InspiredResearch Summer 2018 Issue 12

News from the Department of Computer Science, University of Oxford

What's happening to your data?

Investigating how to make respectful virtual assistants – p26

Oxford's prowess in artificial intelligence:
University-wide expertise on show at expo – p14

Hip surgery guide:
Software to improve diagnosis and modelling for arthritis – p19

Simulations for drug testing:
Heart models that could replace animal trials – p20

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

Contributors

Alex Rogers, Alina Petrova, Anthony Lin, Blanca Rodriguez, Cas Cremers, Despoina Magka, Elisa Passini, Gregory Farquhar, Irina Voiculescu, İsmail İlkan Ceylan, Jakob N. Foerster, Jeremy Gibbons, Jun Zhao, Leanne Carveth, Margherita Nulli, Marina Jirotka, Martin Kraemer, Matt Smith, Michael Wooldridge, Philip Inglesant, Sam Staton, Stefan Kiefer, Thomas Lukasiewicz, Will Seymour.

CONTENTS

Feature articles

Letter from the Head of Department

At some point in the coming few years, I expect to write an introduction to *Inspired Research* that doesn't feel the need to mention artificial intelligence (AI), but right now it seems that time is rather remote.

AI has featured prominently in the life of the department over the past academic year, and I am delighted to report that our recent AI@Oxford event gave a dazzling exposé of the work we do on AI across the whole collegiate university.

We welcomed 150 guests to the event, from government, industry and academia, and showed them the very best of our work in AI, from the philosophical foundations of the subject, through the basic science, to programs that help doctors diagnose heart disease, and driverless cars that can navigate across the frozen wastes of Iceland. It was a wonderful day, and a resounding affirmation of the University of Oxford's primacy in this most exciting of areas. You can find out more about the event on pages 14–15.

Events elsewhere (I don't need to name them) have made all of us that work in technology education acutely aware of the issues surrounding the ethics of what we do, and just how fragile the concept of privacy is in the online world. A decade ago, many students were uninterested in learning about ethics – they regarded it as an unnecessary distraction from the core of their studies. The situation now could hardly be more different: students nowadays positively *demand* ethical training.

What has changed? In short, the internet. Everyone is connected to everyone else, and as a consequence, a student working from their flat on Cowley Road in Oxford can affect the outcome of an election on the other side of the world. The consequences of what we do in the online world are global. We are still finding a path through this brave new world, and I am pleased to say Oxford researchers are leading the charge for an ethical online future.

Professor Marina Jirotka has led the way on responsible innovation, and on understanding the dynamics of social media 'wildfires' (see articles on p7, p24 and p25). And Professor Sir Nigel Shadbolt's recent book *The Digital Ape* gives a fascinating perspective on what it means to be human in the age of social media and AI.

The issues raised by the recent scandal (no, I'm *not* going to name it!) are difficult and complex, and we will continue to wrestle with them in the years to come. Oxford is well placed to rise to the many challenges that we will face.

Professor Michael Wooldridge
June 2018

News in brief

Head of Department, Professor Mike Wooldridge, has contributed to the '2017 AI Index Report'. Created and launched as a project of the 'One Hundred Year Study on AI' at Stanford University, the AI Index is an open, not-for-profit project to track activity and progress in artificial intelligence.

Professor Tom Melham has been appointed to The Alan Turing Institute's Board as the non-executive director and trustee for the University of Oxford. The Alan Turing Institute is the national institute for data science, and its board of trustees is made up of independent members and nominated trustees from each of its six founder members (the universities of Oxford, Cambridge, Edinburgh and Warwick, UCL and the EPSRC). Tom's research is focussed on mathematical methods for assuring quality and correctness of computer systems. Further information about the Alan Turing Institute: www.turing.ac.uk

Professor Marina Jirotka has been elected to the UK Computing Research Community (UKCRC), a highly selective grouping of the UK's leading computing academics. The UKCRC aims to promote the vitality, quality and impact of computing research in the UK. Its members are internationally leading computing researchers drawn from both academia and industry. Marina is Professor of Human Centred Computing, and her expertise involves co-producing user and community requirements and human computer interaction, particularly for collaborative systems. She is also co-founder of ORBIT (see p25) which provides training and consultancy in responsible innovation principles.

Department sweeps up at student-led awards

The quality of the department's teaching and support staff has been recognised at the Oxford SU Student-Led Teaching Awards. Our department had three winners – Lead Administrator Maureen York [pictured above left], Professor Andrew Ker and Professor Ian Horrocks [pictured above right] – in the awards, which are based on student nominations from across the University. Four other members of the department – Professor Cas Cremers, Postdoctoral Researcher Andreas Galanis, Professor Andrew Martin and Senior Lecturer Irina Voiculescu – were shortlisted.

Maureen was a winner in the category for the best member of support staff. She is the lead administrator in the Cyber Security Centre for Doctoral Training (CDT) and assisted in setting it up in 2012/13. She has worked in the University since 1994.

Ian won in the Most Acclaimed Lecturer category. Ian has been working in the area of knowledge representation and reasoning (KRR) for more than 20 years, and leads the department's KRR research group. He has been teaching computational complexity in the department for several years, and this year taught a course on KRR, which is partly based on his recent book, *An Introduction to Description Logic*.

Andrew Ker was a winner in the Outstanding Tutor category. He has been an associate professor in the department since 2009, lecturing on steganography and steganalysis in a course on advanced security, as well as computer security, discrete mathematics, and lambda calculus and types. He has also been a Tutorial Fellow in Computer Science at University College since 2009.

Two other members of the department – Andreas and Irina – were shortlisted for the Outstanding Tutor category, while Cas and Andrew Martin were nominated for the Outstanding Graduate Supervisor award.

Cas' research focuses on information security, and he currently supervises six DPhil students. He has been at Oxford since 2013, and is a fellow of Kellogg College.

Andrew directs the CDT in Cyber Security, and supervises DPhil students there. He has been with the department since 1999, and is also a Governing Body Fellow at Kellogg College.

Andreas teaches the Probability and Algorithms courses at Hertford College.

Irina has been in the department since 1999, specialising in 3D Geometric Modelling, and is a Stipendiary Lecturer at Keble College where she tutors first and second year students.

The winners of the awards were announced at a ceremony in Oxford Town Hall on 10 May.

Rewards for engaging with commerce and non-specialists

Two of the department's professors – Ursula Martin and Blanca Rodriguez [below] – have been chosen to receive Impact Awards by the Mathematical, Physical, Engineering and Life Sciences (MPLS) Division at the University of Oxford.

The MPLS awards aim to foster and raise awareness of impact by rewarding it at a local level, and were presented at a reception in February.

Blanca won an award in the commercial impact category. Her team has actively engaged with industry to demonstrate the power of computer simulations to evaluate the safety and efficacy of medicines. The team developed the Virtual Assay software to predict the response of populations of human cardiac cells to drugs, using computer simulations (see p20–21).

Ursula's award was for her work engaging non-specialists, particularly women, with Mathematics and Computer Science, through new research on Ada Lovelace's science. Her work to mark Ada's 200th anniversary included exhibits at many museums, including the National Museum of Computing, Bletchley Park, the Science Museum and the Computer History Museum in Silicon Valley.

Project to help realise potential of quantum

Professor Bob Coecke [left] and Niel de Beaudrap [right] have received funding from the EPSRC Impact Acceleration Account for developing a quantum computing compiler.

The one-year project, entitled 'Compilation and cost-reduction of quantum computations via ZX-calculus', looks to improve both the efficiency and the flexibility of the compiler, which translates high-level code to machine code. Its success would play a pivotal role in quantum computers becoming part of the 'real world', and realising their potential as one of the most promising future transformational technologies.

As in traditional computing, the efficiency of the compiler in quantum computing is vital, since otherwise the algorithmic gain of quantum computation could get lost in the translation from theory to practical implementation. However, unlike conventional computers which have a well-established hardware architecture, quantum computers may take many different forms, which substantially stretches the role of the compiler, having to mediate between entirely different physical devices.

The research team will use ZX-calculus, developed in 2008 by Bob in Oxford, to fill the gap between quantum hardware and quantum software. They will design ZX-calculus based algorithms that efficiently reduce the cost of implementing quantum circuits on quantum hardware, and are flexible enough to be used with any possible hardware architecture. The team will partner with Cambridge Quantum Computing.

Big data tools aid industry

Five trial platforms that apply the latest big data research findings to solving problems in industry have been successfully rolled out by the ALIGNED project.

The three-year ALIGNED project tackled problems in a wide range of areas, showing how the latest semantic technologies can help create smarter legal information systems, provide better management of health data, and help construct high-quality archaeological and historical datasets. In all of these areas, ALIGNED's semantic and model-driven technologies were able to help people who were struggling to harness the available data. The project, which was a collaboration between industry and academia, received €4 million in funding from the European Commission's Horizon 2020 programme.

'Oxford's contribution has concentrated on two such tools,' explains James Welch, researcher on the ALIGNED project in Oxford's Department of Computer Science. 'The Metadata Catalogue is an online platform for documenting element-by-element descriptions of large datasets or software models. Semantic Booster is an adaption of our existing Booster tool for the automatic generation of information systems from precise specifications.'

ALIGNED's research is explained in videos on YouTube at: goo.gl/cB6jTH

News in brief

Head of Department, Professor Mike Wooldridge, has won the 2018 IFAAMAS Influential Paper Award, which recognises publications that have made influential and long-lasting contributions to the field of autonomous agents and multiagent systems. He was recognised by the International Foundation for Autonomous Agents and Multiagent Systems for two papers: 'The Gaia Methodology for Agent-Oriented Analysis and Design' and 'Developing Mulitagent systems: The Gaia Methodology'. Together these papers have 4,270 citations in Google Scholar.

The 2018 QS World University Rankings for Computer Science and Information Systems put the University of Oxford in second place in Europe. It ranked seventh in the world. The QS World University Rankings are designed to assess universities in four areas: research, teaching, employability and internationalisation. Full details: goo.gl/PpU6Xp

Professor Georg Gottlob had the honour of giving this year's BCS Lovelace Lecture in March at the Royal Society in London. He was invited to give the presentation, which was entitled 'Swift Logic for Big Data and Knowledge Graphs', after being awarded the Lovelace Medal. Also in March this year, a special event was held in Vienna in honour of Georg's influential and long-lasting research contributions in the fields of database systems and artificial intelligence. It was organised as a prelude to the International Conference on Extending Database Technology by Michael Benedikt from our department with TU Vienna, where Georg holds an Adjunct Professor position.

