



NEWS FROM THE DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF OXFORD

BEST IN THE WORLD

FOR COMPUTER SCIENCE TEACHING

according to *Times Higher Education*



ARE SELF-DRIVING CARS SAFE?

Developing techniques to check neural networks – p6



ARTIFICIAL INTELLIGENCE @ OXFORD

A university-wide expo to show AI reality – p8



WHEN TO TRUST ONLINE RATINGS

Research into the robustness of rating systems – p27



DEPARTMENT OF
**COMPUTER
SCIENCE**

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



Editorial board

Suzanna Marsh (Managing Editor)
Kiri Walden (Editor)
Helen Wilcox (Editor)
Suvarna Designs (Designer)

Benjamin Neal
Emma Dunlop
Faris Abou-Saleh
Jason Nurse
Kelly Ryan
Maureen York
Morteza Lahijanian
Peter Jeavons
Rosanna Cretney
Sarah Baldwin
Sharon Lloyd
Stefano Gogioso
Varun Kanade

Contributors

Boris Motik, Elizabeth Phillips, Helena Webb, Jamie Vicary, Janpreet Khabra, Jeremy Gibbons, Jo Ponting, Jolie de Miranda, Jordan Summers-Young, Katherine Fletcher, Mariam Nouh, Marina Jirotko, Marta Kwiatkowska, Martin Dehnel-Wild, Michael Wooldridge, Monica Wilson, Paula Boddington, Reuben Binns, Tim Muller, Vincent Taylor.

Letter from the Head of Department

One of the fascinating aspects of working in Computer Science is witnessing the relentless waves of innovation that sweep over the technology sector. I belong to the first generation that had computers in their home (a Sinclair ZX-80 in January 1981). This was the very beginning of the desktop computing revolution: the 1980s was all about getting computers onto desks, while the 1990s was all about the internet, and the past decade has been all about mobile computing.



The latest wave of innovation to sweep over the sector is artificial intelligence (AI), and in particular, a branch of AI called machine learning (ML). The level of interest in AI and ML right now is truly astonishing. Scarcely a day goes by without a breathless news article about some new AI startup being sold for dizzying sums of money, or a new AI application that will apparently change our world.

Although it is important not to get too carried away with the AI hype that fills the popular technology press, the truth is that AI and ML have seen impressive advances over the past decade. That doesn't mean that we will see robot butlers anytime soon, or that conscious machines are on the horizon. But AI and ML techniques are proving themselves in a huge range of new application areas, and this is genuine cause for excitement.

In this issue of *Inspired Research*, you will read a lot about AI and ML, because Oxford is at the heart of the current AI explosion. Researchers from the Department of Computer Science are at the very forefront of AI development, just as much in classic AI areas such as reasoning as in breakthrough new areas such as deep learning. One activity you might like to take note of is the planned 'AI @ Oxford' day (see p8), where we aim to showcase Oxford's expertise in this area across the whole collegiate University.

Of course it's not just AI in which the University and our department excels, and both have done tremendously well recently in various league tables. The *Times Higher Education Supplement* publishes an annual ranking of world universities, in which, for two years running, Oxford has been placed first. For us in the department, the more exciting news is that Oxford was ranked third overall in the world for Computer Science, and first in the UK, for the second year running.

Any given ranking of universities or university departments will have flaws, and for this reason it is wise not to take individual rankings too seriously. But, when all the rankings point the same way, then I think you are safe to draw some conclusions. And right now, all the indicators are that Computer Science at Oxford is in rude good health.

Professor Michael Wooldridge
December 2017

Professor Rahul Santhanam,
tutor at Magdalen College



Oxford's Computer Science teaching is best in world

Teaching in Computer Science at the University of Oxford is the best in the world, according to the *Times Higher Education's* latest ranking.

Oxford came out top overall for Computer Science in the UK for the second year in a row in the rankings which consider research, teaching, international outlook, industry outcome and citations.

As well as scoring highest for teaching in the world in the 2018 table, Computer Science at Oxford maintained its position as third overall.

Last year was the first time that the *Times Higher Education* had a dedicated league table for Computer Science. This year's world university ranking has expanded to include 300 institutions, up from 100 last year.

The full rankings are at: goo.gl/kpcPSJ

Oxford achieves gold for teaching

The University of Oxford has been awarded the top rating of gold in the first ever Teaching Excellence Framework (TEF), a UK government assessment of teaching and learning in the higher education sector.

The TEF was introduced in June this year to build evidence about higher education performance, and to complement the existing Research Excellence Framework. Gold was awarded to 59 of the 259 organisations who took part in the TEF. The award is valid for up to three years. Undergraduate teaching was assessed against ten criteria covering three areas: teaching

quality, learning environment and student outcomes. Gold was awarded for organisations which delivered consistently outstanding teaching, learning and outcomes for their students, according to the Higher Education Funding Council for England (HEFCE), which implemented the TEF.

For this first TEF, HEFCE assessed teaching excellence at the level of the overall university, rather than for specific subjects. HEFCE is now running subject-level pilots.

The results are at: goo.gl/Z6iJFu

Marta receives Milner Award

The Royal Society has awarded Professor Marta Kwiatkowska the Milner Award for her contribution to the theoretical and practical development of stochastic and quantitative model checking. The award, supported by Microsoft Research, is given annually for outstanding achievement in Computer Science by a European researcher. Marta, the first female winner of the award, will be presented with a medal and a gift of £5,000 at the Milner Award Lecture in November 2018.

Our computer scientist graduates are top earners

Oxford computer scientists have come top of a *Sunday Times* league table of graduate salaries in the UK for 2017. The mean annual salary of Oxford's Computer Science graduates is £45,000 six months after graduation, higher than for graduates of any other course at a UK university. Many Oxford Computer Science recent graduates go on to work for banks and technology companies, often finding roles with companies such as Google and Amazon.

Head of the Complete University Guide table

Computer Science at the University of Oxford came first in the 2018 UK university subject league tables compiled by the *Complete University Guide*.

In making its calculations, the *Complete University Guide* takes into account ten measurements: entry standards, student

satisfaction, research quality, research intensity, graduate prospects, student-staff ratio, academic services spend, facilities, good honours and degree completion. The *Complete University Guide* first started publishing league tables in 2007.

The full table is at: goo.gl/kh1BLX

News in brief

DiffBlue, developers of AI for code, announced in June that it had raised US\$22 million in funding, one year after its spin-out from the University of Oxford. Founded by the department's Professor Daniel Kroening, DiffBlue automates many traditional coding tasks, and is launching three products that are built on its core engine.


Professor Bjarne Stroustrup of Columbia University delivered a Strachey Lecture in November on 'The Continuing Evolution of C++'. Since the development of C++ started in 1979, it has grown to be one of the most widely used programming languages ever. The termly Strachey Lectures are supported by OxFORD Asset Management. The lecture can be watched at: goo.gl/N2vfuZ

Professor Ian Horrocks is a new honorary doctor at University of Oslo's Faculty of Mathematics and Natural Sciences. He received his award on 1 September in a ceremony at the University of Oslo.

In recognition of her research and service, Professor Ursula Martin CBE was awarded a Doctor of Science *honoris causa* of the University of London, by the Chancellor, Her Royal Highness The Princess Royal. Around 500 such degrees have been awarded since its foundation in 1836, including Winston Churchill and Alfred Einstein. Ursula has also been elected to Fellowship of the Royal Academy of Engineering, joining 1,500 world-leading engineers from both industry and academia in the UK's national academy for engineering.



Alonzo Church Award success



Congratulations to Professor Samson Abramsky [pictured above left], Lecturer Hanno Nickau, and Professor Luke Ong [pictured above right] (together with their co-authors) for winning the 2017 Alonzo Church Award. The award is bestowed by The Association for Computer machinery SIGLOG for 'Outstanding Contributions to Logic and Computation'.



The award was given to the group 'for providing a fully-abstract semantics for higher-order computation through the introduction of game models, thereby fundamentally revolutionising the field of programming language semantics, and for the applied impact of these models'.

More information: goo.gl/95LvJw

Double bill: Strachey and Ada Lovelace lectures

There was a rare double treat in early June as two eminent female computer scientists delivered the Ada Lovelace Lecture and a Strachey Lecture on the same day.

Professor Eva Tardos of Cornell University delivered the annual Ada Lovelace Lecture, launched in 2015 in honour of the bicentenary of Ada's birth, on 'Learning and Efficiency Outcomes in Games'. The central question explored in her lecture was 'What learning guarantees high social welfare in games, when the game or the population of players is dynamically changing?'

Professor Sarit Kraus of Bar-Ilan University and University of Maryland delivered a fascinating Strachey Lecture on 'Computer Agents that Interact Proficiently with People', which included video footage of robots deployed in



different scenarios. In her lecture, Sarit discussed the various ways in which automated agents that interact proficiently with people can be useful in supporting, training or replacing people in complex tasks. The Strachey Lectures are a termly series that are sponsored by OxFORD Asset Management.

A vodcast of Eva's lecture is at: goo.gl/5sCh9w, and Sarit's at: goo.gl/N2vfuZ



Oxford to host FLoC next July

The Federated Logic Conference (FLoC), a four-yearly event of epic proportions, is to be hosted by Oxford in July 2018.

The 2018 FLoC event is expected to attract more than 1,000 participants from around the world, and brings together nine major international conferences related to mathematical logic and Computer Science. It boasts an exciting variety of scientific, social and public events.

The Department of Computer Science's Professors Daniel Kroening and Marta Kwiatkowska are co-chairs for the event in Oxford, while the general chair is Professor Moshe Vardi of Rice University.

The Mathematical Institute and the Blavatnik School of Government will house the nine conferences over the course of two weeks, and events of interest to a wider audience will also take place: on 10 July, AI pioneer Stuart Russell (UC Berkeley) will give a public lecture at the Sheldonian Lecture Theatre as part of the Strachey Lecture series, and there will be a public debate in the Oxford Union Debating Chamber on 16 July, on the 'Ethics and Morality of Robots'. The panel will be chaired by London School of Economics Professor Judy Wajcman.

There will also be an event sponsored by the Alan Turing Institute – the Summit on Machine Learning Meets Formal Methods – during which academic and industrial leaders will discuss the potential and future challenges at the intersection of the two areas.

