to ecology and sustainability.

The Consortium's initial goal, and one in which Oxford Computer Science researchers are playing a leading role, is to bring together the existing disparate data sources within a coherent data management and computational science infrastructure. This will then allow this great wealth of data to be analysed through data-intensive techniques, and predictive computational models to be built that consider how human activities are interacting with ecosystem processes, ultimately focusing on how whole communities of organisms interact with human activities and the physical environment at the scale of landscapes or catchments. The resulting improvements in our understanding will then inform our assessment of the impact of both natural and human-induced change, and allow us to support policy-makers in determining management and conservation actions aimed at preserving ecosystem sustainability.

The Moorea Avatar

The consortium is beginning to address these questions by building a place-based data-science infrastructure and computational platform through which the research community can gain access to the wealth of data available. We have termed this infrastructure the 'Moorea Avatar' since it provides a three-dimensional visualization of Moorea which extends the Google Earth-style representation by including the dimension of time, and also enables researchers to zoom into a location, access data, and run (predictive) simulations based on that data. The Avatar is therefore a 'model of place' for Moorea. Since physical structure is at the centre of our understanding of place, this detailed, three-dimensional physical model of the island and the surrounding underwater environment will be at the centre of the model, serving to ground users and provide a natural context for their exploration.

With this 3D Moorea Map in place we can begin to build a virtual environment to provide the infrastructure for the IDEA project as a whole. This map will give the user a sense of 'place' at the outset that is populated with existing resources – research publications, data sets and models relevant for understanding all aspects of the system. This combined resource, which we have termed the 'Moorea Commons' (see figure), will provide a hub allowing exploration and sharing of existing knowledge, nurturing collaborations across disciplines. It will also be a repository for the new outputs – data, models, model outputs, research notes and publications. New research will exist within the context of the whole research effort and be immediately related to previous and on-going efforts.

The longer term goals are to use these resources to build predictive computational models to understand how biodiversity, ecosystem services, and society will co-evolve over the next several decades on Moorea depending upon what actions are taken. The starting point is the physical, biological, and social state of the island system today, and understanding how it has reached this point. This will then form the basis for understanding the future of Moorea under alternative scenarios of environmental change and human activity, including conservation efforts. These computational models could then be used to support scenario-based planning, helping local communities to adapt to environmental change and maximizing ecological resilience.
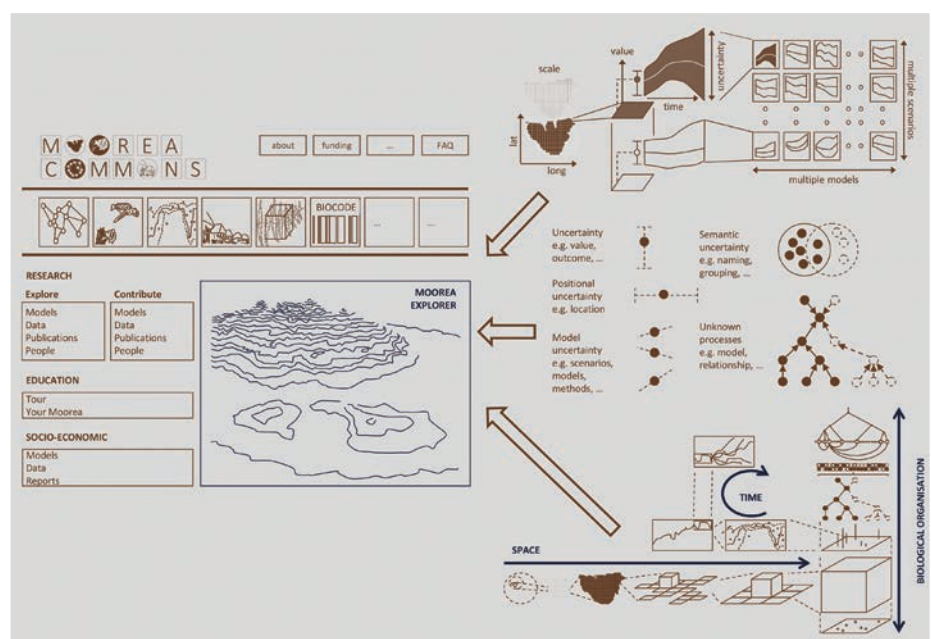
Building our Future

The Department of Computer Science has changed dramatically from the days when it had 20 academics and a handful of research students housed in Victorian houses in Oxford. Over the last five years, we have experienced unprecedented levels of expansion and now have nearly 1,000 people associated with the department as staff, full and part-time students and long-term visitors.

While this growth is certainly positive and necessary in becoming the world-class centre that we are, the department has already outgrown our original building in Parks Road, our more recent e-science building extension in Keble Road, and has now taken on three further locations.

We have acquired teaching space in the Department of Engineering Science, the Human Centred Computing Group based on St Giles, and the Robert Hooke Building in Parks Road. In the past three years, we have refurbished this building for additional teaching space, room for an extensive team in innovative security research, a Cyber Security Doctoral Training Centre, and space for researchers in machine learning.

In 2016 we will open a further bespoke space in the Robert Hooke Building for our newly formed Sensor Network Group, including room for academics, researchers and students as well as experimental and seminar space.

The department has invested substantially in converting these additional spaces to ensure that our research groups are provided with suitable environments in which to work, but our pace of growth and our ambitious plans for the future are best served in a consolidated space with room to expand.

This vision for the future would incorporate a new state-of-the-art dedicated Computer Science Building within Oxford.

We will be pursuing this goal in conjunction with the University and hope to launch a formal campaign later this year. If you would like to donate to help secure the future of innovation and growth in Computer Science at Oxford, please visit this web page: www.alumniweb.ox.ac.uk/cs/make-a-donation
If you would like more information on the plans for the Building the Future of Computer Science at Oxford, please contact: Megan Schaible at megan.schaible@cs.ox.ac.uk

Oxford partners Bebras
in kid's computational thinking challenge

The Bebras Computational Thinking Challenge is primarily an online event, now organised in over 30 countries, that is designed to introduce school students as young as eight to computational problem solving. It reveals how computational thinking can be applied to everyday situations in a way that's fun.

Over a weekend in January, the Department of Computer Science and Hertford College teamed up with Bebras UK, inviting just over 100 of the top-scoring students from the first online round, to attend a second round of competition in Oxford. The invitees in the top four age groups came from schools as far away as Aberdeen, Dundee, Exeter and Cambridge.

In addition to the competition final, the students and their parents enjoyed computing lectures from Oxford academics including

Professors Peter Millican and Mike Wooldridge, a college tour led by current Oxford undergraduates, and a lunch in Hertford College. While the students tackled some more in-depth Computer Science sessions, the parents were treated to a magic show by current Computer Science and Philosophy undergraduate Paavan Buddhdev.

Bebras Coordinator Chris Roffey commented, 'The second round celebration hosted by Hertford College and the Computer Science Department at Oxford was a fantastic event, balancing a celebration of the students' excellence in computational thinking with a friendly enjoyable day for all the competitors and their guests.

The success of the two days was in no small part down to the excellent organisation that was so thorough behind the scenes and yet seemed so relaxed, friendly and effortless to the visitors.'

Bebras is kindly supported by the Raspberry Pi Foundation, ARM Holdings PLC and Computing at School. The winners' hall of fame is available at: www.beaver-comp.org.uk/hall-of-fame.html

The Bebras competition days were a new addition to the department's outreach provision for 2016. Other regular events to introduce young people to Computer Science include the annual 'Women in Computer Science Taster Day', 'Further Maths – What Next?' and the 'Computing at School Conference' in association with BCS. Information on upcoming events for young people and schools is available at: http://goo.gl/qEWkjc

© Khalid Risvi

Alumni Profile

Dr Eoin Hyde (University 2008) & Dr Lorenz Berger (Keble 2010)

Eoin and Lorenz joined the University of Oxford from University College Cork and Bristol University, respectively, with technical backgrounds based in mathematics, applied mathematics, and Computer Science. They are co-founders of a medical technology startup company, Innersight Labs, based in London.

What did you do in Oxford?

We both studied on the same programme here at Oxford, namely the Life Science Interface Doctoral Training Centre (EH:2008-2013; LB:2010-2015).

