InspiredResearch Summer 2015 Issue 6

News from the Department of Computer Science, University of Oxford

Fiat Lux

Eclipse sheds light on solar panel modelling – p12

Quantum computers to leap ahead

Hub to create the world's most advanced quantum computer – p5

Consortium sets up Alan Turing Institute

UK striving to be world leader in big data and algorithms – p8

Security under scrutiny as flights abound

Oxford helps secure the new protocol for air traffic control – p17

University of Oxford

Department of Computer Science

CompSciOxford

DepartmentOfComputerScienceUniversityOfOxford

Inspired Research

is a twice-yearly newsletter published by the Department of Computer Science at the University of Oxford.

If you would like to learn more about anything you read in these pages, please get in touch: editorial@cs.ox.ac.uk

To subscribe to future issues, e-mail: editorial@cs.ox.ac.uk

To download previous issues, visit www.cs.ox.ac.uk/inspiredresearch

Printed on recycled paper

Letter from the Head of Department

My first eight months as Head of Department have been truly exhausting, with a dizzying series of developments in research, teaching, and outreach.

I was less than a week into my new role when two professors knocked on my door with an interesting proposal: a major collaboration with Google DeepMind, the London-based artificial intelligence arm of Google. Nando de Freitas and Phil Blunsom work in machine learning, a subfield of AI that has witnessed huge advances over the past decade, such as reliable automated translation systems. (Read more about recent developments on p18–19.) As part of the collaboration with Google DeepMind, our department received a substantial contribution from Google to fund DPhil students (p4). It will be exciting to see what Oxford and Google can achieve together.

I am delighted to report a major series of successes in grant income, many covered in this newsletter: Oxford is to lead a huge UK quantum initiative (p5); Georg Gottlob received £2 million Engineering and Physical Sciences Research Council (EPSRC) funding for data wrangling innovation (p7); Marta Kwiatkowska was jointly awarded a major EPSRC project focused around the problem of checking software in autonomous vehicles; Edith Elkind and Joël Ouaknine received multi-million pound grants (p6) from the European Research Council, taking the department's tally of these prestigious projects to nine; and Michael Benedikt was awarded an EPSRC fellowship (p3).

Another coup for Oxford (announced this January) is that it is to be one of five founding partners in the Alan Turing Institute, the UK's national focus for data science research. The Institute will provide huge opportunities for staff and students at Oxford. For more information, see p8.

In December 2014, we received the results of the UK's 'Research Excellence Framework' (REF), the quinquennial national exercise under which the research of every university department in the country is assessed and graded. The importance of the REF should not be underestimated: as well as its reputational aspects, the results are used to inform government funding for universities. Oxford did tremendously well, by any possible measure. In Computer Science, no less than 53% of our activity was judged to be of the highest possible standard, a substantial improvement since the last REF in 2008.

This all points to a department in an excellent position for research, but I am pleased also to report exciting developments in teaching. Two colleges – St Catherine's and St Anne's – have agreed to take additional tutorial fellows in Computer Science, meaning that we will be able to take more undergraduates. This may not sound like many tutorial fellows, but prior to this there were only 11 in Computer Science across the whole university. We hope our undergraduate numbers will continue to increase.

We continue to expand our links with industry. If you are interested in working with the department, further information is on p15. And finally, I was delighted to meet so many of our former students at the London alumni event in December. Our alumni activities continue to expand, and I very much hope to meet more of our former students at future events – see p10.

Professor Michael Wooldridge
March 2015

News in brief

Congratulations to Marta Kwiatkowska, Professor of Computing Systems and a Fellow of Trinity College, who was awarded an honorary doctorate degree by KTH Royal Institute of Technology on 21 November in the Stockholm Concert Hall. Marta is a driving force for the development of probabilistic and quantitative methods within Computer Science, and has collaborated with researchers at KTH in various ways.

Dr Vojtěch Forejt was recently awarded a First Grant from the EPSRC on the topic of verification of stochastic systems. His project is concerned with developing algorithms for the verification of a class of models called Markov decision processes, in which the objectives are specified as multiple, possibly conflicting, properties.

Dr Andrew Markham was awarded a £200,000 grant after presenting to the EPSRC's 'Big Pitch' (adventurous ideas presented to a dragon's den style panel) in structural and ground engineering. Andrew's 'MiSixthSense' project seeks to develop a low-cost, wireless, embeddable sensing technology that can measure structural deformations in 3-D from deep within a structure. The key to the technology's success is the use of low frequency magnetic fields that can penetrate rock, soil, concrete and water with minimal loss of signal.

Luke Ong, Professor of Computer Science, has been awarded an EPSRC grant for his work on Compositional Higher-Order Model Checking: Logic, Models and Algorithms.

Computer Science study rewarded by high earnings

Recent Computer Science graduates from the University of Oxford have topped *The Sunday Times* salaries league table. In late 2014, they ranked as the top earners from universities six months after graduating, with an annual average salary of £43,895.

Many of them have been recruited for lucrative posts with technology companies such as Google and Amazon, whilst others have gone into careers as diverse as teaching and social work. Data from the University's careers service shows that Department of Computer Science leavers, including postgraduates, who go into work (or a combination of work and study) have found jobs in: IT and Computing (62.8%); banking and investment (17.7%), roughly equally split between IT support and analyst posts); academia and higher education (13.0%); and Government, Order and Policy (6.5%).

The article also featured recent graduate Nicole Williams *(below)*, who was based at Oriel College, and is currently a software engineer at a small software consultancy in London. She told *The Sunday Times*: 'The tech industry in general seems keen to recruit CompSci grads, and to spend plenty of money doing so.'

Novel systems to improve web data methods

Oxford researchers in Software Engineering – Jim Davies, Jeremy Gibbons, and James Welch – have just been awarded €4 million by the EU through the Horizon 2020 programme for a new project called ALIGNED: Quality-Centric Software and Data Engineering.

ALIGNED, a three-year project starting from 1 February 2015, will develop novel ways to build and maintain software systems that use big data on the web. ALIGNED brings together world class Computer Science researchers (University of Oxford, Trinity College Dublin, University of Leipzig), software companies specialising in data-intensive systems (Semantic Web Company) and in information publishing (Wolters Kluwer), and academic curators of the large Seshat: Global History Databank datasets describing world history and archaeology (Social and Cultural Anthropology at University of Oxford, and Adam Mickiewicz University in Poznañ).

Together they will create more efficient methods for extracting, processing, publishing and sharing web data, laying the foundations for the next generation of big data systems which will lower costs and deal with the web data challenges of dynamism, complexity, scale and inconsistency.

Staff appointments and fellowship

- Tim Furche has been appointed as Department Lecturer from 1 April, having been a Senior Researcher with the department since 2010.
- Sam Staton has been appointed as Research Lecturer on a Royal Society fellowship and commenced on 1 March 2015.
- Michael Benedikt, currently a Professor of Computer Science with the department, has been awarded an EPSRC Established Career Fellowship to start from summer 2015.

Oxford teams up with Google DeepMind on AI

Members of the Department of Computer Science both past and present launched a major collaboration with Google DeepMind, a London-based centre for artificial intelligence research, in November 2014.

As part of the collaboration, which also involves three members of the Department of Engineering Science, the University received a significant seven-figure contribution from Google to be split between the two departments. The Department of Computer Science will use their share to fund DPhil students.

Oxford Computer Science graduates, Dr Edward Grefenstette and Dr Karl Moritz Hermann, together with Professors Phil Blunsom and Nando de Freitas, co-founded Dark Blue Labs. Their goal was to make natural languages, such as English, French or Chinese, understandable by computers, allowing machines and humans to better interact and collaboratively solve problems. Phil and Nando now divide their time between Google DeepMind and the University of Oxford. Edward and Karl Moritz have moved to become full-time employees at Google DeepMind.

Professor Michael Wooldridge, Head of the Department of Computer Science, said: 'Machine Learning is a technology whose time has come. We have invested heavily in this area and we are truly excited at the prospect of what we can achieve together with Google DeepMind. We are extremely proud of Phil, Nando, Ed, and Karl Moritz, and truly grateful for their efforts in securing such a fantastic donation to the department, and for paving the way for future collaboration with one of the world's leading computing companies.'