FLoC will encompass 79 workshops, and the FoPSS summer school, as well as nine conferences: CAV (International Conference on Computer Aided Verification), CSF (IEEE Computer Security Foundations Symposium), FM (International Symposium on Formal Methods), FSCD (International Conference on Formal Structures for Computation and Deduction), ICLP (International Conference on Logic Programming), IJCAR (International Joint Conference on Automated Reasoning), ITP (International Conference on Interactive Theorem Proving), LICS (ACM/IEEE Symposium on Logic in Computer Science) and SAT (International Conference on Theory and Applications of Satisfiability Testing). The call for papers has gone live, with eight out of the nine conferences having a deadline of 31 January for submissions.

For further information about FLoC and these conferences, visit floc2018.org and follow [floc2018](#) on Facebook or Twitter for regular updates.

New academics

The department's new academics in the last six months:

- Giulio Chiribella joined Oxford as a Professor of Computer Science, and a Fellow of St Hilda's College. His research activities are in quantum information, foundations and technologies.
- Yariv Gal joined the department as Associate Professor of Machine Learning, taking up the new post of Tutorial Fellow in Computer Science at Christ Church.
- Nicholas Lane joined the department as Associate Professor in Cyber-Physical Systems, and became a Non-Tutorial Fellow at Kellogg College.
- Andrzej Murawski joined the department as an Associate Professor and a Tutor in Computer Science at Worcester College. His research focuses on programming language semantics and software verification.
- Max Van Kleek was promoted to be an Associate Professor of Human-Computer Interaction and joined Kellogg College as a Non-Tutorial Fellow. He is a co-investigator on the EPSRC PETRAS project: Respectful Things in Private Spaces.
- Standa Živný has become Associate Professor of Theoretical Computer Science, and a Tutorial Fellow of Jesus College.
- Stefan Kiefer was promoted to be Associate Professor of Automated Verification and a Tutorial Fellow at St John's College.

Science family

Professors Leslie Ann Goldberg and Standa Živný from our department are featured in the Royal Society's web pages of parent-scientist profiles. The page aims to increase the visibility of people combining a career in science with a family life, showing the diversity of work life patterns across the UK.

See: ow.ly/zz94308ZBvQ

Extra security for end-to-end encryption

A multi-university team including the department's Professor Cas Cremers has developed a concept to address an issue in end-to-end encryption. A protocol, called DECIM (Detecting Endpoint Compromise in Messaging), solves the problem of when attackers have managed to compromise recipients' devices and can then, unbeknown to the sender or recipient, read and alter future communications. DECIM can detect where messaging has been compromised, and allow users to detect unauthorised usage. The new protocol was published by the IEEE: goo.gl/k3SYH9

Are we safe in self-driving cars?



The safety of self-driving cars relies on the accuracy of the deep neural networks that control them. To verify the performance of these networks, the department is developing model checking techniques, as Professor Marta Kwiatkowska explains.

Deep neural networks (DNNs) have achieved impressive results, matching the cognitive ability of humans in complex tasks such as image classification, speech recognition or Go playing. Many new applications are envisaged, including cancer diagnosis and self-driving cars, where DNN solutions have been proposed for end-to-end steering, road segmentation and traffic sign classification.

However, the reasons behind DNNs' superior performance are not yet understood. They are also unstable with respect to so called adversarial perturbations that can fool the network into misclassifying the input, even when the perturbation is minor or imperceptible to a human.

To illustrate the problem consider the figure below, which shows an image of a traffic light taken in Oxford (on the left) that is correctly recognised as a red light by the YOLO Real-Time Object Detection network, but after the image is perturbed by adding 11 white pixels the image is classified as an oven. To make matters worse, when these images are printed and photographed at several different angles and scales and then fed to the network, they remain misclassified.




Moreover, adversarial examples are transferable, in the sense that an example misclassified by one network is also misclassified by a network with another architecture, even if it is trained on different data.

Adversarial examples raise potential safety concerns for self-driving cars: in the context of steering and road segmentation, an adversarial example may cause a car to steer off the road or drive into barriers, and misclassifying traffic signs may lead to a vehicle driving into oncoming traffic.

As an illustration, let us consider the Nexar traffic light challenge (details at: goo.gl/MnF3vx), which made over 18,000 dashboard camera images publicly available, as part of a challenge to researchers to teach deep networks to identify traffic lights. Each image, **such as that above**, is labelled either green, or red, if the traffic light appearing in the image is green, or red, if the traffic light appearing in the image is red, or null if there is no traffic light appearing in the image.

continued on next page ►



In the figure below, one can demonstrate that, for the deep network that won the challenge, a change of just one pixel is enough to force it into classifying the image that any human would perceive as a red light into a green light, and do so with high confidence.

Safety of AI systems is the most pertinent challenge facing us today, in view of their autonomy and the potential to cause harm in safety-critical situations. Safety assurance methodologies are therefore called for. One of the most successful techniques to ensure system safety and reliability is model checking, now used in industry, for example at Intel, Microsoft and Facebook. Model checking employs techniques derived from logic and automata to verify – automatically, that is, with the help of a computer program – that the software or hardware system meets the desired specifications.

Model checking techniques for machine learning, and neural networks in particular, are little studied and hampered not only by poor understanding of their theoretical foundations, but also their huge scale, with DNNs routinely containing millions of parameters. In this direction, with support from the five-year EPSRC Programme Grant on Mobile Autonomy, which began in 2015, I am developing techniques for safety verification of deep neural networks.

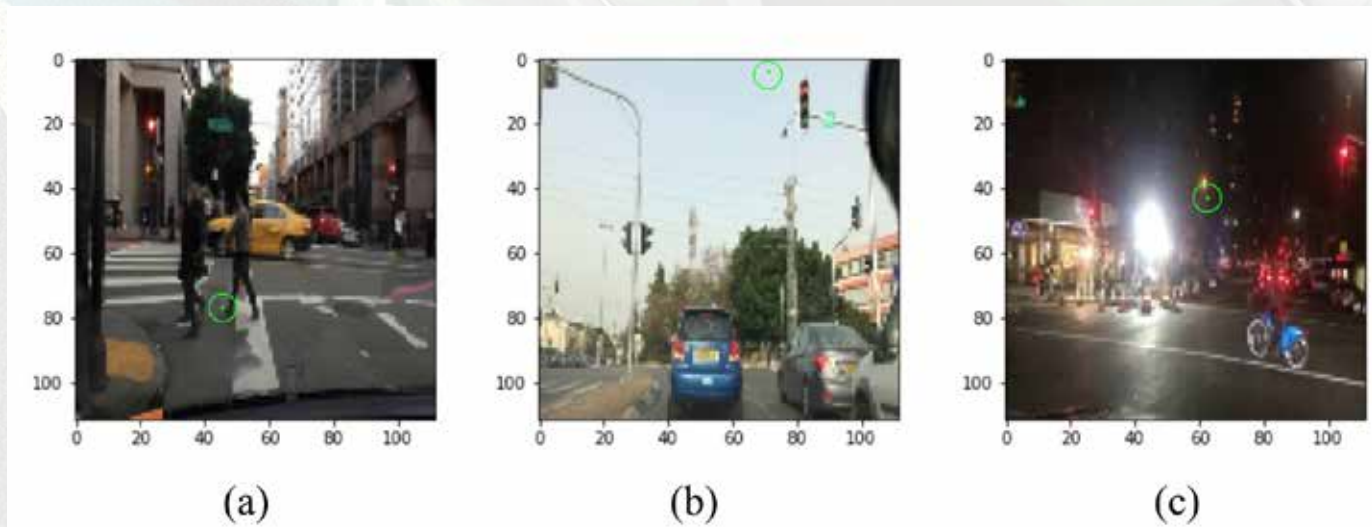
In contrast to existing work, which employs optimisation or stochastic search, the idea here is to automatically prove that no adversarial examples exist in a neighbourhood of an image and therefore the network can be deemed safe for a decision based on this image. I presented the results for several state-of-the-art networks and future challenges for this work in a keynote presentation at the 2017 Computer

Aided Verification (CAV) conference in Heidelberg. The presentation can be watched at: goo.gl/5zinuD

In a further paper just released on this topic (which can be found at goo.gl/yEKLDD), a software package called SafeCV has been developed to test the robustness of a DNN image classifier to adversarial examples. The main contribution is a method to search for adversarial examples guided by features extracted from the image, which uses computer vision techniques and relies on a Monte Carlo tree search algorithm. The software has been evaluated on state-of-the-art neural networks, demonstrating its efficiency in evaluating the robustness of image classifier networks used in safety-critical applications, which is sufficiently fast to enable real-time decision support.

The adversarial images included in this article were, in fact, found automatically by SafeCV. Images of Oxford traffic lights were taken by Matthew Wicker, studying at the University of Georgia and co-author of the paper, who visited Oxford for a summer internship.

Despite some progress, many challenges have to be overcome before we are able to prove that self-driving cars are safe. Chief among them is scalability of the verification techniques to large-scale images, where symbolic methods employed in conventional model checking may offer a solution. Another challenge is how to correct a flawed network, which can be addressed through retraining, fixing errors or even redesigning. Finally, for neural networks that learn as the system executes, for example reinforcement learning, we will need to move beyond design-time safety assurance towards monitoring and enforcement of decision safety at run time.



[Above] SafeCV applied to Nexar Traffic Light Challenge. (a) Red light classified as green with 68% confidence after one pixel is changed. (b) Red light classified as green with 95% confidence after one pixel is changed. (c) Red light classified as green with 78% confidence after one pixel is changed.



Oxford to hold one-day expo on artificial intelligence

The University of Oxford is one of the world's leading centres for Artificial Intelligence (AI) research, and over the past few years, we've witnessed a sharp increase in interest in our AI research from industry, commerce and government.

On 27 March 2018, the University of Oxford will showcase AI research across the whole University with a one-day expo. To be held at Worcester College's stunning new Sultan Nazrin Shah Conference Centre, the event will bring together Oxford's AI and machine learning thought leaders, providing a unique opportunity to hear from them about their work and their vision for the future.

The day will include briefings by the experts, panel question and answer sessions, demos, poster presentations by students, and an opportunity to talk one-to-one with professors at the heart of the AI revolution.

Topics to be covered include robotics, driverless cars, medicine and healthcare, scientific discovery, employment, finance, privacy, and ethical issues.