What attracted you to studying Computer Science as a subject?

Probably keeping in common with most Computer Science people, we are driven by the goal of applying mathematics in the real world to solve today's pressing problems. This is particularly true within the field of mathematical biology – using efficient software and high performance computing to unlock the secrets of complex biological phenomena. We are fortunate enough to have worked on some really fascinating projects during our time at the Department of Computer Science, such as dynamic poroelastic lung and myocardial perfusion modelling.

What aspects of the course you studied here did you particularly enjoy?

We both really enjoyed our time at Oxford. It goes without saying that completing a DPhil is a tough but rewarding challenge. The support network that one finds at Oxford is really second to none and a collegiate system which provide a ready-made and diverse group of friends. The department itself with its world-leading supervisors and energising students also made the course enjoyable for us as students. We were both keen footballers as well, sometimes as team-mates and sometimes as rivals!

What did you do when you left Oxford?

Eoin went on to take up a position as a Research Associate at King's College London working on a pacemaker device design and cardiac electrophysiology basic science research. When Lorenz finished his DPhil he worked at a startup in Silicon Valley. Upon his return to the UK, we met in London and discussed various ideas and technology gaps that we had encountered, eventually deciding that we could best use our skills by providing badly-needed virtual surgical planning tools for surgeons, hence Innersight Labs was born. We are particularly focused on helping surgeons who practice minimally invasive (keyhole) surgery, due to the extra challenges involved in operating in reduced space and without direct sight or touch.

How has the course you studied here helped you in your current profession?

We apply the technical skills that we gained from our DPhils on a daily basis. The course also gave us the confidence to get done whatever needs to be done – in essence a DPhil is the ultimate problem-solving badge of honour. It's that confidence, backed by ability that leads to the necessary bold decisions one must take when starting a brand new company from the ground up.

What advice would you give to current students on applying their knowledge in the workplace, when they leave university?

In our opinion it's important to pick a challenging project that really requires technical knowledge. Not only knowledge already accrued, but crucially also knowledge yet to be gained. One should be constantly seeking ways to upskill in life as, after all, a rolling stone gathers no moss!

If we went back in time and asked, what would the student 'you' have thought about what you are currently doing – would you have been surprised, proud, amazed?

Definitely proud. People tend towards the known path for understandable reasons. The UK, and the EU to a wider extent, is quite risk averse culturally, particularly relative to the US. Rather than fixating on a fear of failure or opportunity costs, society could support people with various backgrounds to come together to start finding solutions to whatever is currently lacking. As students, we always talked about starting our own company, now we're simply doing it.

Preventing lunchtime attacks:

Fighting insider threats with eye movement biometrics

Simon Eberz, and Professors Kasper Rasmussen and Ivan Martinovic have been researching an interesting new application for eye-movement recognition.

Passwords are the most widely used mechanism to authenticate users, both on local systems (eg workstations or laptops), and on the web. Despite their widespread use, they suffer from a number of problems, most notably bad memorability, frequent re-use across different accounts, and users choosing weak, easily predictable passwords. Efforts to motivate users to choose stronger passwords have been largely fruitless, as evidenced by a number of recent password database leaks. For example, the 2013 Adobe password database leak revealed '123456' to be the most commonly chosen password.

In addition, an often overlooked property of passwords is that they can be easily shared with others. While this may seem convenient at times, it poses a major security threat in the context of insider threats. Insiders are in a particularly advantageous position when it comes to acquiring user login credentials, for example through shoulder-surfing, persuasion or coercion (and many other social engineering techniques). In the same context, so-called lunchtime attacks also constitute a threat. In a lunchtime attack, the attacker uses an already unlocked workstation while the legitimate user is away. This attack exploits the fact that passwords only provide login-time authentication, but can not provide continuous authentication.

Biometrics have become a popular approach seeking to remedy many of these issues. Biometric authentication is based on either physiological characteristics (such as fingerprints) or distinctive behaviour (such as typing patterns). Behavioural biometrics in particular offer the advantage of not requiring any significant conscious effort on the user's part, thereby making them particularly suitable for continuous authentication.

Inspired by the wide-spread use of eye movements for market research and medical diagnosis, researchers have investigated the feasibility of exploiting their distinctiveness to provide continuous authentication. Such a system will detect when another user starts using a workstation or web service by observing a sudden change in eye movement patterns. Since eyetracking is becoming more commonplace in different types of consumer hardware (such as smartphones, gaming consoles, and virtual reality headsets), it is believed that eye movements are an excellent candidate biometric to provide additional security at little extra cost.

Harvesting insights from neuroscientific research, we designed a set of 20 biometric features. These capture the distinctiveness of different short-term eye movements, specifically fixations, saccades, and microsaccades. These biometric features identify differences in eye movement speed and acceleration (reflecting the properties of microsaccades), spatial dispersion of gaze points (measuring the steadiness of a person's gaze), and the pupil diameter. For the latter, we use both the static diameter, which largely results from eye shape and size, and the degree to which it changes over the course of a fixation.

Eye movements are a response to a visual stimulus. As different stimuli might lead to different responses (a property that is used in medical diagnosis), the experimental design must account for a number of different everyday tasks. In our experiment, we investigate web browsing, reading, typing, and watching videos. Thirty participants from various age groups and both genders performed our experiment while their eye movements were recorded. In order to be as unobtrusive as possible, we used an optical, video-based eyetracker. The tracking works by projecting an infrared pattern and capturing the reflection of the cornea.

Our initial results are very promising and suggest that eye movements can support reliable continuous authentication, even over a two-week period. We achieve equal error rates between 0.04% for web browsing, and 4.9% for watching a video, while imposters are detected in an matter of seconds. In the future, we plan to investigate whether these error rates can be further decreased through a multimodal approach (ie combining different biometrics into a single system), and how the biometric holds up under a number of active attacks. In addition, we will investigate if our approach can be applied to exciting upcoming devices, such as VR headsets and Google Glass.

Global Cyber Security Capacity Centre:

Focuses on emerging economies

The Global Cyber Security Capacity Centre (GCSCC) is a leading international centre for research into effective cyber security, promoting an increase in the scale, pace, quality and impact of cyber security capacity-building around the world. Its key aim is to develop a framework to understand cyber security capacity through creation of an innovative new capacity maturity model (CMM).

In 2015 the centre began applying this research with a particular focus on emerging economies. The CMM was initially piloted in partnership with the Organization of American States (OAS) in Jamaica and Columbia, before the approach was rolled out across Latin America and the Caribbean region by the OAS in partnership with the Inter-American Development Bank (IDB). In 2015-16 these CMM reviews were undertaken in a further nine countries: Kosovo, Bhutan, Montenegro, Fiji and Armenia with the World Bank, Uganda with the Commonwealth Telecommunications Organisation (CTO), Senegal with the Government of the Netherlands and in Indonesia and the UK at the invitation of the nations themselves.

These reviews enable governments to understand cyber security capacity in-country before developing new strategies, for nations and donors to prioritise investments, and for all partners to address cyber security capacity comprehensively and holistically. The GCSCC will be building on this work in developing its new regional approach over the next few years.

Agrees new partnership in Australia

The GCSCC has recently signed a new agreement with the State Government of Victoria, Australia which paves the way for the centre to create its first regional partnership. Over the next two years the GCSCC aims to create a new network of regional centres worldwide to more effectively disseminate its innovative approach to cyber security capacity-building on a global scale.

The Australian partnership will be based at the Oceania Cyber Security Centre (OCSC), a new facility in Melbourne bringing together eight Victorian universities, the Melbourne-based Defence Science Institute and major private sector partners to be a key focus point for cyber security capacity-building.

The GCSCC will work closely with the OCSC, enabling nations in the region to better understand what good cyber security looks like in order to focus investment and national cyber security strategies. GCSCC Director Professor Sadie Creese said, 'We are delighted to be part of this exciting new partnership. Our aim is for our work to be used internationally as a common framework for effective cyber security, and working with key partners worldwide such as the State Government of Victoria is the only effective way to tackle this truly global issue.'