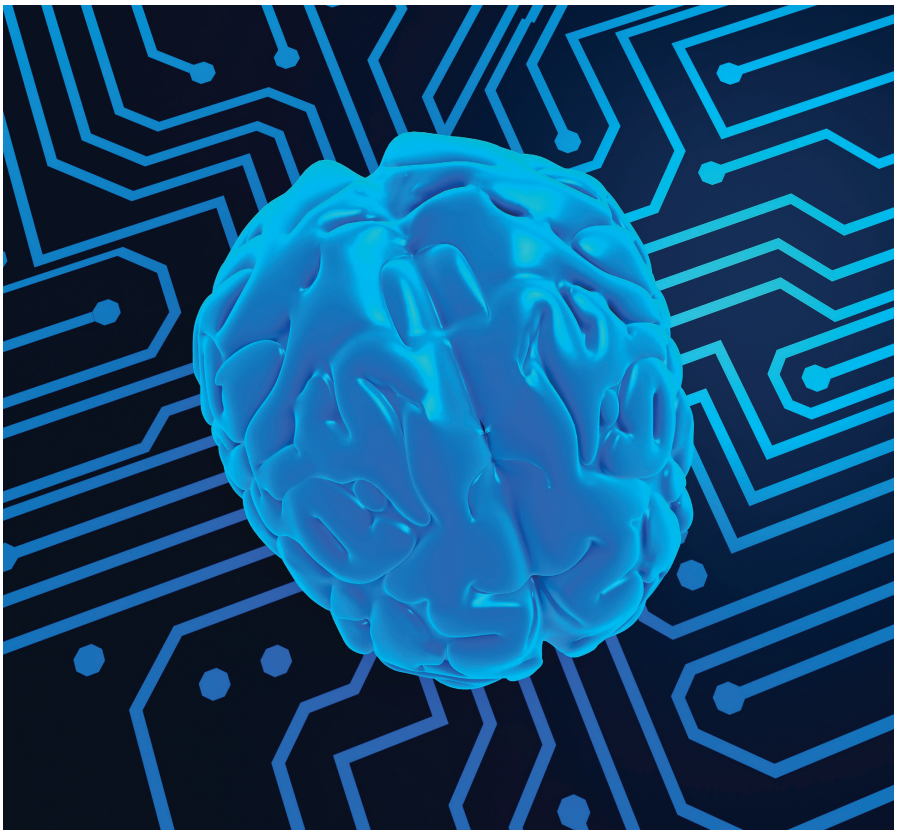

From academia to industry

Dr Edward Grefenstette and Dr Karl Moritz Hermann talk about their move to Google DeepMind:

There are a few points in the history of our field where industry offered a work and research environment similar to that of academia, such as the heyday of Xerox PARC or MSR. Google DeepMind currently feels like it has reached a similar point. Its existence itself is an aptly-named 'moonshot', best illustrated by its stated goal: to 'solve intelligence'. Within the framework of this goal, scientists are encouraged to think about the big questions, the 10-year plans, the blue-sky ideas, and given the computational and engineering resources to make this dream a reality.

The diversity of researchers that we interact with here is anything but minimal. This interaction is both in the form of regular meetings with colleagues, and the steady stream of academic and industrial visitors that come through the doors of our office to give tech-talks, from universities and Google Research offices near and far, visiting to discuss and debate topics from neuroscience to ethical concerns about artificial intelligence, via the latest trends in machine learning methods.

This has made the transition from academia to industry a pleasant one. More than an abrupt stop of our academic lives, we have found this to be a continuation of our work in Oxford, making it easy to look back with fondness at our memorable years at Oxford's Computer Science Department which have so aptly prepared us for this next step in our careers.

Quantum computers to leap ahead

Following an investment of almost £38 million from the EPSRC, the NQIT Hub, a consortium of universities and industrial partners led by Oxford, aims to build the most advanced quantum computer in the world. The experimental physicists and engineers in Oxford and partner institutions will bring world-leading expertise in trapped ions and quantum optics. Meanwhile, researchers from the Department of Computer Science will develop the secure protocols and algorithms that a quantum computer needs to run. University Lecturer in Computer Science, Jonathan Barrett, explains.

Since the early days of quantum theory, it has been known that quantum systems can exhibit phenomena that defy any intuitive understanding. But it is only over the last 20 years or so that researchers have realised how to harness these strange phenomena for powerful kinds of information processing.

An early success was the discovery by Peter Shor of a quantum algorithm for the efficient factorisation of a number. Using this algorithm, a quantum computer could, for example, crack the encryption that keeps our credit card details safe when we shop online. But the potential applications of quantum information science are much broader than this. They include the simulation of classically intractable physical and biological systems, the use of quantum systems for secure communications, improvements in sensing and metrology, and applications to machine learning.

The future importance of quantum science for the information economy was recognised in a £270 million investment in the UK National Quantum Technologies Programme announced by the Chancellor of the Exchequer, George Osborne, in his Autumn Statement of 2013. From this investment, the EPSRC has announced funding for four Quantum Hubs, led by the Universities of Oxford, Birmingham, Glasgow and York, with each Hub focused on a different aspect of quantum technologies. The Oxford-led Hub – Networked Quantum Information Technologies (NQIT) – has received almost £38 million for the development of quantum computing.

Building a quantum computer is difficult, because it requires manipulation of matter at the atomic scale with a very fine degree of control. To meet this challenge, the NQIT Hub will deploy two of the most promising technologies for quantum computation: ion traps and photonics. Given a handful of ions in a trap, each ion can serve as a single quantum bit. Moreover, given more than one trap, quantum information can be passed from one trap to another via single photons of light.

The Q20:20 machine

The flagship goal is the Q20:20 machine: an ambitious vision of 20 ion traps, each containing 20 individual ions, with traps connected by photonic links. The machine can be thought of in two ways: with short links, the whole constitutes a 400 qubit quantum computer; alternatively, if the links are extended to kilometres, the separated nodes form a prototype quantum internet.

Crucially, the technology is designed to be fully scalable, meaning that although 400 qubits is a relatively small-scale quantum computer, there is no fundamental barrier to prevent future extension of the device to an arbitrary size.

As the experimental scientists work to meet these challenges, theorists including a team within the Department of Computer Science, led by Samson Abramsky and myself, will develop abstract languages, along with concrete protocols and algorithms, that can run on a Q20:20 device when it is built. Although the machine will not be capable of factorising large numbers (so we won't be stealing anyone's credit card details anytime soon), small scale algorithms can be developed for specific problems such as the simulation of quantum systems.

We will also be developing protocols for cryptographic tasks such as secure key distribution, and the generation of trusted randomness, which can be implemented on the Q20:20 hardware in its expanded quantum internet form. An exciting prospect is that of 'device-independent' protocols, which use the mysterious phenomenon of quantum nonlocality to ensure that honest users are secure, even in the case that they do not trust the quantum devices that they are using.

It is comparatively early days for quantum technologies, and this investment represents the first big push towards bridging the gap between academic discoveries and industrial exploitation.

More information:
goo.gl/OqwdTI
nqit.ox.ac.uk/index.html

Algorithms to aid group decision-making

Procedures to help collective decision-making are to be designed at Oxford as part of a project that has received a grant of almost €1.4 million from the European Research Council (ERC).

Dr Edith Elkind's project ACCORD (Algorithms for Making Complex Collective Decisions on Structured Domains) aims to design novel procedures that help people make group decisions in a principled and efficient manner. The five-year project, due to start on 1 May, received funding from the ERC under the European Union's Horizon 2020 research and innovation programme*.

The project has its roots in voting theory, which studies methods and principles of collective decision-making, and combines the ideas from this field with algorithmic insights. Its approach is two-pronged: on the one hand, Edith aims to design algorithms for identifying patterns in preference data, using publicly available anonymised datasets from many different sources (varying from parliamentary elections to movie rankings), and, on the other hand, she will search for ways to implement preference aggregation methods so as to take advantage of these patterns.

Another important feature of the project is that it will consider tasks that are more complicated than choosing a single best alternative, such as choosing a committee (a fixed-size set of winners) or a partial ranking of the alternatives.

**Grant number 639945*

European Research Council
Established by the European Commission

Model-checking funding may unravel unsolved maths

A project that could improve software model checking and provide answers to long-standing open mathematical questions, has received funding of almost €2 million in the form of an ERC Consolidator Grant*.

Professor Joël Ouaknine will be leading the five-year project and taking on two Postdocs and two DPhil students to work with him. He hopes that the project will not only result in fundamental theoretical contributions and insights in their own right – potentially answering mathematical questions that have been open for years or even decades – but will also impact the practice of formal verification and lead to new and more powerful methods and tools for the use of engineers and programmers. The project will build on a series of recent advances and breakthroughs (some from within Oxford) to attack a range of specific algorithmic problems.

The central objective is to investigate key algorithmic verification questions concerning two fundamental kinds of mathematical structures used to model and analyse infinite-state systems: linear dynamical systems and counter automata, in both ordinary and parametric form. The team will focus on a series of model-checking and synthesis problems for such infinite-state systems. The aim is to deliver novel verification algorithms along with a map of the complexity landscape. A second objective is then to transfer algorithmic insights into practical verification methodologies and tools, in collaboration with academia and industrial research laboratories.

**Grant number 648701*

Grant supports big data reasoning

The Knowledge Representation and Reasoning (KRR) Group led by Ian Horrocks has secured a joint grant of more than £90,000 equally financed by the American Health Care (HC) provider Kaiser Permanente and the DBOnto Platform Grant.

The project will investigate whether the 'Healthcare Effectiveness Data and Information Set' (HEDIS) for Diabetic Care can be captured with KRR techniques. HEDIS is a benchmark used by more than 90 per cent of America's health plans to measure performance on care and service. Healthcare providers have to supply HEDIS data to be eligible to access social healthcare programs such as Medicare and Medicaid.

