InspiredResearch

Summer 2017 Issue 10

News from the Department of Computer Science, University of Oxford

60 YEARS OF SUCCESS!

We celebrate our anniversary – p6

LipNet:
Lip-reading AI – p12

The Open Data Institute:
Sir Nigel Shadbolt writes about open data – p18

Mathematical Oncology:
Lifting the lid on research in mathematical oncology – p20

UNIVERSITY OF OXFORD

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

Ivo Sluganovic, Ivan Martinovic, Yannis Assael, Sir Nigel Shadbolt, Ursula Martin, Marta Kwiatkowska, Polina Mamoshina, Bill Roscoe, Niel de Beaudrap, Julian Hedges, Artem Kaznatcheev, Rahul Santhanam, Emma Dunlop, Piers O'Hanlon, Michael Wooldridge, Misha Denil, Leanne Carveth, Helena Webb.

Letter from the Head of Department

By the start of Trinity term, there is always a palpable sense that the academic year is ebbing away. The bulk of lectures and tutorials are behind us, with only the hurdle of exams to jump before the vista of the long vacation opens up before us. Trinity term is, naturally enough, a time for reflection, and I am pleased to report that this academic year feels like it has been a very fruitful one for the department.

While we have had much to celebrate over the past year, surely the biggest news of all is that Oxford's Sir Tim Berners-Lee was announced as the recipient of the ACM A.M. Turing Award. The ACM A.M. Turing Award is the ultimate recognition that can be given to a computer scientist. Initiated in 1966, the award is given annually to an individual or group that has made fundamental contributions that have been of lasting significance to the whole field of computing. With $1m prize money, the award is often referred to as the 'Nobel prize for computing', and prior to Sir Tim, just two Oxford professors had received the ACM A.M. Turing Award: Dana Scott in 1976, and Sir Tony Hoare in 1980. And it is more than a quarter of a century since a member of a British university won the award. Tim's achievement, of course, speaks for itself. He gave us the World Wide Web, and in doing so, he changed our world more profoundly than perhaps any other person alive today. I am simply delighted that Tim received the award - I can think of no other person more deserving of it.

It is perhaps hard to compete with the ultimate glory of a ACM A.M. Turing Award, but over the past year, many other of our staff have received honours of various kinds. Marta Kwiatkowska was elected a Fellow of the ACM (FACM); Ursula Martin was made a Fellow of the Royal Society of Edinburgh (FRSE); and Georg Gottlob received the Lovelace Medal, the most prestigious award given by the British Computer Society. And the department itself received honours, being ranked first for Computer Science in the UK in the Complete University Guide, and also first in the UK in the Times Higher Education Supplement's first ever ranking of UK Computer Science departments. And this is not to mention the many best prizes, awards, and distinctions that you will read about in this edition of Inspired Research.

Events have quickly unfolded with respect to our desire to expand undergraduate teaching. We are in the process of appointing tutorial fellows for three colleges that are admitting our undergraduates for the first time. Christ Church, Hertford, and Jesus colleges will all have their first tutorial fellow in Computer Science in October, and we are in discussions with a number of other colleges.

I will sign off with a reminder that this year is the 60th anniversary of the founding of the Department of Computer Science (or Computing Laboratory, as it was originally). We are planning a series of events to mark the occasion, with a celebration in September - we hope to see you there!

Professor Michael Wooldridge
June 2017

Professor Sir Tim Berners-Lee awarded 'Nobel prize for Computing'

The Association for Computing Machinery has named Sir Tim Berners-Lee, a Professor at Massachusetts Institute of Technology and the University of Oxford, as the recipient of the 2016 ACM A.M. Turing Award. Sir Tim was cited for inventing the World Wide Web, the first web browser, and the fundamental protocols and algorithms allowing the Web to scale. Considered one of the most influential computing innovations in history, the World Wide Web is the primary tool used by billions of people every day to communicate, access information, engage in commerce, and perform many other important activities.

The award is named for Alan M. Turing, the British mathematician who articulated the mathematical foundation and limits of computing. ACM is marking '50 Years of the ACM A.M. Turing Award' with special events this year. This celebration will culminate in June 2017 in San Francisco, CA with an ACM conference on the future of computing featuring 20 Turing laureates and other distinguished members of the computing field. The conference will be followed by the ACM Awards Banquet, at which Sir Tim will formally receive the 2016 ACM A.M. Turing Award.

'The first-ever World Wide Web site went online in 1991,' said ACM President Vicki L. Hanson. 'Although this doesn't seem that long ago, it is hard to imagine the world before Sir Tim Berners-Lee's invention. In many ways, the colossal impact of the World Wide Web is obvious. Many people, however, may not fully appreciate the underlying technical contributions that make the Web possible. Sir Tim Berners-Lee not only developed the key components, such as URLs and web browsers that allow us to use the Web, but offered a coherent vision of how each of these elements would work together as part of an integrated whole.'

'The ACM A.M. Turing Award is the greatest academic honour that can be bestowed upon a computer scientist', said Professor Mike Wooldridge, Head of Department of Computer Science at the University of Oxford. 'It is entirely fitting that Sir Tim should receive it now: I can think of no other living individual who has changed our modern world more profoundly. Sir Tim is the third University of Oxford professor to receive the ACM A.M. Turing Award' said Professor Michael Wooldridge, Head of the Department of Computer Science at the University of Oxford, 'we are immensely proud of this illustrious heritage.'

Professor Sir Nigel Shadbolt, commented, 'It has been my enormous privilege to work with Tim over the years and witness at first hand his single-minded dedication to ensuring that the web he invented serves all of humanity and not just the few. The ACM A.M. Turing Award is the Nobel prize for Computer Science and it is a fantastic recognition of Tim's achievements.'

Sir Tim is a graduate of Oxford University, where he received a first-class Bachelor of Arts degree in Physics. He is the 3Com Founders Professor of Engineering in the School of Engineering with a joint appointment in the Department of Electrical Engineering and Computer Science and the Computer Science and Artificial Intelligence Laboratory (CSAIL) at Massachusetts Institute of Technology (MIT), where he also heads the Decentralized Information

Group (DIG). He is also a Fellow at Christ Church and a Professorial Research Fellow at the Department of Computer Science, University of Oxford.

Sir Tim founded the World Wide Web Consortium (W3C) in 1994, where he continues to serve as Director. W3C is an international community that develops open standards to ensure the interoperability and long-term growth of the Web. In 2009, he established the World Wide Web Foundation, which works to advance the Open Web as a public good and a basic human right. He is the President of the Open Data Institute (ODI) in London.

News in brief

Oxford is number one according to the 2018 Computer Science subject league table published by The Complete University Guide. Universities were judged on entry standards, student satisfaction, research quality and graduate prospects. Our overall score is a perfect 100%!
Read more: goo.gl/eeiD65

NEWS

News in brief

In Spring 2017, Professor Marta Kwiatkowska took over as Deputy Head of Department of Computer Science, with responsibility for research. Marta's research is concerned with modelling and analysis methods for complex systems, and spans the whole spectrum, from theory and algorithms, to software implementation and applications. Marta has also just been elected as a Fellow of the European Association for Theoretical Computer Science (EATCS). Marta's fellowship was awarded for her pioneering work in the development of model checking for quantitative systems and in its application to a wide range of areas, as well as for outstanding mentorship, and serving as a role model for female researchers in Computer Science. Read more here: goo.gl/S5y1OS

Professor Ursula Martin has been elected a 2017 Fellow of the Royal Society of Edinburgh, one of 59 distinguished individuals across the arts, business, science, technology, and academia joining the 1600 members of Scotland's national academy.

Dr Maria Bruna has won the 2016 Women of the Future Award for Science, for her work on stochastic models of interacting particles and its application to industry. Maria is a member of the 2020 Science programme for the Mathematical and Computational Modelling of Complex Natural Systems, where her research seeks to explain how individual-level mechanisms give rise to population-level behaviour in biology and ecology. Website: awards.womenofthefuture.co.uk

Georg Gottlob honoured with Lovelace Medal

Professor Georg Gottlob has been selected by the Awards Panel of British Computer Society, The Chartered Institute for IT as the 2017 winner of the Lovelace Medal.

Georg's research has transformed our understanding of database systems, and in particular, the logical foundations of databases and the computational complexity of associated reasoning problems. He has developed deep theoretical results of lasting significance, and developed techniques that have huge value in domains such as web data extraction. He has won a string of awards and distinctions for his research, most recently being elected FRS (Fellow of the Royal Society) in 2010. His research brings together techniques from logic and model theory, complexity theory, database systems, and knowledge representation/AI, and the fact that he is regarded as a leader in each of these domains is an indication of the significance and influence of his work

The BCS Lovelace Medal is the top award in computing in the UK. The award is presented annually to individuals who, in the opinion of BCS, have made a significant contribution to the advancement of Information Systems. The BCS Lovelace Medal was established in 1998 in honour of Lady Augusta Ada Byron, Countess of Lovelace and daughter of Lord Byron. She was born in 1815.