Robot morality to be debated at FLoC

The Federated Logic Conference (FLoC) 2018 taking place in Oxford on 6–19 July 2018 will include some events that will appeal to the wider, non-specialist, community. A debate on the morality of robots is among the fringe events that will form part of the two-week conference. These fringe events (described below) are available to book on an individual basis, and will be recorded for those who cannot make it.

Much of the rest of the conference will tackle specialist subjects under the guise of several international conferences related to mathematical logic and Computer Science. Logic underpins many computing subjects, such as programming, machine learning and algorithms.

FloC this year is being organised by and hosted at Oxford with Professor Moshe Vardi of Rice University as General Chair. FloC, which has been running since 1996, now operates on a four-year rotation, and is this year being co-chaired by Professors Daniel Kroening and Marta Kwiatkowska from Oxford's Department of Computer Science.
FLoC fringe events include:

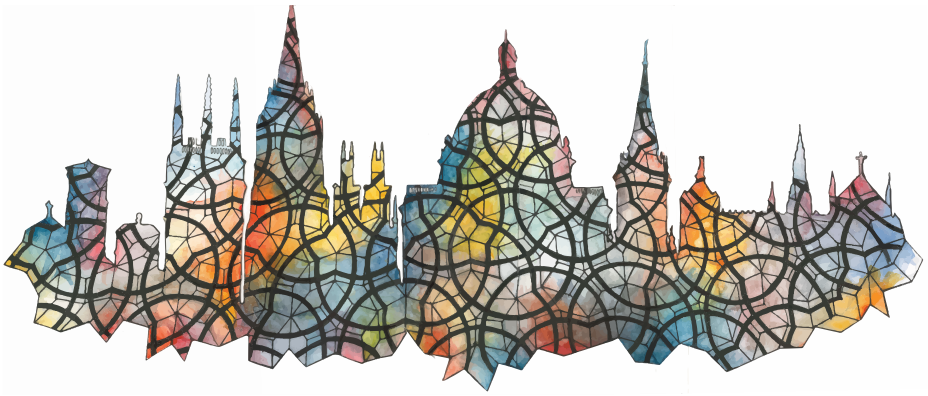

The Summit

The Summit on Machine Learning Meets Formal Methods on 13 July will bring together 13 academic and industrial leaders discussing the benefits and risks of machine learning solutions. The overall aim is to identify promising future directions for research and innovation of interest to The Alan Turing Institute and UK research councils and government agencies, which will be summarised in a written report that will be made public.

The Debate

The Debate questions the ethics and morality of robotics, such as: Will robots be able to act as agents in their own right and make moral and ethical decisions? Are robots capable of replacing human beings? What are the ethical implications of such a future?

A panel specialising in ethics, law, computer science, data security and privacy will debate these issues in the Oxford Union Debating Chamber on 16 July.

Public Lecture

Professor Stuart Russell will give a talk on 10 July at the Sheldonian Theatre, Oxford, entitled 'Unifying Logic and probability: The Blog language'. In this public lecture, Russell will describe the Bayesian Logic (Blog) language mainly through examples and cover its application to monitoring the Comprehensive Nuclear-Test-Ban Treaty. The talk is part of the termly Strachey Lectures which are generously supported by OxFORD Asset Management.

More information on FLoC, including the programme with session timings and abstracts, is at:
www.floc2018.org

A second paper has long-term impact

The department's Professor Samson Abramsky is to receive the LICS Test-of-Time Award for a paper that he wrote in 1998 with Kohei Honda and Professor Guy McCusker. The paper is entitled 'A fully abstract game semantics for general references'.

The award, bestowed by the IEEE Symposium on Logic in Computer Science (LICS), recognises a small number of papers from the LICS proceedings from 20 years prior that have best met the 'test of time'.

Samson, who is Christopher Strachey Professor of Computing and a Fellow of Wolfson College, previously received a LICS Test-of-Time Award in 2007 for the paper 'Domain theory in logical form' published in 1987. The papers are selected by an awards committee that is appointed by the LICS General Chair and consists of between three and five members. In selecting these papers, the awards committee consider the influence that the papers have had since publication; because of the foundational nature of LICS' work, impact is often not fully felt immediately, hence a 20-year perspective.

The award will be presented at LICS 2018, which will be held in July in Oxford, as part of the Federated Logic Conference (FLoC) 2018.

The paper is at: goo.gl/ao56Gk

Tool to build trust in AI

In order to gain the public's trust in artificial intelligence (AI) online systems, the University of Oxford will be leading a project to develop a tool that allows users to evaluate and critique algorithms.

The 'ReEnTrust: Rebuilding and Enhancing Trust in Algorithms' project has received a grant of almost £1 million from the EPSRC. As well as setting up an online tool to help users assess major web-based platforms, the initiative aims to develop techniques to allow all parties to explain their views and suggest possible compromise solutions when trust in algorithms has been lost.

The researchers, who are from the Universities of Oxford, Nottingham and Edinburgh, also plan to develop a methodology for deriving a 'trust index' for online platforms that would allow users to easily assess the trustworthiness of web-based platforms.

'We will be working with industrial partners to make sure that our research relates to real-world examples, and that it fairly represents all those affected by this work,' said Principal Investigator Marina Jirotka from Oxford's Department of Computer Science. 'We also believe that by talking extensively to users, platform service providers and

other stakeholders while developing this tool, we will gain a deeper understanding of what makes AI algorithms trustable.'

Marina's team will be using their expertise in responsible research and innovation to help create methodologies that support responsible development of AI online systems. The project will also consider to what extent user trust can be regained through technological solutions and if it might be necessary and appropriate to seek other ways to rebuild trust, for instance through policy, regulation or education.

A new video animation by the UnBias project explains what algorithms are, how they shape online browsing and how they can create risks of bias: goo.gl/dhKVTn

The project summary is at: goo.gl/pZU4yC

Report calls for environment fit for childhood

Professor Marina Jirotka was one of the contributors to a major report by academics that is calling for a digital environment that is 'fit for childhood'.

The authors believe that the digital environment has created welcome opportunities for children and young people and excluding them from it is neither an option, nor desirable. However, mundane and prevalent risks, including insomnia, obesity, low self-esteem and oversharing, present real harm to childhood development and tend to be overlooked. For example, sleep deprivation caused by extended use of digital devices can affect concentration, performance at school and general wellbeing.

Key recommendations from the Digital Childhood report include: children receiving smartphones for the first time being educated about them; design standards meeting childhood development milestones; and government taking into account children's views when creating policy. The report was written by a group of academics led by Angharad Rudkin from the University of Southampton.

The full report is at: www.5rightsframework.com

Ethics and design challenge combined

The ethical hackathon, which is a novel twist on the hackathon that highlights ethical issues alongside design features, made its debut in February.

A group of Oxford researchers organised its first ethical hackathon at the University of Nottingham as part of the ongoing UnBias project, which is examining the responsible and fair use of algorithms. They challenged teams to design a social networking platform that uses algorithms responsibly. The teams presented their ideas to an expert panel, who gave prizes for the best designs. In an ethical hackathon, teams carry out a design task but in addition to thinking about technical features, as would be the case for a traditional hackathon, they are also required to identify the ethical implications of the particular technology involved. For instance, does the technology risk causing harm? Will it decrease user autonomy or undermine privacy? Will users have the opportunity to determine for themselves how the technology will fit into their day-to-day lives?

Ethical hackathon participants are challenged to identify novel and creative solutions to these potential problems and to find ways to embed ethical considerations into their designs. The concept has been developed by Professor Marina Jirotka, Helena Webb and Menisha Patel [below] from the department's Human Centred Computing theme.

Literary feast focuses on AI

Artificial intelligence (AI) has been a popular topic in recent books published by the department's academics.

Head of Department Professor Mike Wooldridge [left] wrote *Artificial Intelligence: A Ladybird Expert Book*. Aimed at adults, the book is an introduction to intelligent machines and the humans that program them.

The Digital Ape, co-authored by Professor Sir Nigel Shadbolt, looks at how smart machines are transforming our lives and societies, and aims to dispel misconception and confusion about them.

Senior Researcher Paula Boddington wrote *Towards a Code of Ethics for Artificial Intelligence*, which investigates how to produce realistic and workable ethical codes or regulations to address immediate and longer-term issues.

Mike and Nigel spoke at the Oxford Literary Festival in March. Mike talked about the reality of AI, and Nigel chaired a discussion on robot sex.

Also at the Oxford Literary Festival, Dame Wendy Hall [pictured below], Regius Professor of Computer Science at the University of Southampton, gave a talk entitled 'Artificial Intelligence through the Looking Glass' as an Ada Lovelace Lecture held jointly by the University of Oxford's Bodleian Libraries and Department of Computer Science.

Dame Wendy, who is executive director of the Web Science Institute, and co-chair of the UK government's AI review, explained in the lecture how AI can already outperform humans on some tasks. She also looked at the implications for human society.

Hundreds listen at festivals

Nearly 1,000 festival goers attended events run by our academics at the Hay Festival this June. Professor Sir Nigel Shadbolt talked about machine learning and his book *The Digital Ape*. Elisa Passini led her audience on a journey through the virtual heart. Jamie Vicary and David Reutter ran a 'Build your own quantum computer' workshop and Bushra AlAhmadi, Klaudia Krawiecka and Saumya Jetley led a workshop on robot programming.

A week later, the department's Professor Sadie Creese and Ana Mincholé were on stage at the Cheltenham Science Festival. Sadie talked about the catastrophic effects that would result from an attack on a nation's critical infrastructure, and work to defend against hostile actors in the virtual world. Ana took part in the 'Building a Virtual Human' talk with the Science Museum's Roger Highfield and other members of the European CompBioMed project.

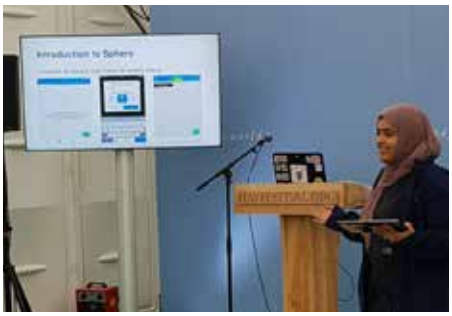

Sixty glorious years animated

A new three-minute video to celebrate 60 years of Computer Science by the department has been put together by the University.

The video starts when the Oxford University Computer Laboratory (as it was known until 2011) was set up with one hand-me-down computer known as 'the monster in the basement' and a lot of paper. It then runs through the department's expansion in research and development, namechecking important figures such as Christopher Strachey who became the first professor of computation.