HEDIS is based on a set of indicators, and, to determine these, patients need to be classified according to their patient records into cohorts, a generic task in the field of KRR. The ongoing

investigations have shown that the HEDIS specifications can straightforwardly be translated into if-then-relationships which can be expressed by SWRL (Semantic Web Rule Language – pronounced: swirl) rules.

The data together with these rules are to be evaluated by the KRR Group's newly developed RDF triple store and parallel SWRL reasoner called RDFox. The project will thus put RDFox to a crucial test, showing whether it can potentially handle data from millions of patient records, and therefore scale to a level suitable for production.

Data, data everywhere, but wrangling tools soon there

A project that aims to help data scientists effectively wrangle data has received a five-year EPSRC Programme Grant of £5.7 million. It was awarded jointly to the Universities of Edinburgh, Manchester and Oxford.

Data is everywhere, generated by increasing numbers of applications, devices and users, with few or no guarantees on the format, semantics and quality. The economic potential of data-driven innovation is enormous. To realise it, and to provide meaningful data analyses, data scientists must first spend a significant portion of their time (estimated as 50% to 80%) on data wrangling – the process of collecting, reorganising, and cleaning data.

There is thus an urgent need to provide data scientists with a new generation of tools that will unlock the potential of data assets and significantly reduce the data wrangling component. This is where the Value Added Data Systems (VADA) project comes into play. VADA investigates principles and solutions for supporting users in discovering, extracting, integrating, accessing and interpreting the data of relevance to their questions, and for automating these tasks. The project is directed by Georg Gottlob, Professor of Computer Science. The Oxford co-Investigators are Dr Tim Furche, Professor Thomas Lukasiewicz, Professor Dan Olteanu and Dr Giorgio Orsi.

Additional contributions come from industrial partners who are both developers of data management technologies (LogicBlox, Microsoft, Neo) and data user organisations in healthcare (The Christie), e-commerce (LambdaTek, PricePanda), finance (AllianceBernstein), social networks (Facebook), security (Horus), smart cities (FutureEverything) and telecommunications (Huawei).

Project supports new school Computing curriculum

A new project to use Oxford-developed systems to help school children learn Computer Science began this academic year, designed to support the new compulsory Computing component of the National Curriculum. It is financed by the Department of Computer Science and Faculty of Philosophy, plus a generous private donation, together with matched funding from the UK Department for Education.

The project involves two teaching systems developed by members of the Department of Computer Science: Professor Peter Millican's 'Turtle System' (which introduces conventional 'imperative' programming in any of BASIC, Java, Pascal, or Python) and Dr Michael Spivey's 'GeomLab' (which introduces functional programming, starting from the creation of Escher-style recursive images). Both systems have been used extensively in Oxford's outreach events, and are now being supplemented with teaching materials and web resources to enable them to provide the basis for school lessons and qualifications.

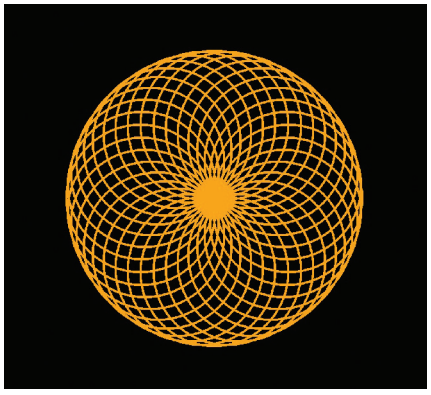

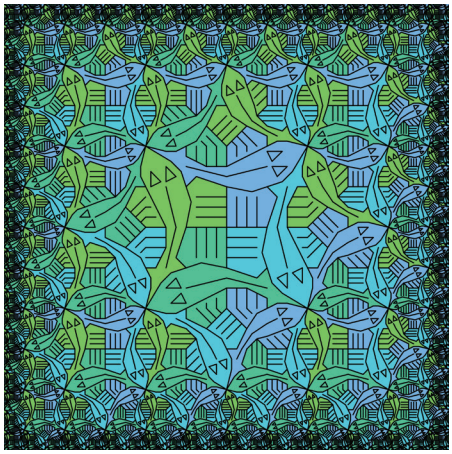

Peter and Mike have been employing schoolteachers and students to write and test these materials, and also developing a new online web forum which can be used to organise teaching and facilitate the submission and assessment of coursework.

To get a flavour of what they are doing, visit www.turtle.ox.ac.uk and take a look at the various example programs under 'TurtleOnline': here the basic concepts are illustrated in three (soon to be four) different languages, with access to the underlying machine code and memory, giving an insight into Computer Science. For more information on GeomLab, see: www.cs.ox.ac.uk/geomlab.

By the time the project is completed, Peter and Mike hope it will attract increasing numbers of schoolchildren towards this fascinating subject, which offers such wonderful prospects for interesting employment and is so important for the UK economy.

Alan Turing Institute to break new ground in big data

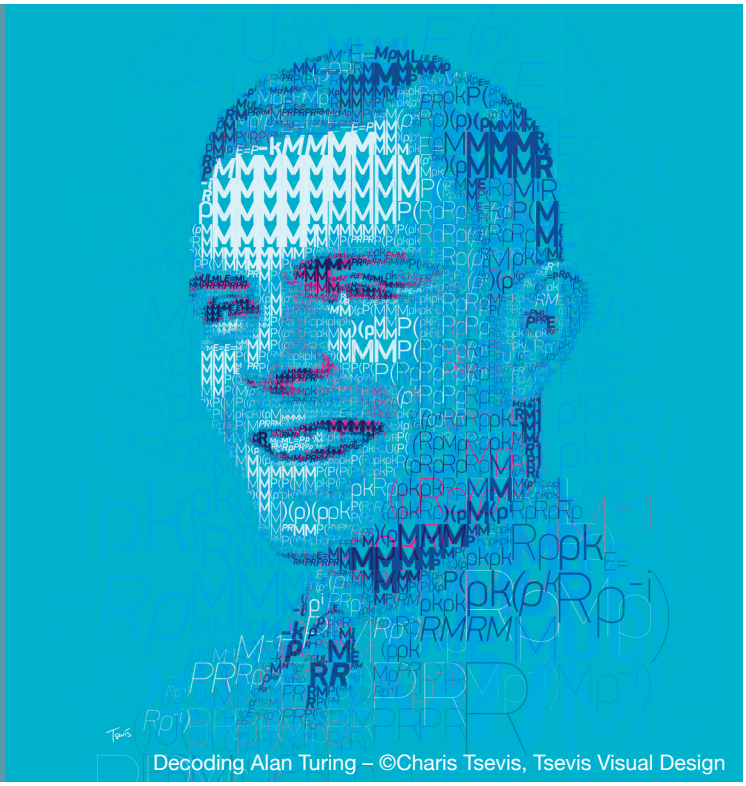

Decoding Alan Turing – ©Charis Tsevis, Tsevis Visual Design

Oxford has been chosen as one of five universities to lead the Alan Turing Institute. The aim behind setting up the new Institute is to build on the UK's existing academic strengths and help position the country as a world leader in the analysis and application of big data and algorithm research.

The Department of Computer Science will be one of five departments from the University of Oxford helping create and run the activities of the Institute, which will have its headquarters in the British Library in London.

Business Secretary, Vince Cable, who announced which universities would be involved this January, said: 'Alan Turing's genius played a pivotal role in cracking the codes that helped us win the Second World War… Headed by the universities of Cambridge, Edinburgh, Oxford, Warwick and UCL, the Alan Turing Institute will attract the best data scientists and mathematicians from the UK and across the globe to break new boundaries in how we use big data in a fast moving, competitive world.'

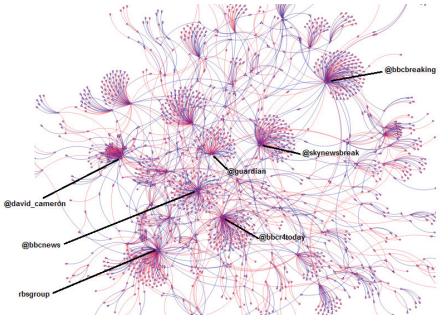

The delivery of the Institute is being coordinated by the EPSRC. The Institute is being funded over five years with £42 million from the UK government, with the selected university partners contributing further funding. In addition, the Institute will seek to partner with other business and government bodies.

Professor Philip Nelson, EPSRC's Chief Executive, said: 'It will use the power of mathematics, statistics, and computer science to analyse big data in many ways, including the ability to improve online security. Big data is going to play a central role in how we run our industries, businesses and services. Economies that invest in research are more likely to be strong and resilient; the Alan Turing Institute will help us be both.'

Researchers across the University of Oxford are already conducting world-class research in data science and analytics. The new Institute will tap into world-leading strengths and achievements across these scientific disciplines, for example the Department of Computer Science's research into machine learning and natural language processing.

The other departments from Oxford which are to be involved in the Alan Turing Institute are the Mathematical Institute, Department of Statistics, Department of Engineering Science, and the Oxford Internet Institute.