Previous winners include World Wide Web inventor Sir Tim Berners-Lee, Linux creator Linus Torvalds, ARM designer Steve Furber, information retrieval pioneer Karen Spärck Jones, and Doug Engelbart who developed the computer mouse and the modern style of computer interface.

On learning of this award Georg commented, 'The Lovelace Medal came as a great surprise to me, and I feel deeply honoured, in particular considering the previous winners. I see this as an award not only for myself, but for the Database and Information Systems research community in the UK, which has gained significant momentum in the recent decade.'

Read more: www.ox.ac.uk or academy.bcs.org

Federated Logic Conference coming to Oxford

Professors Marta Kwiatkowska and Daniel Kroening will be co-chairing the Federated Logic Conference (FLoC'18), in Oxford in July 2018, to be held in the Mathematical Institute and the Blavatnik School of Government. This is an extremely large event that only occurs once every four years, and encompasses nine major logic-related computer science conferences.

Keynote speakers will include Byron Cook, Shafi Goldwasser, Georges Gonthier, and Peter O'Hearn. There will also be an exciting programme of public events, including a lecture by Stuart Russell at the Sheldonian and an Oxford Union style debate on ethics for robotics.

For more information you can visit floc2018.org

Andrew Martin appointed as Bletchley Park Trustee

The Bletchley Park Trust has appointed Professor Andrew Martin to its board of Trustees. Andrew is Professor of Systems Security, Director of the Oxford Centre for Doctoral Training in Cyber Security, and Deputy Director of the Software Engineering programme. He was instrumental in setting up the University's Cyber Security Network and helps to lead it, heading Oxford's EPSRC/GCHQ-recognised Academic Centre of Excellence in Cyber Security Research.

'The story of Bletchley Park is crucial to the development of our academic discipline, as well as to the national life,' said Andrew. 'It is a great privilege to be involved in this way, and through my connection with Kellogg College, I hope we can bring that story to an even wider audience.'

Professors Ivan Martinovic and Cas Cremers scoop MPLS Impact Awards

We are delighted that Professors Ivan Martinovic and Cas Cremers have been recognised with MPLS Impact Awards. These Impact Awards aim to foster and raise awareness of impact by rewarding it at a local level.

Ivan won a 'category 1' award that recognises 'research that has had substantial impact since 1 August 2013'.

Cas was given a 'category 2' award, which recognises 'excellence in generating broad user interactions that achieved impact in the past year'. He was nominated for his recent work on the Transport Layer Security Protocol (TLS), which has led to significant improvements in the next generation of Internet security.

New Year honours

We are immensely proud that two of our alumni were recognised in the New Year Honours list. Following his Olympics success Paul Bennett was made an MBE for Services to Rowing, while Anne-Marie Imafidon [pictured left] was made an MBE for Services to Young Women within STEM Sectors.

News in brief

Each year Computer Weekly compiles a list of the top 50 most influential people in UK IT. This year our department was well represented, with Professor Sir Tim Berners-Lee and Professor Jim Davies both included. Read more: goo.gl/a5EeaJ

Congratulations to our two winners in the 2017 Oxford University Students' Union Teaching Awards. Dr Quentin Miller won the Mathematical, Physical and Life Sciences 'Outstanding Tutor' award and David Hobbs, Programme Administrator for the CDT in Cyber Security, won in the 'Best Support Staff' category.

Many congratulations to our former Head of Department Professor Sir Tony Hoare, who has been elected as a foreign associate of the National Academy of Sciences (NAS). He is one of 21 new foreign associates recognised for their distinguished and continuing achievements in original research. goo.gl/jZZuWU

Congratulations to Samson Abramsky, Hanno Nickau, and Luke Ong (together with their co-authors) for the 2017 Alonzo Church Award. The award is given by ACM SIGLOG for 'Outstanding Contributions to Logic and Computation'. This year the award has been given jointly to Samson Abramsky, Radha Jagadeesan, Pasquale Malacaria, Martin Hyland, Luke Ong, and Hanno Nickau 'for providing a fully-abstract semantics for higher-order computation through the introduction of game models, thereby fundamentally revolutionising the field of programming language semantics, and for the applied impact of these models.'

Department celebrates its 60th Anniversary

2017 marks the 60th anniversary of the Department of Computer Science (formerly the Computing Laboratory) at the University of Oxford. This year we will be celebrating our success and our achievements through a series of events and activities for members of the department, alumni and industry partners. Throughout the year we will be holding themed seminars in honour of our anniversary. Three highlights of our anniversary year will be:

In June we will be holding our 'Research Showcase' event in the department where research groups demonstrate their work to visitors through short talks, posters and demonstrations. This event has historically attracted people from the department, the university and industry to explore the exciting cutting-edge research being performed here. There will also be demonstrations from spinout companies who are now commercialising research outputs and showing the benefits of the work to society and industry. Through this event we hope to communicate details of the innovation we are fostering and explore potential partnerships with new organisations.

Over the summer we will be working with the Knowledge Exchange and Impact team to develop a publication on the achievements of the department over the last 60 years, from the foundational techniques developed by Sir Tony Hoare and Professor Christopher Strachey, to the applied work resulting in industrial benefits. This will be a great collection, showcasing how our department has changed the landscape for Computer Science over the decades and the benefits to society and technological advancement.

In September we will be hosting a 1950s themed ball at the fabulous Oxford Town Hall in celebration of our anniversary. This ticketed event will be open to members of the department and alumni and promises to be a lively evening. We will have a Jive/Rock/Swing band so we expect a night of jiving and hip swinging. Smart dress from the era, 50's style decorations and special lighting will add to the ambience.

To find out more, visit: www.cs.ox.ac.uk/60ball

Concurrency, Security, and Puzzles: A Festschrift for Professor Bill Roscoe

In honour of Professor Bill Roscoe's 60th birthday, in January 2017, members of the department organised a two-day symposium, bringing together collaborators, colleagues and friends from around the world to celebrate.

Bill is one of the world's leading minds in the theory of concurrency and is most notable for developing large portions of the Communicating Sequential Processes (CSP) theory, leading the development of the CSP-based verification Failures Divergences Refinement (FDR) tool and making major contributions in the field of computer security. The Festschrift celebrated his work with keynote speakers Sir Tony Hoare and Professor Stephen Brookes, and an international selection of speakers drawn from Bill's past and present collaborators.

In an after-dinner speech, Bernard Sufrin (Emeritus Fellow) said of Bill 'His mathematical intrepidity and engineering inventiveness have helped transform the science of computing as well as the craft of programming. It would be hard to count the numbers of teachers, engineers, and programmers his work has influenced – directly, transitively, or by osmosis.'

Fun facts

As it's our anniversary year we thought we would ask some searching questions to find out more about our department. This is the result!

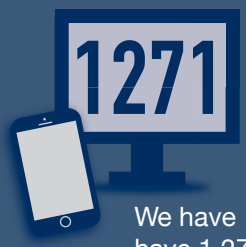

We have have 1,271 computers (including phones and tablets).

Average number of emails per week processed through the CS Departmental mail server: 62,700

There are 7000 books in the Department library.

The number of post-its it takes to create Mario on the wall...

There are 12 grotesques in our buildings.

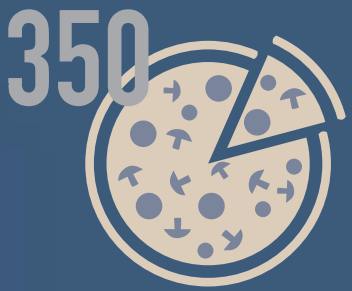

CompSoc ate 350 pizzas and spent £3,600 at Dominoes in a year!

The longest serving member of support staff is Julie Sheppard, who joined us on 2 July 1984!

Someone gets locked out at night and has to be 'rescued' by security: once a month.

We drink 140 kilos of instant and 2,500 pouches of ground filter coffee in a year.

We have created a map [below] that shows every nation represented here at Computer Science. It's in response to the news that Oxford is the 6th most International University in the world (according to *Times Higher Education*).



Women in computing at Oxford, from the pioneers to now

Walk into Professor Blanca Rodriguez's office and you'll see a striking image of the physiology of the heart. She's part of a long Oxford tradition of using computation to unravel complex systems: Dorothy Hodgkin, winner of the 1964 Nobel Prize for chemistry, was one of Oxford's computing pioneers, and her beautiful computer-drawn images of the structure of insulin now have pride of place in the Bodleian Library.

Campaigning by Dorothy and others led to the founding of the Oxford University Computing Laboratory in 1957, supporting research and teaching in numerical analysis, and a computing service. In the early days many women were employed as operators and programmers. In 1960 Joan Walsh was awarded the very first DPhil in the Computing Laboratory, and in 1970, as a Professor in Manchester, set up (with Linda Hayes and Shirley Lill), the group that became the highly successful software company NAG. Susan Hockey, later a Professor at University College London, was a pioneer of what is now called digital humanities, studying computer representation of non-Western languages.

The Masters programme in Computation started in 1979: then, as now, Oxford was an international melting pot, and Leonor Barocca recalls how, excluded from education in Portugal by political events, taking the course in 1982 shaped the rest of her life. Another overseas student recalled that, when she came top of the class in her undergraduate exams, entitling her to a scholarship to study in Europe, it was initially offered to the second placed (male) student. When she protested, funding was swiftly found for both of them!