Oxford – Singapore link–up for trusted computing and verification

A new collaboration with the National University of Singapore, led in Oxford by Professor Andrew Martin, will explore applications of a trustworthy remote entity (TRE) and formal verification of its operation. The idea of the TRE was developed in the recent Oxford DPhil thesis of Dr Andrew Paverd using trusted computing technologies to build a network proxy. This device can undertake privacy-sensitive operations, using the guarantees of modern hardware to provide evidence to all parties that the calculations have not been subverted. The first example scenario is the smart power grid. The TRE allows usage data from several households to be aggregated (to preserve privacy but give an accurate instantaneous load figure for the neighbourhood), usage data from a single user to be aggregated over time (so an accurate bill is produced, without disclosing the detailed usage pattern), and to allow signals (such as for load-shedding) to be sent from the supplier back to the individual households, even though these are not identified to the supplier directly.

The TRE is designed to have the smallest possible trusted computing base so as to be amenable to verification both of the protocols used and the code implementing the system. This project will concentrate on approaches to that verification task, using tools developed in Oxford and by the group of Professor Jin Song Dong at the National University of Singapore. The project will run for three years, and will employ staff in both Oxford and Singapore. It is funded under an initiative established jointly between the UK's Engineering and Physical Sciences Research Council (EPSRC) and the Singapore National Research Foundation.

Accelerating the impact of Computer Science research

Made available by the UK's main Research Councils, Impact Acceleration Account awards (IAAs) are designed to 'do what they say on the tin' – ie to speed up or amplify the impact arising from research that is already funded. In many cases, it is recognised that there will be a long lead time between research 'discovery' and impact in the real world. In the case of the Engineering and Physical Sciences Research Council (EPSRC), Oxford University is one of 31 UK institutions to have been awarded an Impact Acceleration Account (IAA) award. The grant is available to support activities which will reduce this lead time and assist in accelerating the impact from past, current or future EPSRC-funded research. Grants fund either the development of the technology itself to the point where it is suitable for venture capital, or the secondment of people inward (eg industrial partners into Oxford) or outward (eg Oxford researchers into industry).

Under the IAA secondment scheme, Professors Boris Motik and Bernardo Cuenca Grau were recently awarded IAA funding to support a postdoctoral researcher, Robert Piro working on secondment on a project applying semantic technology to the broadcast industry. This award builds on technology developed within two distinct EPSRC-funded projects – RDFox and SemFacet. The ultimate goal of the project is to develop TV channels which adapt to viewers' preferences.

Content selection poses a significant challenge because the selected content must satisfy both the users preferences and the broadcasters' policies regarding what content can be shown to whom and at what times of day, etc. Existing solutions are based on ad hoc rule systems of limited expressivity, and the project aims to explore whether semantic technology and automated reasoning can provide a more principled and flexible solution to selection.

In the case of this award, the secondment partner is Imagine Communications, a leading broadcasting technology provider who had been approached by Isis Innovation seeking to partner with researchers working in this field. Having made the link between Boris and Imagine, the spinout company Covatic evolved. Covatic is now developing a product which will enable established broadcasters such as the BBC to retain viewers and maintain their brand given the shift to view-on-demand (VOD) services by recognising the benefits opened up by using broader, up-front data, linked data and semantic reasoning.

Boris says that in addition to the spin-out opportunities, the IAA secondment has provided his research team with useful hands-on experience of what the technology they are developing is capable of, which in turn will help to inform the direction of their ongoing research.

Spin-out news

Members of the department continue to be involved in spin-out activity, with current academics as well as alumni linked to a number of exciting companies which originated at Oxford. Mi6Sense, TheySay, OxCept, Semmle and Wrapidity are just a few of our spin-out successes, ensuring that our department continues to make a real-world impact by encouraging commercialisation across a wide range of research areas. In the past few months OxCEPT was a Finalist of The 'Anti-Fraud Security Strategy of the Year' title at the FSTech Awards 2016, members of the department were involved with DeepMind's AlphaGo success (read more about that on page 26) and sentiment analysis spin-out TheySay was referred to in the media during the Corbyn Labour Leadership win.

Isis Innovation has linked up with the department to provide a 'hot-desk'. The hot-desk is staffed by Daniel Stachowiak enabling members of the department to gain easy access to advice on the commercialisation of research. Isis Innovation is the research and technology commercialisation company of the University of Oxford. Isis is the highest university patent filer in the UK and is ranked 1st in the UK for university spin-outs.

News in brief

NVIDIA, a world-leading specialist in visual computing, has recognised the department by naming two academics as 'AI Pioneers' and giving us our very own DGX-1 supercomputer. Professors Phil Blunsom and Nando de Freitas have been named by the company as 'Pioneers in Artificial Intelligence Research'. NVIDIA has chosen AI pioneers at ten different universities to receive their own ground-breaking supercomputer, the DGX-1.

Semantic technologies in health care

Modern health care crucially builds on data and its effective analysis, particularly in the scope of quality control. In the US, Health Maintenance Organisations (HMOs) are required to regularly report quality measures specified by the National Committee for Quality Assurance (NCQA); such reporting is a mandatory condition for participation in government funded health care schemes, which cover more than 40 Million patients across the US.

Traditionally, the computation of these quality measures involves a series of SAS programs or SQL queries. Such solutions tend to be large and involved so validating and maintaining them is difficult, requiring a high level of IT-skills. In contrast, semantic technologies provide tool support for modelling domain knowledge with high-level languages in so called ontologies, which are declarative, succinct and much closer to natural language.

The Knowledge Representation and Reasoning (KRR) group in the Department of Computer Science, together with the US HMO Kaiser Permanente (KP), conducted a project to demonstrate the advantages of using semantic technologies in health care.

In the project we computed a particularly difficult set of NCQA quality measures, namely HEDIS Diabetic Care, which we encoded in a Datalog ontology. The ontology consists of rules, which are intuitive if-then statements, and are in close correspondence with the text of the HEDIS specification. Instead of 3,000 lines of low-level SQL code, we ended up with 174 Datalog rules, each of which is easily legible and easily checked for correctness.

Rules were evaluated using RDFox, an in-memory RDF triple store and parallel Datalog reasoner. RDFox has been developed in the KRR group, and scales exceptionally well on modern multi-core processor architecture. Running on a relatively modest commodity server, RDFox computed the quality measures over the medical history of an entire KP regional branch with more than 465,000 patients in less than 30 minutes. KP's data analyst in charge acknowledged this to be fast in comparison to the currently used NCQA approved vendor product. they were similarly impressed with the small number of iteration cycles needed to reconcile discrepancies between the results delivered by our solution and the vendor product. Due to the intelligibility of rules, we started with a low number of discrepancies from the outset, and using so called explanations provided by RDFox we quickly identified modelling errors and even discovered problems with the vendor's solution.

With this project we have successfully demonstrated the advantages of semantic technologies over traditional solutions in the context of health care. Further projects with KP are envisaged, and KP is now considering the deployment of semantic technologies, and in particular RDFox.

5G–ENSURE to make 5G networks and systems secure and resilient

The 5G-ENSURE project has received funding of just over €7.5m from Horizon 2020, the largest ever EU Research and Innovation funding programme. For Europe to truly benefit from 5G services across sectors key to its economy, it is essential that trust and privacy are design priorities. 5G-ENSURE follows an innovative approach where 5G's inherent security needs are understood and addressed by design. A 5G testbed will be used for validating operational capacity in response to security and privacy issues.

The project involves Professor of Systems Security, Andrew Martin and two new Research Assistants Piers O'Hanlon and Ravi Borgaonkar.

Read more: www.5gensure.eu

Strachey anniversary

Professor Christopher Strachey (1916–1975) was a pioneering Computer Scientist and the founder of the Programming Research Group, now part of the Department of Computer Science at Oxford University. We will be marking the occasion of 100 years since Christopher's birth on Saturday 19 November 2016, three days after his birthday, with a symposium of invited speakers. The morning will look back at Christopher's life and works from a historical and technical perspective, and the afternoon will concern the future of Strachey-inspired theoretical Computer Science at Oxford University. There will also be a display of related archival material at the Bodleian Library on Friday 18 November.