Professor Michael Wooldridge, Head of the Department of Computer Science, said: 'The UK has rightly come to recognise Alan Turing's essential contributions to code breaking at Bletchley Park, which played such an important role in bringing the second world war to an end. But for computer scientists, he is also recognised as one of the founders of our field. His seminal work during the 1930s laid the foundations for the digital computer. We are therefore delighted to be part of the institute bearing Turing's name.'

Cyber projects flourish

The University of Oxford Cyber Security Centre is moving from strength to strength, with a variety of new projects about to start and established research being put to use in the wider world.

The Cyber Security Capacity Portal (www.sbs.ox.ac.uk/cybersecurity-capacity) is a new resource for delivering effective cyber security both within the UK and internationally. It contains information for policy-makers and those with responsibility in this area and has been created by the Global Cyber Security Capacity Centre. The

Capacity Centre is now putting the Capacity Maturity Model to the test, with pilots in Jamaica, Colombia, Montenegro, Bhutan, Armenia and Kosovo in the past few weeks, and more pilots planned. The Centre also advised the Government of Jamaica on its new National Cyber Security Strategy, launched at the end of January 2015.

The first-year students in the Centre for Doctoral Training (CDT) in Cyber Security are about to start their mini-projects, working with supervisors across the University, and many with external sponsors too. The CDT will hold an open showcase day on 1 October for the students to show off what they have been doing.

The Cyber Studies Programme at the Department of Politics and International Relations is a major new research and teaching initiative on various aspects of cyber security and e-government; Cyber CDT student Florian Egloff has just published the first working paper for the Cyber Studies Programme, available at goo.gl/4tRQoP.

Joint research and teaching between the Department of Computer Science and the Saïd Business School has led to the creation of several new teaching resources for the Business School: Oxford provides the only executive education offering that looks at cyber security issues from a leadership rather than a technical perspective. Cyber security is now included in the core MBA syllabus, and the new Cyber Risk for Leaders Programme is an intensive two-day programme covering all the key questions necessary to understand and manage cyber risks in a large organisation. Enrolment is now open for the first programme, 8-10 June 2015. Contact: dean.worthington@sbs.ox.ac.uk, +44 (0)1865 422 767.

For the latest updates, see the (brand new) website at www.cybersecurity.ox.ac.uk, or follow @Ox_CyberSec on Twitter.

Celebrating the 200th anniversary of Ada Lovelace

This year the University of Oxford is celebrating the 200th anniversary of the birth of computer visionary Ada Lovelace. The centrepieces of the celebrations will be a display at the Bodleian Library and a symposium presenting Lovelace's life and work in the light of contemporary thinking on computing and artificial intelligence.

Ada, Countess of Lovelace (1815 – 1852), is best known for her remarkable article on Charles Babbage's Analytical Engine, which presented the first documented computer program. Lovelace had wide scientific and intellectual interests and studied with the scientist Mary Somerville and the mathematician Augustus De Morgan. The display in the Bodleian's new Weston Library (29 October – 18 December) will offer a chance to see Lovelace's correspondence with Babbage, De Morgan, Somerville and others, and her childhood exercises and mathematical notes.

The symposium on 9 and 10 December is aimed at a broad audience interested in the history and culture of mathematics and computer science, presenting current scholarship on Lovelace's life and work, and linking her ideas to contemporary thinking about computing, artificial intelligence and the brain. For full details and to register for further information, visit the website below.

Oxford's celebration is being led by the Bodleian Libraries and the Department of Computer Science, (organised by Professor Ursula Martin) working with colleagues in the Mathematics Institute, Oxford e-Research Centre, Somerville College, the Faculty of English and TORCH, The Oxford Research Centre in the Humanities.

blogs.bodleian.ox.ac.uk/adalovelace

Alumni round-up

It's now been nearly a year since we launched the Department of Computer Science Alumni Network, and I'm delighted to bring you this short round up of activity over the past months. I also hope that you enjoy reading about Simon Hay (Lincoln, 2006), this edition's alumni profile – do get in touch if you'd like to be featured next time.

December drinks event

Thank you to all of you who joined us at the Royal Society on 4 December to celebrate the launch of the alumni network. Mike Wooldridge and I really enjoyed meeting so many of our alumni, and we got a lot of useful feedback on the development of the network. We've included some photos from the evening here and the full selection is on our Flickr page. www.flickr.com/photos/computerscienceoxford/

Forthcoming events

Make sure you save the date for the 2015 Oxford Meeting Minds: Alumni Weekend, 18–20 September. Two members of the Department of Computer Science, Professors Nando de Freitas and Stephen Pulman, will be participating in a panel discussion on artificial intelligence on Saturday 19 September. This event will be open to all guests but we will be holding a lunch before the lecture exclusively for our alumni with the opportunity to meet Nando and Stephen.

The full programme will be announced in early June and tickets go on sale on Monday 29 June. I do hope that you'll be able to join us.

Get involved

Another thank you – this time to those of you who have given up their time and expertise to help the department. Over the last few months, our alumni have organised careers talks, led group design practicals, participated in outreach events and sat on committees. If you'd like to become more involved with your old department, then do get in touch.

And finally – keep in touch

We'd love to hear your news and find out what you've been up to. Please also let us know if you have any suggestions or comments about the network – we really want to make sure that the network works for you, so your feedback will make all the difference.

Frances Wheare,
Development and Alumni Manager

Mentoring scheme launched

The Oxford Women in Computer Science Society (OxWoCS) and the Department of Computer Science Alumni Network have joined forces to launch the Computer Science Mentoring Scheme, and we would like to invite you to participate.

The scheme has been developed to offer support for current and former members of the department by introducing them to people who can help them to navigate their educational and professional development. This can range from offering insight and guidance to more practical solutions such as a work placement or making introductions. Additionally, the act of sharing and articulating one's situation can have valuable benefits of its own. Mentors can also benefit from participating in the scheme by developing leadership expertise and seeing how they can apply their skills and experience to a range of scenarios.

continued on next page ▶

Alumni Profile

Simon Hay (Lincoln, 2004)

Simon Hay read for a degree in Computer Science at Lincoln College. Together with fellow alumnus Joe Mathewson (PPE, 2004), he founded educational technology start-up Firefly which was named by Forbes as one of five British digital businesses to watch in 2015. Here he talks about what inspired him to start his own business as well as his experiences at Oxford.

What inspired you and your co-founder to start Firefly?

My co-founder Joe and I wrote the first version of Firefly together as 14-year-old school friends. We were frustrated that the Internet was making it easy to access a world of information from anywhere – except for information we needed about school. We'd been tinkering with old computers and learning to code and thought we could make it easier for teachers who weren't computer experts to create and share things online – whether that was sharing teaching resources with each other, homework with us or news with our parents. So the business was really born out of our own frustration. We hoped it would make our own learning easier but we never imagined it would become a serious business.

How did your time at Oxford help to prepare you for starting your own business?

As well as a solid technical foundation, of course, which is crucial to the success of any software start-up, I think my time at Oxford really helped me with confidence, organisation, leadership and other, softer skills which are equally important. Joe and I both went to Lincoln College together, and tried to fit in growing our nascent business amongst lectures, tutorials and other commitments, which really teaches you the importance of managing your time efficiently, and prioritising ruthlessly!

What did you enjoy most about your time at Oxford?

I did a lot at Oxford apart from my degree, and I'm really proud to have rowed for the university lightweights against Cambridge, and learned to fly with the University Air Squadron, for instance. I met many of my closest friends – and my wife! – through getting involved with clubs and societies outside my college. But I think my degree also convinced me that computing really is the thing I'm most passionate about – in any event, I enjoyed it enough to go back for more and do a PhD in Computer Science, too.

What advice would you give to someone who is unsure of where their degree in Computer Science might take them?

First of all, I think it's OK not to know. Plenty of people aren't sure what they want to do when they leave university, and I don't think that's necessarily a problem. I certainly wouldn't have predicted 10 years ago that I'd be doing this now, and I think that's true for many of my colleagues and friends as well. A Computer Science degree definitely keeps your options open: people with analytical and quantitative minds will always be in demand, whether you use the skills you learn on the course directly or not. It's worth looking around, and casting your net a bit wider than the traditional milk round. If you can, try a few internships in different fields to see what you enjoy most, and workinstartups.com and siliconmilkroundabout.com are good places to start if you think you might enjoy being part of a new business.

from previous page ▶

The scheme is open to people of any gender who have a connection to the department, from current students and alumni to teaching, research and support staff, both past and present.

Although support and facilitation will be provided by OxWoCS and the Computer Science Alumni Network, once a mentor-mentee partnership has been established it will be up to the participants to set the frequency and format of their meetings. The scheme is intended to be flexible but we do ask that you make a commitment to meet three times a year (once a term if you are within the academic system).

If you would like to participate, please send a brief email to mentoring@cs.ox.ac.uk stating your name, connection with the department and whether you are interested in being a mentor or a mentee. We will then send you a link to register for an online service which allows mentees to contact mentors directly.

The scheme is intended to be shaped by its participants, so please don't hesitate to get in touch using the email address above if you have any comments, feedback or questions.