Early students did pioneering work. MSc student Hanan Mohamed wrote an expert system for intricate Sudanese inheritance law. Mary Sheeran, now a Professor in Sweden, developed new programming languages for computer hardware in her DPhil, and was one of the first academics and college fellows hired to teach on the undergraduate Computer Science programme, launched in 1984.

A major benefactor of internet research in Oxford, Dame Stephanie Shirley, founded a highly successful software company which employed mostly women, but in general, the numbers of women in computing declined in the later twentieth century. It was not until 2007 that Oxford appointed a female Statutory Professor of Computer Science, Marta Kwiatkowska. Recently elected a Fellow of the ACM, Marta now acts as mentor and passionate advocate for OxWoCS, our energetic group supporting the next generation of women in computing. These days our department is home to eight female professors and many more women at other levels within the Faculty.

Go girls!

The department is always keen to encourage more young women to think of Computer Science as a career – and study – option.

On 13 February we welcomed 16 year 7 girls, and their mums to the department for our first 'InspireHer!' event. InspireHer! aims to increase girls' interest and curiosity in Computer Science by introducing coding in a fun and interactive way. On the day girls, their mums (and one brave Dad) got stuck in to an exciting range of activities and challenges: an introduction to Sphero SPRK+ programming, learning to program a Sphero robot, and visits to Oxford colleges, Somerville and St Anne's.

InspireHer! was supported by Google, with DPhil student Bushra Alahmadi and fellow members of the OxWoCS (Oxford Women in Computer Science) society.

On 14 February 48 Guides from across Oxfordshire took part in a busy programme of talks and activities that covered multiple sciences including Computer Science, Maths, Physics, Material Sciences and Earth Sciences.

The girls, ranging in age from 10-16, took part in a computer programming practical and went on college tours of St John's and Somerville. Taking part in new experiences and building new skills are key areas of the Guide Programme. One Guide commented, 'I didn't know anything, and thought that it would be really boring – you have changed my mind and now I know that it is exceedingly fun.'

Coming up this summer, a series of three multi-science taster days for young women based around the themes of technology, health and energy & environment. Plus Oxford's annual Women in Computer Science Day. More information here: goo.gl/bhqdpF

Bebras competition

Two weekends in January and February once again saw a group of talented school and college students from across the UK descend upon the department to take part in a final challenge as part of the Bebras Computational Thinking Challenge. Finalists attended lunch and a prize-giving ceremony at Hertford College. The standard was remarkably high, with several of the top scorers achieving full marks. In the Senior Finals Alex Darby narrowly beat his twin Stephen to 2nd place.

4th Oxbridge Women in Computer Science conference

On 16 March 2017 Oxford hosted the 4th Oxbridge Women in Computer Science Conference, an annual event organised jointly by OxWoCS (Oxford Women in Computer Science) and its Cambridge counterpart Women@CL.

The conference is a great opportunity for Oxbridge women at all stages of their academic career to network and collaborate.

This year, the busy programme included eight presentations, ten posters, a career panel and an interview workshop. The talks included a keynote from Professor Blanca Rodriguez and a talk on Android Development by Robin Bennett, from Conference sponsor Google.

News in brief

Andreea Marzoca was the joint winner of the 2017 Worshipful Company of Information Technologists Student Award. The awards recognise outstanding undergraduate and postgraduate IT students within the UK, and were created in 2015 by The Worshipful Company of Information Technologists Charity (WCIT Charity). Criteria for the award included academic excellence, overcoming adversity, entrepreneurial skills and contribution to charity or community. Andreea, and joint winner Joanna Joss (of Brunel University, London) are the first female winners of this award.

Oxford students take first place at Deloitte CTF final

Left to right: Chris, Alastair, Ilias, Cecile, Emma.

Having placed third in the Deloitte CTF (Capture the Flag) qualifier, the Ox002147 team composed of Chris Bellchambers (3rd year Maths & Computer Science), Cecile Berillon (MSc), Emma Espinosa (3rd year Computer Science), Alastair Janse van Rensburg and Ilias Giechaskiel (both of the Centre for Doctoral Training in Cyber Security) headed off to the final round at the Deloitte offices in London.

The competition featured 17 security-related challenges, and included problems in cryptography, forensics, web application exploitation, binary reverse engineering, and some physical security challenges such as lock picking and hardware hacking.

The competition was very closely fought – towards the end, Edinburgh and Oxford were tied in points, with Edinburgh in the lead. But in a dramatic turn of events, a final flag submission four minutes before the end gave Oxford first place!

Programming competition success

The Northwestern Europe Regional Contest (NWERC) of the ACM International Collegiate Programming Contests took place at the University of Bath in November. Oxford entered three teams, two of which were placed in the top five, out of more than 100 competitors.

A gold medal, and third place overall was won by team 'Catz CS Society (1)'. Team 'Curious Coding' also won a gold medal, achieving fourth place in the competition.

NWERC 2016 is an official regional contest in the ACM International Collegiate Programming Contest. The goal of each team is to solve as many algorithmic problems as possible within the 5-hour time limit. Potential solutions are submitted then scored by an automated judging system. The regional competition draws students from colleges and universities throughout Belgium, Luxembourg, Great Britain, Ireland, Iceland, Norway, Sweden, Finland, Denmark, Germany and the Netherlands.

The competition scoreboard is available here: goo.gl/pPWfSs

The top three teams from this round of the completion – Oxford's 'Catz CS Society (1)', alongside the University of Helsinki's 'Game of Nolife', and 'Omogen Heap', the team from Sweden's KTH – Royal Institute of Technology – advanced to the ACM ICPC world finals.

The winning teams consisted of:

Catz CS Society (1):

* Hristo Venev (1st year undergraduate Computer Science, St Catherine's College)

* Mihail Jianu (1st year undergraduate Computer Science, St Catherine's College)

* Sauyon Jason Lee (2nd year undergraduate Computer Science, St Catherine's College)

Curious Coding:

* Alexandru Valeanu (2nd year undergraduate Computer Science, Oriel College)

* Gabriel-Robert Inelus (2nd year undergraduate Computer Science, Keble College)

* Ragnar Groot Koerkamp (Masters student, Mansfield College)

Meltwater
acquires Wrapidity

Meltwater, a global leader in media intelligence, has announced the acquisition of Wrapidity Limited, an artificial intelligence (AI) spinout from this department which uses innovative technology to automate data extraction from content on the web. Wrapidity's underlying technology will help Meltwater to improve and scale its industry-leading analytics software and address the volume of data coming from new sources as required for future product generations.

Wrapidity helps businesses get the right data from the web, turning billions of webpages into a highly-structured database of relevant data and enabling better search, recommendations and analytics for that business. The company was launched by Oxford University Innovation, the research commercialisation company of Oxford University, in 2015.

Department of Computer Science Professor Georg Gottlob, Co-Founder of Wrapidity, said: 'Instant access to products, places, people and news has changed our lives in the last decade. The same access, but at a much larger scale, is now changing business in ways we can't even imagine yet. At Wrapidity, we have responded to this by developing a completely new AI-based technology for extracting massive amounts of relevant data from websites.'

Tim Furche, Co-Founder and Chief Technology Officer added: 'Meltwater already monitors and analyses millions of articles per day across several languages. Combining Meltwater's industry leadership and global footprint with Wrapidity's advances in AI technology, we will be able to surface more accurate, timely and insightful content for Meltwater's customers.'

Zoubin Ghahramani Strachey Lecture

On 7 March we welcomed Professor Zoubin Ghahramani to deliver a Strachey Lecture on 'Probabilistic machine learning: foundations and frontiers'.

Probabilistic modelling provides a mathematical framework for understanding what learning is, and has therefore emerged as one of the principal approaches for designing computer algorithms that learn from data acquired through experience. In his lecture, Professor Ghahramani reviewed the foundations of this field, from basics to Bayesian nonparametric models and scalable inference. He talked through some current areas of research at the frontiers of machine learning, leading up to topics such as probabilistic programming, Bayesian optimisation, the rational allocation of computational resources, and the Automatic Statistician.

Zoubin Ghahramani is Professor of Information Engineering at the University of Cambridge, Co-Director of Uber AI Labs, and the Cambridge Director of the Alan Turing Institute.

The Strachey Lectures are generously supported by OxFORD Asset Management. Watch here: www.cs.ox.ac.uk/mediawall

The world outside

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

Peter Turner, a facilities support technician, attended the Pickleball championship in Amsterdam in November. Peter and his wife played a tough doubles match but came away finishing in fourth place. Pickleball is a racket sport played on a badminton court, which combines various net sports and supports two teams of one, two or three per match.

Every year, BBC Radio 3 challenges amateur composers to create a piece of music. This year, the challenge was to create music for a modern version of a medieval English poem. Ian Collier's entry wasn't shortlisted, however he was invited to conduct his piece at the local St Frideswide's Church, which he did on 22 December 2016. Ian works as a Computing Officer within the IT team.