Although Christopher was keenly interested in the practical aspects of computing it is on the theoretical side that he most indelibly left his mark – notably by creating, with Professor Dana Scott, the denotational (or as he called it, 'mathematical') approach to defining the semantics of programming languages. Strachey also spent time writing complex programs and puzzles for various computers, such as a draughts-playing program for the Pilot ACE in 1951. During his career he developed some fundamental concepts of machine-independent operating systems, including an early suggestion for time-sharing, and was a prime mover in the influential CPL programming language. Christopher came from a notable family of intellectuals and artists, perhaps most famous for his uncle, Lytton Strachey, a writer and member of the Bloomsbury Group.

Book here https://goo.gl/6wcLpK

News in brief

The Conversation is a website collaboration between editors and academics to provide informed news analysis and commentary that's free to read and republish. It is an innovative way of sharing research on a 'creative commons' basis. Recently two of our academics have contributed articles. Professor Jeremy Gibbons wrote an article on what happens when computer modelling goes wrong, and Researcher Reuben Binns wrote about how data can be used negatively to reinforce discrimination. Read these articles and more at www.theconversation.com

....

The Careers in Computing fair was held in association with the Oxford Careers Service in November. It attracted exhibitors ranging from small local companies, to multinationals including American Express, Google, and IBM. Interested in exhibiting this year? Contact fairs@careers.ox.ac.uk

ERC Consolidator Grant for Scalable Data Management

Professor Dan Olteanu has been awarded a €1.98 million grant from the European Research Council to conduct a five-year research project on 'Foundations of Factorised Data Management Systems'.

The project has two main objectives: to investigate fundamental challenges in scalable data management at the confluence of compression, distribution, and approximation for both data and computation; and to build a practical large-scale data management system that incorporates the techniques and algorithmic insights from the first objective.
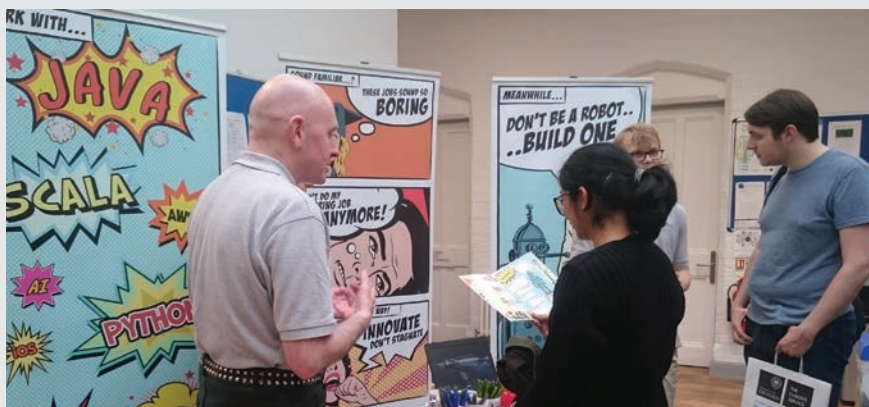
A key scalability ingredient put forward in this project is a lossless compression technique called factorisation. This is conceptually the same as factoring numbers and algebraic expressions, which is taught in primary school and dates back to Euclid. For example, given two sets, consider counting the number of possible pairs consisting of one element from each set. Instead of explicitly creating and counting the pairs, one can multiply the sizes of the two sets. Such simple ideas can be applied to many computational tasks underlying query answering and the construction of prediction and classification models. By avoiding redundancy in data and computation, factorisation becomes prerequisite to optimality of data processing algorithms. In practice, factorisation can speed up data processing by orders of magnitude.

Student Profile

Samuel Littley, (Computer Science, Keble)

Many Oxford students choose to undertake internships during their undergraduate degrees. We talk to fourth-year student Samuel Littley about his experiences, since joining us from a grammar school in Lancashire.
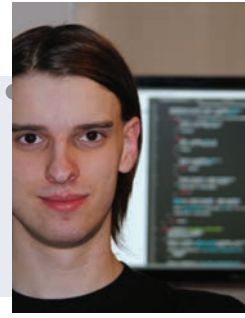
Tell us about your internships?

During my time at Oxford I've had the opportunity to do three separate internships.

My first internship was with Fatsoma, an events ticketing company based in Manchester. I worked on WordPress based promotional sites for clubs and venues, building the sites from designs and also on managing the WordPress system. I also wrote custom plugins in PHP for interaction between the sites and Fatsoma's own ticketing system, Facebook, Twitter and other social networks. For this internship I worked from home mostly, using Skype or email to communicate and Git to manage code. This also allowed me to spend three weeks at the Edinburgh Festival Fringe, working for Fatsoma during the day, running lights and sound for a show in the evening.

My second internship was with Google in London. I worked with the Site Reliability Engineering (SRE) team for Android and Play Services, completing a Python project to assist with planning and managing capacity for the various services the SRE team maintained. The Google offices were right in the heart of London by Victoria station – I spent three months living in Vauxhall with a 30 minute walk into work.

My final internship was with Semmle in Oxford, a company started by one of the Department of Computer Science lecturers in 2006, who do analysis of source code to improve code quality and correctness. I worked on a Python-based plugin for Sublime Text to display analysis results in the text editor and also a Java-based backend for a system allowing users to interact with the analysis in a web browser. It was great living in Oxford all summer with a 10 minute cycle ride to the office each day.

News in brief

A Google PhD Fellowship 2016 has been awarded to Oxford Computer Science student Marcelo Sousa. In 2009 Google created its annual PhD Fellowship programme to support outstanding Computer Science students across the globe. This year 39 students from North America, Europe and the Middle East were recognised. Oxford's Marcelo Sousa was the only person awarded a Fellowship in the category of Programming Languages and Software Engineering.

How did you hear about the placements?

For Fatsoma, I had previously done a week's work experience with the company, having known one of their software engineers through Scouting.

For Google, it was a company I was considering applying to anyway, having looked up internship opportunities on their website. I was then fortunate enough to know a graduate who had recently started at Google and was able to recommend me before going into the normal recruitment process

Semmle was a surprising experience of the company reaching out to me (somewhat unexpectedly) inviting me to join them for dinner! After spending the evening talking about the company (and plenty of other things) I was invited to an interview a few days later.

How have the internships helped you?

This has been a fantastic way to get industry experience and has really helped me decide where I want to work once I graduate. I was originally concerned as to whether I wanted to work for a large multinational corporation, such as Google, or somewhere smaller where I'd be able to make more of an impact. The experience both at Google and at Semmle reassured me that there wasn't that much difference between big and small.

The other huge benefit of doing internships has been the sheer improvement in the quality of my work. Writing code and doing projects at university and in my free time is quite different to working in a business with other people. I learnt a huge amount about how to write good maintainable code and how to work alongside my colleagues.

What are your plans for when you finish your degree?

I decided at the start of my 4th year that I would like to work for Google again – I spoke to my recruiter at Google who arranged a few more interviews for me. Fortunately I had the advantage of being able to refer to my work done during my internship and to the interviews I did for my internship. After a few agonising weeks of waiting I finally heard that Google had a job for me, in the same role as my internship. I'll be starting with them in July in Dublin – a huge, but very exciting step!

Research in Conversation: Machine learning and linguistics

Professor Stephen Pulman was recently interviewed by Georgina Brooke as part of the Oxford University 'Research in Conversation' series. The interview is reproduced in an abbreviated version below.

Stephen, what problems are you working on right now?

One of the problems that's really interested me from the start and that I'm working on at the moment is the problem of inference and natural language.

I'll give you an example: if you read somewhere that 'John's taller than Bill,' you can conclude that Bill is shorter than John. But you can't actually conclude that either of them is tall. If I give you an extra bit of information. If I say, 'Bill is tall, and John is taller than Bill,' then I can conclude that John is tall. It follows logically from the other statements. So that inference is an example of a process where you've deduced something from the grammatical and semantic structure of language.