The event is aimed at industry, commerce and government, and we welcome representatives from the press. This event will debunk the myths about AI, which sometimes appear in the popular press, and show you the reality of AI today: what is possible and where the technology is going.

Enquiries are welcome: aioxford@cs.ox.ac.uk
Further information and registration: www.ml.ox.ac.uk



Provisional speakers include:

Professor Paul Newman, leader in robotic technology and autonomous systems, and co-founder of driverless car company Oxbotica (www.oxbotica.ai)

Professor Steve Roberts, machine learning expert and co-founder of Mind Foundry

Professor Mike Osborne, co-author of the celebrated report on AI and employment, and co-founder of Mind Foundry

Professor Andrew Zisserman, leading computer vision researcher and scientist at DeepMind (deepmind.com)

Professor Mihaela van der Schaar, leading expert on AI in medicine and healthcare

Professor Shimon Whiteson, machine learning expert and founder of Morpheus Labs (morpheuslabs.co.uk)

Professor Michael Wooldridge, head of Department of Computer Science and leading expert on cooperating AI systems.





Mike Wooldridge speaks on AI at The House of Lords

Professor Michael Wooldridge describes his experience giving evidence to the House of Lords Select Committee on Artificial Intelligence.

Artificial Intelligence (AI), and in particular the subfield of AI known as machine learning (ML), are very big news these days.

There has been genuinely rapid development in ML over the past decade, prompted in part by scientific breakthroughs, but also by the availability of large amounts of data with which to ‘train’ ML programs, and the availability of cheap computer processing power. Just as important to the growth in interest has been the dire

‘The UK is well placed to benefit from the AI boom’

predictions made in the press from high profile commentators about the risks that AI might ultimately pose to humanity.

In 2017, the UK government responded to the AI boom by organising an All-Party Parliamentary Group on AI starting in April, and launching in October a House of Lords Select Committee on AI led by Lord Tim Clement-Jones. Select committees work by first calling for written evidence (which anyone can submit), and then using these as the basis to structure a series of evidence sessions, which take place over several months.

On the basis of my written evidence, I was called to the first session on 10 October 2017. The other witnesses in my session were Professor Dame Wendy Hall, Regius Professor of Computer Science at the University of Southampton, and Oxford’s Professor Nick Bostrom,

author of the *New York Times* bestseller, *Superintelligence*.

The chair, Lord Clement-Jones, was friendly and welcoming, but as the session began, I discovered to my horror that a pen I had been nervously playing with had exploded, covering my hands in black ink. (If you watch the video of the session you can see my inky hands and my attempts to keep them out of sight.)

The session began with brief statements from the witnesses, and then moved briskly into the question and answer session. I was enjoying myself enormously: it is tremendously rewarding talking about a subject that you are passionate about in such a high-profile venue, and having senior members of our government listen attentively to your answers. The questions ranged from highly practical to rather philosophical.

I aimed at being enthusiastic about AI, but cautiously measured about what is likely to be possible, and to dispel some of the crazier ideas circulating. The main messages I wanted to put across were:

- There have been genuine and exciting developments in AI recently, which will provide opportunities for new AI-based products and services, but these advances have been confined to very narrow domains – the Hollywood dream of conscious machines is not imminent, and indeed I see no path taking us there.
- The UK is well placed to benefit from the AI boom with truly

world-class AI expertise at many UK universities (notably Oxford), a vibrant start-up scene in London, and the presence of DeepMind. However, the UK’s position is fragile, made more so by the uncertainty of Brexit, because it is essential to attract the best talent internationally to succeed in this field.

- To nurture the field, we need advanced skills – not just in programming, but MScs/DPhils. Advanced skills in AI are in desperately short supply – people with them are able to command startling salaries. A concerted investment in MSc/DPhil training would be highly desirable.
- Knee-jerk legislation about AI is not likely to be appropriate. I think it would be more useful to focus on specific areas such as health, insurance and finance, and think about whether specific legislation is needed in those areas to cope not just with AI, but with our digital future in general.
- My main concerns about AI in the short to medium term are related to unemployment and privacy, plus the implications of autonomous weaponry.

The event attracted a lot of press coverage. Some reports presented my contributions as uncritical flag waving for AI, which I didn’t intend and find hard to see in the transcript. Don’t take my word for it. Why not watch the video yourself. Feel free to enjoy the sight of me squirming as I realise my hands are covered in black ink, just as I am about to give evidence in the House or Lords. The video is at: tinyurl.com/yba4punx

Rocky ethical territory to traverse to reach AI code

How people relate to intelligent machines is one of the complicated questions to solve in work to develop ethical codes for Artificial Intelligence (AI) research. Senior Researcher Paula Boddington describes work currently in progress, potential pitfalls and opportunities.

There is growing attention to ethical issues in AI, with many calling for the development of codes of ethics, or for ethical watchdogs to oversee research and development. There are debates about why precisely we need such ethics guidance. Some cite the threat from a looming singularity (when intelligent machines achieve unstoppable growth, unpredictably changing the nature of human existence), while others express scepticism about this, but urge attention to ethical issues in AI that are already upon us, such as the impact upon employment, or unchecked bias in algorithms.

Indeed, there are several initial sets of broad principles already in existence, such as the Future of Life Institute's AI Asilomar Principles formulated in January 2017, and the Conference Toward AI Network Society's set of eight principles in 'AI R&D Guideline'. By their nature, such sets of principles tend to be aspirational; they may be a useful starting point for discussion. Some common overlapping themes can be found, such as calls to ensure AI is aligned with human values, and calls for transparency and accountability.

Unsurprisingly, such sets of principles share a great deal with other sets of ethical principles. But what is most useful in such broad statements is to stress those values that AI is likely to challenge the most, and to explain why. For example, the opaque nature of the decision-making of many AI systems provides a very strong reason to emphasise values of transparency and auditability. Codes of professional ethics rest on the broad assumption that professionals have the power to control their products or services, in order to provide benefit and to prevent or mitigate harms.

The control problem in AI therefore presents a significant difficulty, which gives a strong reason to believe that in AI, we are facing particularly rocky ethical territory. Ethical issues in AI tend to arise from its use to replace or supplement human agency, and this means we have to work through questions about the nature of persons, and how they relate to intelligent machines, which go to the heart of philosophical thinking about ethics.

Many ethical issues, such as privacy in the use of personal data, may be highlighted by AI, yet are not unique to it. We should think about what can be learned

from discussions of ethics elsewhere, paying close attention to similarities and differences. The ethical regulation of social science research suffered greatly from being squeezed into a model of regulation derived from medical research; there is ample reason to think this has hampered much valid research.


Likewise, with AI, we need to make sure that the shoe fits the foot. Ethics regulation is not and never should become an end in itself, and needs to be focused on enabling beneficial research and development, not on 'banning' things. This is particularly the case given that ill-judged brakes on the development of AI can in some instances mean that attempts to combat bad actors may be hampered. For instance when OpenAI, which

'Ethical issues in AI tend to arise from its use to replace or supplement human agency'

aims to create open source AI, was launched, a stated aim was precisely to mitigate the threat of malevolent superintelligence by making the technology freely available. Others, however, disagree with the premise of such an approach.

Enabling beneficial R&D means that principles must be implemented in practical contexts. The Institute for Electrical and Electronic Engineers is engaged in an ambitious project, the Global Initiative for Ethical Considerations in Autonomous Systems, which includes the development of standards for engineers in a large variety of practical areas to try to ensure that broad ethical aspirations can be achieved in practice. The standards currently in development include ethical considerations in system design, bias in algorithms, transparency in autonomous systems, and the control of personal data by AI.

Paula's book *Towards a Code of Ethics for Artificial Intelligence*, was recently published by Springer, and was written while working on a project 'Towards a Code of Ethics for Artificial Intelligence Research', with Oxford Professors Mike Wooldridge and Peter Millican. Funding for the project was generously provided by the Future of Life Institute.



Black boxes for robots to reassure public

If robots were fitted with black boxes, it would improve their safety and help curb public mistrust about their use in daily life. Senior Researcher Helena Webb describes research by two professors in this area.

As the 21st century progresses, the possibility that robots might become an everyday presence in our lives becomes greater. There are more and more news stories about what robots and other autonomous systems are now capable of doing and what roles they might take in our workplaces, homes and public spaces. Robots are seen by proponents as offering the opportunity to free us from mundane and repetitive tasks as well as increasing precision in technical or delicate ones.

However, amongst optimism over the ways that robots can benefit society, there are also concerns that they might be harmful. These include worries that the use of robots in employment will lead to significant loss of jobs and that reliance on automated systems to make decisions will strip us of our own autonomy. Furthermore, there are safety concerns over our current and potential future use of robots. In 2016, a human test driver was involved in a fatal crash whilst piloting a semi-autonomous car and we can imagine scenarios in which, for instance, the malfunctioning of care or work assistance robots could cause massive harm.

So when a robot behaves in an unexpected way and causes damage, how can we take steps to understand what happened and make sure it doesn't happen again? The potential future ubiquity of robots and concerns over the damage they might cause when they malfunction are issues considered in a recent article by Professor Alan Winfield of the Bristol Robotics Lab at the University of the West of England, and Professor Marina Jirotko of the Human Centred Computing Group at Oxford.

In recent years, Marina has played a key role in the development of the field known as 'Responsible Research and Innovation', which explores and develops the means through which societal and ethical concerns can be recognised and embedded into processes of research and innovation.

In their research paper, 'The case for an ethical black box', Alan and Marina combine technical understandings of robotics with a responsible research and innovation approach. They observe that new heights of machine autonomy create fear as well as optimism and put forward an innovative idea that can enhance safety in automated systems whilst also advancing public trust. They suggest that robots should be equipped with a 'black box', equivalent to the flight data recorders used in aviation. The black box continuously records sensor and relevant internal status data and can be extended in scope to also capture the AI decision-making process and environmental factors occurring before an adverse incident.

Just as black boxes in aviation can be drawn on to provide crucial

evidence following an accident, they can also be used in incidents involving robots. The information provided can help us to understand why a robot behaved in the way it did and then make recommendations for changes to prevent similar incidents or limit the potential damage caused.

The black box therefore offers to significantly advance the safety of robots. In addition, it can also foster the societal acceptability of these innovations. The presence of the black box, and associated professional groups making use of it in the course of an investigation, would demonstrate that the robots and their developers are responsible and accountable for their behaviours. It would also provide transparency so that members of the public can see these processes of responsibility and accountability in action.