New book co-authored by Bob Coecke

The unique features of the quantum world are explained, entirely through diagrams, in 'Picturing Quantum Process: A First Course in Quantum Theory and Diagrammatic Reasoning', by Bob Coecke (and co-author Aleks Kissinger).' Requiring only basic mathematical literacy, this entirely diagrammatic presentation of quantum theory represents the culmination of ten years of research.

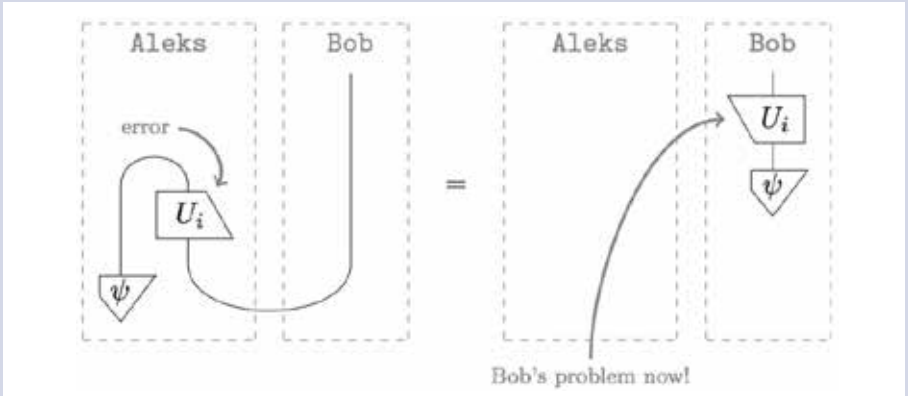

LipNet: The first end-to-end sentence level lip-reading model

Doctoral Student Yannis Assael writes about the important work that he, Brendan Shillingford, and the rest of the LipNet team have been doing.

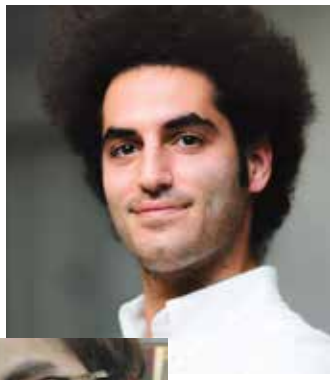

[Above: Yannis Assael, Left: Brendan Shillingford]

Lipreading plays a crucial role in human communication and speech understanding, as highlighted by the McGurk effect, where one phoneme's audio dubbed on top of a video of someone speaking a different phoneme results in a third phoneme being perceived. A phoneme is the smallest unit of sound in speech.

Lipreading is a notoriously difficult task for humans, especially in the absence of context (goo.gl/u3DDJ3). Consequently, human lipreading performance is poor. Hearing-impaired people achieve an accuracy of only 17% even for a limited subset of 30 monosyllabic words and 21% for 30 compound words, based on previous literature. An important goal, therefore, is to automate lipreading.

Furthermore, machine lipreaders have enormous practical potential, with applications in improved hearing aids, speech recognition in noisy environments eg cars (goo.gl/jvb0s4), silent dictation in public spaces, biometric identification, and silent-movie processing.

LipNet performs lip-reading at the sentence-level using machine learning. It is the first model to perform lip-reading at the sentence level. More specifically, it is able to process a sequence of video frames using neural networks, then map to a sentence.

At the heart of LipNet lies a spatio-temporal convolutional neural network architecture, able to take as input a sequence of frames and output a sequence of characters of the sentence spoken.

continued on next page ▶

[Above] Evaluation of LipNet on the GRID corpus.

News in brief

In February Professor Daniel Kroening published 'AI Will Save Us (Or At Least Correct Our Mistakes)' in the Huffington Post. He dispelled some myths about the potential negative impact on our lives and said that, 'I would argue that the most dangerous area in modern technology is not software that's too clever, but rather software that isn't clever enough.' Read the full article here: goo.gl/U8hGwK

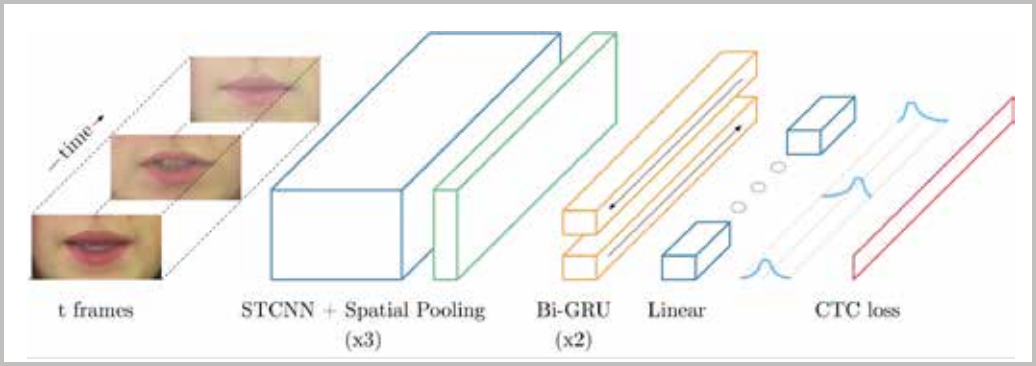

[Left] At the heart of LipNet lies a spatio-temporal convolutional neural network architecture, able to take as input a sequence of frames and output a sequence of characters of the sentence spoken.

from previous page ▶

We compare the performance of LipNet (93%), against human lipreading experts (52%), in the largest publicly available sentence-level lipreading dataset called the GRID corpus. The GRID corpus has a fixed grammar and limited vocabulary, which is why LipNet performs so well there. However, there are no signs that LipNet wouldn't perform well when trained on larger amounts of real-world data. In collaboration with NVIDIA we also tested the performance of LipNet for potential usage in their autonomous vehicles. A demonstration of that can be found here (goo.gl/j1QQzy).

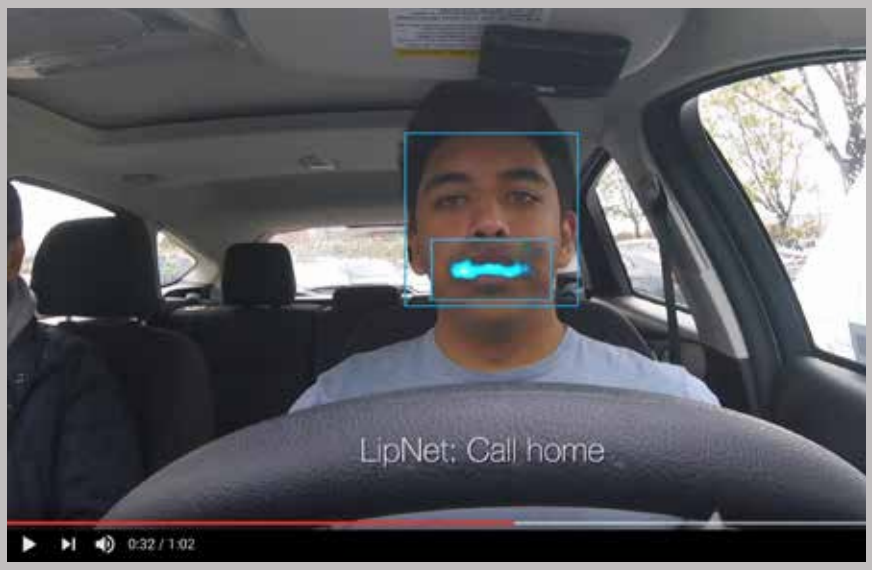

[Above] Evaluation of LipNet in collaboration with the NVIDIA autonomous vehicle team.

What does LipNet see?

One of the most interesting parts of this research was the analysis of the learned representations of LipNet from a phonological perspective. To do that we apply saliency visualisation techniques to interpret LipNet's learned behaviour, showing that the model attends to phonologically important regions in the video. In particular, in the figure above we analyse two saliency visualisations for the words 'please' and 'lay' from a speaker of the GRID corpus. Using the phonetics expertise and guidance of Mitko Sabev, DPhil in General Linguistics and Comparative Philology, University of Oxford, we saw that LipNet learned by itself to pay attention to the crucial parts of the mouth that a lipreading person would look at.

The LipNet team consists of Yannis Assael, Brendan Shillingford, Shimon Whiteson and Nando de Freitas; a major part of the model was developed during Yannis' summer internship, and was finalised in Oxford in November 2016.