There's a separate question of how I know whether that's actually true in the world. Well, I have some idea of what the standards of tallness are for people, which can vary: what's tall for a basketball player is quite different from what's tall for you and me – and that's different from what's tall for a tree or a building. The process of trying to write programmes that will automate those inferences is really interesting to me, and it also has lots of practical applications. If you take Siri, at the moment the way that works is by trying to find some sentences or some text that contain the answer to your question directly – they don't use much by way of inference. It would be much more useful if they did! If I were to ask: 'Is President Obama tall?', and somewhere or other it could work out that Bill Clinton is tall but Obama is taller than Clinton, then it can tell me that, yes, Obama is

tall even by the standards of US presidents – that would be really exciting.

Computers are really good at some things and really bad at others. What they're really good at is putting together bits of information from different sources such as numerical information or stuff from databases. What I'd really like to be able to do is to have a machine that could go off and read normal human text, understand it, and then be able to explain things about what it had read to me when questioned. It would have to be able to figure out the basics of what it had read without me necessarily having understood it. So getting a computer to tell me something I don't know would be really exciting. And not just to tell me, but to be able to explain it: to teach me.

What needs to happen for that to be a reality?

Lots of things! There's the basic ability of machines to understand English, which is very far from perfect. We can do it to a certain extent. But to actually have a machine go off and read, say, a biology textbook and then understand the content of it properly – we're a very long way away from that.

What is really difficult are the areas where there isn't that predefined logical structure, and where you and I would use our common sense and experience of the world. In some ways getting computers to understand language about particular technical domains (like, say, maths or logic) is less challenging because those domains do have a logical structure

which you can use to guide the understanding of the text. What is really difficult are the areas where there isn't that predefined logical structure, and where you and I would use our common sense and experience of the world to fill in the gaps. Computers are really bad at that and we just don't know how we do that.

Some simple examples would be things like resolving linguistic ambiguities. For example: 'The porters refused admission to the students, because they advocated violence.' So the question is, who does 'they' refer to there? 'They' would, very likely, be the students. If I change it to 'The porters refused admission to the students because they feared violence,' in that instance, 'they' is much more likely to be the porters. And we can deduce that from the sentence. But that deduction process is – at the moment – completely mysterious as we just don't know how to make a computer simulate it. That common sense reasoning that goes into resolving ambiguities and almost every aspect of language processing is the most challenging aspect of language understanding.

You're also a co-founder of TheySay, could you tell me more about that?

TheySay came out of work done with a DPhil student called Karo Moilanen, who's another co-founder. He came to me because he'd been working in a company on 'sentiment analysis', which is the process of discovering positive and negative attitudes towards things in text. Most people working in sentiment analysis treat it as a machine learning problem: you get

a corpus of texts, some of which are labelled positive, some which are labelled negative, and then you train the machine to distinguish the two. Karo was interested in some of the fine-grained linguistic properties that those approaches got wrong. So a few examples would be something like 'kill,' in isolation that has negative connotations. 'Bacteria,' again negative connotations. 'Killing bacteria' that's positive however 'failing to kill bacteria' that's negative again whilst 'never failing to kill bacteria' that's positive again.

It wasn't just that there were more positive words than negative words, it was to do with the grammatical structure and the semantic structure of the sentence. So that's interesting because you've got 'never' (which on its own is probably negative), 'fail' (negative), 'kill' (negative), 'bacteria' (negative) but put together they become positive. We got interested in how that happened. It wasn't just that there were more positive words than negative words, it was to do with the grammatical structure and the semantic structure of the sentence.

We have two basic products. There's an API which you send text to and it analyses it and you get the result of that analysis back in a report. The second product, 'MoodRaker', is a dashboard. MoodRaker allows you to choose topics you're interested in. For example, 'Oxford University,' you decide what text sources

you're interested in (blogs, RSS feeds, tweets, news sites etc) and the system then starts monitoring mentions of that and categorises and annotates them in various ways such as by sentiment.

We also tried to look for other things such as emotion: fear, anger, surprise. It's particularly interesting to track emotion in political discussions. We also try to look out for sarcasm. We try to detect advertisements because, of course, they're always positive which is going to skew your results and we try to do some 'topic detection,' is this a text about sport, finance, or politics? We're also beginning to look at demographic properties: gender, age and political orientation of the text authors.

What type of content are most of your clients interested in analysing with your tools?

It varies. Sometimes it's social media. A lot of our clients are in healthcare monitoring so, in these cases, there are various websites, for example the NHS, which people use to record their experience as a patient. However it can also be company reports and press reports.

A lot of the things we've done recently for fun are around politics. We were tracking what people were saying about the parties or people in, for example, the Corbyn campaign, the Scottish independence referendum and the general election.

It's very interesting, but we're never trying to predict the outcome. We can't assume that the people on social media are representative of the voting public (in fact quite the opposite in the case of the Scottish referendum – the independence campaign absolutely dominated social media, despite losing in the actual vote).

Do you have to analyse British English differently to American English ('not bad' might mean really pretty good in the UK but is more likely to be a more indifferent sentiment in the US)?

We do calibrate for whether we're processing American or British text, but the more difficult thing to take into account is the context in which particular words are being used. So if you're talking about beer, then 'cold' is good but if you're talking about 'coffee' or 'food', it's not. Then there can be differences depending on who's producing the text. So if I describe something as 'sick', that's not good. But if my son describes it as 'sick' – that's a good thing.

It's also very difficult to determine who the target audience is… the intended positivity or negativity can depend on the interests of the readers. If I read about oil prices coming down, I'm probably pleased about that, but if I'm an investor or the Government (who want to get taxes from oil) then that could be perceived as negative. We've only got the text, we don't know about the readers' preferences or prejudices.

What would you ultimately like to consider the final legacy of your research to be?

It's so nice to see your students all over the place. I've got former students in universities here, in London, in the US and in companies. It's getting to the point where I see students of my students!

In terms of research, I'd hope people continue to use and develop my work, particularly where I've been able to combine linguistics with computing or logic, and I hope TheySay continues to be successful.

The world outside

The department's students and academics are not only accomplished computer scientists. Below are some other achievements by members of the department.

Photo © Spurbrick

Katherine Fletcher, Coordinator of Cyber Security, is a member of the Oxford-based Brazilian dance and percussion group Sol Samba. This colourful crew brought some South American fizz to the department BBQ last summer.

···

Simon Eberz, a DPhil student studying Software and System Security, won Silver with his team at the British University Team Championships for Archery. The BUTC is a national competition that features 32 teams coming from all across the United Kingdom. The tournament was won by Warwick University.

···

Claire Rugg, a member of the finance team, had 10 inches of hair cut off for charity. She raised £419 in sponsorship for The Little Princess Trust, which provides real hair wigs to boys and girls across the UK who have lost their hair through cancer treatment.

···

Carlo Maria Scandolo, a DPhil student in the Quantum Group, recently restored some previously hard to decipher works written by Nicola Matteis (fl. 1650 – 1700), an Italian composer active in the latter half of the 17th century. Carlo Maria transcribed the works into a modern notation playable today, and the transcripts have now been published by Edition Walhall, a German music publisher.

···

Martin Lester, a Senior Teaching Assistant, organised the Great Scholars Riichi Mahjong Tournament. It took place in the department on 2 April and attracted 20 players from around Europe. Mahjong is a game with elements similar to bridge and poker that is popular throughout East Asia. Martin used a constraint solver to generate as fair a tournament schedule as possible. He has been invited to help organise the European Championship later this year.

···

Three teams from the department (two mixed teams, one men's team) took part in the annual Teddy Hall Relay running event, which took place on 9 March. Each leg of the race was 3.6mi, beginning and ending at Iffley track. Some of the happy team members can be seen in the photo [below].

···

The 3rd Oxbridge Women in Computer Science Conference

© Luana Bulat

Oxford and Cambridge female computer scientists forged many connections at the 3rd annual Oxbridge Women in Computer Science Conference held at the University of Cambridge's Computer Laboratory on 10 March 2016. The conference is co-organised by Oxford Women in Computer Science (OxWoCS) and Cambridge women@CL, and is now an established annual event.

University College London's Professor Ann Blandford presented the keynote talk entitled 'Empowering people through computing'. It reflected on her education, impressive career and experience as a female computer scientist. The keynote was followed by student presentations, and a poster session, with topics ranging from social networks to computational linguistics. This was followed by our 60 second lightning talks and undergraduate presentations. The academic portion of the conference was closed by Dr Alyson Pitt, from Google.