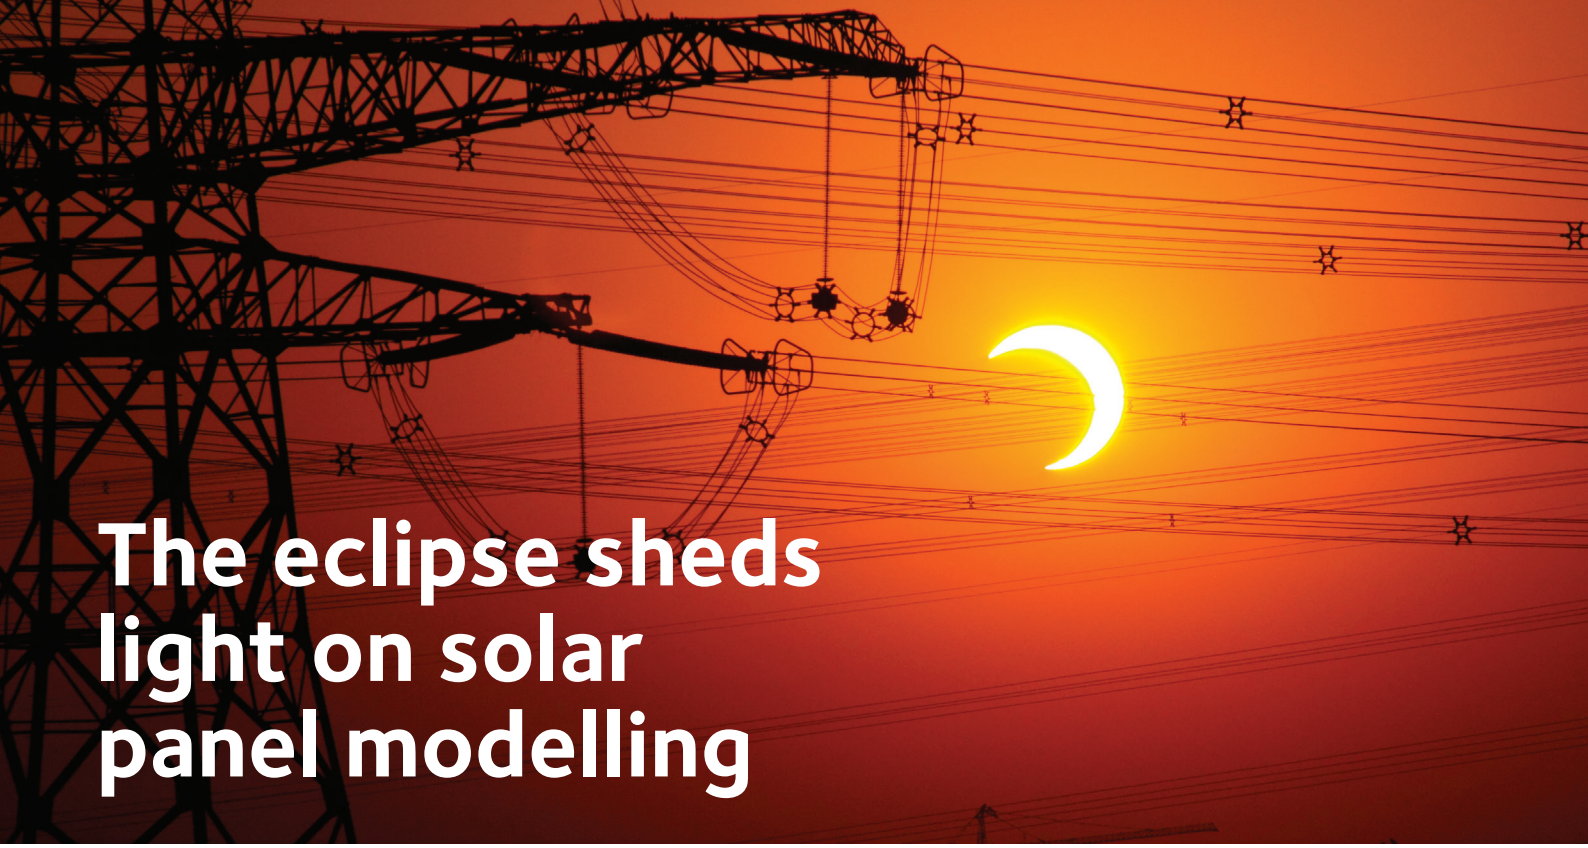

The eclipse sheds light on solar panel modelling

Mathematical models being developed at the Department of Computer Science to help predict the behaviour of renewable energy sources are going to be tested with data gathered from the partial solar eclipse. Professor Alessandro Abate, who is leading the work, describes how such models should result in a more robust and reliable electricity grid.

As the sun was disguised by the moon on 20 March, it led hundreds of thousands of solar panels to tune down as they experienced sudden partial darkness, and then jointly and rapidly turn on again as sunlight fully returned. Data from how solar panels reacted during the eclipse will inform mathematical models, which are developed to help ensure the stability of the power grid when energy sources don't perform steadily.

Solar panels now feed substantial amounts of electricity into the power grids of countries including Germany, Italy, and France: in the summer up to 40% of Germany's electricity comes from its solar farms. Whilst reducing greenhouse gas emissions, renewable energy sources, such as solar and wind, introduce uncertainty into power grids because of the oscillation in power production levels due to local weather variations.

The last European solar eclipse was in 1999, before the pervasive proliferation of solar energy sources feeding into grids. It was therefore not known what impact an eclipse, and the subsequent synchronised behaviour of large populations of solar panels, would have on the stability of European power grids, which were engineered in the pre-renewable era. Abrupt events such as a solar eclipse are thus relevant to the study of less infrequent grid instabilities, such as those related to local power failures, which can potentially result in a cascade of blackouts.

This eclipse provided a rare opportunity to challenge in a worst-case scenario the mathematical models developed to predict the behaviour of large populations of solar panels. With very many solar panels of different types, brands, and set-ups sited in different locations, current models have to cope with a lot of variables: the team in the Department of Computer Science believes that new, simpler, yet precise, models can be obtained to reliably predict the overall behaviour and its likely impact on the stability of the power grid. In the longer term, the quantitative prediction of the behaviour of populations of solar panels could enable the mitigation of grid instability.

The team hopes to contribute to helping design strategies to ensure that power grids are more predictable, robust, and reliable. The models under study are not only relevant to solar power, but should also be useful for predicting the behaviour of other renewable energy sources such as wind power. As we move towards ever more interconnected grids in which renewables are providing more and more power, there is a great need to understand how these variable energy sources can engage with electricity grids that are improved to fully accommodate them.

Where on earth...?

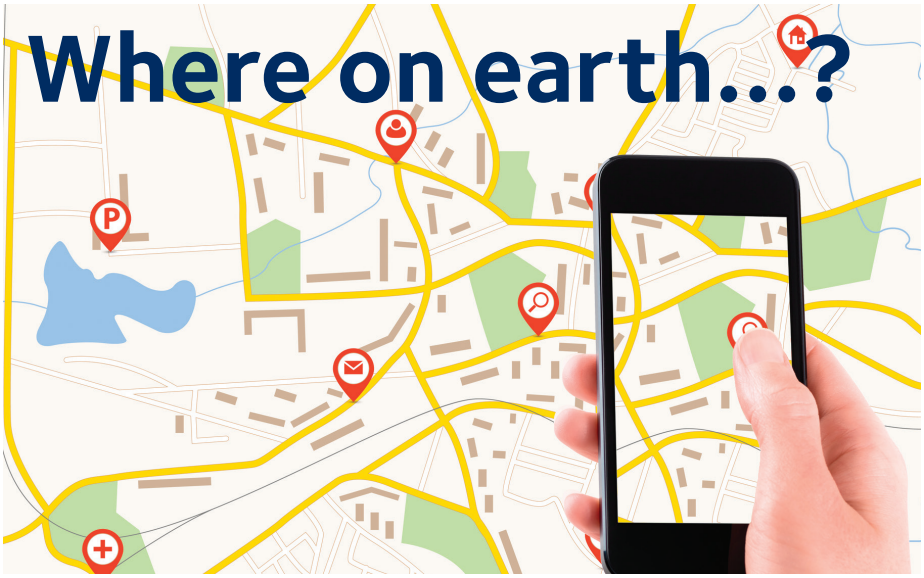

A revolutionary approach to GPS-free navigation that works reliably regardless of network connections is described by Dr Fred Kemp.

There is increasing demand for smartphone applications and services that are location aware. Whilst many phones are GPS-capable, satellite reception is notoriously inconsistent in urban environments and frequently non-existent inside buildings.

To combat this, many phones and fitness trackers contain a range of other sensors, such as a compass, accelerometer and gyroscope, and use these, often in combination with WiFi-based systems, to provide position estimates. However, such 'dead-reckoning' systems are particularly error-prone, and are thus ineffective without constraining the user's track within the context of an indoor map.

This leads to another problem; while indoor maps are becoming increasingly available, until now there has not been a practical and reliable indoor 'map-matching' solution, which is also capable of running directly on a mobile device.