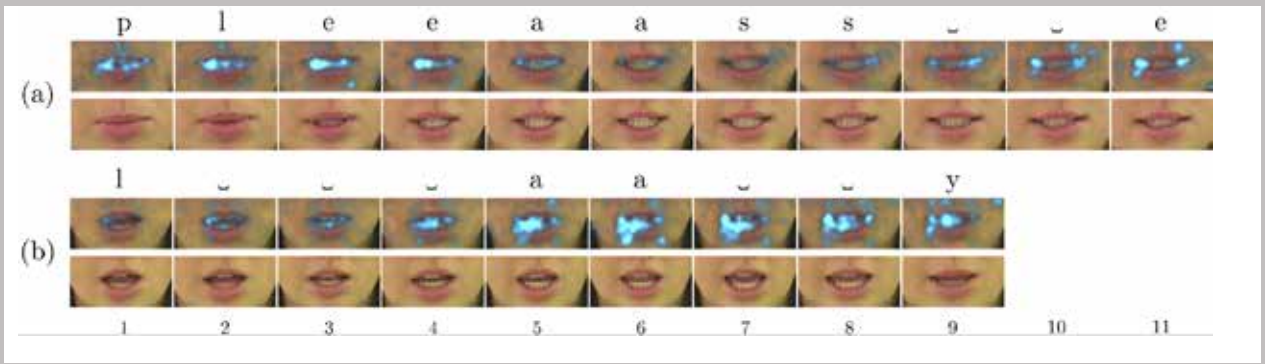

[Above] Saliency maps for the words (a) please and (b) lay, produced by backpropagation to the input, showing the places where LipNet has learned to attend.

Alumni Profile

Amelia Gould: a journey from Computer Science to Naval Engineering

Amelia used her degree in Engineering & Computer Science as the basis for a career in the Royal Navy which eventually led to her current role as head of engineering on BAE Systems' Naval Ships Combat Systems team.

What course did you study here and when?
I studied an undergraduate degree in Engineering & Computer Science between 1996 - 2000.

What was your background before that (where you come from, which school, other University or college etc prior to Oxford)?
I lived in East London. I went to my local comprehensive to do my GCSEs, and was then awarded an Assisted Place to a private school for my A-levels (City of London School for Girls).

What attracted you to studying Computer Science as a subject?
I originally applied to study Engineering Science and in my first year that is what I studied. Sadly in that time I discovered that no matter how hard I tried I could not understand Mechanics. My tutor suggested that I dropped the Mechanics part of the course and as I was good at Maths and logic I took up Computer Science instead. I had never programmed a computer before, and did not know much about the course, but I was willing to give it a go. It is a decision I have never regretted!

I loved the course, and though I found some bits really hard the guys on my course were great and helped me when I needed it (I was the only girl in my year doing Computer Science). I wish I had been exposed to Computer Science earlier as I would have definitely chosen it straight away.

What aspects of the course you studied here did you particularly enjoy?
My project, (which was related to 3G Mobile Phones) and the open way the course was taught. It was taught very well, and I am very proud to have been awarded a First class degree, one of only two that year in Engineering & Computer Science.

What did you do when you left Oxford?
I joined the Royal Navy. They had sponsored me through university, so once I graduated I joined properly and served for a further seven years.

How has the course you studied here helped you in your current profession?
The course taught me how to think, and how to approach problems logically. I learnt the principles of programming and Computer Science, which meant that I could interpret what may be going wrong when I was working as a Test & Trials Engineer. I can also now understand how algorithms are constructed and a system architecture when it is being described by my software engineers.

What advice would you give to current students on applying their knowledge in the workplace, when they leave university?
Try to keep an open mind about the career options available to you with your new-found knowledge. Computer Science is about much more than programming, I have worked on other parts of the engineering lifecycle – as a Business Analyst writing requirements and as a Tests & Trials Engineer breaking software for a living.

If we went back in time and asked, what would the student you have thought about what you are currently doing – would you have been surprised, proud, amazed?
All three! I never thought I would end up being Head of Engineering of a cutting-edge business, but I am very proud to have got here and amazed by the great people and interesting things I did along the way.

Alumnus Dr Alastair Donaldson awarded the BCS Roger Needham Award 2017

The British Computer Society award (sponsored by Microsoft Research Cambridge) has been made in recognition of outstanding work in the area of many-core programming. Alastair (now at the Department of Computing at Imperial College London), has made a distinguished contribution through his design and application of rigorous program analysis methods to the emerging field of many-core programming. His techniques and case studies have also made a major contribution to fundamental Computer Science research.

The Roger Needham Award was established in memory of Microsoft's first director of research outside the US. It is awarded for a distinguished research contribution in Computer Science by a UK based researcher within ten years of their PhD.

Happy Birthday ERC

In March 2017 the European Research Council (ERC) celebrated its 10th anniversary. We marked the occasion in the Department of Computer Science with a drinks reception recognising the projects and academics who have received ERC funding.

The ERC is a flagship component of Horizon 2020, the European Union's programme for research and innovation. It was set up by the EU in 2007 to fund excellent scientists and their most creative ideas and to help European institutions attract and retain the best researchers of any nationality. Since its inception, more than 7,000 projects have been funded and as each ERC grantee employs on average six team members this represents a huge boost to research training, particularly for doctoral students and postdoctoral researchers. The total budget allocated to the ERC for the period 2014-2020 is €13.1 billion (17% of the overall H2020 budget of €77 billion).

Over the past decade, the Department of Computer Science has received over £16 million (€19m) from the ERC. A total of 14 projects, with an average award size of approximately £1.14m, have been funded. 12 of our academics have received ERC awards and their projects represent much of the core research conducted in the department.

At the departmental celebration held on 8 March Professor Michael Wooldridge congratulated recipients on their success and acknowledged the importance ERC funding continues to play in the department's research portfolio.

Starting Grants (for researchers with 2-7 years' experience post PhD, up to €1.5 million per grant) have been awarded to Edith Elkind, Daniel Kroening, Shimon Whiteson and Stanislav Zivny.

Consolidator Grants (for researchers with 7-12 years' post PhD, up to €2 million) have gone to Professors Dan Olteanu, Joel Ouaknine and Rahul Santhanam.

Five academics have received Advanced Grants (for established researchers, up to €2.5 million) which is testament to the excellent advanced research taking place in the department. Recipients are Professors Leslie Goldberg, Georg Gottlob, Elias Koutsoupias, Marta Kwiatkowska, and Michael Wooldridge. Marta Kwiatkowska and Georg Gottlob have also been awarded Proof of Concept grants providing top-up funding to bridge the gap between research and marketable innovation.

Faster generation of prime numbers

by Professor Rahul Santhanam

A large part of mathematics is driven by the desire to better understand prime numbers. Prime numbers are important not just in number theory but also in more applied areas such as cryptography.

A fundamental result known as the Prime Number Theorem tells us that prime numbers abound – for large enough N, about a $1/\log(N)$ fraction of integers between 1 and N are prime. However, the abundance of primes does not seem to help us with efficiently computing a fixed prime number of a given length, which can be important in applications. The fastest known algorithm to generate a prime that is n bits long was only guaranteed to halt in time $2^{n/2}$, which grows astronomically with n. Somehow, primes manage to be both abundant and elusive!

In recent work with Igor Carboni Oliveira (currently a postdoc at Oxford), to be presented at the Symposium on Theory of Computing (STOC) conference this June in Montreal, I significantly improve the state of the art, giving a randomised algorithm that halts in time subexponential in n for infinitely many n. This suggests that a polynomial-time algorithm might be within reach, but more remains to be done.

News in brief

Associate Professor Anthony Lin is the recipient of a Google Faculty Research Award, for his research on Cascading Style Sheet (CSS) Minification via Graph Optimisation and Constraint Solving. Read more here: goo.gl/rUUoOJ

Professor Shimon Whiteson wins Google Research Award

Shimon Whiteson has been announced as a recipient of a Google Research Award. Google puts out an annual open call for proposals on Computer Science and related topics including machine learning, machine perception, natural language processing, and security. The grants cover tuition for a graduate student and provide both faculty and students the opportunity to work directly with Google researchers and engineers.

For this round of funding Google received 876 proposals covering 44 countries and over 300 universities. Shimon's proposal on, 'End to End Deep Model Based Reinforcement Learning' was one of only 143 successful proposals.

Shimon explains, 'Deep reinforcement learning has had some exciting successes in the past few years, but so far all those successes have been "model-free" ie they work without building an explicit model of the world, even though such a model has large potential advantages. We have a hypothesis as to why model-based approaches to deep reinforcement learning are failing, and in this research we will build a new method that, if our hypothesis is correct, will enable model-based approaches to deep reinforcement learning to succeed.'

Julian Hedges awarded EPSRC funding for game-theoretic modelling research

Jules Hedges is the principle investigator of the 3-year EPSRC-funded project 'A compositional approach to game-theoretic economic modelling'. The project aims to apply mathematical techniques from theoretical computer science to economic theory, with a special emphasis on scalable reasoning. This is based on recently-discovered similarities between game theory and the algebraic structures that underlie diagrammatic reasoning, that have been previously applied in quantum computing and linguistics by Oxford's Quantum group. This makes it possible to draw formal but intuitive pictures of economic processes, and reason about them graphically. Another surprising but deep connection is to 'lenses', an interest of Oxford's Algebra of Programming group. Lenses, also known as functional references, are a powerful way of looking at, constructing, and using functions on complex data types.

Besides understanding the hidden mathematical structures underlying economics, the aims of the project include making the theory practical by developing prototype software tools for economists. Just as modern programming languages aid collaboration between groups of programmers on large software projects, an 'economic programming language' would enable economists to work together on a previously unimagined scale. Another aim is to investigate problem areas – such as the economics of smart energy grids, where sheer scale makes existing theory impractical to apply. Collaborators on the project include Neil Ghani (University of Strathclyde), Philipp Zahn (University of St. Gallen) and Viktor Winschel (OICOS GmbH).