The conference was free for all, thanks to the generous sponsorship of Google. Starting with only 50 attendees in 2014, the conference has gone from strength to strength with over 150 attendees this year. This allowed the forging of many lasting connections between the two societies, and the creation of new working relationships. The 4th Oxbridge Women in Computer Science Conference will be held at Oxford's Department of Computer Science in March 2017.

Satisfied students push up the Barometer

The latest (2015) Student Barometer results are out, with Computer Science students at Oxford reporting a 95% satisfaction rate with their university experience – above the University's overall rating of 93%.

Students on the undergraduate programme in Computer Science particularly valued the organization and smooth running of the course, the prompt feedback on their work, the quality of tutorial teaching, plus the library and laboratory facilities, all of which were rated at over 98%.

The department's part-time graduate programmes – the MScs in Software Engineering and in Software & Systems Security – both scored particularly highly in the quality of lectures measure: at 98% both courses came in a full 10% higher than the average score across the University. The Software Engineering programme also scored over 98% for class size, and subject expertise of academic staff.

The overall feedback from the department's doctoral programme students was also above the University average. There was a 95% satisfaction rate with the level of research activity. The ability to work and study alongside people from a range of other cultures was another particularly positive point, along with the subject expertise of the staff.

The Student Barometer is a yearly questionnaire that is sent out to all students within a university in mid- to late autumn, collecting data from newly integrated students as well as from more seasoned students. A single authority, i-graduate, collects and publishes the results, allowing comparisons across departments, colleges and even universities.

© Gerald Sharp Photography

Student judged 'outstanding' by WCIT

Second-year student Will Bell has been awarded an Outstanding Information Technology Student award by the Worshipful Company of Information Technologists. Will was presented with a certificate and prize at a ceremony held in the historic Apothecaries' Hall in London. He is pictured above with Master Alderman Sir David Wootton (left) and Master Educator John Leighfield CBE (right).

Going back to work with a Returning Carers' Award

Ana Minchole is a Postdoctoral Researcher working in the computational biology research team. Ana was selected to receive a Returning Carers' Award from the University of Oxford. These awards operate as a small grants scheme intended to support the return to research of women and men who have taken a break of at least six months for caring responsibilities.

Ana's research interests include cardiac modelling and simulation and signal processing. The aim is to provide selective biomarkers from the electrocardiograph signal which will help to predict cardiac arrhythmias occurring either due to drugs or pathological conditions. Cardiac modelling provides a unique tool to understand the specific ionic mechanisms of each disease condition and simulate how these mechanisms are reflected on an electrocardiogram.

Ana says, 'this award has helped me to re-establish my research by funding two stays in the Department of Experimental Cardiology at the University of Amsterdam, which possesses a large set of clinical and experimental data that will allow me to connect the knowledge and data from the ionic level to whole organ and body surface ECG.'

Staff appointments and changes

Since the last edition of *Inspired Research*, there have been some new additions to the department:

Babis Nikolaou – Research Assistant on the DBONTO project.
Christiana Panayiotou – Marie Curie Fellow.
Diederik Roijers – Research Assistant on the TERESA project. (See p28.)
Dmitry Chistikov – Research Assistant on Infinite-State Systems and Dynamical Systems Verification and Synthesis: Algorithms and Complexity.
Evgeny Sherkhamov – Researcher/ Software Engineer on the Knowledge, Representation, and Reasoning project.
Menisha Patel – Researcher on the ZOON project.
Ohad Kammar – Researcher in Programming Language Theory.
Philippa Hopcroft – Senior Research Fellow.
Piotr Skowron – Research Assistant on the ACCORD project.
Rahul Santhanam – Professor of Computing Science & Tutorial Fellow, Magdalen College.
Ravishankar Borgaonkar – Research Fellow on the 5G Security Architecture project. (See p15.)
Ruiwen Chen – working on a unified approach to algorithms and lower bounds.
Ruslan Fayzrakhmanov – Research Assistant on the VADA project.
Varun Kanade – Research Lecturer in Machine Learning.
Xin Zhou – Research Assistant on the Computational Cardiovascular Science project.

Digital wildfires: Showcase workshop

Baroness Beeban Kidron gave a powerful keynote speech arguing for 'iRights' for young people, including the right to be forgotten online.

Senior Researcher Helena Webb describes a recent Digital Wildfire project event that promoted debate over how to limit the spread of harmful content on social media.

On 12 January 2016 the Digital Wildfire project co-hosted a showcase workshop that brought together researchers with stakeholders from policy, the police, civil society, commerce and education. The aim of the workshop was to discuss the spread of harmful content on social media. Opening the event, Digital Wildfire project leader Professor Marina Jirotka noted the many benefits brought about by platforms such as Facebook, Twitter, and Instagram, describing their ability to enhance communication and societal cohesion. However she also highlighted the capacity for social media to facilitate the rapid spread of false information, inflammatory content and hate speech, all of which can cause significant damage across society. Marina argued for a collective debate to determine effective and responsible ways to limit this spread of harmful content.

The workshop welcomed a diverse range of speakers to share their perspectives on this debate. We heard from groups such as Kick It Out – football's inclusion and equality organisation – who deal with large amounts of harmful social media content on a daily basis. We also heard from academic researchers who have studied social media phenomena such as the spread of rumour, (lack of) privacy, and online 'shaming'. Iain Bourne from the Information Commissioner's Office described data protection issues surrounding social media content. Paul Giannasi, Head of the Cross-Government Hate Crime Programme at the Ministry of Justice, was joined by Gabrielle Guillemin from freedom of speech organisation Article 19 and academics from the Universities of Middlesex and Cardiff in a roundtable discussion about social media regulation.

In a powerful keynote speech Baroness Beeban Kidron, founder of the iRights campaign, argued that young people should have the right to remove their online content in order to allow them to forge their own identities, learn from mistakes, and reduce their risks of suffering harm. Alongside our co-hosts at the University of Nottingham, the Digital Wildfire team also ran a youth panel in which six young people presented their views on digital citizenship. These highly articulate young people were winners of a competition to answer the question 'What makes a good digital citizen on social media?' and each received a £100 prize voucher thanks to sponsorship from Santander Universities. More information about the workshop is available at www.digitalwildfire.org and clips of the presentations can be seen on our 'Digital Wildfire project' YouTube channel.

Our six youth panel winners were introduced by Elvira Perez Vallejos, University of Nottingham (far right of picture) and received prizes from Santander's Mary Ridgway.

The ethics of AI

by Paula Boddington

The research project 'Towards a code of ethics for Artificial Intelligence (AI) research' is one of 37 projects in an initiative, backed by a $10m donation from Elon Musk, to ensure that AI 'remains beneficial'. It is funded by the Future of Life Institute, which aims to support research 'safeguarding life and developing optimistic visions of the future'.

This project is cross-disciplinary, Professors Michael Wooldridge of Computer Science and Peter Millican of Philosophy, are working with Paula Boddington, a philosopher with a background in the ethics of emerging technologies, who joined the Department of Computer Science in October to work on this project.

The initial aim is to clarify the challenges which regulation in this field presents, drawing on lessons from other areas of rapidly developing technology. 'Ethical regulation' often conjures up the idea that the aim is to curtail or ban research and development but ethics should also be interpreted much more positively, as assisting with attempts to ensure that technology is developed in ways which have a positive impact. There is an ever-present need for dialogue between disciplines in asking questions in ethics, and this approach is far more promising than one which sees a stand-off between science and technology on the one hand, and codes and regulations on the other. Just as there are some risks in technology, there are also risks in codes of ethics – for instance, not only does a code in itself fail to guarantee good behaviour, but complacently thinking that 'we're the good guys, we've done the ethics' may even discourage careful examination of the hard issues.

Early work on the project has examined the particular challenges AI gives for producing general ethical guidelines. For instance, codes of professional ethics are generally based upon certain starting points. One assumption is that the need for a code arises from the relative vulnerabilities of the layperson in relation to the professional, who has skills and knowledge the layperson lacks; the professional is assumed to be in control. However, with AI, there are fears that in some circumstances, the ability to keep control might be compromised. So, should we only develop technologies where human control is retained at all times?