New standalone approach

To overcome this, the Sensor Networks group in the University of Oxford's Department of Computer Science have taken a revolutionary new approach to indoor positioning that is robust, reliable and requires no user input or infrastructure (such as WiFi or Bluetooth beacons).

Using pattern recognition and machine learning techniques that have had great success in the field of natural language processing, but have previously not been applied to GPS-free positioning, they created a flexible and lightweight method to combine data from mobile phone sensors with a digital map. The map then acts as a 'sanity filter', thus avoiding 'impossible' situations such as a user appearing to walk through a wall.

MapCraft is a smartphone app combining this map-matching technique with a number of unique refinements to the sensor-based dead-reckoning process and, importantly, the ability for the system to learn. Not only does MapCraft learn about the user (e.g. typical step length, gait, how they tend to carry the phone, etc), but also about the environments in which the user is navigating. Information about corrections derived from one user can be fed back into the metadata associated with the map, and thus improve navigational accuracy not just for the returning user, but also for new visitors, who would automatically benefit from this crowd-sourced data when they download a map.

Outstanding test results

The system has already been tested in a number of environments and outperforms all competing solutions, working reliably across different users, motion modes and devices. Consequently the system has received a lot of positive interest, winning awards at a number of conferences and attracting praise from industry leaders.

MapCraft could have significant impact throughout our daily lives. Imagine visiting a museum or exhibition centre, and being able to use your smartphone for turn-by-turn navigation; to receive notifications of other exhibits to see, based on where you have been so far; or in the shopping centre, to quickly locate the item you viewed online earlier (along with price matching of course!); or in a crowded bar, to order a drink without leaving your seat, and for the waiter to know exactly where to bring your order; or to know, as you enter the building, where exactly 'ward 37f' is, when visiting friends and family in hospital.

We might still take the same (small) steps through our urban spaces, but with Mapcraft's giant leap in GPS-free navigation, the journey can be easier and perhaps a little more interesting.

This article first appeared in Isis Insights*, the quarterly magazine of Isis Innovation, the University's technology commercialisation company.*

> ***'Indoor localisation through consumer mobile devices will revolutionise how people interact with their physical environments and with each other. Among more than a decade of research efforts in indoor localisation, MapCraft is one of the most elegant solutions, especially for its low sensing and computation cost and high robustness. It has clear value proposition and the potential to be widely adopted.'***
>
> **Jie Liu, Principal Research Manager, Microsoft.**

News in brief

The city of Oxford has been listed as one of the top 10 UK digital clusters in the 2015 Tech Nation Report 'Powering the Digital Economy'. The interactive data project, developed by government-backed organisation, Tech City UK, shows the growth of digital businesses within specific UK clusters. The Department of Computer Science, the Capacity Centre, the Oxford e-Research Centre and the Oxford Cyber Security Centre are all namechecked in the Oxford cluster.

The FDB research project led by Dan Olteanu has received a US$10,287 Amazon Web Services research grant. The support from Amazon will enable benchmarking of distributed techniques for query processing and machine learning that exploit compact, factorised representation of data to mitigate network communication cost and offer horizontal scalability. It is the next milestone in the FDB (Factorised DataBases) project, which addresses the problem of processing large-scale relational data via compression.

Head of Department Michael Wooldridge and his co-authors Shaheen Fatima and Sarit Kraus have published a new book: *Principles of Automated Negotiation*. The book explores key issues in the design of negotiating agents, and discusses the potential benefits of automated negotiation, its unique challenges, possible applications, and current deployments.

Three teams representing the Department of Computer Science took place in the 55th Annual Teddy Hall Relays in March. Nearly 160 teams and 600 athletes competed. The race begins from the track where Roger Bannister broke the 4 minute barrier. Each leg of the relay is 3.6 miles.

Our pick of the latest podcasts and vodcasts that feature computer science research at Oxford

- **Deep Learning**
  A lecture series on deep learning by Nando De Freitas.
  goo.gl/y5Ey9e
- **Oxford and Cybersecurity**
  Andrew Martin, Sadie Creese, and Ian Brown offer an in-depth look at the issues facing Oxford's Cybersecurity researchers, across many of their disciplines.
  goo.gl/E3C99W
- **Adventures in Mathematics**
  Helen Byrne talks about her work in Mathematics and Medicine using Computer Science
  goo.gl/7oDu4y
- **Multilingual Distributed Representations without Word Alignment**
  Karl Moritz Hermann and Phil Blunsom – A short lecture
  goo.gl/rbo0Nq

Student recognised for outstanding efforts

Andrew Paverd received one of the 2015 Outstanding Student Awards from the Worshipful Company of Information Technologists (WCIT) this February.

WCIT is the 100th City of London Livery Company and focuses on four key areas: Charity, Education, Industry and Commerce, and Fellowship.

Andrew, who is in the final year of his DPhil in the Systems Security Group, was nominated by the department for academic excellence and contribution to community. As a voluntary Computer Science Student Ambassador, Andrew regularly gives presentations to visiting school groups with the aim of inspiring students to pursue degrees in Computer Science.

Earlier this year, Andrew was selected as one of five young researchers from Oxford, and the only one from his department, to participate in the multidisciplinary Global Young Scientists Summit in Singapore where participants from around the world met and interacted with Nobel Laureates, Fields Medallists and Turing Award winners.

Pictured above: Andrew Paverd receiving the WCIT award.

Fry's work explored with technology

Stephen Fry was at the Bodleian Libraries on 30 October for a *YourFry Hackathon* which explored his work through technology.

'Groups of artists, writers, filmmakers, scientists and programmers worked on projects as diverse as short plays, videos, websites and Twitter bots, all based around the newest instalment of Stephen Fry's autobiography, said doctoral student Beth McMillan. 'My team used cluster analysis, Markov chains and a lot of Perl to make a website pulling out different pieces of information from the text.'

Industry links in the Department of Computer Science

The Department of Computer Science has a long history of working with industry across a range of sectors, allowing companies to raise their profile within our research and student community. Here are just some of the ways in which companies can get involved.

Tech Talks on Tuesday are lunchtime lectures or seminars given by invited companies to students during term time. Previous talks have covered a range of subjects, from Bitcoin to the challenges of a 24-hour news cycle, and offer students an insight into the practical applications of their degree.

Companies are also invited to set challenges for second year students who undertake group design practicals. Industry partners are encouraged to actively participate in the process, scheduling regular progress updates and offering talks that cover a range of 'soft skills' such as project and time management, team working and leadership skills.

Some group design practical participants also choose to sponsor a prize, based on their own criteria. Other companies elect to sponsor a student prize for best performances in the end-of-year examinations.

The Careers in Computing Fair, held every winter in association with The Careers Service, offers companies the opportunity to target Computer Science students from Oxford and Oxford Brookes, as well as those studying other relevant degree subjects such as Engineering, Physics and Statistics.

The department has had a number of successful industry research partnerships over the years, responding to key issues faced by industry, government and society. We are always happy to discuss in-depth research collaborations, as well as opportunities for philanthropic support to meet corporate social responsibility objectives.

For more information or to discuss your requirements in more detail, please contact industry@cs.ox.ac.uk.

News in brief

Benjamin Neal, the Publicity and Schools Liaison Apprentice in the Department of Computer Science, was joint winner of this year's apprenticeship prize for an outstanding contribution to his department. Ben and co-winner Kathryn Scott (Engineering Science) are students at Abingdon and Witney College.

Junior Research Fellow, Thomas Gibson-Robinson, has won an award for his contribution to science outreach from Oxford's Mathematical and Life Sciences Division. He, along with other student ambassadors from across the division, received their awards at a reception at the Natural History Museum in October.

Oxford's Professor Michael Benedikt, along with his two co-authors, have won an award for the impact of their paper 'XPath satisfiability in the presence of DTDs' over the decade since it was published in the Principles of Database Systems (PODS) proceedings. The 2015 ACM PODS Alberto O. Mendelzon Test-of-Time Award recognised the paper as having a large number of citations, influencing many researchers, and today being considered the standard reference for complexity results on the satisfiability problem for XPath expressions.

Dr Irina Voiculescu has received a John Fell OUP Research Fund award. This will enable her to employ a postdoctoral researcher for a year to extend some of the Spatial Reasoning research group's work on segmentation and 3D modelling from medical images. This work will be carried out in collaboration with the Nuffield Department of Orthopaedics, Rheumatology and Musculoskeletal Sciences.

Professor Tuomas Sandholm from Carnegie Mellon University gave the Hilary Term Strachey Lecture, entitled 'Modern Organ Exchanges: Generalized Matching under Dynamics, Edge Failures, and Fairness'.

Prize-winning model predicts cardiac variability

A new approach to computer modelling has won the 2014 3Rs Prize awarded by The National Centre for the Replacement, Refinement and Reduction of Animals in Research (NC3Rs). The work has been recognised for its potential to reduce the number of animals used in research, particularly in the safety assessment of new drugs.