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

- Dr Jason Nurse explains his research in cyber security, and the risks surrounding security and privacy in cyberspace with the use of digital technologies. The research is focused specifically on the risks of social media and information divulged unknowingly via emails and the devices we use. You can find his talk here: goo.gl/T9g3iR

- DPhil student Maaike Zwart has produced a video based on the linguistic research of the Quantum Group. It gives a brief introduction on the basic techniques of how quantum theory can help understand natural language. goo.gl/VrOIkU

- OxfordSparks have produced a podcast asking the question of whether we can protect buildings from earthquakes. They spoke to the department's Sensor Networks group to see how best to achieve this. goo.gl/aHQ0Hw

These videos, and many others, are available on the department's Media Wall: www.cs.ox.ac.uk/mediawall

Chao Charity Mbogo: alumna in Africa

At the Department of Computer Science, we try to keep track of our alumni and we often hear of success stories as our graduates apply what they have learnt here in the 'real world'. Chao Charity Mbogo did an MSc in Computer Science at Oxford before going on to complete a PhD at the University of Cape Town.

She writes of her time at Oxford, 'One lost application form, a new form and a gruelling scholarship interview later, I was admitted for an MSc in Computer Science at Oxford's St Catherine's College.

I was on cloud nine for several months, but then reality hit: during orientation and the first weeks of class, I had to learn UNIX from scratch, complete practical lab assignments within short periods, and adapt to a faster and more dynamic learning process than I was accustomed to.

My interactions with other friends attending Oxford and similar institutions, who had completed undergraduate degrees in Kenya, revealed that most of us had to work twice as hard to bring ourselves on par with our classmates.'

Chao now works as a Researcher and Lecturer in the department of Computer Science at Kenya Methodist University. She (like many other women around the world) is very much aware of the lack of women pursuing studies and careers in Computer Science. In common with some current students, she is involved in initiatives to encourage more women to get involved in computing.

You can read an article she wrote for *The Conversation*, 'Africa needs more women computer scientists. How to make it happen' here: goo.gl/KI7MSD

AI child's play

Doctoral student Misha Denil has been part of a team of researchers teaching Artificial Intelligence (AI) to play with blocks, just as a baby does – learning about the physical world through experimentation.

Working at Google DeepMind, and also with researchers at the University of California, Misha used reinforcement learning techniques to enable the AI to learn using experiments within two virtual environments.

The first experiment used five blocks that looked the same but weighed a different amount. Just as a baby would, the AI had to 'poke' each block to identify which was the heaviest. The AI had to interact with all the blocks to come to the right conclusion.

In the second experiment five blocks were piled into a tower. Some blocks were stuck together and others were single. The AI had to pull the blocks in order to identify which ones were stuck together.

The AI was able to learn how to solve the tasks without having any prior knowledge of the laws of physics, or of the physical properties of the blocks used in the experiments.

Although this research is at an early stage, it has important implications, particularly for robot development. It proves that AI has the potential to learn how to solve problems when clear instructions are not available, acquiring an understanding of the world that exceeds passive perception.

Read more here: goo.gl/B3SHH7

The Open Data Institute: Unlocking the power and promise of open data

Professor Sir Nigel Shadbolt is co-founder and Chairman of the Open Data Institute – here he writes about the work of the Institute, and the importance of open data in general.

Five years ago I co-founded the Open Data Institute (ODI) alongside Sir Tim Berners-Lee. The founding ambition was to realise the value latent in open data. Open data is data which anyone can access, use and share for whatever purpose. Since 2009 Tim and I have been working with successive UK governments to open up non-personal public sector data for wider reuse.

All governments collect, generate and commission huge amounts of data on behalf of their citizens. Everything from how taxes are spent, to data about the weather, transport, health outcomes, reported crime, companies, agriculture, education, the list goes on. Since we began our work tens of thousands of data sets have been released by governments both in the UK and around the world. The UK is now widely regarded as internationally leading in its open data work.

Data exists of all types, not just open data, but personal data about an individual, sensitive corporate or public sector data. There is a whole spectrum of types of data whose interactions need to be understood. The ODI now embraces the wider challenge of finding the value, both economic and societal, in this mass of data. Societies around the world are amassing data at an unprecedented rate, but the awareness of how to make best use of it, and the skills to do just that, haven't kept pace. At the ODI, we're doing something about that.

So the ODI's wider mission is to connect, equip and inspire people around the world to innovate with data. We do this in a variety of ways, from helping businesses and government to shape data and digital policy; to training thousands of people around the world in basic data skills; to supporting startups and entrepreneurs to make game changing products and services from data.

Since its launch in 2012, the organisation has developed a global team to help innovate around data in areas from banking and transport, healthcare, art, sport, and many others.

Reaching out around the globe

The ODI has reached millions of people through its international network of 28 franchised ODI 'Nodes' across 18 countries. In addition to our International work, the organisation has developed a network of regional centres throughout the UK, including Aberdeen, Belfast, Birmingham, Cardiff, and Leeds, to provide training and networking support.

The ODI has incubated over 60 startups, connecting organisations and governments with young companies who provide innovative solutions to complex problems using open data. New ideas are crucial to transforming businesses, and from the beginning, the ODI startup programme has been key in forging these vital connections.

Building the UK's data infrastructure

Much of the ODI's work encourages and supports the creation of a strong data infrastructure in the UK. We are used to thinking of areas like transport and energy requiring physical infrastructure. From roads and rail networks to the national grid and power stations, we understand that investment and management of these vital parts of an infrastructure is essential to our economic well-being and future prosperity.

This is no less true of key data assets. A data infrastructure consists of data assets and the technology to store, share and use them; the organisations that operate and maintain them; the processes by which they are maintained; and guides describing how to use and manage them. A trustworthy data infrastructure should be sustainably funded and managed so it can be used in ways that benefit all of society.

We have some excellent examples of infrastructure data from the likes of Companies House, Land Registry, Ordnance Survey and Defra, but core parts of the data

continued on next page ▶

from previous page ▶

infrastructure that we need within the UK are missing, unreliable, or of a low quality. The ODI encourages the government to invest here just as it invests in our other traditional infrastructures.

Supporting data learning and culture change

Improving data literacy is another core area of focus for the ODI. In order to take advantage of the data revolution, policymakers, businesses and citizens need to understand how to use data. In other words, they must become data literate. The ODI generates content and courses to build data literacy amongst a wide range of groups and have so far delivered training to over 6,600 people from around the world.

Creating an environment where openly sharing ideas and collaborating is encouraged, is as much about the culture of an organisation, as it is about the data. The ODI has developed a culture change framework and open data readiness assessment based on research on organisational change and insights from civil servants leading change in governments around the world, from Mexico to Macedonia, Kenya to Chile.

Data opportunities in transport

Data innovation is transforming all parts of industry. For example, the ODI has recently created a report alongside Deloitte for the Transport Systems Catapult (TSC), to identify how government can incentivise data sharing in the UK 'Intelligent Mobility' sector. Currently, fears around cyber security, lack of data literacy skills and a legacy of viewing transport modes such as rail and road in isolation are restricting the free flow of information. Overcoming these barriers could unlock £14bn of benefits from new innovations by 2025.

Technologies such as driverless cars, journey planning apps and smart ticketing are all identified as opportunities which can be fully exploited when data is published and shared as much as possible, while respecting privacy.

The report shows that investment in data could lead to faster journeys, lower emissions, improved regional connections and opportunities for job creation in an emerging technology sector – without the need for massive infrastructure building projects.

Data and culture

Open data isn't just a powerful tool for public services and the private sector, but for culture too. The Open Inheritance Art project, funded by the ODI, invited the public to discover and visit private collections of Britain's most important cultural heritage artworks in a bid to increase access to art. An online platform was built to enable the public to easily find the location of the artworks and arrange viewings with the owners.

The artworks are made available under the HMRC Conditional Exemption Incentive scheme. It aims to keep national assets of outstanding beauty, historical or scientific interest in the UK. Open Inheritance Art improved HMRC's data, by allowing people to flag missing or outdated information. Visitors can photograph and upload digital copies of the works onto the website. In many cases, these photos are the first digital representations of the works, displayed on the website and uploaded to WikiMedia, a freely usable media source under open licences.

Changing the banking industry with data

In 2015, the ODI was invited to help set up the initial Open Banking Working Group (OBWG), which called for an open banking standard and the introduction of open banking APIs, enabling services to be built using open and shared data from banks and customers, based on an open standard.

The work and outputs of the OBWG were largely adopted by the Competition & Markets Authority (CMA) as they launched Open Banking in the UK in September 2016. I am proud that the ODI played such a key role in this important regulatory shift. The ODI continues to be involved in Open Banking conferences, reports and government consultations. The organisation has been advising governments, treasuries and banking associations in New Zealand, Australia, USA, Canada as they make their own move towards open banking.