Likewise, codes of professional ethics talk of producing benefit and avoiding harm. Yet, where a far-reaching technology such as AI is rapidly developing, it is harder to assess what changes would count as beneficial, to assess who would benefit most, and who might be harmed.

The project team is organising a workshop on Ethics for Artificial Intelligence at IJCAI-16 in New York this July, which will also be a forum to discuss wider issues concerning AI and society. We are working with the journal *Minds and Machines* to produce a special edition arising from our workshop.

The IEEE has embarked upon a programme of 'values by design' which will produce a joint declaration on AI, autonomous agents, and robotics, for its 400,000 members as well as for wider propagation. Our project is participating in an IEEE subcommittee examining lethal autonomous weapons. We'll be contributing to an IEEE meeting on Ethical Considerations in the Design of Autonomous Systems, collocated at the ECAI16 conference in The Hague in August this year.

Meanwhile, our project has also contributed to a workshop at the (British) Society for the Study of Artificial Intelligence and the Simulation of Behaviour (AISB), examining the five-year-old EPSRC Principles of Robotics to consider their continuing relevance. A report will be delivered to the EPSRC and more widely.

We're very excited to be undertaking this work at a time when so much is occurring in the field.

© Photos: Jennifer Balakrishnan

Celebrating Ada Lovelace in Oxford

In 2015 we saw the bicentenary of computer pioneer Ada, Countess of Lovelace, famous for the prescient paper she wrote about Charles Babbage's unbuilt analytical engine, which led to her being called `the first programmer'. Celebrations were led by Oxford's Department of Computer Science, in partnership with some of the world's major computing museums.

At the heart of the Oxford celebrations were the archives of Lovelace family papers, deposited by their owners in the Oxford's Bodleian Library, and thus kindly made available to scholars globally. The 200th anniversary celebrations saw a particular focus on Lovelace's mathematics, and, remarkably, its first study by professional historians of mathematics, including Christopher Hollings, of Oxford's world-leading History of Mathematics group, and Adrian Rice of Ralph Macon University, working with Professor Ursula Martin. They've debunked many of the myths about Lovelace, to show the advanced level at which she was studying, and her deep and perceptive way of thinking. Thanks to the generosity of the Clay Mathematics Institute, and the descendants of Ada Lovelace, a digital edition of Lovelace's

mathematical papers has been made available online.

A display in the Bodleian Library, curated by Dr Mary Clapinson of the Bodleian, and Ursula Martin, brought to the forefront Lovelace's childhood; her correspondence with Babbage; her mathematical notes and speculations, and the only known photographs of Lovelace, kindly loaned by a private collector. The display has moved stateside for 2016, to the Computer History Museum in Mountain View, California, the world's leading institution exploring the history of computing and its ongoing impact on society.

A Symposium in Oxford in December 2015 attracted around 300 participants, including 50 students, who heard interdisciplinary presentations from computer scientists, mathematicians, historians and a graphic artist.

The Symposium had an amazing and positive response in thank you notes, articles and tweets. Valerie Barr, Union College USA, summed up the interdisciplinary awareness: 'As a computer scientist, I was humbled by the extent of research the scholars from the humanities disciplines have done. They have faced a significant big data problem, reading letters and diaries from multiple people that spanned decades, drawing connections between them, using calendars and newspapers to confirm and clarify details.' Others wrote: 'incredible experience, dramatic, moving, educational, cathartic.

Will never forget it', 'in terms of interdisciplinary engagement across very different disciplines, it was one of the best conferences I have been to' and 'I have never attended a Symposium of such multi-disciplinary interest, and such constant intellectual and imaginative stimulation, provocation, nudging'.

With Engineering and Physical Sciences Research Council (EPSRC) support, we partnered with Queen Mary University of London, whose cs4fn (Computer Science for Fun) is a global campaign to enthuse and teach both students and others about inter-disciplinary Computer Science research. They have created a special 'Ada Lovelace' edition of their cs4fn magazine, with around 20,000 copies distributed to UK schools, and linked web activities. Working with The National Museum of Computing in Bletchley, UK, and the Computer History Museum, we developed a competition for young women to 'write a letter to Ada', which attracted over 1,000 entries. The National Museum of Computing houses a rebuilt Colossus, the world's first electronic computer. We were delighted that Colossus veterans Irene Dixon and Betty O'Connell were able to present the UK prizes at the Ada Lovelace Symposium.

All this was made possible through generous sponsors and partners: http://goo.gl/itwf5a

Digitised versions of Ada's letters: http://goo.gl/wqVCe3

Update from the Networked Quantum Information Technologies project (NQIT)

Christopher Strachey, Oxford's first professor of Computer Science, felt that those working in the theoretical and practical disciplines of computer technologies should work closely together. Following in this tradition, the Quantum Group are involved in the Networked Quantum Information Technologies project, or NQIT (pronounced 'En-Kit'). Members of the Quantum Group discuss how they are working with colleagues in Physics and Materials Science to design a prototype quantum computer.

The approach being taken is to design simple components to control small physical systems, which can store the somewhat delicate quantum states required if one wishes to access more computational power than randomised algorithms. These modules are then connected together using 'entanglement' — the subtle phenomenon described by Einstein as 'spooky action at a distance'. By sharing entanglement, the small modules can perform more complicated computations, working in concert on their fragile quantum data by bridging the unforgiving and choppy straits of the noisy classical world between them. Provided with a few thousand of these modules, one could, in principle, perform computations that would stymie a conventional 'classical' computer the size of a planet.

Under the guidance of Professor Samson Abramsky, and in consultation with Professor Simon Benjamin of Material Sciences, postdoctoral researchers Dr Niel de Beaudrap and former department member Dr Dominic Horsman approached the question of abstract notations for the NQIT engine by digging down to the physical construction of the modules. As quantum systems engineering is still the preserve of experimental physicists, Niel and Dominic consulted ion trap experts Dr Chris Ballance and Dr Thomas Harty (Oxford's Department of Physics) to determine what physical limitations (such as sources of noise) could be expected of the small memory modules, and outlined approaches to protecting quantum memories in a quantum network. They then investigated the common features of those protocols, and identified bialgebras as a promising mathematical formalism for the analysis and design of the NQIT project.

In another research program overseen by Dr Jonathan Barrett, Postdoctoral Researchers Dr Matty Hoban and Dr Nathan Walk are working on protocols that exploit quantum phenomena for their security as well as their computational power. In collaboration with optical physicists, they are designing prototype devices to securely generate random numbers and transmit secret keys. Another major application is the development of so-called blind computing protocols, in which an encrypted problem may be safely solved by an untrusted quantum computer. This will be particularly important for ensuring equitable access to quantum computing technology in its first generations — likely to be expensive devices controlled by a small number of companies and governments.

Quantum computation has transformed in the time since it was first proposed by American physicist Professor Richard Feynman and pursued by Oxford Professor David Deutsch, from what many viewed as a whimsical theoretical model, to a highly speculative proposal for what computations can be realised, and finally to a challenging engineering problem. While the theoretical underpinnings of quantum computation may now be settled, Professor Christopher Strachey would be heartened by the fact that theoretical computer scientists at Oxford are still engaged in the problem of designing and describing the next wave of computing technology.

Inaugural Lovelace Lecture by Professor Barbara Liskov

The inaugural Ada Lovelace lecture took place on 27 October 2015 and was given by Barbara Liskov [pictured right], Institute Professor at the Department of Electrical Engineering and Computer Science at MIT, and winner of the 2008 ACM A.M. Turing Award for Computing.

This talk discussed how the abstraction mechanisms we use today came to be and how they are supported in programming languages. The lecture was generously supported by Facebook, and introduced by the company's Head of Global Diversity, Maxine Williams.

Google DeepMind's Demis Hassabis gives Strachey Lecture

Dr Demis Hassabis, an Artificial Intelligence researcher and CEO of Google DeepMind gave a talk at the Sheldonian Theatre on 24 February 2016 as part of the distinguished Strachey Lecture series (generously sponsored by Oxford Asset Management). With reference to his position at the cutting-edge of Artificial Intelligence (AI) research, Demis discussed the developments achieved by DeepMind since the company was founded in 2010 (renamed once it was acquired by Google in 2014 to Google DeepMind). Google DeepMind's mission is to 'solve intelligence, and then use that to solve everything else.'