The black box provides a pathway to greater public trust in robots. In fact, the authors warn that without it, public fear and mistrust over the use of robots in daily life is likely to remain.



Our first 60 years in photos

In 2017 the department celebrated its 60th anniversary. More photo montages are at www.cs.ox.ac.uk/CShistory

Photos from left to right, top to bottom: Emeritus Fellow Bernard Sufrin; MSc students, 2005. Photo by Xi Cheng; A gaggle of CompSoc (former) presidents. Photo by Ganesh Sittampalam; InspireHer! mother/daughter coding day 2017: Paula Fiddi and Bushra AlAhmadi; Ian Lynagh, finals 2002. Photo by Ganesh Sittampalam; Coffee: Carroll Morgan, Bernard Sufrin and Ian Page; Queen's Award for Technical Achievement: Tony Hoare, Bill Roscoe, Michael Goldsmith and Jim Woodcock; Maureen York, Administrator, CDT in Cyber Security; Wolfson building, 1989; Dragon boat racing team, 2013; Richard Bird and Tony Hey, 2003; Oxbridge Women in Computer Science Conference, 2017; Departmental Administrator Christine O'Connor; Bill Morton, Professor of Numerical Analysis; Jim Davies, Director, Software Engineering, 2011; Professor Stephen Cameron, 2013; MSc Computer Science students, 2014-2015. Photo by Aziza Bekniyazova; Ian Page. Photo by Guillaume Puyo; first iteration of the 'Building Information Governance' course in the Software Engineering Programme; Programming Research Group 1980-81; Inaugural Lovelace Lecture: Joël Ouaknine and Donald Knuth, 2015; Departmental photo, 1986.



Alumni Profile

Jolie de Miranda

Jolie completed both her BA and DPhil at the Oxford University Computer Laboratory (the department's former name). Her time here armed her with the core skills for her current role in systematic trading in the quantitative finance industry.



What course did you study here and when?

I studied Mathematics and Computation as an undergraduate (1999-2002). This was followed by a DPhil in Computer Science under the supervision of Professor Luke Ong (2002-2006).

What was your background before Oxford?

I am half Brazilian and half Chinese and moved around a good deal as a child. I had lived in six countries and studied at five different schools by the time I reached university. My working definition of 'home' was wherever my parents happened to be at the time. It's much the same now – home is where my husband and daughter are!

What attracted you to studying Computer Science as a subject?

I fell in love with computers at the age of 11 when I first saw a PC running Windows 3.1 in my aunt's office. I was blown away: I had never seen a colour screen PC before, nor a mouse! After that day, I campaigned relentlessly for my parents to buy me a computer and on my 12th birthday my dream came true. Within a few weeks I had written my first few lines of code in BASIC and I was hooked! My fate was sealed from that moment on; it was inevitable that I was going to end up studying Computer Science.

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

Oxford places a great emphasis on the mathematical foundations of Computer Science. I strongly believe in this approach. For my undergraduate degree, I originally enrolled to study Computation and I am embarrassed to admit that I was initially a little dismayed at the prospect of having to spend half my first year on maths courses.

However, my opinion of this changed rapidly: not only did I enjoy the maths courses but I started to realise that, above all, I was learning how to think in a mathematically precise and rigorous way. I was becoming a much better computer scientist as a result of this. At the end of my first year, I changed to the joint degree of Mathematics and Computation.

What did you do when you left Oxford?

I left Oxford to join the world of quantitative finance. After experimenting with a few different roles at both banks and hedge funds (the quantitative finance industry is enormous), I finally found my niche in systematic trading. Over the last five years I have been running systematic trading strategies that trade a wide variety of financial instruments including commodities, currencies, bonds and global equity indices.

My work combines developing mathematical models to forecast security prices and the risks involved, but it also involves coding up the systems and infrastructure to generate the trades and execute these trades electronically.

Although the work is occasionally stressful (particularly if I'm going through a losing streak!), it's a huge amount of fun and I feel incredibly fortunate to be able to work on such interesting problems day in and day out. I currently work for Balyasny Asset Management as an associate portfolio manager, and am based in London.

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

The joint degree of Mathematics and Computation has armed me with the core skills I need for my day-to-day work: mathematical modelling and development skills to turn these models into robust real-time trading systems. Although the topic of my DPhil is unrelated to my current line of work, it taught me a great deal about the highs and lows of research and how to keep my motivation up in times of adversity. In other words, it taught me to persevere.

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

Don't be put off if the end application doesn't immediately grab you or fall within your immediate area of expertise. My knowledge of financial markets was virtually non-existent when I first entered the world of finance. But that didn't seem to matter to anyone: what mattered was that I had been equipped with a good tool set to analyse and tackle problems in a logical and efficient manner. The domain-specific knowledge will come.

News in brief

The Oxford Women in Computer Science Society (OxWoCS) has won an equality and diversity award. The University of Oxford's Mathematics, Physical and Life Sciences Division (MPLS) chose OxWoCS for its inaugural 'Equality and Diversity Award for Team Effort' in recognition of its development of a successful portfolio of academic, industrial and social events, together with travel scholarships and collaborative activities with other institutions, all aimed at supporting and promoting women in Computer Science.

Lecturers Mark Kaminski and Egor Kostylev, and Professors Bernardo Cuenca Grau, Boris Motik and Ian Horrocks won the distinguished paper award at the International Joint Conference on Artificial Intelligence (IJCAI-17) for their paper 'Foundations of Declarative Data Analysis Using Limit Datalog Programs'. It was selected from more than 2,500 papers that were submitted to IJCAI this year.

Graduate student Daniel Nichol attended the 5th Heidelberg Laureate Forum (HLF) in late September. His trip was funded by a Romberg grant awarded by the HLF Foundation, and a departmental travel grant. The HLF brings 200 of the brightest young researchers in mathematics and Computer Science together with winners of the ACM A.M Turing Prize, the ACM Prize in Computing, Fields Medal, Nevanlinna Prize and Abel Prize, for a week of informal scientific discussion.

DPhil students Martin Dehnel-Wild, Louise Axon and Elizabeth Phillips participated in the second iteration of the HutZero programme: an intensive five-day cyber-security startup bootcamp at CyLon's London premises, jam-packed with workshops on technical development, business fundamentals and team building.

Booking.com offers women MSc scholarships

With the support of Booking.com, the university is offering ten scholarships to female UK/EU students studying MScs in mathematics, statistics and Computer Science in 2018-19.

The scholarships will cover both fees and a stipend at the level of the national minimum doctoral stipend as set by the research councils. Scholarships are available to women applying to MScs in Computer Science, Mathematics and the Foundations

of Computer Science, Software Engineering (part-time), Software and Systems Security (part-time), Mathematical Sciences, Statistical Science, Mathematical Modelling and Scientific Computing, and Mathematical and Theoretical Physics. Interested applicants should apply by the normal route by the relevant January deadline; women fulfilling the eligibility criteria will automatically be considered for a Booking.com scholarship.

More details: goo.gl/w7t3NT



Familiar faces and places appear in AlphaGo film

The AlphaGo film, which was released in the UK in October 2017, includes footage of a Strachey Lecture in the Sheldonian Theatre and of Professor Mike Wooldridge walking through Oxford with AlphaGo's CEO Demis Hassabis.

The film relays the story behind the Google DeepMind Challenge, a modern battle in the ancient Chinese game of Go. The competition in March 2016 pitted AlphaGo, an artificially intelligent computing

system, against Lee Sedol, the greatest Go player of the past decade. The film chronicles the journey from the back streets of Bordeaux, past the coding terminals of Google Deep Mind in London, to the competition in Seoul.

Students and staff of the department were invited to attend a private screening of AlphaGo in Oxford in December.

The trailer is at: alphagomovie.com

Oxford hosts ICFP 2017

The annual premier conference in functional programming (FP) was held in Oxford in September 2017 and chaired by the department's Professor Jeremy Gibbons. The 22nd Association of Computing Machinery (ACM) SIGPLAN International Conference on FP (ICFP 2017) took place in the Mathematical Institute.

At ICFP, researchers and developers heard about the latest work on the design, implementation, principles and uses of FP, from theory to practice and from academia to industry.

The conference has been steadily growing, and set new records again this year: 444 registered participants for the main conference, and 594 for whole week of events, which also



Photo: Alexandra Cárdenas performing at the Evening of Algorithmic Arts, the closing event of the Workshop on Functional Art, Music, Modelling and Design, and of ICFP itself.

included the Haskell Symposium and Commercial Users of FP, ten smaller workshops, and ten half-day tutorials. ICFP 2017 was also co-located with the International Conference on Formal Structures for Computation and Deduction, which held 11 workshops.

In addition to two keynotes, the main conference featured 44 accepted papers. The proceedings appear in the new journal: *Proceedings of the ACM on*

Programming Languages, PACMPL. More information: icfp17.sigplan.org

Cristina Matache won first prize in the undergraduate division of the ICFP Student Research Competition, for her work 'Formalisation of the Lambda-Mu-T Calculus in Isabelle/HOL'. Cristina is now studying for her MSc in Computer Science at St Cross College. Further information: goo.gl/bdHhB5

Student prizes

First year prizes

Alexandru Strimbu (St Anne's College) won the BCS Prize in Computer Science for the best performance in Computer Science Prelims papers. Denitsa Markova (St Catherine's College) won the Gibbs Prize for Mathematics and Computer Science Prelims (paying particular regard to Computer Science papers), and Vladimir Mikulik (Hertford College) was awarded the Gibbs Prize for Computer Science and Philosophy Prelims (paying particular regard to Computer Science papers).

Third year prizes

Mantas Pajarskas (St Hugh's College) won the Hoare Prize for the best overall performance in Computer Science and the G-Research Prize for best project

Congratulations to students who won prizes for their Prelims, the exams at the end of the first year of undergraduate study, and to third and fourth year students who have won prizes for their work.

in Computer Science. Declan Manning (New College) received the British Telecom Research and Technology Prize for Mathematics and Computer Science for best overall performance (paying special regard for Computer Science papers). Wenkai Lei (St John's College) won the Junior Mathematics Prize for Mathematics and Computer Science for outstanding performance in the mathematical papers in Part B. Richard Ngo (Hertford College) received the Gibbs Prize for Computer Science and Philosophy in Parts A and B (paying particular regard to Computer Science papers) and Nicholas Buckner (Balliol College) was awarded the Gibbs Prize for Computer Science and Philosophy for Parts A and B (paying particular regard to Philosophy papers).