The video highlights important landmarks in the department's research, such as work with Intel on chips, and a virtual model of the heart being made available to drug companies. It also traces the introduction of the various teaching programmes – doctorates, masters' and undergraduate degrees – through to today's exciting research projects, such as those in the fields of quantum, artificial intelligence, ethics and cryptography.

Tweeting influences in UK election analysed

The news media's role in the flow of news content on Twitter during the 2017 UK General Election has been analysed by a team from the Reuters Institute for the Study of Journalism (RISJ) and the Department of Computer Science.

The researchers – including Jun Zhao and Jason R C Nurse [pictured above right] from Oxford's Department of Computer Science, and Sílvia Majó-Vázquez and Rasmus Kleis Nielsen from the Reuters Institute – examined the role of long-established brands (legacy media) and the more recently created digital-born outlets on Twitter, by looking at their tweeting activities and influences. To do so, they analysed 4.26 million tweets between May and June 2017, starting by assessing the distribution of news content throughout the electoral campaign cycle.

They then ranked the providers of news-related tweets according to their tweeting activities and retweet counts. These measures of news production were combined with audience engagement (the number of mentions and replies to each tweet) to offer a more nuanced understanding of the outlets' influence over audiences during the electoral event.

The analysis showed that UK broadcasters and national newspapers were dominant in both the overall production of news content on Twitter and in attracting the highest levels of engagement from audiences. The case was not totally clear cut, however, with some small purely digital outlets exhibiting a very high level of influence on Twitter conversations around the election, some of which are at the extremes of the ideological spectrum. These new media outlets, such as *Breitbart* and the *Lad Bible*, gained much higher levels of audience engagement than would have been anticipated given their overall audience reach, as did brands such as *The Economist.*

The research is online as a RISJ fact sheet: goo.gl/GHGh9x

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

- A short film about future virtual humans incorporating material from research by Professor Blanca Rodriguez and her group has been released on YouTube. It premiered at the London Science Museum IMAX in September 2017, and shows the possibilities of using High Performance Computing (HPC) to delve deeper into the processes taking place within the body.

- The latest Strachey Lecture, entitled 'Privacy-preserving analytics in, or out of the cloud', was delivered by Professor Jon Crowcroft from the University of Cambridge on 6 March. Jon's talk was about how to provide privacy when running analytics on users' personal data. His talk was recorded and posted on the Media Wall as part of the Strachey Lecture series, which is generously supported by OxFORD Asset Management.

- Professor Alex Rogers answers the question 'Where have all the cicadas gone?' in a podcast produced by Oxford Sparks, a University portal used to engage the public with exciting science. Alex is leading an experiment with the hope of rediscovering The New Forest Cicada, for which the last unconfirmed sighting was in the year 2000.

The above videos, and many others, are available at: www.cs.ox.ac.uk/mediawall

Four £500 prizes awarded for project work

This year's Group Design Practical culminated in an exhibition and formal presentation held in the department on 9 May. Teams of five to six second year undergraduates had worked on their chosen projects since January, battling it out for four £500 prizes. Many of the challenges were set, or sponsored by, industry partners.

The judging panel consisted of one representative from each of the sponsors – GResearch, Palantir, Vodafone and Oxford's Department of Computer Science. Each industry partner awarded a prize based upon their own criteria. The Department Prize was awarded for the best presentation.

Teams chose a challenge from 14 different project briefs proposed by industry partners and department colleagues. Part of the work was to undertake a proper requirements analysis, working with their project mentor.

Team 2 won the Palantir Group Project Prize with its 'Cryptocurrency Correlations' project. The team members were Marilena Bescuca, Andreea Stanciu, Joohyun Yoo, Debjit Mandal, Matthew Hillman and John Mewes. Joe Pitt-Francis was their mentor, and this project was supported by the investment company OxFORD Asset Management.

Team 8 was awarded the GResearch Prize for its 'Fake News' project. The team consisted of Zihan Ye, Eleanor Williams, Barbara Bilakiewicz, Max Timbs, Thomas Foster and Andrei-Alexandru Radoi. Their mentors were Helena Webb, Jason Nurse and Ross Gales.

Team 10 received the Presentation Prize for its project 'Carbon Intensity and Smart Devices' (see below). Team members were Denitsa Markova, Alexandru Strimbu, Zihang Lai, Zhenxian Yang, Jan Bialas and Fanny Duneau. The group was mentored by Professor Alex Rogers.

Finally, the Vodafone Prize went to Team 11 for a project in collaboration with GResearch. The team consisted of Edward Salkield, Hristo Venev, Xinyue Tong, Oliver Weller-Davies, Anna Whitmore and Annazita Barry. They were mentored by Senior Researcher Jiaoyan Chen.

All second year undergraduates reading Computer Science, Computer Science & Philosophy or Mathematics & Computer Science took part in the Group Design Practical. It is designed to allow the students to practise their skills learnt previously on the course, and helps them develop and apply theory already learnt. It also helps develop team-working abilities, and project and time management skills.

Companies interested in sponsoring projects or prizes in future years, or wanting to learn more, contact Leanne.Carveth@cs.ox.ac.uk

Project goes to House of Lords

Undergraduates from the department showcased their project work at the House of Lords in May, as part of a launch event for a tool that predicts regional carbon intensity (www.carbonintensity.org.uk).

The event, which was held by the tool's developers – National Grid, the Environmental Defense Fund Europe and the World Wide Fund for Nature, also showcased new approaches to de-carbonising the electricity grid. The Oxford student team was invited to demonstrate its second year Computer Science Group Design Practical work, which National Grid had sponsored.

The undergraduates have built a web-based system to allow consumers to optimise and schedule significant electrical loads (such as electric vehicle charging) to minimise their carbon footprint.

Denitsa Markova, Alexandru Strimbu and Tiffany Duneau gave the demonstration in London, accompanied by the department's Professor Alex Rogers. He helped verify the methodology of the partner's API tool, and maintains GridCarbon – a smartphone app that presents live carbon intensity information about the UK electricity grid: www.gridcarbon.uk

WCIT prize for Marius

Marius Gavrilescu, a third year undergraduate student in Computer Science at St John's, has won a WCIT Outstanding Information Technology Student Prize 2018. The Worshipful Company of Information Technologists (WCIT) advises charities on IT issues, provides mentoring for IT professionals and supports schools. Award criteria for its prizes are academic excellence, overcoming adversity, entrepreneurial skills, and contribution to charity and community. Marius was presented with his certificate and cheque in the Armourers' Hall in London.

Oxford team makes semi-finals at cyber policy competition

A team of four cyber security DPhil students made the semi-finals at Cyber 9/12, an Atlantic Council policy competition held in London in February. Erin Chapman, Mariam Nouh, Jantje Silomon and Andrew Dwyer, all from the Centre for Doctoral Training in Cyber Security, also won the award for the best policy brief which outlined cyber security policy recommendations.

Fifteen teams from universities across the UK, law enforcement, and the Ministry of Defence were briefed at BT Tower in London on a crisis scenario, and had to deliver policy recommendations to a mixed panel of decision makers including expert representatives from academia, government and industry.

At the end of two gruelling days, the final few teams had just 20 minutes to present their final ideas to the judges. The event saw a gathering of industry professionals across sectors who acted as judges in competition panels and a fantastic pool of talent from across the country.

This event would not have been possible without a strong organisation team, including four Oxford DPhil students from the Centre for Doctoral Training in Cyber Security: Arianna Schuler Scott, Monica Kaminska, Jack Sturgess and Martin Kraemer.

'It was inspirational to see such diverse and enthusiastic talent coming together to tackle issues in our world where technologies continue to evolve, and the speed at which we have to regulate for and respond to cyber incidents increases,' said Arianna, Deputy Director of Cyber 9/12.

Beginner coding sessions recorded for YouTube

Oxford's Computer Society (CompSoc) recorded its packed-out 'Learn to Code' course for the first time this year. Thomas Denney, [pictured above] an undergraduate in Computer Science and President of CompSoc, delivered and recorded the set of five introductory Python sessions for non-computer scientists, which he also presented last year.

More than 200 students attended the course held in Hilary Term. This was a new record for CompSoc which has been running the sessions for a few years.

The students were from across the university, with around two-thirds from STEM subjects. The sessions packed out both lecture theatres in the Wolfson Building, with fellow undergraduate Sauyon Lee [pictured right] presenting in the second lecture theatre.

Around another dozen Oxford Computer Science students and CompSoc members dedicated their Thursday evenings to help run the course this year, including registration and assisting students with exercises.

The videos are available on YouTube at: goo.gl/Q63nfP
Lecture materials: goo.gl/YVDUJL
Thomas's blog: goo.gl/hCDCgx

Oxbridge women join up for conference

The 5th Oxbridge Women in Computer Science conference on 15 March attracted more than 100 participants – undergraduate and graduate students and postdocs – from the Universities of Oxford and Cambridge.

The conference, sponsored by Google, was organised by the Oxford Women in Computer Science Society (OxWoCS) together with a sister society from the University of Cambridge, women@CL.

The conference consisted of two keynote talks by Amanda Prorok from Cambridge and Anna Ukhanova from Google, 12 student presentations, a poster session, a panel discussion, and a minute madness session (lightning talks presenting early research ideas and work in progress as well as completed work, in a short and concise way).

Nathalie Cauchi [pictured above], a doctoral student, from Oxford's Department of Computer Science won the best presentation award.
Conference website: goo.gl/TdCeWw

Alumni Profile

Despoina Magka: *works as a software engineer with Facebook's Account Integrity team. The fundamentals of Computer Science and core skills that she learned during her MSc and DPhil at the Univerdity of Oxford have stood her in good stead in her career.*

After studying Electrical and Computer Engineering in Athens, I arrived in Oxford in 2008. I completed an MSc in the Department of Computer Science, where I delved into a wide range of exciting subjects, from computational linguistics and quantum computer science to dynamic epistemic logic.

That year was one of the most intense and enjoyable of my life, but way too short, so I was thrilled when I received an offer to pursue a DPhil in the same department with generous EPSRC support. I had always been fascinated by logic and also had a strong interest in its practical applications, so Professor Ian Horrocks' Knowledge Representation and Reasoning group was the perfect match for me. The group gave me a precious opportunity to work alongside world-leading researchers and learn from the best.

My DPhil years were a formative experience that helped me hone some of the core skills I still use today, such as approaching problems in a structured and disciplined way and shaping my own research agenda. They also gave me a bedrock of Computer Science fundamentals on which I have been relying on ever since. For my doctorate, I designed and implemented ontology languages for the classification of complex objects; I applied those to taxonomies of biochemical structures from the European Bioinformatics Institute in Cambridge. I travelled to present my research at academic conferences in various places, including Canada, Germany, Italy, China, France and Scotland.