In his introduction, Head of Department Professor Michael Wooldridge described Google DeepMind as 'a company doing remarkable things at a remarkable point in their trajectory' currently at the forefront of the international press for its achievements in playing Go. For those unfamiliar with AI research, Google DeepMind's CEO describes the research conducted by the company as the 'Apollo programme for Artificial Intelligence'. The company hopes to solve what is termed internally by them as 'artificial general intelligence' enabling the programme to think outside the box rather than through a strict regimen of programming for specific tasks, termed by them as 'narrow Artificial Intelligence'.

Most of the AI we interact with on a daily basis is 'narrow Artificial Intelligence' which is only able to deal with specific preprogrammed

templates, for example the Apple programme Siri. Artificial general Intelligence has the ability to be adaptive and inventive in relation to all tasks. Demis identified Deep Blue (a chess computer designed and produced by IBM in 1996) as the first breakthrough for AI. Although this was the first time a chess computer had beaten a reigning champion under regular time controls, Demis said that he came away being more impressed by the opponent Garry Kasparov, who was able to do more than just play chess. He could also speak several languages as well as tie his shoelaces!

Evolving from this idea, Google DeepMind has since developed AlphaGo, an agent based on artificial general intelligence, learning automatically from raw inputs, using reinforcement learning in order to be able to choose the best possible action for a task.

AlphaGo was recently pitted against the three-times European Champion Fan Hui in a game of Go, to demonstrate the capabilities of AlphaGo. The computer won 5-0. The programme was able to combine intuitive pattern recognition rather than brute calculation. We usually think of computers as using brute calculation to achieve a goal, such as the logic involved in a game of chess, however the purpose of AlphaGo was to have the computer 'think' more like a human using intuitive pattern recognition. In the same way as a human player, the programme replicates making a move because it feels right (ie intuitive pattern recognition) rather than using calculation to make a move based on a logical conclusion.

The game between AlphaGo and Lee Sedol, one of the world's top Go players, was played (after the Hassabis lecture took place) during March in Seoul, Korea, resulting in a 4-1 win for AlphaGo, the first time an AI has beaten a top-ranked Go professional. See p27.

The lecture concluded with Demis's plans for the future — in summary to use advances in Artificial Intelligence to solve real-world problems in healthcare, robotics and smart assistance.

Information about future Strachey Lectures can be found here: www.cs.ox.ac.uk/seminars/strachey/
The Hassabis lecture is available to view in full here: podcasts.ox.ac.uk/artificial-intelligence-and-future

AI solves a version of the 100 hat riddle

A team of Department of Computer Science academics has published a paper 'Learning to Communicate to Solve Riddles with Deep Distributed Recurrent Q-Networks' which has been reported on by the *New Scientist* and *The Independent*.

Working with Google DeepMind, the team of researchers created an AI which is able to solve a variation of the '100 hat riddle', a puzzle used in Google job interviews. The research, which extends existing algorithms, could be developed to solve other multi-agent communication problems.

Professor Nando de Freitas, a leading member of the research team explains, 'We wanted to investigate how AI individuals with limited perception and reasoning can learn to communicate so as to solve very complex problems, so that the group as a whole benefits. Riddles and brainteasers are a fun and challenging environment for both humans and AI agents.

We observed that the AI agents learned communication protocols, and also learned to use objects in the environment to communicate — eg they learned to use the light switch in one of the riddles (the light bulb riddle) to signal to other prisoners that they had already been in the interrogation room, and hence in the end they all tasted sweet freedom.

This research is interesting to us for several scientific and engineering reasons. First, the behaviour of the AI agents mirrors that of some animals like vervet monkeys and honey bees, except that in our case communication is learned rapidly instead of via a slow natural selection process. Second, this paper provides engineers with guidance on how to build simple and cheap machines that could learn to act as a group so as to solve difficult problems — eg self-driving cars with limited perception that learn how to communicate so as to improve traffic flow and safety, distributed medical sensors that learn how to communicate so as to discover new treatments by making sense of the global picture. Third, it is encouraging that the same methods — deep neural networks and reinforcement learning — can be used to solve riddles, play Atari or reach human level performance in Go; this re-assures us that we are on the right scientific track to understand all aspects of intelligence.'

DPhil student Jakob Foerster comments, 'The results show that we can reformulate tasks, which are made to be challenging for humans, as communication-based AI problems and that our extension of existing algorithms can successfully solve these kind of challenges. For example, the AIs mange to discover the known optimal strategy for the hats riddle and invent novel solutions for the light bulb riddle. This research also illustrates that AI research is a lot of fun and it's great to see the interest from the broader public!'

Another DPhil student, Yannis Assael concludes, 'Our work paves the way towards a new class of multi-agent communication problems. We believe that the challenge now is to gain deeper understanding of the communication protocols that arise and to see what other real applications we can solve with this approach.'

Oxford and AlphaGo

It has been an exciting time for the Department of Computer Science, as a Go-playing program developed by DeepMind, Google's London-based artificial intelligence lab, triumphed in a five match tournament against Lee Sedol, one of the World's top Go players. Sedol won just one game, against AlphaGo's four victories.

The tournament is the first time a Go-playing program has beaten a grand master, and has been likened to the 1997 Chess match between the computer Deep Blue and Chess Master Garry Kasparov — it marks a huge step in the development of artificial intelligence. AlphaGo is the latest in a series of impressive achievements by DeepMind, and Oxford has been at the heart of it: Professor of Computer Science Nando de Freitas and Oxford Computer Science graduates Karl Moritz Hermann and Ed Grefenstette are all part of Google DeepMind. DPhil student Nal Kalchbrenner is part of the AlphaGo team.

The TERESA robotics project counters isolation of elderly

Associate Professor Shimon Whiteson has recently joined the department. Here Shimon discusses a research project he is leading, that takes robotics in an unexpected direction, helping housebound elderly people to continue social interaction remotely.

The TERESA project aims to develop a telepresence robot of unprecedented social intelligence — helping to pave the way for the deployment of robots in settings that require substantial human interaction. In telepresence systems, a human controller remotely interacts with people by guiding a remotely located robot, allowing the controller to be more physically present than with standard teleconferencing. So, for example, a person with mobility issues unable to attend a social event in person. They will be able to remotely control a robot that could move around the room, allowing the controller to participate in conversations and feel less isolated. TERESA robots will semi-autonomously navigate among groups, maintaining face-to-face contact during conversations, and displaying appropriate body-pose behaviour.

The new telepresence system being developed by project researchers frees the controller from low-level decisions regarding navigation and body pose in social settings. Instead, TERESA will have the social intelligence to perform these functions automatically. This means that the controller does not have to be someone with advanced computer operation skills — making the technology accessible to a wider range of people.

The project is generating new insights into socially normative robot behavior, and is also producing new algorithms for interpreting social behavior, navigating in human-inhabited environments, and controlling body poses in a socially intelligent way. Machine learning plays a key role in the project, from processing audiovisual data to learning to quantify the 'social value' of different behaviours.

The project is conducting annual experiments in a day centre for the elderly in France, where elderly people go to participate in physical and intellectual activities designed to promote healthy ageing. When the elderly are sick and unable to travel to the day centre they are cut off from their social network, something that improved telepresence systems could help address. Three experiments involving dozens of elderly subjects have already been successfully completed and a final experiment, evaluating the complete semi-autonomous system, is planned for October of this year.

In a recent interview (*Business Insider UK*, October 2015) Shimon explained his vision of human-robot interaction, 'I really think in the future we are all going to be cyborgs […People] have a tendency to think there's us and then there's computers. Maybe the computers will be our friends and maybe they'll be our enemies, but we'll be separate from them. But I think that's not true at all, I think the human and the computer are really, really quickly becoming one tightly coupled cognitive unit.'

TERESA has received funding from the European Community's Seventh Framework Programme (FP7-ICT-2013-10) under grant agreement n° 611153. The project partners are Oxford University, Imperial College London, University of Amsterdam, University of Twente, Universidad Pablo de Olavide and MADoPA.