Fourth year prizes

Matthew Gripton (St Catherine's College) received the Hoare Prize for best overall performance in Computer Science in Part C. Dan-Andrei Gheorghe (Somerville College) won the Microsoft Prize for best Computer Science project in Part C. Florian Jaeckle (Worcester College) was awarded the Hoare Prize for best overall performance in Mathematics and Computer Science in Part C. Tomas Vaskevicius (St Anne's College) was awarded the G-Research Prize for best Computer Science project by a Mathematics and Computer Science candidate in Part C and Alexander Bridgland (Hertford College) received the Hoare Prize for best overall performance in Computer Science and Philosophy in Part C.

The world outside

The department's students, academics and support staff are not only accomplished in their respective fields, but have many other talents too, ranging from triathlons to Mahjong. Here are a few of their recent activities.

Researcher Martin Lester [pictured below] entered the World Riichi Mahjong Championship 2017, which was hosted in Las Vegas. He sailed through to the knockout stage, and came 24th of 224 competitors.



Ivo Sluganovic, a DPhil student in the Software and Systems Security group, competed in his first IronMan 70.3 – a triathlon with a 1.2 mile swim, a 56 mile bike ride, and a 13.1 mile run. He finished 809th overall in the Weymouth race, or 77th in his divisional rank, putting him approximately a third of the way down the field.

HR Manager Laura Jones took part in an overnight marathon walk for Cancer Research in September. She finished in 9 hours and 44 minutes.

Researcher Michaël Cadillac took part in the 43rd Annual British barbershop convention with the Oxford Harmony Barbershop Choir. Although the group didn't qualify for a prize, they fared slightly better than in previous years.



DPhil student Elizabeth Phillips and undergraduate Janpreet Khabra were sponsored to attend the Grace Hopper Celebration (GHC) in October. The GHC aims to recognise, celebrate and encourage women in technical fields.

The department and OxWoCS have committed for six years to send two students to the GHC, which was held in Florida in 2017 and hosted over 18,000 attendees. It was produced by the Anita Borg Institute.

Elizabeth said: 'The conference included a wide variety of keynote talks ranging from philanthropists such as Melinda Gates to BCSWomen founder Susan Black. There were over 250 professional development workshops and technical sessions. These ranged

from career development to research-focused presentations on security and data science.

'I was a mentor at the Open Source Hackathon, which enabled attendees to contribute code to the open source community on a variety of projects including Cloudera's one on the prevention of child abductions.'

Janpreet said: 'The key theme throughout the event was celebration. Every keynote talk was a true celebration of women in tech, with lights and glow sticks and music and so many awards. What stood out to me was the commitment to the representation of diversity. There were speakers from all academic, socioeconomic and ethnic backgrounds from all over the world, and specialised social sessions for minority groups.'

Plaudits for DPhil Mariam Nouh

Mariam Nouh won third place in the Association for Computing Machinery Student Research Competition (graduate category) at the GHC for her research 'CCINT: The cyber-crime intelligence framework for detecting online radical content'. Mariam studied the characteristics of existing online radical groups and defined automated approaches to detect them. She was able to analyse the content they posted, as well as their behaviour, using natural language

processing and machine learning techniques. She is a DPhil student in Cyber Security, focusing on research topics related to detecting and analysing cyber-criminal behaviour online.



OxWoCS extends a large welcome to all

In October 2017, the Oxford Women in Computer Science Society (OxWoCS) expanded its annual welcome event for all members of our department by joining up with other societies at the Oxford Town Hall.

The event was for the first time jointly hosted with five other societies for women in STEM: OxFEST, the Mirzakhani Society, Women in Physics, Women in Materials Science Society, and Women in Engineering. Sponsor representatives and committee members from all the societies mingled with attendees while they enjoyed light refreshments.



Programming teams excel in European competition

Three student teams from the department showed off their skills in a programming teamwork challenge in November.

In the Northwestern European Regional Contest (NWERC) of the ACM International Collegiate Programming Contest (ICPC), the Tractor_Specialists team came 3rd, the Assemblers 10th, and the Hen 114th. Teams of up to three students had to solve as many programming

problems as possible from a given problem set, using only one computer.

Having come third at NWERC, the Tractor_Specialists team (Andrei-Costin Constantinescu, George Chichirim, and Tamio-Vesa Nakajima) have qualified for the ACM-ICPC World Finals held in Beijing in April.

The teams previously competed at UK-IEPC, the contest to help universities choose which teams to send to NWERC. Nine teams entered from Oxford.

The NWERC leaderboard is at goo.gl/HM2rSw



In the first round of University Challenge, Winston Wright, an undergraduate in the department, was part of the winning Magdalen College team. ©BBC

News in brief

Oana Tifrea-Marcuska has been awarded an EPSRC Doctoral Prize 2017. During her DPhil project, Oana has explored several different approaches to personalised search on the social semantic web, in the form of preference-based ontological query answering. For the EPSRC prize project, Oana will explore and develop suitable techniques for mining preferences for personalised search. The aim of EPSRC Doctoral Prizes are to support students to increase the impact of their PhD research and improve retention of doctoral students in research careers.

DPhil student Elizabeth Phillips was awarded one of five scholarships to attend the 2017 VizSec conference as part of IEEE VIS in Arizona in October. She was also selected as one of 18 global recipients of the first Codess Microsoft Professional Program Sponsorship and awarded a full scholarship to complete the six-month Data Science Professional Program.

Research assistant Elisa Passini received the Technological Innovation Award at the Safety Pharmacology Society Meeting 2017 for 'Virtual Assay: a User-Friendly Framework for *In Silico* Drug Trials in Populations of Human Cardiomyocyte Models'. The Virtual Assay software was developed by our Computational Cardiovascular Science team in collaboration with Oxford Computer Consultants, and supported by an EPSRC Impact Acceleration Award. It is used for *In Silico* drug trials, using populations of human cardiac cellular models based on well-understood human cardiac physiology. The human cell populations are calibrated against experimental data and used to predict the effects of different pharmaceutical agents in humans at the population level.

App users need protecting from aggregation of permissions

To help speed up code writing, many app developers use common libraries, but this can result in such libraries getting access to significant privileges on a device. DPhil student Vincent Taylor describes the problem and his research in the area.

Smartphones have rapidly become a leading part of our daily lives. Smartphone usage is fuelled predominantly by apps, small pieces of software ready to be downloaded from app stores at the touch of a button. Android is the most popular smartphone operating system, with an official app store (Google Play) containing over two million apps.

Each Android app is delivered as a single archive, which contains all app code and resources needed for the app to function. Just as on traditional computers, third-party libraries are available for app developers to use to rapidly deploy advanced functionality to their apps. Common libraries include those for advertising, analytics, and social networking.

The libraries used in Android apps are tightly integrated into the binary code of the app itself and the Android operating system treats an app and its embedded libraries as a single



entity. An undesired consequence of this is that embedded libraries obtain the same permissions that have been granted to their host app. This means that granting permissions to apps often grants permissions to other third parties as well.

While this permission leakage from apps to libraries is well understood by Android security and privacy researchers, a far more insidious problem has been left unaddressed. Popular libraries are likely to be used in more than one app on a device. The problem stems from the fact that apps typically have different sets of permissions granted to them. This means that an instance of a library in one app may have several permissions, and another instance of the same library in another app

on the same device may have several different permissions. Thus, if the library were to aggregate its permissions, it would achieve significantly more privileges on a device than it would seem at first glance. This is the problem of 'intra-library collusion'.

Professor Ivan Martinovic and I from Oxford (with data and insights from Alastair Beresford of Cambridge University) studied the potential for intra-library collusion in the real-world using data from 30,000 actual smartphones. Over 57% of devices were susceptible to intra-library collusion. By performing a historical study, we showed that the risks from intra-library collusion have increased significantly over the past two-and-a-half years.

Having investigated the problem, we are now conducting research into providing mitigations for it as part of my DPhil research. The next steps are to extract URLs and network traffic from libraries to see what private data is being sent, and to where. This work fits within the security research theme of the department, which has the aim of enabling users to use technology with confidence that their privacy will not be breached.

Cyber expertise on show

Cyber Security Oxford, the network for people across the University with an interest in cyber security, hosted its first ever research showcase on 9 June. The event was attended by industry and government experts from 28 different organisations, spanning a wide range of

interests, and gave our students and academics a chance to discuss the real-world needs of various sectors.

Feedback was very positive: guests appreciated the opportunity to see research at the cutting edge, and to gain an overview of the wide range of expertise in Oxford, as well as to get to know our researchers in a sociable setting where they could discuss ideas.

Based on discussions at the event, two new collaborations are already under way, and two graduate students are expected to take up part-time industry roles alongside their studies this year. A follow-up meeting will be held in October to discuss further options for joint research.

To learn more, contact enquiries@cybersecurity.ox.ac.uk.

Framework to guide banking on mobile biometrics

In response to developments taking us towards a password-free future, our department teamed up with Mastercard to create a mobile biometrics framework for banking executives.

Banks see biometric technology – using fingerprints, iris scans or ‘selfies’ – as a more convenient and secure way for customers to prove their digital identity. However, only 36% of relevant banking executives feel they have adequate experience to deliver mobile biometrics, according to the report ‘Mobile Biometrics in Financial Services: A Five Factor Framework’.

Oxford and Mastercard worked together on the report, which set out guidelines for banking executives to follow. It advises them to focus on five factors: performance, usability, interoperability, security and privacy.

‘Biometric authentication has a lot of potential, but it is important to address the objectives of each of the five factors when designing solutions. Working together with Mastercard enables us to solve realistic threats to the industry with the best technical and scientific ideas. Users will need consistency, quality and assured security for this technology to thrive,’ said Professor Ivan Martinovic, who was a joint author of the report, along with Oxford DPhil students Giulio Lovisotto, Ivo Sluganovic, Marc Roeschlin, plus Mastercard’s Raghav Malik and Paul Trueman.