It is the first time that the 3Rs prize has been awarded to a Computer Science paper since its launch in 2005, and the first time that the highest accolade has gone to a PhD student: Oliver Britton from the Department of Computer Science.

The winning paper by Oliver and others is entitled 'Experimentally calibrated population of models predicts and explains inter-subject variability in cardiac cellular electrophysiology'. The paper, a collaboration between the Department of Computer Science, Oxford's Centre for Collaborative Applied Mathematics, and Janssen Pharmaceutica, Belgium, was authored by Oliver Britton, along with his supervisors, Alfonso Bueno–Orovio and Blanca Rodriguez, plus Karel Van Ammel, Hua Rong Lu, Rob Towart and David J. Gallacher.
The authors have built a computer model of cardiac electrophysiology that incorporates natural variability. Normally when using a computer model to test how a drug might affect the heart, its effect is compared to an average profile of electrophysiology. But this average profile is not really representative of the whole population, where natural variations in heart properties occur from person to person.

This new approach has the potential to make computer models that are far more powerful and more predictive of human response, and therefore a more viable alternative to using animals in research. This is the first time that natural variability has successfully been considered in such a model, and the methodology could be applied to other diseases. The authors are already planning to use the same methodology to build computer models for understanding pain and diabetes.

Professor Ian Kimber OBE, 3Rs Prize panel chair, said: 'Mr Britton's paper really stood out to the panel because of the model's potential to be used to replace early-stage animal tests in drug safety studies, across a broad range of disciplines. The model has also been developed into a piece of user-friendly software, encouraging uptake and use by industry, which could have large scale impacts on the reduction of animals in research.'

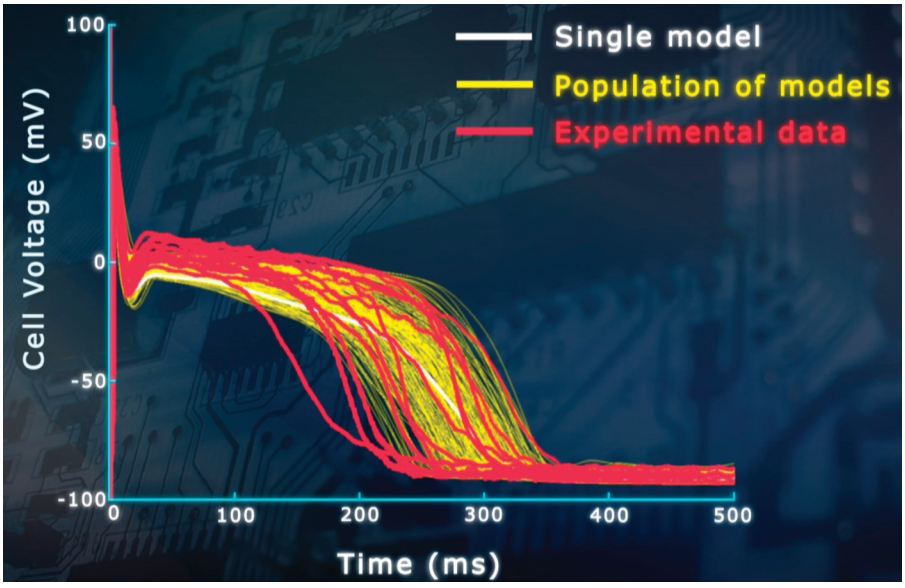

Industry uptake

The computer model has been developed into a user-friendly software package called Virtual Assay, which should increase industry uptake for use in drug safety testing, as it can be used without the need for specialist programming and modelling experience. It is possible to use the modelling method to study variability in any biological system, increasing the potential of the model to reduce animal use. The software Virtual Assay has been developed by Oxford Computer Consultants, funded by the EPSRC Impact Acceleration Account awarded to the University of Oxford, and supported by Isis Innovation.

The 3Rs prize, presented at an evening celebration in London in March, consists of £18,000 to be used to further the 3Rs potential of the study and £2,000 for personal use. It is awarded for an outstanding contribution to scientific and technological advances in the 3Rs in medical, biological or veterinary sciences.

Oliver said: 'It's great to have our research recognised as having a potential impact on the 3Rs, and winning the prize helps with publicising our work to the audience who might be interested in it. We plan to use the prize grant to apply our methodology in neuroscience, specifically pain research in humans, through a collaboration we have developed with an industrial partner. We will investigate how variability between individuals alters the response of different types of pain-sensitive neurons to drugs.'

Further information: www.cs.ox.ac.uk/ccs/tools
Podcast in which Oliver explains his work: goo.gl/ktS6og

Security under scrutiny as flights abound

With flights growing rapidly in number, air traffic control needs new technologies to cope with the increase. A novel protocol has therefore been mandated for use in Europe from 2017, and the University of Oxford is working on authentication techniques to improve its security. In this article, Martin Strohmeier, a DPhil student, describes the various issues and the department's work on resolving them.

Air traffic control is the backbone of what is arguably the key means of transport in the modern world, and it has to deal with ever more aircraft as the traffic load is growing tremendously.

Large European airports such as London Heathrow experience spikes of more than 1,500 daily take-offs and landings and industry forecasts predict a doubling of worldwide flight movements between now and 2030. With the growing adoption of unmanned aerial vehicle (UAV) technology for civil applications, we can even expect a further boost in air traffic over the coming years.

Traditional air traffic control relies on radar systems originally developed for use in military applications. The requirements of the increased traffic density exceed the physical limits of such conventional radar systems and new technologies are needed. The future of air traffic control is moving away from mere object detection systems towards flexible modern data communication networks that vastly improve the situational awareness of pilots and ground controllers.

The Automatic Dependent Surveillance Broadcast system (ADS-B) is an example of this future. ADS-B uses satellite based positioning to continuously and automatically broadcast location updates, velocity and other data to both ground control and nearby aircraft. It is mandated for use from 2017 in Europe (2020 in the United States) as part of the next generation of air transportation systems. As of today, most airlines have started to upgrade their fleet, and about 80% of aircraft in Europe are already sending out ADS-B signals.

Unfortunately, the protocol and many related ones are also highly insecure. Since it does not offer any type of authentication, an attacker can exercise full control over the messages sent over the wireless channel. The easy manipulation of ADS-B radar signals allows an adversary to inject fake ghost aircraft, change the identity of an aircraft, make it disappear to ground control, and execute many other potentially fatal attacks. Consequently, recent high profile cases of aircraft incidents such as the disappearance of Malaysian aircraft MH370 and the hijacked emergency signals created a lot of speculation in the media about the potential impact of insecure protocols on the safety of air traffic.

In the Department of Computer Science at the University of Oxford we have been working on the security aspects of current and future air traffic communication over the last three years. Initially, we analysed the practical feasibility of attacks on ADS-B, and have since been researching the best ways to improve the security of the protocol before its widespread operational use.

Since existing legacy technology and other real world constraints found in the aviation industry impede the use of cryptography in air traffic communication for the foreseeable future, we are currently focusing on cyber-physical countermeasures. We use physical layer information to, for example, securely verify the location of aircraft or provide an alternative way to authenticate ADS-B messages. Such techniques are able to instantly improve the security of the air traffic system without requiring protracted and costly changes to existing hardware and software.

Research requires data

For this research, we require large amounts of real ADS-B data from commercial aircraft in Europe. To facilitate the collection of such data, we are a founding member of the OpenSky initiative (www.opensky-network.org), together with our European partners. Overall, OpenSky currently has captured and stored about 15 billion messages for further research, providing surveillance of an area of more than 1 million km^2, and observing ca. 30% of Europe's daily flight movements.

Finally, we are also working to raise awareness of cyber security issues with air traffic control stakeholders such as ground controllers, pilots, and system engineers. We hope that our continued efforts will lead to measurable improvements in air traffic security as the next generation of protocols takes over.

Has artificial intelligence become a threat?

Some commentators have recently suggested that artificial intelligence may pose an existential risk to the human race. In this article, two of our professors set out their views on the development of AI and how they think it could affect society.

Michael Wooldridge,

Professor of Computer Science

Since the field was founded in the mid-1950s, artificial intelligence (AI) has enjoyed sharply mixed fortunes. Early results on programs that could undertake primitive forms of learning and planning gave reason for optimism. By the mid-1970s, the dominant paradigm for AI, often called *symbolic* AI, tried to directly model reasoning and explicit decision-making processes. However, it soon became clear that symbolic AI had fundamental – and so far insuperable – problems.

These difficulties led some researchers to reject the symbolic paradigm and look for fundamentally new approaches. One idea was to use computers to simulate evolutionary processes, in order to artificially evolve intelligent systems, but this hasn't provided a route to AI yet. Another idea was *connectionism*, which tried to construct neural nets to model brain-like processing. Early results were not encouraging, and the field languished until the mid-1980s when real progress began to be made. But it is this century that has seen the area explode, and much of the current interest in AI arises from the huge progress in neural nets over the past 10 years.