Conclusion

Open data, available to use and reuse, will fuel the public and private services of the future. Businesses will flourish or struggle depending on their ability to realise the value inherent in their data by improving data literacy skills, and allowing others, as well as themselves, to innovate with their data. Every organisation should consider which of their data assets might best serve them and others as open data.

Artem Kaznatcheev lifts the lid on research in mathematical oncology

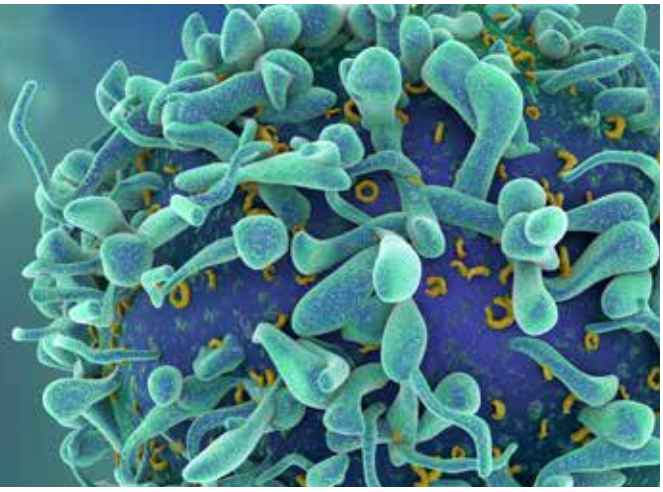

It was a classic hack-a-thon, except for one key difference: my team wasn't just the usual mathematicians, programmers, computer and physical scientists. A central part of it were biologists and clinicians specializing in blood cancers. And we weren't prototyping a new app, we were trying to predict the risk of relapse after stopping Imatinib for patients with chronic myeloid leukemia. This was 2013 and I was at the 3rd annual integrated mathematical oncology workshop. It was one of my first exposures to using mathematical and computational tools to study cancer — the field of mathematical oncology.

Most are familiar with the use of mathematical (and computational) models in the physical sciences but over the last few decades, mathematical modelling has started to play a bigger and bigger role in medicine, including oncology.

It is always hard to find firsts, but a popular candidate for the first mathematical model in oncology is Armitage & Doll's 1954 multi-stage theory of cancer. They modeled the age-distribution of cancer incidence and developed a theory that the prevalence of cancer will increase with a power of age that is one less than the number of changes needed for its progression.

But the biggest push for mathematical models of cancer started in the 1970s. Greenspan introduced mechanistic models of growth by modeling tumour spheroids. These masses of pure tumour tissue, although not overly realistic — the lack of blood vessels in these avascular tumours made many fundamental questions inaccessible — were consistent and reliable in the lab and allowed for a closer coupling of theory and experiment. Around the same time, looking instead to the clinic, Wheldon et al. introduced the linear-quadratic model for cell death under radiotherapy — a model that is now used to calibrate treatment protocols.

In the late 80s, the first mathematical models of angiogenesis — the recruitment and co-opting of blood-vessels by tumours — were proposed by Balding & McElwain. These were still models based on differential equations and continuum approximations of large numbers of cells. But unlike avascular tumours, the experimental systems for studying angiogenesis were much less consistent and much more chaotic in their dynamics. Whereas avascular tumours looked relatively similar to each other the vascularized tumours could look completely different from one to the next. And this heterogeneity between tumours was central to clinically relevant difference like the rate at which a tumour could invade other tissues (metastasize). It is one of the reasons why your doctor might be more worried about a new mole with irregular boundaries than a circular one.

To address this challenge, the community moved towards explicitly stochastic and discrete agent-based models in the 1990s. These models allowed for easier communication with experimental biologists and parametrising from existing measurements in the literature. But they also forced mathematical oncologists to rely heavily on simulations and parameter sweeps. The many, often implicit, modeling choices that researchers could now make (researcher degrees of freedom) also allowed for conceptually different models to have very similar observable dynamics. This made the task of model selection very important and encouraged mathematicians and computer scientists to work more closely with biologists and clinicians.

By the early 2000s, ARA Anderson and colleagues introduced multi-scale hybrid models that combined discrete cells with continuous diffusion of chemical species like oxygen or growth-factors. To parametrise, these models needed to combine data collected under varying conditions and in different experiments. Breaking through silos, reconciling data, and serving as connectors between different areas of biology became a more integral parts of the modelers' work. This was made possible in part by mathematicians moving to cancer centers and research institutes but also by hybrid-modelers who wear both the hats of clinical doctors and mathematicians.

The integration continues today: as a community we are able to quickly absorb mathematical and computational tools and techniques

continued on next page ▶

from previous page ▶

from other disciplines. Through this, mathematical oncologists are helping to budge the ingrained paradigm that cancer is a cell-autonomous process: that the defining properties of a cancer are just properties of individual cells (like their genotypes). By bringing in tools from ecology and evolutionary biology, we can start to view tumours as heterogeneous evolving ecosystems where the cell's microenvironment can play a determining role in disease progression. For example, David Basanta and colleagues have used evolutionary game theory to reason about the strategies that cancer cells adopt for propagation and survival.

But the biggest contribution of mathematical oncology isn't any particular model or technique. Our biggest contribution is the dialogue with experimentalists. Theorists help motivate new experiments and hypotheses, propose new measurements. Experimentalists give theorists an endless supply of subtle and complex problems, to challenge our tools and understanding. Mathematicians and computer scientists can clarify and challenge the paradigms in cancer research, and cancer researchers can give us challenges and motivations for new kinds of mathematics and refinements of existing tools.

The field is young — and most experimental biologists are too uncomfortable with equations — to allow us to rely exclusively on journals for communication. Instead, we flourish when we are brought together into integrated teams that combine experiments, clinical care, and modeling. For this, hack-athons like the Integrated Mathematical Oncology Workshop and inter-disciplinary groups are an invaluable resource. And although a week-long hack-a-thon or even dozens of carefully-calibrated and validated models won't cure cancer, the discussions started and contributed to by mathematicians and computer scientists contribute to our progress against this malady.

Group Design Practical presentation day

The Group Design Practical, runs from January, and sees teams of 5-6 undergraduate students battling it out with their chosen project. Many of the challenges are set or sponsored by industry partners. The students' work culminates in an exhibition and formal presentation, which this year was held in the department on 14 May.

The judging panel consisted of one each from sponsors - GResearch, Metaswitch, Palantir and Winton Capital. Each industry partner awarded the teams a prize based upon their own criteria.

Teams chose a challenge from a list of 14 different project briefs proposed by industry partners and colleagues from within the department. Project topics are presented in the form of an outline design brief. Part of the work is to undertake a proper requirements analysis for the chosen project, working with the project mentor.

Winning teams:

Team 8 was awarded the GResearch Prize for their Musical Machines project. The challenge was to develop a system that could generate music. The team consisted of Ruofeng Yuan, Szymon Kowara, Nicholas Sale, Michael Wang, Jo Lee and Robert Mitchell, with Departmental Lecturer Dr Joe Pitt-Francis as their mentor. The project was in collaboration with Winton Capital.

Team 2 won the Metaswitch Prize with their Proof of Location project. Their task was to build a prototype time-location verification system using low-cost hardware such as Raspberry Pis to prove that someone was in a certain place at a certain time. Their team consisted of Samuel Hindmarch, Jesse Sigal, Andrew Kenyon-Roberts, Ioana Vasile, Sauyon Lee and Catherine Vlasov. Dr Reuben Binns was their mentor.

Team 5 was awarded the Palantir Prize for their project 'AI Racing Market'. This challenge was set in collaboration with GResearch who challenged the team to develop a competitive market in which members of the public could submit algorithms to an interactive game to see which is the best. Team members were Aleksandar Monev, Michael Agius, Janpreet Khabra, Paul Horvath-Bojan, Calin Tataru and Daniel Mroz. The group were mentored by Dr Peter Minary.

The final prize, sponsored by Winton Capital, was awarded to Team 7 for their 'Virtual Reality Finance' project. The project asked the team to implement visualisations that would allow users to immerse themselves in real-world financial datasets, allowing them to explore and interact with the data from within. The project was sponsored by Oxford Asset Management and the team was mentored by DPhil student Ivo Sluganovic.

All second year undergraduates reading Computer Science, Computer Science & Philosophy or Mathematics & Computer Science take part in the Group Design Practical. This area of work is designed to allow the students to practise skills learnt on the course, and helps them develop and apply the theory learnt in the previous year and a half. It also helps develop team-working abilities and project and time-management skills.

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

Bill Roscoe remembers: Com Lab 1977–2000

Professor Bill Roscoe shares early memories of his time in the department

I first set foot in the Computing Laboratory exactly forty years ago, in Trinity Term 1977, when Steve Brookes (my tutorial partner: we were in our second year reading for a degree in Maths) and I went to the Jenkin Building for tutorials in Numerical Analysis (NA) with David Mayers. In those days NA made up most of the Computing Laboratory. The Programming Research Group (PRG) were in 45 Banbury Road, a Victorian house. The two of us both had some programming experience and decided that in our third year we would do the Scott-Strachey course on Denotational Semantics and Domain Theory, given by Joe Stoy. My favourite memory of my NA tutorials is going punting with Steve and some other friends, being dropped off in the parks so we could go to our tutorial and then rejoin the punt afterwards.