‘This framework is fundamental to accelerating the deployment of mobile biometrics for consumers and industry alike, but collaboration is key,’ commented Ajay Bhalla, Mastercard’s president of global enterprise risk & security. ‘We

can only achieve this if industry, academia, governments and technology vendors understand and contribute to the evolution of the Five Factor Framework for mobile biometrics.’

Opus Research has produced a synopsis of the research, with a breakdown of the critical issues for financial services companies. It is available at: goo.gl/69hQa8



Stingray detectors provide little protection

Some freely available tools to detect threats to mobile device users from fake mobile phone towers, also called stingrays and International Mobile Subscriber Identity (IMSI) catchers, have been shown to have major shortcomings in a recently presented paper.

A research team from the University of Oxford and the Technical University of Berlin evaluated the capabilities of five IMSI catcher detector apps for Android devices, which claim to protect mobile device users from spying by fake base stations. The researchers found that none of the popular

tested apps were able to detect the IMSI catchers effectively. As part of the study, which was presented in August to the USENIX Workshop on Offensive Technologies, the team developed their own White-Stingray, a framework with various attacking capabilities in 2G and 3G, using techniques described in the patent database for IMSI catchers and commercial product brochures. The framework managed to circumvent many of the apps’ detection techniques.

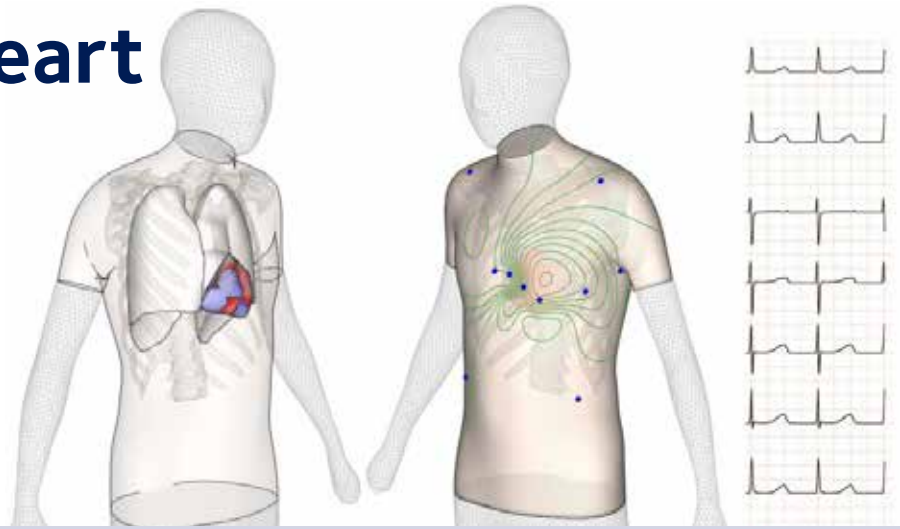
‘People incorrectly think that IMSI-catcher detection apps can protect you from tracking, but they can in fact trick people into handing over their private data,’ said Ravishankar Borgaonkar, a research fellow from Oxford’s Department of Computer Science who led the study. ‘The root cause of the ineffectiveness of catcher apps lies in both the architecture of the GSM/3G system and the limited access to baseband chip data to app developers which

means that spies can always stay one step ahead. The GSM/3G network design gives all the power to the base station, while the phone is just a dumb device that listens for and accepts commands.’

The team looked at ways to help solve the architectural issues, and how to improve the detection capabilities of the IMSI-catchers on mobile devices. ‘We found that not only do free Android apps need to be improved, but phone manufacturers, baseband chip providers, and carriers need to actively coordinate their efforts to tackle this fake base station problem,’ said Ravishankar.

The research was partly performed within the EU’s 5G-ENSURE project, which is addressing priorities for security and resilience in 5G networks. Ravishankar’s co-researchers were Oxford’s Professor Andrew Martin, and Berlin’s Shinjo Park, Altaf Shaik and Professor Jean-Pierre Seifert. Their paper can be read at: goo.gl/J15kuE

Computer heart simulations feature on big screen



A short film about future virtual humans has been made that incorporates material from research by Professor Blanca Rodriguez and her team.

The film, which was made by CompBioMed (a user-driven centre of excellence in computational biomedicine in which Oxford is a

core partner), was first screened at the London Science Museum IMAX in September. Afterwards Blanca gave a short presentation about virtual hearts.

The film includes material on high-performance computing simulations of the human heart based on clinical data provided by Research Fellow

Ana Mincholé, DPhil student Hector Martinez-Navarro and Researcher Francesc Levrero.

The film has since been shown at the Imagine Science Film Festival in New York City, and CompBioMed hopes other screenings will follow in the coming months

Carnival demystifies research

Members of the department helped to take real research projects out to the public as part of Oxford's Curiosity Carnival on 29 October. Oxford ran a city-wide programme of activities across its universities, libraries, gardens and woods for members of the public.

The Human Centred Computing group – with other colleagues – ran a stand in the Pitt Rivers Museum that was based on two of their projects, UnBias and Digital Wildfire, giving information about common challenges to online safety and fairness in everyday life such as trolling and hate speech, fake news and bias in search results. Visitors were invited to reflect on the challenges and vote on how best to resolve them.

In the Weston Library, Research Fellow Jason Nurse participated in the Living Library, where members of the public could 'borrow' researchers from a range of domains for face-

to-face conversations. Jason was advising on keeping safe online, but with a special focus on social media, voice assistants and smartphones. He spoke to students, families, teenagers, and adults of all ages.

Meanwhile, on Broad Street, DPhil student Maaïke Zwart performed aerial acrobatics with a group of students from the Oxford Aerial Dance Society. The main goal was to entertain the audience, but also to show that the research community has others interests too.


The University's YouTube video of the event is at: goo.gl/MtF4Tj



News in brief

DPhil student Maarten Scholl won the Best Quantitative Finance MSc Thesis Award for his thesis entitled 'The Role of Central Clearing Parties in Over-the-Counter Derivatives Markets' on 3 November. This annual award is open to students from the BeNeLux and Scandinavia. Maarten is a graduate of the joint University of Amsterdam/Free University Amsterdam MSc program in Computational Science. Maarten is now reading for the DPhil in Computer Science at the University of Oxford and the Institute for New Economic Thinking.

DPhil student Klaudia Krawiecka won a national competition in Finland for the best information security thesis. She was the first woman to receive the award. Klaudia graduated with joint degrees from Aalto University, and the Norwegian University of Science and Technology, and is now at the Centre for Doctoral Training in Cyber Security in Oxford.



Jason Nurse

Festivals reach the masses

Our research teams have been explaining their work to the public at a number of science festivals this year: Hay, Cheltenham and Oxford.

At Hay Festival current members (and one alumna) of the department spoke to more than 1,500 people 'face to face' in a series of lectures and workshops. Members of the public heard about the latest research in quantum computing, had a go at a computational thinking challenge, and even built their own quantum computer.

Professor Marta Kwiatkowska gave a talk about her research entitled 'When to trust a robot'. While at the festival she recorded a three-minute video on computing with DNA for Hay's series of masterclasses answering big questions for A-level students, watch at goo.gl/usY7tk. Oxford alumna Anne-Marie Imafidon recorded another video for the Hay Levels series on the way in which binary numbers operate in computing, which you can see at goo.gl/B5MUA7

Anne-Marie and Senior Research

Fellow Jamie Vicary took part in some BBC programmes made at Hay. Anne-Marie and DPhil student Linde Wester were invited to be part of the Hay's 30th anniversary celebrations, named as part of Hay's '30 under 30' group of academics: hayfestival.com/hay30/

At the Cheltenham Science Festival in June, Jamie ran the 'Build your own quantum computer workshop' again, and Research Fellow Jason Nurse explained how voice assistant technology works and considered the implications for privacy and security. Jason's talk was covered in an interview with the BBC. Watch at goo.gl/iGyBmG

Jason took to the stage again at the Oxford Science Festival in June when he spoke about the advantages and risks of maintaining a digital presence, especially as platforms begin to integrate with the Internet of Things.

Linde gave another talk on quantum computing at the Malvern Festival of Innovation in October.

Oxford Sparks is a portal used to engage the public with exciting science taking place across the University. The day-in-the-life videos are aimed at an audience of 12 or older: goo.gl/5L9iCB Other videos are available on the department's Media Wall: www.cs.ox.ac.uk/mediawall

A day in the life of a researcher

Research Fellow Jason Nurse (photo above) is the star of the latest Oxford Sparks' video showing what it's like to work as a researcher at the Department of Computer Science.

News in brief

Martin Strohmeier's DPhil dissertation 'Security in Next Generation Air Traffic Communication Networks' has received a commendation and been published on the BCS website. During his DPhil at Oxford, Martin extensively analysed the security and privacy of wireless aviation technologies and focused on the development of novel cyber-physical approaches to improve their resilience quickly and efficiently. Read it here: goo.gl/fLwzLy

Professor Leslie Fox has been honoured by his home town of Dewsbury with the mounting of a blue plaque, a permanent sign installed in a public place, in the railway station to commemorate his achievements. Leslie was a founding member of our department as the Oxford University Computing Laboratory in 1957. During this initial period the research was almost entirely towards numerical analysis, and Leslie became the University's first Professor of Numerical Analysis in 1963.

Professor Cas Cremers and DPhil students Martin Dehnel-Wild and Kevin Milner won best paper award at ESORICS 2017 in September. This research considered the security of the world's power grids: these are often controlled remotely, so they need to be confident that received commands are from a legitimate control station, and not an attacker. The authentication mechanism is defined by an IEEE standard, 'DNP3: Secure Authentication v5', and the research team conducted the first full security analysis of this protocol (using the Tamarin Prover), concluding that it meets its required security properties. This four-month project shows that real-world standards can be formally analysed in an achievable amount of time. Read the paper at: goo.gl/ZEDo86

Defining trust helps robots behave



An important step in making sure that robots act safely and effectively in our society is to be able to formally express what trust is in a human-robot partnership. The University of Oxford is working on how to express the properties of trust, writes researcher Morteza Lahijanani.

Robots are becoming members of our society, which consists of convoluted social rules, relationships and expectations. Complex algorithms have rendered robots increasingly sophisticated machines with rising levels of autonomy, enabling them to leave behind their traditionally simple work places in factories for a more complicated world. Driverless cars, home assistive robots, and unmanned aerial vehicles are just a few examples.