I immersed myself in Oxford's wonderfully diverse community, attending and organising events of the various international societies, socialising in the premises of centuries-old colleges and meeting and chatting with talented people in some of Oxford's numerous pubs. Looking back now, I realise how privileged I am to have savoured this unique experience in one of the most magical cities in the world.

After defending my DPhil thesis in 2013 and spending a summer interning in Cisco, I joined Yahoo as a developer working with backend Hadoop technologies. Following a one-year stint there, I joined Facebook's vibrant London office as a software engineer. During my first year at Facebook, I was part of the Entities team, where I contributed to building and understanding the Graph Of Things that people are interested in and connected to.

Since then, I have been part of Facebook's Account Integrity team, where my job is to protect billions of people from online abuse, a role which I find tremendously rewarding. As a software engineer leading large scope efforts, I feel empowered by the skill set acquired during my doctoral studies. I am comfortable working on risky research projects, clearly communicating technical content and confidently asking the right questions.

The advice I would give to students considering their next career choice is to follow their passion, never stop learning new things and always seek opportunities that are fulfilling to them.

Portrait depicts role model

Department alumna Anne-Marie Imafidon (2006, Mathematics and Computer Science) won the Role Model of the Year prize at Booking.com's inaugural Technology Playmaker Awards on 8 March. The prize was aimed at women or men who acted as role models in a technology field, driving change through that role, and encouraging the next generation of leaders.

Anne-Marie is the co-founder of the STEMettes, a social initiative set up five years ago to inspire and promote the next generation of young women in the STEM sectors.

A new portrait of Anne-Marie has been installed in the Exam Schools, after being on display in Oxford's Weston Library in 'The Full Picture: Oxford in Portraits' exhibition. It was commissioned as part of the Diversifying Portraiture project, which features more than 20 paintings, drawings and photographs.

[Right] Anne-Marie Imafidon © Sarah Muirhead, the Diversifying Portraiture project.

Flight simulator evaluates cyber-attack reactions

Did you know that there's a flight simulator in the department's Robert Hooke Building? DPhil student Matt Smith describes how the computer scientist team is using it to evaluate flight crew response to cyber-attacks.

As system security researchers, we spend a lot of time on the outside of systems, looking in. This is particularly the case with our research into the security of avionic communications. Recent years have seen researchers in the department take an outsider look at vital systems such as ADS-B (Automatic Dependent Surveillance – Broadcast) and ACARS (Aircraft Communications Addressing and Reporting System), which are important in the smooth and safe operation of airspaces. We know a lot about how attackers might try to attack these systems, but we did not know how existing flight crew training would handle such attacks.

To investigate this, our team (Professor Ivan Martinovic as Principal Investigator, Martin Strohmeier and I) constructed a flight simulator using commercial and off-the-shelf components and software, and invited Airbus A320 pilots to take part in our experiments. As part of each session, a pilot flew scenarios in which we had simulated the effect of a cyber-attack on a different system that we consider to be technically possible and plausible.

The team carried out scenarios for three systems: glideslope, Traffic Collision Avoidance System (TCAS) and Ground Proximity Warning System (GPWS). Glideslope systems are part of the Instrument Landing System (ILS), which enables the autopilot to guide the aircraft to the runway at a predetermined angle and rate of descent – in situations where the aircraft uses auto-land capability, it relies on ILS to the point of touchdown. TCAS provides automatic resolution when aircraft become, or are about to become, too close, instructing each on how to separate themselves. GPWS is a top-level safety system which monitors a number of instruments to check whether the aircraft is getting too close to the ground.

Through this, we assessed two main parts of the response. Firstly, we could see which aspects of their training were called upon to handle different attacks, and the workload which such an attack added. Managing workload in the cockpit is crucial to keeping the aircraft safe, so attacks which significantly add to this are problematic. Next, we were able to measure the potential disruption of different attacks, both with regards to the level of trust the pilots had in the systems but also from a wider commercial and airspace system point-of-view.

Based on preliminary analysis, our research, which is funded by the Department Impact Fund, has

The team's flight simulator

demonstrated the need to work on security solutions to prevent these attacks happening. Each attack caused most pilots to have an increase in workload, as well as them having to carry out manoeuvers to resolve the problem. Typically, this would result in missed approaches or diversions, thus causing delay, passenger discomfort and inconvenience, and reputational harm to the airlines.

Beyond this, the work will inform future simulation-based experiments in aviation. The team will develop further simulations for pilots based on these results, as well as exploring other aspects of aviation which could benefit from cyber security simulations. More immediately, the results of the work are being shared with a number of industry bodies with the aim of understanding and acting on the data.

Research website:
www.avsec-oxford.com

Oxford showcases its strengths in AI

The first ever expo on artificial intelligence (AI) at the University of Oxford showcased the tremendous advances being made in the field, and looked at where Oxford's research is likely to lead. The AI@Oxford event brought together researchers from across the collegiate University, highlighting their expertise through a series of talks, debates, and demonstrations. Head of the Department of Computer Science Professor Mike Wooldridge reviews the event.

Mike Wooldridge gives the opening speech at the expo.

Reaching the point where a machine can translate Proust's beautifully written French into equally expressive English sentences is still some way in the future – if it is even ever possible. The issues with translating a classic tome by Proust are manifold: computers do not have knowledge of human relations and reactions, and they do not understand the nuances of French life in the period in which Proust was writing – and we see no easy way to give them this knowledge. On the other hand, automatic translation for non-literary works has become a daily reality for many, which 20 years ago seemed in the league of science fiction.

This is an example that I gave during my opening presentation at the AI@Oxford expo to help illustrate both the impressive reality of AI today, and the many challenges that remain for the future. During the event, delegates were given many other examples of the current state of AI – from driverless vehicles to health assistants to banking systems. Expert panels also discussed how to ensure that the technology leads to positive developments for society, minimising any possible downsides.

Another point I emphasised in my talk is that the AI community is currently focusing on developing AI in a relatively narrow way, which is a long way from the dream of general purpose intelligent systems. Current AI systems may be able to perform at super-human levels in certain narrow tasks, but scientists still have no idea how to make self-aware, conscious machines. The suggestion of robots being able to take over the world has no basis whatsover in current scientific reality.

A lot of the excitement in AI at the moment is being generated by advances in machine learning: the idea of machines that can learn for themselves how to conduct tasks, rather than be explicitly programmed to carry them out.
This may sound like an alarming idea, but again, although progress has been rapid this century, machine learning currently works only on certain narrow tasks.

AI systems are proving their worth in carrying out routine tasks more quickly and precisely than humans in various industries, and the AI event explored various possible uses, from automating ports to supporting healthcare. A port, for example, is ideal for automated vehicles and cranes because its borders are limited geographically, there are few people on site for machines to avoid, and the tasks performed are narrowly defined.

It is much more complex to use self-driving vehicles on public roads, due to the presence of more road users, and many more unexpected events, for example from pedestrians. Nevertheless Professor Paul Newman, who heads up the University's Oxford Robotics Institute (ORI), said he thinks it is a question of when, not if, driverless cars become the norm.

Oxford is conducting various pieces of research relating to self-driving vehicles. For example, Professor Shimon Whiteson from

continued on next page ▶

from previous page ▶

the Department of Computer Science talked at the expo about how his team is developing smarter simulators with the aim of making autonomous driving safer, and thus speeding up the process of getting self-driving cars onto public roads.

Meanwhile, the ORI is looking beyond driverless cars on the road to solving how to deploy them in more difficult conditions, for example off-road on a glacier, in the desert, and in poor visibility. To do this research, Paul's team has fitted out an ordinary Land Rover with a range of sensors including cameras, radars and lasers. The Land Rover was on show at the Sultan Nazrin Shah Centre in Worcester College, where the expo was held.

A host of system demonstrations were also on show at the expo, including: visual search technologies that enable users to search BBC footage for a specific face or scene; a machine learning system that induces logic (Prolog) programs from data; and a machine-learning image analysis solution which makes it easier for anyone anywhere to use ultrasound technology.

The delegates also received a flavour of student research being conducted at the University in a poster session. The posters covered subjects such as: the application of machine learning to drug research; the use of low-cost smartphones to detect mosquitos; and the application of machine learning to detect evidence of online radical behaviour.

As well as exploring the technological possibilities, the University is conducting research into the potential impact of AI on society. How to ensure that the future of AI has a positive effect on the population at large was a theme

The ethics panel: (from left) Professor Marina Jirotka (Department of Computer Science), Mariarosaria Taddeo (Research Fellow, Oxford Internet Institute), Allan Dafoe (Co-director, Future of Humanity Institute), Paula Boddington (Senior Researcher, Department of Computer Science).

that cropped up time and again at the AI Expo, and was discussed in depth by a set of panels on employment, ethics and security.

Ethics, explained Senior Researcher Paula Boddington from the Department of Computer Science, is all about making sure that AI is developed to be positive by setting and enforcing good ethical codes. The field can, on many points, learn lessons from codes developed in other areas, such as medicine. What is troubling about AI, in Paula's view, is that there is a guiding assumption in ethics that the professionals (in this case, machines) know what they are doing.

While the introduction of any emerging technology is never without risk, many speakers felt that the enormous potential benefits of AI outweighed the possible downsides. An example, given by the security panel, was: while a rogue group may use AI in social media to recruit new members, a defensive AI system could catch a vulnerability and repair it before anyone else finds out. Most speakers were in agreement that the greater danger in AI use is from malevolent people, rather than from glitches in the systems.

Another area for public concern, which was discussed by a panel at the expo, is the potential impact of AI on employment, and how the resulting changes might affect society. The employment panel, while agreeing that some jobs might be affected negatively (particularly those which could be easily automated), also believed that AI had an upside, for instance in increased productivity. Sandra Wachter, a Lawyer and Research Fellow at the Oxford Internet Institute, said that machines have a tendency to enhance the capability of humans, for example, working in tandem with a doctor on treatment plans, or helping banks make responsible lending decisions. The advantage of AI systems in making judgements is that they are not swayed by emotion or prejudice in the way that humans can be.

Feedback from the day has been overwhelmingly positive, as you can hear from delegates (from industry and government) themselves in clips from the day – link below.

It's heartening to hear that they were impressed by the research being conducted at the University of Oxford, and that we conveyed the central role that Oxford is playing in this most exciting of research areas.

AI@Oxford delegates' comments: goo.gl/ymoQDp
The Ethics panel debate: https://goo.gl/5RQBnv
Facebook Live interview with Paul on the future of autonomous cars: https://goo.gl/q9rPZi

Unravelling the complexities of AI in a team

The key to enabling several artificial intelligence (AI) entities to collaborate is for them to act independently but be trained centrally, according to a University of Oxford team. DPhil students Jakob N Foerster and Gregory Farquhar describe a method developed by computer scientists and engineers which could make it possible to deploy learning multi-agent systems in the real world.