I should put my cards on the table right now and say that, as an AI researcher of 25 years' standing, I truly believe that the neural net systems I have seen over the past year represent a real breakthrough. Things have become possible in the past few years that would have been unthinkable a decade ago. Fundamental advances in the algorithms to handle neural nets led to a new area called *deep learning*, which makes it possible for neural structures to handle much more challenging tasks, and we now have the necessary data and processing power to develop such neural nets. We have only just begun to understand what we can do with this incredible new capability.

Earlier this year, deep learning research hit the headlines when Google DeepMind (the London-based AI research arm of Google) demonstrated a system that taught itself to play 49 video games from the 1980s Atari console. After about eight hours of standard desktop time, the program could play these games better than a human, having taught itself to play *unsupervised*, using only the video feed from the game as input.

So where next? Some pundits have made very public statements about how AI might affect society, even claiming that AI may pose an existential risk to the human race itself. What happens when we have constructed systems that are smarter than we are?

I don't see any immediate cause for concern here. Impressive though the capabilities so far demonstrated are, they are mainly focused on very specific, narrow tasks, and are a long way from human-level intelligence. One issue with deep learning seems to be that it is hard to build programs that can do long term reasoning – tasks that require thinking about the future. The 'Skynet' scenario from the *Terminator* films, where a machine becomes sentient just hours after being switched on, makes a thrilling storyline but doesn't seem likely – although for the record, I don't recommend giving control of the world's nuclear arsenal to an AI program!

So how will these AI technologies affect us? I think the exciting thing is that *intelligence is going to be embedded in everything we build*. This is going to change society fundamentally. Perhaps the most obvious and immediate concern is that AI techniques are going to make a huge number of workers redundant across the globe, just as the automation of factory production lines did in the 1970s and '80s. This time, the complexity of tasks being automated will be higher, so the effects of the changes will not be limited to unskilled labour. I don't think these changes imply any obvious existential risk for the human race, but we do need to think about how society will respond.

continued on next page ▶

from previous page ▶

Peter Millican,

Professor of Philosophy

Mike gives a nice summary of the historical development of AI, with periodic bursts of optimism followed by a gradual realisation, as hopes failed to materialise, that things were much more difficult than expected. But the recent dramatic rise of deep learning may point towards a real paradigm shift in this mixed history. If AI is about to transform our ability to analyse and predict over a wide range of new domains, then this naturally raises fears about the implications for our society.

Some of these fears are over-hyped, and, like Mike, I see no need to worry about machines suddenly becoming self-aware and deciding to destroy us. Conscious thought and feeling remain deeply mysterious, but we have no reason to suppose that mere processing power will produce them. We know, at least in outline, why the machines we create are able to display intelligence, and it has nothing to do with the sort of conscious awareness that we experience. So there is no risk here – unless, just possibly, we move to the stage of creating self-reproducing machines able to evolve in novel ways beyond what we have anticipated.

Of far more concern is the potential imbalance of power amongst *humankind* that these new developments might bring, not only in respect of employment (as Mike highlights), but also in economic life and society more generally. To take just one example, if dedicated AI share-traders, programmed with proprietary algorithms that learn from (and take advantage of) our foibles, can wipe the floor with the rest of us, then we could see yet further concentration of wealth in the world while our pension funds – and much else – suffer as a result.

The economic structure of our society is predicated on a basic assumption that it is not fundamentally biased or open to exploitation, and that market trading will tend to benefit all participants. But if this is wrong – if the dice can be systematically loaded by those in the know in a way that makes it impossible for those without AI superpowers to compete – then the way we organise these things will have to be re-thought. Such re-thinking of principles is difficult, especially when our politics is so heavily influenced by established financial and media power, and while economic 'common sense' is dominated by the complacent dogma that the market is both irresistible and legitimate, wherever it might lead.

Major technological changes, whether in warfare, transport, manufacture or medicine, typically require social change, but the latter always lags well behind. Unlike previous technological leaps, moreover, AI has the potential to transform all of these fields, more comprehensively and much faster than other historical changes.

Fortunately we live in a democracy, where it is at least possible in principle to adapt our systems so that our developing technologies serve the interests of society as a whole rather than a small elite. To achieve this, however, our population as a whole has to be engaged and educated in ethical and political thought, rather than seduced into frittering its intellectual energies in entertainment (where again AI has the potential to be a game-changer). People have to be able to think through the basis of how our world is organised, and to consider alternatives beyond the current orthodoxies.

If the new wave in AI proves to be as disruptive as some anticipate, then education in philosophical and critical thinking, historically informed by the views of other societies and by thinkers outside the current paradigms, has never been more vital to our collective well-being.

Sentiments spin out of control

In the recent Scottish referendum, emotions ran high. It was the perfect opportunity for TheySay, an Oxford spin-out, to demonstrate the potential of its technology to analyse sentiments. Professor Stephen Pulman, co-founder of the company, and David Morgan, CEO, describe how their machine monitored voters' emotions.

In the early hours of 19 September 2014, emotions in the United Kingdom, and more specifically Scotland, ran wild. A potential turning point in the history of the UK constitution had been reached, determined by a huge turnout to vote 'yes' or 'no' to Scottish independence.

In the run-up to the vote, gladiatorial television debates between Alex Salmond and Alastair Darling resulted in minute scrutiny of their personalities, and issues like 'currency' became pitched battles.

The strident rhetoric from both sides was also matched by vitriolic exchanges across social media. Instant emotion, discussion and opinion were being expressed on a huge scale. Social media, news, blogs and chat rooms were seething with outrage, delight, despondency and joy.

Between 2 and 19 September alone, we classified 1.225 million tweets as showing positivity towards the 'yes' campaign and 0.525 million towards 'no' using a sentiment analysis algorithm; the volume of other types of text was equally huge. Note that we immediately learn that voting outcomes cannot be predicted on the basis of social media volume: the 'no' tweeters were clearly a small subset of the 'no' voters.

It is a challenge both for people and machines to track and understand those emotions and opinions. No-one could possibly read everything, let alone aggregate insights into how people are feeling, across all the multi-dimensional loyalties, biases and events expressed via thousands of messages an hour. But TheySay, a spin-out from Oxford's Department of Computer Science (set up in 2011),

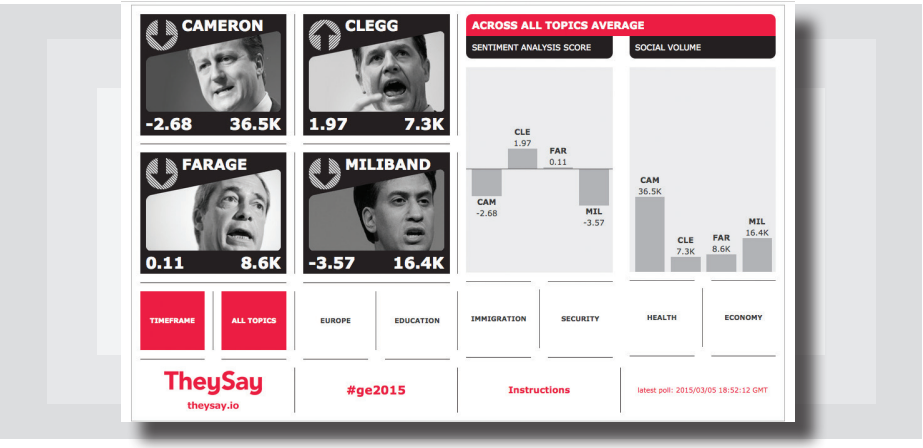

has built a sophisticated sentiment analysis algorithm that analyses huge amounts of data in near real time. Using a combination of machine learning and more traditional rule-based techniques the system builds a picture of the relationships between the words in a phrase, sentence and document in a very similar way to which humans understand text. Then it uses this analysis to detect sentiment (positive or negative), emotion (fear, anger, happiness and sadness), and sureness (certainty and confidence). The company sells its services to businesses who wish to know what consumers are saying about their brand or a particular topic on social media sites, as well as to customers in healthcare monitoring and financial services.

Can we trust the machine's interpretations? In the Scottish referendum, there were peaks of activity around the televised debates, but the real activity started to ramp up in the seven days up to the day of the vote, and for hours after. TheySay analysed all of the Twitter traffic from 2 to 19 September, and in that last week, things really got hot.

Instinctively, we would guess that there would be euphoria from the 'no' voters when the result became a foregone conclusion, matched by despair (with a smattering of anger) from the 'yes' voters. TheySay's technology detected exactly that. It also showed that during the week before the vote, 'yes' voters showed high levels of certainty and confidence, and more extreme levels of anger, fear and happiness. With 'no' voters, there was less certainty and confidence (with noticeable drops and peaks in fearfulness), and flat readings for happiness and sadness. In contrast, in the early hours of the day after the vote, the 'yes' voters descended into abject misery whereas the 'no' voters were, not surprisingly, bouncing with positivity.

TheySay is repeating this exercise for the 2015 General Election: ge2015.theysay.io

This time, four party leaders, Cameron, Miliband, Clegg and Farage, are tracked against key issues like the economy, education and Europe. We display time slices: e.g. today, last week and last month. We display numbers for Twitter and other message volume, as well as sentiment scores for each party leader in the aggregate, and for each specific issue.