As the level of involvement of such systems increases in our daily lives, their decisions affect us more directly. We instinctively expect robots to behave morally and make ethical decisions. For instance, we expect a firefighter robot to follow ethical principles when it is faced with a choice of saving one person's life over another in a rescue mission, and we expect an eldercare robot to take a moral stance in following the instructions of its owner when they are in conflict with the interest of others. Spoiler alert: This is in conflict with the robot in the movie 'Robot & Frank' which partakes in a robbery with Frank, its owner, to achieve its goal: Frank's wellbeing.

Such expectations give rise to the notion of trust in the context of human-robot relationships and to questions such as 'how can I trust a driverless car to take my child to school?' and 'how can I trust a robot to help my elderly parent?' Failing to answer such questions appropriately can cause a major blow to the field of robotics or, more generally, autonomous systems.

In order to design algorithms that can generate trustworthy decisions and hence an ethically reliable system, we need to understand, formalise and express trust. This is a challenging task because it involves

'How do we quantify trust and how do we model its evolution?'

many aspects including sociology, psychology, cognitive reasoning, philosophy, logic and computation.

We believe formal methods, specifically quantitative verification and synthesis, can provide a venue to approach the above questions. In recent years, these methods have been receiving a lot of attention in the robotics community and have been explicitly adapted to provide guarantees for the safety and correctness of robot behaviours.

We have begun a thorough investigation into formalisation of trust and expressing its properties with the aid of collaborators

in the philosophy and human factors communities as part of an ES/PRC-sponsored project. It is a collaboration between the Oxford Robotics Institute and our department entitled 'Mobile Autonomy Programme Grant: Enabling a Pervasive Technology of the Future', which runs from March 2015 to February 2020. Professor Marta Kwiatkowska is leading the Safety, Trust and Integrity theme within the project (goo.gl/uBZctr), working with Research Associate Wenjie Ruan and me. New DPhil student Maciej Olejnik is also joining the project.

The vision of the project is to create, run and exploit the world's leading research programme in mobile autonomy, addressing fundamental technical issues, which impede large-scale commercial and societal adoption of mobile robotics. Understanding trust and being able to evaluate it to inform trust-based decision-making and reliance on mobile robots is key to their widespread adoption. We have organised two workshops on the topic, and the third one will be at FLoC in Oxford. More information: goo.gl/NAkCgL

The study of trust is a cross-disciplinary challenge creating central research topics, notably from the formalisation angle. The immediate technical research questions are how to quantify trust and how to model its evolution. Another key question is how to design a logic that allows the expression of specifications involving trust. From the verification perspective, the questions are how to verify (reason about) such specifications in the context of a given partnership or, even more prominent, how to synthesise (design) an autonomous system such that, in a partnership with a human, these specifications are guaranteed.

In our research group, we have started taking initial steps in investigating these questions and believe that, only by a thorough study of them, one day, we may be able to guarantee the success of robots in our society.

Badgers bridge territorial borders

A clearer understanding of badgers' territorial habits give vital clues to the likely effectiveness of badger culls. By introducing different technology to collect tracking data, an interdisciplinary team from Oxford and Cambridge has gleaned important new insights into badgers' potential interactions.

Strategies for the culling of badgers to help prevent the spread of bovine tuberculosis are typically based on assumptions that badgers mainly move within certain territorial borders. Growing evidence, however, suggests that European badgers may socialise more with other social groups than previously thought.

To demonstrate that is the case, it is important to be able to track where badgers are relative to each other, and how often they are in the same place at the same time (or a different time). This pinpoints the occasions when badgers may transmit disease, directly or indirectly, among themselves.

When aiming to build a continuous picture of the infrequent interactions of badgers, all previously used technologies have had various disadvantages, as well as certain merits. Deficits that hampered analysis included, for example, tag battery life being too short, the difficulty in measuring individual badgers' positions at specific points in time, and data security issues.

In a 13-week proof of concept study a research team set out to test a fully automated active RFID (aRFID) system to better measure badgers' movements. Our computer scientists teamed up with zoology experts from Oxford's Wildlife Conservation Research Unit and Cambridge to fit badgers with 20g aRFID tags, and set up wirelessly networked base stations in setts and near latrines.

For the task, researchers from the Department of Computer Science (led by Professors Niki Trigoni and Andrew Markham) adapted commercially-available technology, which has usually been used by the security industry to protect high value assets in small areas, such as valuable paintings.

The study demonstrated that the aRFID system could collect data on the location of individual badgers when they were within 31.5m of base stations (which could then be used to show when badgers were in the same place at the same time). It also showed that the system was capable of operating continuously over a long period, with the tags predicted to last between two to five years. By storing data off the tag, it could be accessed remotely, and data security was high.

The study also showed that this niche aRFID technology is highly suited to this sort of tracking and analysis. It has an advantage over other technologies because of tag longevity and because it allows infrequent events to be detected in small defined areas, which we can achieve by restricting the detection range. The wireless interconnectivity of base stations also means it could potentially be accessed from anywhere in the world.

The data collected enabled researchers to track how much time

badgers spent within their own social groups, and with other groups, at setts and latrines. It showed that of the time spent by badgers mixing together, they were for the most part with members of their own social groups in their own setts. This was as expected, but proved that analysis of aRFID data can corroborate known location patterns.

Further analysis threw up new findings about badgers' habits: around 16% of the time when badgers were near each other was spent mixing with other social groups near latrines and at setts. In each week of the study some badgers visited others' setts. The analysis suggested that badgers regularly go beyond their own territorial borders, which contradicts previous assumptions, and puts into question conventional strategies for badger culling.

A research article about the findings of the study was published in *Methods in Ecology and Evolution*: [goo.gl/3y2V2M](https://doi.org/10.1111/mee.12222)



Royal Society awards Jamie Vicary competitive fellowship

Jamie Vicary has been named as a 2017 recipient of a Royal Society University Research Fellowship (URF) for a project on 'Higher Structures for Compositional Science'.

The Royal Society's URF scheme provides outstanding early career scientists, who have the potential to become leaders in their chosen fields, with the opportunity to build an independent research career. An initial tranche of funding for five years is provided, with the potential to extend for a further three years.

The scheme is extremely competitive; each year only around 30 applicants are successful.



Boris Motik receives EPSRC fellowship extension

Professor Boris Motik has been awarded a three-year fellowship extension by EPSRC for his work on intelligent information systems. The extension will build on his EPSRC early career fellowship, 'MaSI3: A Massively Scalable Intelligent Information Infrastructure', which resulted in the team developing techniques that are now used by two startups.

Intelligent data management techniques play a key role in many areas. For example, oil producers such as StatOil analyse streaming sensor data to diagnose faults and prevent failures. To simplify the management of the data, intelligent information systems provide services that capture background knowledge about the application domain, and infer information implicit in the data and the background knowledge.

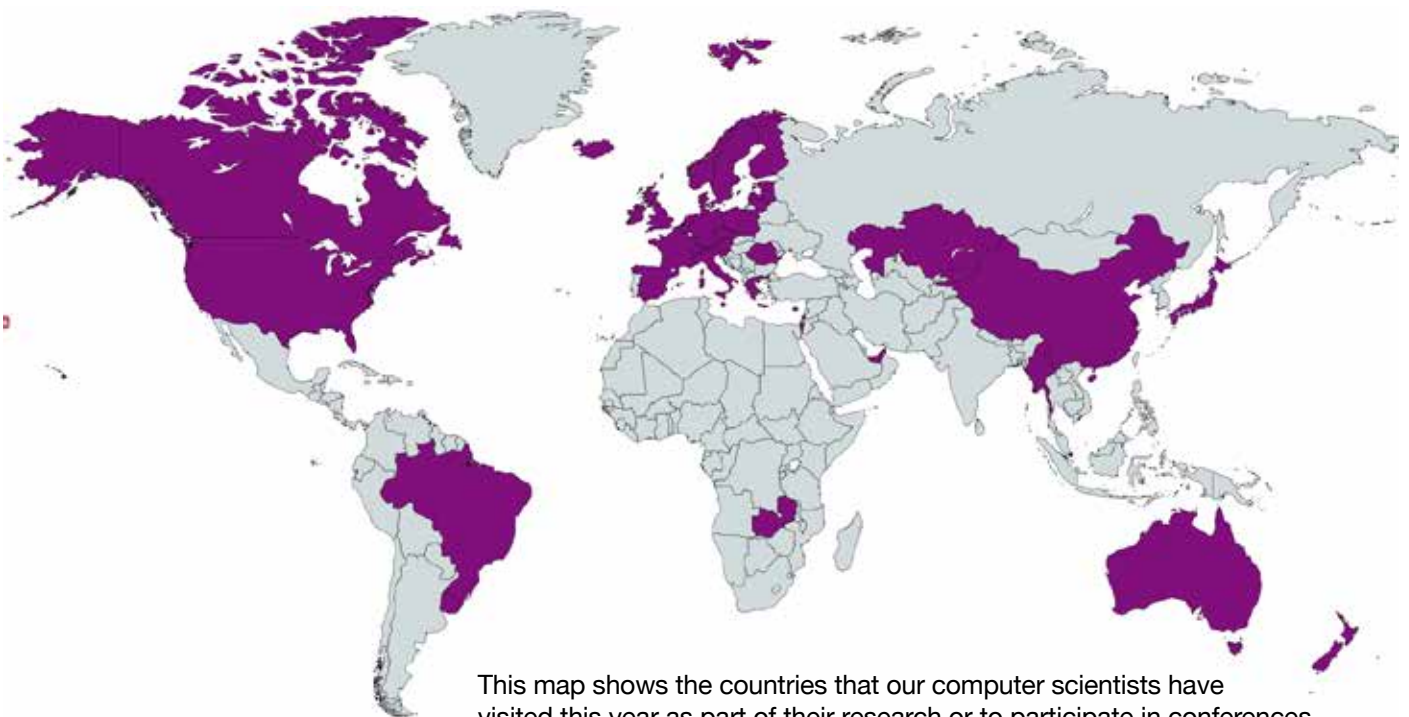
The objective of MaSI3 was to make intelligent information systems

a reality by developing scalable reasoning and query answering techniques. These underpinned an intelligent information system called RDFox, which recently provided the technological foundation for two startups: Covatic (www.covatic.com) and Oxford Semantic Technologies (www.oxfordsemantic.tech).