In a fleet of search-and-rescue drones, each single drone typically needs to be able to decide on its best course of action using only local information. This is commonly referred to as 'decentralised execution'. However, often the design of the policies is carried out in a centralised fashion. For example, policies can be trained using a simulator which has access to the observations and actions of all agents.

A fleet of search-and-rescue drones is an example of multi-agent systems that can use learning AI individually while needing to be aware of one another. Our research differs from a lot of AI research, which often focuses on single agent settings and two player games.

'We're excited about the possibility of using these methods for the training of autonomous drones or cars...'

The world is full of challenging multi-agent problems: these range from self-driving cars to drones and even social interactions. In many of these applications a number of independent entities need to be able to take independent actions based on local observations in order to achieve a common goal. We believe that this domain of centralised training and decentralised execution is one of the key avenues for successfully developing and deploying multi-agent systems in the real world.

One of the great challenges when training multi-agent policies is the credit assignment problem. Just like in a football team, the reward achieved depends on the actions of all of the different agents. Given that all agents are constantly improving their policies, it is difficult for any given agent to evaluate the impact of their individual action on the overall performance of the team.

Above: Starcraft, a science fiction strategy game, was used to demonstrate the model's potential

To address this issue, our team (Computer Science's Jakob Foerster, Gregory Farquhar and Professor Shimon Whiteson, with Engineering's Triantafyllos Afouras and Nantas Nardelli) developed a method called 'Counterfactual Multi-Agent Policy Gradients' (COMA).

We demonstrated the model's potential by using StarCraft, a science fiction strategy game. The problem setting was unit management, which represents a challenging cooperative multi-agent problem.

In this setup, each of the units was one agent with partial observability, illustrated by the red circle [see image above]: We extended the actor-critic architecture with a centralised critic which learned expected returns (total points) given the joint action and complete observation of all agents. We then used this centralised critic in order to calculate a 'counterfactual' advantage for each agent.

This advantage compares the expected return after taking a given action to what would have happened had this agent taken a different (counterfactual) action. While evaluating each counterfactual in turn might seem slow and computationally expensive, we used deep learning to efficiently compute the counterfactual values in parallel.

Importantly, once the training was finished we threw away the critic and deployed the policies, allowing for decentralised execution. Putting everything together we obtained a training method which outperformed existing methods and achieved high win rates against the StarCraft bot.

This COMA method was described in a paper which won the 'Outstanding Student Paper Award' at the Association for the Advancement of Artificial Intelligence's conference, AAAI 2018, in February. The project was funded by the ERC and under the EU's Horizon 2020 Research and Innovation programme.

Our team is continuing to research further challenges in multi-agent coordination and exploration. We're excited about the possibility of using these methods for the training of autonomous drones or cars, which must learn to coordinate in the real world without being able to reliably communicate.

The full paper is at: goo.gl/CZ9n92

AI start-up tackles sophisticated data analysis

Artificial intelligence (AI) systems that can reason about and query large amounts of knowledge provide basic building blocks of many advanced data management applications. Three professors from the department have launched a spin-out, Oxford Semantic Technologies, with the aim of applying this technology to data integration and analysis.

Logic-based knowledge representation is a branch of AI that has an array of potential applications, from the oil to the healthcare industries. At its core are a set of technologies for querying large amounts of data while taking into account background domain knowledge. These technologies provide the basis for the start-up called Oxford Semantic Technologies founded by Professors Ian Horrocks, Bernardo Cuenca Grau and Boris Motik around a year ago. The spin-out is the culmination of a journey which began in 2009 with a piece of consultancy work that sowed the seed in Boris's mind to then pursue academic research in the area.

The consultancy work was for a Spanish travel start-up, ExperienceOn. The company wanted to revolutionise the way customers searched for travel industry information, and Boris advised them how to set up a logic-based knowledge system. It works by capturing background knowledge about a domain, and then inferring information implicit in the data and background knowledge.

'Consider, for example, if a customer wanted to search for accommodation at Easter near the Coliseum in Rome,' explains Boris. 'The system needs to know facts about where the Coliseum is, what the dates for Easter are, and where Rome is. For the system to be able to handle this, information has to be put in, stating for example that hotels and bed and breakfasts are sorts of accommodation, and a query on that has to be applied.'

Unlike other prominent branches of AI such as machine learning, this approach is deductive in that reasoning is performed using logical axioms – statements that explicitly describe what is true in an application domain. With advice from Boris's group, ExperienceOn built an ontology (description of tourism domains, covering items such as accommodation and Spain) using the OWL ontology language and the Resource Description Framework (RDF) to describe their data.

The travel consultancy work gave Boris the idea to embark on a five-year research project called MaSI^3 – a Massively Scalable Intelligent Information Infrastructure, which tackled the problem of scale in the sort of system used by ExperienceOn.

By developing scalable reasoning and query-answering techniques, the group developed an intelligent information system called RDFox, which was optimised for modern computers. It could be used, for example, by oil producers analysing streaming and sensor data to diagnose faults and prevent failures.

The potential of RDFox was demonstrated in a proof of concept pilot with Kaiser Permanente, an American healthcare provider that wished to show it met a set of federal quality measures. The project demonstrated that ontologies and RDF provide a flexible formal language that can capture criteria naturally and precisely, and that systems such as RDFox can apply criteria to healthcare records and analyse the data.

The pilot revealed high potential for industry uses, and, to ease the relationship with customers, the professors set up Oxford Semantic Technologies. 'I think it would be very difficult to run a spin-out and conduct research by yourself,' said Boris. 'The collaboration between the three of us made it so much easier. Everyone brings different strengths. The infrastructure offered by Oxford University Innovation was very useful too.'

One year on and Oxford Semantic Technologies has already built real systems for real customers using RDFox and is conducting sophisticated analysis of data for them. Meanwhile, Boris is also continuing with research on MaSI^3, which has received an ESPRC three-year fellowship extension grant called AnaLOG. The project's aim is to develop extensions of ontology languages that can make systems such as RDFox applicable to complex data analysis tasks. The consultancy and research complement each other, he believes.

'I like seeing the theory being brought into practice,' says Boris. 'It's good to put the components together, make a system, and see it working. I can then bring problems from industry back into my research and translate them into a scientific problem.'

[From left]: Professors Ian Horrocks, Boris Motik and Bernardo Cuenca Grau

Realistic data models for large-scale probabilistic knowledge bases

Systems that crawl the web encountering new sources and adding facts to their databases have a huge amount of potential uses. However, a lack of common-sense knowledge about their stored data is currently limiting their potential in practice. Oxford researchers are working to overcome these constraints, as Professor Thomas Lukasiewicz and co-investigator İsmail İlkan Ceylan explain.

Palaeontology, geology, medical genetics and human movement are domains for which large-scale probabilistic knowledge bases have already been built. There is an endless list of other potential applications for these systems that continuously crawl the web and extract structured information, and thus populate their databases with millions of entities and billions of tuples (structured sets of data). For example, such systems may also be used to enable digital assistants to answer natural language questions in healthcare, such as: 'What are the symptoms of appendicitis in adults?'

In recent years, there has been a strong interest in academia and industry in building these large-scale probabilistic knowledge bases from data in an automated way. This has resulted in a number of systems, such as Microsoft's Probase, Google's Knowledge Vault, and DeepDive (commercialised as Lattice Data and then bought by Apple).

Artificial intelligence research has also joined the quest to build large-scale knowledge bases. Fields such as information extraction, natural language processing (for example, question answering), relational and deep learning, knowledge representation and reasoning, and databases are all moving towards a common goal: the ability to query large-scale probabilistic knowledge bases.

However, these search and extraction systems are still not able to convey some of the valuable knowledge hidden in them to the end user, which seriously limits their potential applications in practice. These problems are rooted in the semantics of (tuple-independent) probabilistic databases, which are used for encoding most probabilistic knowledge bases. To achieve computational efficiency, probabilistic databases are typically based on strong, unrealistic completeness assumptions.

These assumptions not only lead to unwanted consequences, but also put probabilistic databases on a weak footing in terms of knowledge base learning, completion and querying. More specifically, each of the above systems encodes only a portion of the real world, and this description is necessarily incomplete.

'The team plans to produce a prototype implementation and experimentally evaluate the proposed algorithms.'

However, when it comes to querying, most of these systems employ the closed-world assumption, meaning that any fact that is not present in the database is assigned the probability '0', and thus assumed to be impossible. As a closely related problem, by also using the closed-domain assumption, all individuals are assumed to be known, and no new individual can exist. Furthermore, it is common practice to view every extracted fact as an independent Bernoulli variable, which means that any two facts are probabilistically independent.

For example, the fact that a person starred in a movie is often assumed to be independent from the fact that this person is an actor, which is in conflict with the fundamental nature of the knowledge domain. Furthermore, current probabilistic databases lack common-sense knowledge; such knowledge is useful in many reasoning tasks to deduce implicit consequences from data, and is often essential for querying large-scale probabilistic databases in an uncontrolled environment such as the web.

The main goal of our current EPSRC-funded research, which began in December 2017, is to enhance large-scale probabilistic databases (and so to unlock their full data modelling potential) by including more realistic data models, while preserving their computational properties. The team from the Department of Computer Science (Thomas as principal investigator with İsmail and Professors Georg Gottlob and Dan Olteanu as co-investigators) is planning to develop different semantics for the resulting probabilistic databases and analyse their computational properties and sources of intractability.

Over the three-and-a-half years of the project, we are also planning to design practical scalable query-answering algorithms for databases, especially algorithms based on knowledge compilation techniques. We will extend existing knowledge compilation approaches and elaborating new ones, based on tensor factorisation and neural-symbolic knowledge compilation. Once designed, the team plans to produce a prototype implementation and experimentally evaluate the proposed algorithms. These prototypes should help demonstrate the full potential of large-scale probabilistic knowledge bases as data models.

Hip surgery guide

With an ageing population, arthritis is an urgent research topic. The department's Medical Imaging Group led by Irina Voiculescu has been developing software to improve support for early diagnosis and intervention, and to print 3D models to guide surgeons when preparing for hip replacement surgery. An early diagnosis helps clinicians reduce the duration of pain suffered by patients as hip osteoarthritis progresses.

Although surgeons now routinely have access to wonderfully detailed MRI scans, a surprising amount of measurement still has to be made manually to identify and pinpoint the early stages of arthritis. That requires months of training on the job and it is so time-consuming that manual measurements are rarely used in clinic, surgeons relying instead on training and experience. This process could, however, be sped up by the OxMedIS (Oxford Medical Image Segmentation) software developed by Irina's group. The software provides accurate semi-automated measures which allow clinicians to make precise treatment plans for patients.

OxMedIS automates the process of selection and analysis of anatomical features, which allows for greater precision. Crucially, it takes new measurements which are not currently available in clinic. This gives a fuller picture of the state of anatomical features around the joint. The highly developed system converts scan data internally into a complex hierarchical data structure, which gets displayed to the user in an intuitive way allowing clinicians easily to select bones or other tissue.

Reinforcement learning techniques are used to identify the femur bone in the scan, and then this information is extrapolated to infer the presence of hip cartilage damage. OxMedIS then delivers detailed information on the extent of both the damage and its location within the joint. Even when damage is not yet visible, by reading the biochemical signals of the cartilage, the software can predict where in the hip joint arthritis may develop next, potentially allowing surgeons to bring forward the date of a hip replacement operation.

These features help clinicians make evidence-based decisions and develop patient-specific treatment plans, including the kind of surgical intervention to use, and

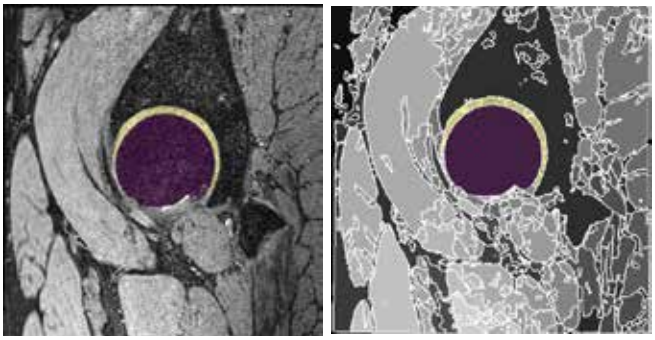

[Above]: Section through MRI hip joint scan with femur and cartilage labelled using OxMedIS

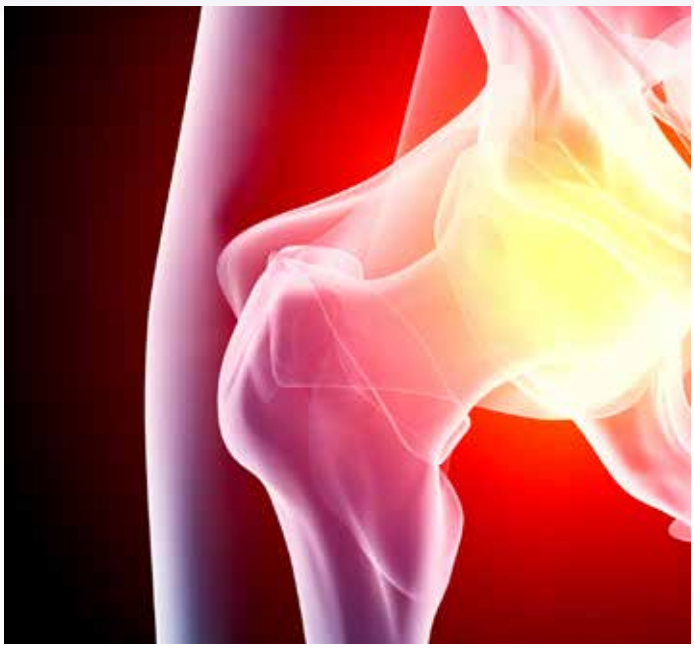

when best to operate. The decision to operate early sometimes removes the need to repeat costly scans and diagnosis as arthritis progresses. It also eradicates the need for long-term physiotherapy and medication, resulting in significant financial savings to the NHS. In real world terms, this also means that many patients are in pain for less time and lose fewer work days.

Irina's group developed the clinical features of the software in collaboration with Professor Siôn Glyn-Jones, who is a consultant orthopaedic surgeon at the Nuffield Orthopaedic Centre and conducts research for the Nuffield Department of Orthopaedics, Rheumatology and Musculoskeletal Sciences (NDORMS). Together they used data from previously treated patient cohorts to demonstrate the capabilities of OxMedIS as a clinical tool.

Until now the work has been funded by a John Fell Award and by the Department Impact Fund. The team is currently applying for funding to further automate the software to classify patients into different risk groups.

Irina says: 'No other automated product is currently available for this type of analysis. We believe this software will make a huge difference to the lives of arthritic hip patients, and the way they are managed.'

The OxMedIS project website is at: www.oxmedis.com

Why computer simulations should replace animal testing for heart drugs

Safety of new medicines is imperative before they are given to patients – which is why drugs are tested on millions of animals worldwide each year to detect possible risks and side effects. However, research shows computer simulations of the heart have the potential to improve drug development for patients and reduce the need for animal testing, as Elisa Passini, Patricia de Benito and Professor Blanca Rodriguez explain.

Animal testing has, to date, been the most accurate and reliable strategy for checking new drugs, but it is expensive, time consuming and – for some – highly controversial.

There is also the potential for some side effects to be missed due to the differences between animals and humans. Drug trials are particularly problematic for this reason and it's clear that new testing methods are needed to enable the development of better and safer medicines.

A variety of species of animals – including rats, mice, rabbits, guinea pigs, dogs and pigs – are used each year in drug development to predict the possible side effects for the heart in humans. While the underlying biology is similar, small differences between animal and human cells are amplified when a patient takes a drug. It means predicting the risk to patients is limited to an accuracy rate of around 75–85%, research shows, and it also leads to drug withdrawals from the market because of cardiovascular safety issues.

However, it's now possible to test a new heart drug in a 'virtual human'. The department's recent research (outlined in the paper 'Human In Silico Drug Trials Demonstrate Higher Accuracy than Animal Models in Predicting Clinical Pro–Arrhythmic Cardiotoxicity') demonstrates that computational models representing human heart cells show higher accuracy (89–96%) than for animals in predicting an adverse drug effect, such as dangerous arrhythmias – where the heart beat becomes irregular and can stop.

It shows that human computational models would bring additional advantages by reducing the use of animal experiments in early stages of drug testing. This would improve drug safety, thereby lowering the risk for patients during clinical trials, and speeding up the development of medicines for patients in urgent need of healthcare.

British biologist Denis Noble first began experimenting with computer models of the heart in Oxford in 1960. Since then, the technology has evolved and it is ready to be integrated into industrial and clinical settings.

'Human computer models are now available at different scales'

Thanks to human experimental data, human computer models are now available at different scales, from single cells to whole hearts, and they can be used to explore the behaviour of the human heart in healthy or diseased conditions, and under drug action.

Instead of a one-model-fits-all method, there are also new population-based approaches. Everyone is different, and some drugs can have harmful side effects only for certain parts of the population, such as people with a specific genetic mutation or disease.

The study by the Computational Cardiovascular Science team demonstrated that computer models of human heart cells are more accurate than animal experiments at predicting the drug-induced side effects for the heart in humans. This research won an international prize [see box on p21] because of its potential to replace animal testing in labs.

We incorporated the technology into software, named 'Virtual Assay', which is easy to use for non-experts in modelling and simulations. The software offers a simple user interface for Microsoft Windows in which a control population of healthy cardiac cells with specific properties, based on human data, can be built. It can then be used to run computer-simulated – known as 'in silico' – drug trials, before analysing the results.

continued on next page ▶

from previous page ▶

The whole process is very quick: it takes under five minutes using a modern laptop to test one drug in a population of 100 human cardiac cell models.

Several pharmaceutical companies are already using and evaluating Virtual Assay, which is available with a free academic licence and can be used by clinicians and pharmaceutical companies.

This research is part of a wider move towards the integration of computer models for drug safety testing which includes 'the Comprehensive in vitro Proarrhythmia Assay initiative', promoted by the US Food and Drug Administration and other organisations.

While simulations of heart cells can run in a few minutes, 3D computer models of the whole heart still require a huge amount of computational power. One heartbeat, for example, can take about three hours in a supercomputer with almost 1,000 processors.

We're now working on 3D simulations of the heart to explore drug cardiac safety and efficacy on a larger scale. It includes an exploration of diseased conditions, such as acute ischemia – where the blood flow in one of the arteries around the heart is obstructed. This research is also part of the European CompBioMed project to build computer models for the whole human body: a virtual human. By bringing together academia, the pharmaceutical industry and regulatory agencies we hope to accelerate the uptake of human-based in silico methodologies for the evaluation of cardiac drug safety and efficacy.

Computer simulations are a faster, cheaper and effective alternative to animal experiments – and they will soon play an important role in the early stages of drug development.

• This article first appeared in *The Conversation*.

Web links

The prize-winning research paper: goo.gl/W4rBPy
The project website: goo.gl/PHtX2B
Virtual Assay in the Oxford University Innovation Store: goo.gl/Ac7966
The NC3Rs website: www.nc3rs.org.uk

Photo [left]: The Computational Cardiovascular Science team testing the Virtual Assay software for in silico drug trials: (from left) Alfonso Bueno-Orovio, Professor Blanca Rodriguez, Elisa Passini and Patricia Benito.

Paper wins international 3Rs prize

The National Centre for the 3Rs (NC3Rs) – an organisation dedicated to replacing, refining and reducing the use of animals in research and testing – awarded its top prize to a research paper by a team from the Department of Computer Science.

The International 3Rs prize was awarded for research which developed original software which predicts the clinical risk of drug-induced side effects for the heart with higher accuracy than animal experiments.

The 3Rs prize, sponsored by GSK, was presented for the research paper 'Human In Silico Drug Trials Demonstrate Higher Accuracy than Animal Models in Predicting Clinical Pro–Arrhythmic Cardiotoxicity' by Elisa Passini and colleagues from the University of Oxford and Janssen Pharmaceutica at a ceremony in March. The prize consists of a £28,000 grant and a £2,000 personal award. The same paper also won the Safety Pharmacology Society Technological Innovation Award 2017.

The research was funded by the Wellcome Trust, Engineering and Physical Sciences Research Council, CompBioMed project (EU), TransQST Project (IMI) and the NC3Rs. It was conducted in Professor Blanca Rodriguez and Alfonso Bueno-Orovio's group. They are recipients of an NC3Rs Infrastructure for Impact grant to promote the profile of in silico human models for the 3Rs.

The winning research builds on work by the same team, which won the 3Rs Prize in 2014. Oliver Britton and colleagues then established a computer model that incorporated variations in 'normal' heart electrophysiological properties based on existing data from rabbit Purkinje fibres (cardiac cells). Oliver has subsequently been awarded an NC3Rs project grant to extend the same principles to pain research.