Boris's extension project, 'AnaLOG: Datalog Extensions for the Analysis of Static and Streaming Data' will develop research in this area, focusing on the key data analysis problem to design a language that can express the relevant tasks.

Datalog has been identified as a natural starting point but has some shortcomings. The objective of AnaLOG is therefore to develop Datalog extensions for data analysis in intelligent information systems, establish links with known problem-solving methods, and evaluate the results in practice.

Oxford Computer Science has global reach



This map shows the countries that our computer scientists have visited this year as part of their research or to participate in conferences.

Not so private devices



One of our research teams is working on ensuring that Internet of Things (IoT) users are aware of the sensitive data that their devices could collect and share. The EPSRC-funded project aims to empower end users with smart tools to help them understand and control devices' data collection activities, writes Researcher Reuben Binns.

Internet-connected devices are increasingly used in highly private spaces. These include sensors, smart voice assistants and other devices in living rooms, bedrooms, bathrooms, and private offices, as well as wearables and implantables worn on or even inside the human body.

While the privacy implications of personal computers and smartphones have been much studied, the very personal nature of the data collected by such devices raises novel challenges. Have you ever wondered whether smart conversation systems such as Amazon's Alexa or Apple's Siri are secretly judging you? Or what your internet bathroom scales might be saying behind your back?

Just as digital advertising technology has evolved to track users of the web and smartphones,

it is also increasingly embedded within IoT devices. Most consumers lack any awareness of who their devices may be talking to, and what kind of data they might be sharing.

The 'Respectful Things in Private Places' project (part of the EPSRC-funded PETRAS IoT Hub) looks at empowering end users with smart tools that will help them to understand, make informed decisions about, and exert effective control over the data collection activities of such devices. It will also identify methods to enable more ethical and preference-respecting data processing.

The project began in October and runs until February 2019. It is led by Principal and Co-Investigators Professors Nigel Shadbolt and Max Van Kleek. It aims to empower people, not by making choices for them, but by helping them to make

sense of the flow of their data and the entities who collect it and giving them meaningful controls. Inspired by different theories of privacy, we hope to give end users new ways to think about their concerns and the choices available to them.

The project aims to increase our understanding of end users' perceptions of data collection practice of devices in highly private settings, as well as an understanding of the spectrum of concerns, and preferences regarding such practices. We will examine whether different privacy framings, conceptualisations and theories influence the ways people think about privacy trade-offs in these spaces, and identify those that most effectively reflect and resonate with end users' concerns.

More project information: goo.gl/VbLQp5

Location tracking to benefit emergency responders

Professors Niki Trigoni and Andrew Markham have been awarded US\$1.1 million for a research project entitled 'Pervasive, Accurate and Reliable Location Based Services for Emergency Responders'.

The project aims to develop fit-for-purpose location-based services for emergency responders in

environments without GPS by addressing the specific challenges they face such as lack of floorplan information, poor visibility and communication blind spots. The team hope their research will translate into a key safety tool for emergency teams, enabling better coordination and faster incident resolution, and preventing fatalities and injuries.

The award was made by the US Commerce Department's National Institute of Standards and Technology (NIST). It is funding 33 research and development projects

to the tune of US\$38.5 million, and is aimed at advancing broadband communications technologies for first responders.

The grants are intended to help modernise public safety communications and operations by supporting the migration of data, video and voice communications from mobile radio to a nationwide public safety broadband network, as well as accelerating critical technologies related to indoor location tracking and public safety analytics.



Tackling the blight of fake news

Professor Marina Jirotko and Senior Researcher Helena Webb consider what can be done by government and social media platforms to tackle the problem of fake news.

As campaigning in the UK General Election gained momentum in April 2017, the Chairman of the Government's Culture, Media and Sport Select Committee called for Facebook to improve its handling of fake news on the platform. Referencing concerns that the spread of false stories across social media had influenced the results of the 2016 US Presidential election, Damian Collins MP suggested that the propagation of content of this kind could threaten the 'integrity of democracy'.

Worries over the apparent prevalence of false content online and its capacity to have significant offline effects have grown rapidly over the past year and fake news has become an established social problem. Whilst the spread of rumour has always been a feature of social life, we can observe certain novel dynamics in the fake news phenomenon.

Firstly, the hyperconnectivity brought about by the popularity of social media means that online content of any kind can spread on an unprecedented rapid scale. Combined with the apparent growing user reliance on social media as a news source – in particular amongst young people – this creates a vulnerability where false stories can easily propagate.

They might then take hold if users conduct themselves online in a 'filter bubble' in which they surround themselves with only similar

viewpoints and are not exposed to alternative or conflicting versions of the 'truth'. These filter bubbles are in turn reinforced by social media platforms' own algorithmic processes as users are presented with personalised content that complements what they have already looked at and liked and are less likely to be shown counter-content.

User behaviour and the nature of social media thereby appear to provide fertile ground for the spread of fake news. A further key concern is that this vulnerability can be exploited so that false content is propagated in an organised way for the purposes of profit (gained via online advertising) or political interference.

Inevitably questions arise over how fake news can be addressed, with much attention – including from the UK Government – focusing on the suggestion that social media companies should take more responsibility for resolving the problem. Research studies, including our UnBias project, explore how changes to the regulation of social media might prevent or limit spread of fake news.

One of the most radical changes could involve a shift in the legal status of social media organisations so that they become more comparable to traditional publishers such as newspapers in terms of the responsibility they must take for content posted. Less radical, and perhaps more technically and politically likely, is the development of Codes of Conduct for social media platforms. Platforms could sign up to undertake various practices in response to potential

fake news. This would not necessarily involve the removal of content – something which would lead to strong objections on the grounds of freedom of speech.

Other practices might include the use of 'kite marks' to display the trustworthiness of certain news stories (based on features such as the provenance of the story and the existence of counter stories), feedback functions through which users can vote on the likely truthfulness or otherwise of what they have read, or the use of algorithms designed to pierce filter bubbles via the presentation of alternative content. However as the fake news phenomenon is not simply a technology-based one but is also grounded in social practices, solutions to the problem may also need to look beyond changes to the regulation of platforms.

Particularly important might be education and user self-governance practices in which individuals identify potentially false content and act to stop themselves and others spreading it. Such practices can also help mitigate the spread of false information and also develop more critical faculties in consumers, particularly the young who may often accept news online as the truth. Public and political debates about fake news seem set to continue in the future and efforts to address the apparent problem will benefit from careful research scrutiny.

This article was first published in the summer 2017 issue of ESRC's magazine, *Society Now* (issue 28), and then on the ESRC blog.

UnBias is on Twitter: @UnBias_algos

Research to identify when to trust online ratings

Users justifiably believe that some ratings given to a product or service online are manipulative. Tim Muller from our Security Research Group is embarking on EPSRC-funded research to identify when a rating system is sufficiently robust to overcome such manipulation.

There is a certain wisdom of the crowd. Others may have information that you're lacking, and be prepared to share it with you online. Many online systems capitalise on this. For example, recommender systems take in users' preferences, and try to predict what they would prefer by extrapolating this data. Online reviews allow users to read what others are thinking, and influence their decisions.



A specific form of sharing information is via ratings. Examples are: 1-5 stars, flag a comment, thumbs up/down, or selecting an appropriate emoticon.

Whenever you design a system that uses ratings, there will be attackers trying to manipulate these ratings. Several classes of such attacks exist in theory and in practice. A simple example is attackers using multiple accounts to provide good ratings about themselves. We might be able to detect attacks after they've occurred, and block the attacking accounts. However, there is nothing stopping the attacker from creating more accounts, and masquerading them as honest, up to the point where they perform another attack. So, the question is, can we use these ratings to make the right decision, even if these attacks inevitably do occur?

To define 'the right decision' is not always trivial. Think of a movie recommendation system; is it meaningful to talk about the right movie to watch? If, on the other hand, you want to download an application that has been rated to (not) contain malware, then the right decision is obvious: you want to download it, if, and only if, it's actually malware free. We look at cases where the right decision is obvious, but the issue is only to deduce this from the ratings that are given.

Now, if (and this is a big if!) honest ratings are perfectly accurate and a rating is more likely to be honest than malicious, then the majority of the ratings is expected to be correct. For example, say 70% of the raters are honest, and we ask five people, then we expect three

or four honest ratings. But we may get unlucky and get only two honest ratings. What if we want to be 99% sure that we make the right decision? In this case, it's a matter of asking enough people to rate: namely 29.

For that simple example, we implicitly asserted that attackers always lie. What if they don't? Well, in the example, the amount of correct ratings goes up, so the

strategy of going with the most rated option can only improve. We can therefore say confidently that the probability of making the right choice is at least 99%, no matter what the attackers do. Therefore, we can say that asking 29 people to rate, and following the majority vote, is '1%-robust'. When asking fewer people to rate (27, the largest odd number below 29, avoiding ties), there exists a strategy for the attacker to increase the probability of making the wrong decision to over 1% (1.17% if the attacker always lies). Thus, asking 29 people is the minimum required for 1%-robustness; we call this optimality. Finally, stability is about the fact that, as long as honesty is more likely than maliciousness, following the majority is the best option.

When users interact multiple times, simple majority schemes cease to be optimal. Then, furthermore, 'always lying' is not our worst-case; we must consider attackers who sometimes tell the truth. The upside of multiple interactions is that robustness is achievable, even when honest users are outnumbered. Intuitively, this is because fake ratings have a non-zero probability of being identified as such, and decrease posterior probability of honesty for attackers, and increase it for honest users. Effectively, we slowly start trusting honest raters. Finally, introducing the notion that honest users aren't perfect, and may make mistake lowers the effectiveness of trust.

My goal during my one-year EPSRC First Grant is to find relationships between all these parameters, and effectively identify when robustness of ratings can be achieved.



Artificial Intelligence @ Oxford

- Talks by the experts
- Q&A panels
- Demos
- Meet the professors
- Student posters

Topics include:

- Robotics
- Driverless cars
- AI in healthcare
- AI for scientific discovery
- AI and employment
- AI and finance
- AI and privacy
- Ethical issues for AI

A One Day Expo



A university-wide event
at Worcester College, Oxford

27 March 2018

A unique opportunity to see the state-of-the-art in artificial intelligence and machine learning at one of the world's great universities, and to meet Oxford's AI experts one-to-one.

<http://ml.ox.ac.uk/>