Automated verification of string-manipulating programs

Programs written in scripting languages with heavy use of string variables can easily lead to programming errors, with potentially serious security consequences. A new project will tackle these issues by developing algorithms to analyse how strings are manipulated in a program, and decision procedures to automate this reasoning. Professor Anthony W Lin explains.

Strings have been a fundamental data type in programming languages since the inception of Computer Science. This is true now more than ever with the past decade having witnessed the rapidly growing popularity of scripting languages (for example JavaScript, Python and PHP). JavaScript, for instance, was ranked as the most used programming language worldwide ever by back-end developers in a recent developer survey conducted by Stack Overflow.

Programs written in scripting languages tend to make heavy use of string variables, which are difficult to reason about and could easily lead to programming errors. In some cases, such mistakes could have serious security consequences. For example, in the case of client-side web applications, cross-site scripting (XSS) attacks could lead to a security breach by a malicious user.

The top ten classes of web application security vulnerabilities today include XSS, based on OWASP's (Open Web Application Security Project) most recent studies. These vulnerabilities are typically caused by improper handling of user inputs and submitted requests (in the form of strings, for example 'Donald Duck') by the web applications.

The goal of the project 'Algorithmic Verification of String Manipulating Programs' (for which I am the Principal Investigator) is to develop a constraint language for reasoning about how strings are manipulated in a program, along with decision procedures for automating this reasoning. For instance, the boxed example below requires us to have logic for reasoning about concatenation and finite-state transductions applied to strings. This is a challenging problem in both theory and practice since it requires us to solve some long-standing open problems about constraint solving over strings in the most general case.

The European Research Council is funding the project (November 2017 to October 2022), which aims to make scientific advances towards this in both theory and tool implementation, and apply the results to challenging real-life problems, including the analysis of XSS vulnerabilities in web applications.

Example of an XSS vulnerability

The following JavaScript code snippet is adapted from the paper 'Securing the tangled web' by Christopher Kern in 2014.

```
var x = goog.string.htmlEscape(cat);
var y = goog.string.escapeString(x);
catElem.innerHTML = '<button onclick= "createCatList
(\" + y + '\')">' + x + '</button>';
```

The code creates a DOM element catElem by assigning HTML markup. The value for the category cat is provided by an untrusted third party. The code attempts to first sanitise the value of cat. This is done via the Closure Library, string functions, htmlEscape and escapeString. Inputting the value 'Flora & Fauna' into cat gives the desired HTML markup:

```
<button onclick="createCatList('Flora &amp;amp; Fauna')
">Flora &amp;amp; Fauna</button>
```

On the other hand, inputting the value ');attackScript()
;// into cat, results in the HTML markup:

```
<button onclick="createCatList('&amp;#39;);
attackScript();//')">&amp;#39;);attackScript();//')</button>
```

When this is inserted into the DOM via innerHTML, an implicit browser transduction will take place, ie, first HTML-unescaping the value inside the onclick attribute and invoking the attacker's script attackScript() after createCatList. This subtle XSS bug (a type of mutation XSS) is due to calling the appropriate escaping functions in the wrong order.

Complex statistical methods benefit from probabilistic programming

New ways of understanding and building basic structures in statistics and machine learning is the research area of Professor Sam Staton, a Royal Society University Research Fellow. In this article he explains his work into understanding what a probabilistic program means as a mathematical structure.

Probabilistic programming is emerging as a way of building the kinds of complex statistical models that are used in computational statistics and machine learning. Businesses from Uber to YouGov are investing in it. In brief, a probabilistic program is a program that uses two additional programming features: one for making random choices, and one for recording observations about data. What makes it a statistical model is the idea that a probabilistic program will not simply be run once. Rather it will be treated as a model specification, and typically run thousands of times, so that all the random choices are considered: this is called a 'Monte Carlo' simulation.

For example, a simple 'Bayesian' inference problem involves updating a 'prior' belief according to the likelihood of some observation. This can be coded up as a two-line probabilistic program, in which the first line makes a random choice (incorporating the prior belief) and the second records an observation (incorporating the likelihood). We find the updated 'posterior' belief by running a Monte Carlo simulation.

The interest in this method of statistical modelling comes from the separation between the model (now a program) and the simulation method. Some simulation methods are very efficient and yet they can be applied to a wide variety of programs. This makes it easy to prototype different models without rewriting new simulation methods each time.

My own work in this area (with colleagues in Oxford and elsewhere) is based around understanding the meaning of probabilistic programs. One can always run programs and get results, but I'd argue that it is important to understand what a program means as a mathematical structure. This general point of view was first emphasised in Oxford 50 years ago by Dana Scott and Christopher Strachey. The conventional foundation for probability theory is called measure theory. It turns out that while the measure-theoretic interpretation of probabilistic programs is often quite elegant, there are some useful probabilistic programs that are taxing for this traditional foundation.

As an example, let's consider higher order random functions – random functions that themselves produce random functions. In conventional software engineering, higher order functions are an important way of building modular programs. Probabilistic programming systems, such as the Anglican system developed in Oxford, have no problem with higher-order functions in practice. However, the traditional measure-theoretic foundation of statistics cannot cope with higher-order functions: this was proved by Aumann in 1961. Our new direction in measure theory, based on 'quasi-Borel spaces', can cope with higher-order functions, and so it is a good foundation for probabilistic programming.

This is what I find exciting about my research: that structural analysis from logic and programming can suggest new ways of understanding and building basic structures in statistics and machine learning.

This article is based on the paper 'A Convenient Category for Higher-Order Probability Theory' by Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang: goo.gl/JACiVq

Sam is on the organising committee for The Federated Logic Conference (FLoC) 2018 – see p6 for more details.

Examples of regression problems

A function is something that gives an output for each input, and the problem of regression is to ask: 'given these inputs and outputs, what function generated them?' For example, given observations $f(0)\cong 0.6$, $f(1)\cong 0.7$, $f(2)\cong 1.2$, what is f likely to be?

The answer depends on our prior belief about f. The first graph shows f when our prior belief is that f is linear with unknown slope and base. In the second graph, f is piecewise linear. Both are generated by Monte-Carlo simulation from short Anglican programs. The beauty of probabilistic programming is that it is very easy to vary the model: the program for the second piecewise graph differs from the first only by inserting a higher-order random function 'piecewise'. This higher-order function doesn't exist in traditional measure theory, but it does exist in quasi-Borel spaces, our new foundation.

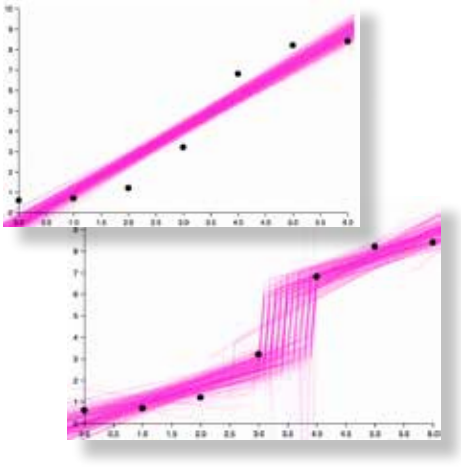

Will quantum computing make the world better?

Quantum computing has the potential to transform a whole host of technologies, from cryptography to artificial intelligence (AI). To help ensure these changes are positive for society, Oxford researchers are aiming to embed responsible research and innovation in quantum projects, as researcher Philip Inglesant explains.

New technologies raise new questions and new opportunities. What are the opportunities and risks? Who will benefit and who will lose out? Who decides? What are the alternatives? What if we're wrong?

Responsible research and innovation (RRI) – also known as responsible innovation – is an emerging set of practices to help ensure that the outcomes from research are beneficial and acceptable to society. This is a transparent, interactive process which encourages multiple stakeholders to work together throughout the research and innovation process to engage with ethical and social issues arising from a technology.

'Quantum methods also provide the potential for theoretically unbreakable, secure communications'

Networked Quantum Information Technologies (NQIT), one of the four Hubs in the UK National Quantum Technologies Programme, has the ambitious objective of building the core components of a scalable, universal quantum computer, placing the UK amongst the world leaders in this exciting new technology. RRI, led by Professor Marina Jirotka from Oxford's Department of Computer Science, has been embedded right from the start of this major public investment, as these technologies emerge from the laboratory to enable a new generation of products and services. NQIT is highly interdisciplinary; researchers in the Department of Computer Science are collaborating with Oxford researchers in Physics, Materials and Engineering as well as nine other universities and over 30 companies.

Some issues with quantum computing are already widely discussed. For example, it is predicted that powerful quantum algorithms may disrupt much of the public key cryptography scheme which underlies most internet security. Conversely, quantum methods also provide the potential for theoretically unbreakable, secure communications – which could protect not only high-value transactions and state secrets but might also provide a hiding place for terrorists and organised crime.

Other technologies that could also be transformed by quantum computers are machine learning and AI. These are technologies which are already having an impact through applications such as targeting of products, language processing, and, on the horizon, self-driving cars. Already, AI has implications for economic activity, for manufacturing and skilled labour as well as for knowledge work and highly expert professions such as law and medicine (see articles in issue 11 of *Inspired Research*). These raise important ethical issues, but also questions of their wider consequences for society, the innovation pathways they will open up or close down, legal and public policy choices, and governance to steward these technologies towards an equitable and sustainable future.

Powerful universal quantum computers are still some years in the future, but other forms of quantum computing, using other paradigms, are already in existence, and there is increasing interest in potential uses for quantum simulation and already achievable quantum applications. Quantum-related spin-off technologies will also make important research contributions and have significant industrial applications. All of these technologies, in various ways, will raise questions for RRI.

Marina is well placed to advise on this area. She has led an EPSRC project to develop a RRI framework for information and communications technology (ICT), and led on RRI for various European projects. She currently co-leads the development of the Observatory for Responsible Research and Innovation in ICT (ORBIT) – see p25.

Marina says: 'Although we can already start to see how quantum computing could change the world, there are many things that we cannot predict. But if we fail to consider these questions, then undesirable as well as beneficial outcomes may become embedded in society. RRI doesn't pretend to have all the answers, but it does provide a flexible and practical framework to anticipate what might happen, so that we are ready to respond to whatever challenges quantum computing brings.'

NQIT's website: nqit.ox.ac.uk
Orbit's website: www.orbit-rri.org
Results of EPSRC's public dialogue on quantum technologies: goo.gl/9cwJxN

Building a better IT future

In a world where technological advances can create undesirable futures, there is an ever stronger need for more research into the ethical and societal consequences of technology. Responsible research and innovation is a way of achieving more desirable futures, as outlined below by Margherita Nulli, ORBIT Project Officer.

Once a technological product is manufactured and out in the wild, it may often be used in ways that were not originally anticipated by the designers. For example, affective devices – robots that can interpret signals to infer a human's mood and respond accordingly – are mainly developed with beneficial purposes in mind, such as helping people to fight loneliness. Even though the original intention is good, what happens to all the data about its owner that the robot stores? Who has access to it and how can it be used? What would happen if people start treating the robot as a real human being?

If we want to build a better IT future we need to ask ourselves a key question: how can we maximise the benefits of a technology while minimising its risks? One of the most prominent approaches to addressing this question is responsible research and innovation (RRI).

RRI is a process that aims to ensure that the outcomes and processes of research are socially desirable, acceptable and sustainable. This highlights important questions such as: what are the benefits of the research? What are the risks? What are the alternatives and who is responsible for the outcomes of the research? RRI states that researchers should take a long-term perspective and consider what will happen to their products once they leave the laboratories. RRI can help researchers to obtain better outputs and minimise the side effects of their research by helping them anticipate the effect of their products, reflect on the purposes of the research, engage with stakeholders and act to influence the direction of their research. These aspects form the AREA 4P framework, which is the core of RRI in the UK. The framework consists of a set of questions that help the researcher to consider all aspects of their research and to evaluate if it is being done responsibly.

It has to be emphasised that RRI is not a silver bullet. RRI cannot help us predict the future or avoid all side effects arising from a technology. It can however help us to reflect on the strengths and weaknesses of a project (or a tool) and stimulate a better informed conversation about possible solutions.

The Observatory for Responsible Research and Innovation in ICT (ORBIT) Project

RRI is a concept embraced by a range of research funders including the EU and EPSRC. ORBIT (see box) was recently created with the support of the UK's biggest funder of ICT research, the EPSRC. ORBIT aims to help put RRI into practice by offering services that can integrate all of its aspects.

In the case of an affective device, ORBIT could, for example, help to identify resources, provide case studies, technology descriptions and other documents to learn from good practice. A self-assessment tool developed by ORBIT can help identify the strengths and weaknesses of a piece of work and guide researchers in their RRI activities. For example how do you decide whether an object could or should have a female or male voice? How can you envisage the consequences of a broad adoption of affective robots in homecare settings? ORBIT can help researchers find ways to answer such questions and gain competences through training and proposal development.

Key to the success of ORBIT will be the development of a community of interested scholars which will be supported by the newly launched ORBIT journal, a platinum open access venue for dissemination and discussion of the role of ICT in modern societies.

Orbit website: www.orbit-rri.org

ORBIT launch event at the Palace of Westminster

ORBIT, a joint venture between Oxford and De Montfort Universities, was launched at the House of Lords on 30 January 2018. ORBIT is led by Professors Marina Jirotka of Oxford University and Bernd Stahl of De Montfort [pictured above], with Martin de Heaver as Managing Director, Margherita Nulli as Project Manager and Carolyn Ten Holter as Marketing Officer.

Who's storing your conversations?

It may feel like chatting with Alexa and Siri is fairly innocuous, but it can be disconcerting to know that your conversations are stored somewhere beyond your control. A research team from the department is investigating how to make virtual assistants more respectful through a combination of technical and social techniques. DPhil student William Seymour explains.

There is a large body of human-computer interaction literature that describes the potential perils of treating a machine as if it was human. At the same time, people over the years have anthropomorphised everything from their kitchen appliances to their Roombas. While this behaviour is mostly harmless, new technology threatens to change that.

Virtual assistants such as Siri and Alexa offer us a means of automating simple and mundane tasks. But more so than the devices that came before them, they invite us to personify them. They have names, speak in human voices, and can, to a limited extent, converse with us and tell jokes. So what's the problem? This behaviour primes us to expect our technology to behave in a human-like manner.

'The data collected by your Echo belongs to Amazon and not you.'

Combined with distinctively human interaction methods (such as speech) we can start expecting our technology to behave in ways that would be obvious to a human, but are simply beyond the current capabilities of machines.

In addition to this, the relationships that we have with these devices are now built on a very different footing than they were in the past. While Cortana might tell you that she likes you, the legal agreement that exists between you and the company that created your virtual assistant's software does not contain a clause obliging them to like or respect you (and when was the last time you read any terms and conditions?) Normally this is not something we notice, but when the divide collapses it can lead to creepy behaviour, and a feeling akin to betrayal. In a 2017 murder trial in Arkansas, recordings from an Alexa owned by the accused were turned over to the prosecution as evidence. This is not something that one expects from a gadget that is primarily used to play music. We expect to enjoy privacy in our own homes, but this is ultimately incompatible with the reality that the data collected by your Echo belongs to Amazon, and not you.

Is the solution to start reading privacy policies? An analysis of 184,000 privacy policies reveals that the reading comprehension needed to understand the average privacy policy is roughly the same as that required to read an article in *Nature*. Simplifying policies might help, but only if they are actually read. In order to find a longer term solution, Professor Max Van Kleek, Reuben Binns and I have been asking what a respectful virtual assistant might look like, as part of the EPSRC-funded Respectful Things In Private Spaces (ReTIPS) project (goo.gl/VkDyoW).

Perhaps a respectful assistant stores and processes the data it collects locally, so that it never leaves your home, or offers different voice recognition options so that users can trade off the accuracy of cloud-based models against the privacy of local ones. We also think that the most respectful devices will incorporate tools which help users make their other devices more respectful if they are unable to do it themselves.

Along the same lines, we are also putting together an 'honest' version of the Amazon Echo. When using a modified Alexa, users are informed when they make a request how that information is going to be used: where it goes, if it's used to profile them, and which third parties will have access to it. By doing this we hope to examine how people use this information to make decisions about the devices in their homes, and particularly how they change their behaviours after interacting with our speculative lab prototypes.

So while virtual assistants themselves are not inherently problematic, they do bring to the fore many of the problems that have built up around our use of centralised data-driven services, and especially so for home Internet of Things devices. We believe that using the notion of respect when thinking about smart devices, while not solving all the problems identified above, will help make the smart home of the future a better place to live.

What do Angry Birds know about your children?

Kids Online Anonymity & Lifelong Autonomy

When evaluating games for children to play on mobile devices, parents tend to focus on content, rather than information gathered. Yet these apps track a notable amount of information, prompting an Oxford team to work on increasing parents' and children's awareness of privacy issues, as Senior Researcher Jun Zhao explains.

Parents are working hard to protect their children's online safety. However, although risks related to social media platforms or social video sharing sites (such as YouTube) are widely known, risks posed by mobile applications or games (ie apps such as Angry Birds) are less known. Behind the cute characters played, apps used by primary age children not only have the possibility of exposing them to age-inappropriate content or excessive in-app promotions, but may also make a large amount of their personal information and online behaviour accessible to third parties in the online marketing and advertising industry.

Such practices are not unique to children's apps, but young children are probably less capable of resisting the resulting personalised advertisements and game promotions. Currently there are no effective ways to stop these tracking behaviours on mobile devices or notify parents and children of these risks. This is a timely challenge given that tablet computers are widely used by young children for both education and entertainment purposes.

Our initial research in our project, Kids Online Anonymity & Lifelong Autonomy (KOALA), showed that one in four apps designed for children can be linked to more than ten different tracking companies. We are using data privacy visualisation tools developed by the Oxford team in KOALA's forerunner, the five-year EPSRC Programme Grant SOCIAM, to raise awareness in schools, and amongst young children and their parents about the risks associated with the use of mobile devices.

'One in four apps designed for children can be linked to more than ten different tracking companies'

KOALA's next step will focus on developing our data privacy visualisation tool into something that is more suitable for the cognitive ability of young children, and help parents to use the tool to mediate their children's use and choice of mobile apps. The team will test the impact of the tool with parents and children during the project, which is funded by Oxford University's EPSRC Impact Acceleration Account and will run until March 2019. It is led by Professor Sir Nigel Shadbolt and myself from the Department of Computer Science.

By partnering with the Anna Freud National Centre for Children and Families, we also aim to explore the impact of these personal data tracking practices of mobile apps upon the general well-being of young children aged six to ten. Interactions with parents and young children have shown so far that parents lack support when choosing safe and appropriate digital content for their young children, and they struggle to manage the stressful moment of stopping their children using risky or inappropriate apps.

At the same time, parents largely believe that their children are too young to understand or discuss privacy risks, and instead they often seek technical restrictions or monitoring to safeguard their children. However, children are already facing the risks of excessive in-app promotions and losing control of their personal information online. We have therefore outlined concrete steps that parents can take – see box below.

The General Data Protection Regulation (GDPR), which came into place on 25 May 2018, requires children's data to have specific additional protection. It demands more transparency of organisations when accessing and using children's personal information, particularly for profiling or marketing purposes.

KOALA is one such project that is endeavouring to increase this transparency for parents and young children especially.

More information about the KOALA project: goo.gl/dpaQ22

Top tips for SHARP mobile apps choice

- Select: Select apps carefully by consulting resources such as Common Sense Media.
- Help: talk to your children about asking you for help when they really need it.
- Avoid: avoid providing any sensitive personal information to the apps.
- Rating: check 'Age Rating' of the apps.
- Privacy Permissions: check 'Privacy Permissions' of the apps.

Facebook collaboration to improve secure group messaging

A team from Oxford's Department of Computer Science led by Professor Cas Cremers has been working with Facebook software engineer Jon Millican to improve the security of group messaging.

As part of the research project, the team designed a mechanism called Asynchronous Ratcheting Tree (ART), a protocol for end-to-end encrypted group messaging. It aims to improve on existing systems' security guarantees for scalable group conversations, which are often weaker than those they provide for one-to-one messages.

'Modern security applications already offer strong security for two-party communications, but these often diminish once more parties are involved. The existing techniques for strong security don't scale to the group sizes required by WhatsApp or Facebook. In our paper we combine techniques from group messaging with strong modern security guarantees,' said Cas.

'The resulting ART protocol offers the advantages of point-to-point security while still being usable on mobile devices, which are often offline. An important point is that it can re-establish a secure connection even if there is a security breach on one device, a property which we call "post-compromise security".'

Facebook has made a proof-of-concept implementation of the new protocol publicly available via GitHub. The researchers hope that the protocol can in future be used to enhance the security of apps such as WhatsApp, Signal, Facebook Messenger, Google Allo and Wire.

Oxford DPhil students Katriel Cohn-Gordon, Luke Garratt and Kevin Milner worked on the research with Cas. The paper is at: ia.cr/2017/666