That summer (1977) saw the arrival of Tony Hoare as professor, Strachey having died two years earlier, meaning that Joe was the only Computer Science academic 1975-7, though Dana Scott was very much around but based in Philosophy. Steve and I enjoyed the course so much that we decided to do our doctorates in Computer Science rather than Maths, and we arrived as fresh doctoral students in the autumn of 1978. Tony and Joe were the only academics in PRG, and there were three support staff including Malcolm Harper who had been Strachey's research assistant and stayed in the department for 40 years until 2010. there were a couple of research assistants including Bernard Sufrin, who had arrived at the same time as me, Steve and Geraint Jones, who was also my undergraduate contemporary though at a different college. Bernard, Geraint and I are jointly the longest serving active members of the department.

The MSc in Computation took its first students in 1979. In those days the work of the department was still heavily influenced by Strachey and his students, the emphasis being on programming language design and semantics, and operating system work. Tony's influence was beginning to be felt, with his work on CSP and promotion of formal methods work in general. I remember many seminars in Tony's office, discussing developments in CSP and related areas, often joined by Dana Scott. Steve and I finished our doctorates in 1982. Peter Henderson, a functional programming academic, was briefly a lecturer before leaving again. In 1982 we moved to 8-11 Keble Road, where I remember sharing an office with (at different times) David Gries and David Park. By then Tony was much involved in inmos, the great hope for a British microprocessor industry, and so I, Geraint and other colleagues became much involved in the Transputer and occam.

The big opportunity for departmental expansion came in 1983 and the following few years with the Alvey Programme. The first tranche brought four new lecturers, trebling the number of Computer Science academics to six. These were me, Bernard, Richard Bird carrying on the functional programming tradition, and Ib Sorensen in industrial formal method. The establishment of the undergraduate degree of Mathematics and Computation in 1985 led to a slow but steady trickle of appointments to Tutorial Fellowships.

By about 1989 the department (PRG and NA) had outgrown 8-11 Keble Road, and so we adopted an annexe (affectionately or maybe not) known as Tasmania, in South Parks Road. This convinced all of us who were around that it was far better to be on one site. This became true again in 1994 with the opening of the Wolfson Building. This extension was built almost like a medieval castle under the supervision of administrator Mike Field. The number of degrees was growing: the Software Engineering Programme grew out of a big formal methods collaboration with IBM on the CICS transaction system that had led to the department winning one of its two Queen's Awards in 1990 and 1991. We set up the Computer Science undergraduate degree in 1994.

The range of topics addressed by our researchers slowly grew so that by the end of the 1990s we had a flourishing group in security, one in requirements, the very beginnings of one in Computational Biology and, with Richard Brent, one in Algorithms. Meanwhile Tony, who had never been seen to touch an actual computer at work for many years (his secretary had printed out all his emails and he had written his replies on them in his inimitable longhand) began to take an interest in a computer Microsoft had given him. It was therefore less of a surprise than it might have been when he retired from Oxford in 1999 and moved to Microsoft.

By this time I was running the Mathematical Sciences faculty as its last-ever faculty board chairman (a semi-executive position) and we appointed Samson to take over from Tony.

That is when what I think of as our modern age began.

AFFECTech: wearable tech for better mental health

Mental health problems affect mood and the way people behave, think and react. Referred to as *affective* or *mood* disorders, this group of psychiatric diseases includes depression, bipolar disorder and anxiety disorder. Affective disorders, which are highly disruptive to one's life, have become a major public health concern.

Personal wearable technologies have become popular for physiological monitoring of physical health, including weight management and chronic illness management. In comparison to physical health, personal technologies for affective health have not yet been explored to a similar extent.

With the advent of sensor-based technologies now available on mobile devices, hybrid techniques are emerging that integrate a variety of sensor input, for example heart rate and accelerometer sensors, with self-report questionnaires for stress monitoring.

A limitation of current personal technologies for mental health is that, while supporting self-monitoring in real life settings, they rely on remote data storage and analysis by professional therapists for diagnosis and medical treatment. What is lacking is support for people to develop own understanding of emotions, their triggers and habitual responses, as well as self-help through training emotion regulation strategies.

The AFFECTech project builds on recent advances in emotion monitoring while enabling greater empowerment for the patients themselves. The overall aim is to support self-understanding and successful adoption of adaptive emotion regulation strategies in daily life. The project will design and develop low-cost self-help technologies for visualizing, exploring and regulating emotions.

The AFFECTech network brings together human-computer interaction researchers, biomedical engineers, computer scientists and clinical psychologists from across Europe. At Oxford, the team is led by Professor Marta Kwiatkowska and includes Andrea Patane, Clarendon scholar at the Centre for Doctoral Training in Autonomous Intelligent Machines Systems (AIMS CDT).

A key objective of the project is to capture emotion regulation strategies across the three main affective disorders. AFFECTech will develop tools capturing moment-by-moment emotional changes based on physiological markers such as changes in heart rate and heart vagal tone.

Another key objective is supporting the understanding of one's emotions and the value of regulation strategies in order to support behaviour change for healthier life choices. This involves high-level processes for making sense of physiological data, identifying the triggering events, remembering one's response in dealing with the arising feelings and the adaptive value of such responses. AFFECTech will develop representations of emotional events based on physiological and contextual aspects of emotional experience, informed by sensor input, to support both recall and sense making.

At the technology level, techniques will be developed for extracting such causal representations, visualising them, and integrating a range of sensor input, eg heart and accelerometer sensors, for management of affective disorders.

The final objective is training of adaptive emotion regulation strategies. With a few exceptions, affective health systems offer limited support for emotion regulation, and, although this can be trained, developing successful interventions for people suffering from affective disorders is challenging. Most such interventions have limited applicability to a single class of disorders, and limited user engagement or adoption in daily life.

Findings agree on the strong benefits of cognitive reappraisal, mindfulness and biofeedback. AFFECTech will focus on developing interactive tools for training these three strategies, which have been successfully employed for depression, anxiety and bipolar disorder. It will design engaging user experiences so that people are motivated to adopt them in daily routines. This will involve developing interventions for cognitive reappraisal, integrating wearable EEG (electroencephalogram) and actuators for real-time feedback, and interactive systems for biofeedback training. A challenging goal is the automation of the delivery of a personalised therapy via a mobile device. The developed applications will be tested across patients with stress, depression and bipolar disorder.

AFFECTech thus focuses on advancing the understanding of how wearable technologies can empower people to understand their emotion and develop adaptive regulation strategies responsible for modifying emotional responses in daily life. This marks significant shifts in personal technologies for mental health: (i) from monitoring emotions to identifying emotion regulation processes; (ii) from representing emotions to enabling understanding emotions and their regulation processes; and (iii) towards self-help by empowering people to successfully learn adaptive emotion regulation strategies in daily life.

News in brief

The Department of Computer Science has donated some used computer equipment which is now on its way to Zimbabwe, where it will be used by Emergination Africa. The charity enables students to access online educational resources and mentoring, and was co-founded by current DPhil student Prince Abudu. The shipment of equipment is likely to take a month to travel to Zimbabwe, where it will be used to set up local 'hub' computer classrooms that students from local areas can access. Read more at emerginationafrica.org

The Building Facilities team and IT department have been working very hard recently to refurbish our lecture theatres. As you can see from these 'before and after' shots [below], they have done an amazing job. We now have lecture theatres that are a pleasure to teach and study in.

€14m cash boost for adverse drug reaction research

The Department of Computer Science is one of the partners of the TransQST project on adverse drug reaction research funded with €14 million by the European Innovative Medicines Initiative and with the participation of Professor Blanca Rodriguez and her team.

The new research project aims to improve the understanding of adverse drug reactions and the approach of systems modelling approaches to drug safety. Adverse drug reactions (ADRs) are the unwanted side effects of medication. They can contribute significantly to patient morbidity, mortality and hospitalisation costs.

Funded by the Innovative Medicines Initiative 2 Joint Undertaking (IMI 2 Joint Undertaking) the five-year project, called Translational Quantitative Systems Toxicology (TransQST), aims to develop novel computational approaches using the best available data from the public and private domains to address the problems of drug safety.

TransQT is a partnership between ten academic institutions, three small and medium-sized enterprises (SMEs) and eight pharmaceutical companies, with a total budget of £14m. The project will be coordinated by the University of Liverpool, and the pharmaceutical company AbbVie is the Project Leader.

One of the main focuses of the project are 'off-target reactions' which cannot be predicted from the known pharmacological properties of the drug. The main organs of concern for such reactions are the liver, the kidney and the cardiovascular and gastrointestinal systems.

The focus of the TransQST project will be to provide innovative methodologies and software tools for systems toxicology modelling.

For more information see: transqst.org

Deep learning in biomedicine

DPhil student Polina Mamoshina writes about how research in deep learning is contributing to developments in healthcare.

After one decade of rapid accumulation of data in public repositories, biomedicine has finally joined the 'big data club'. Large open access repositories such as the Library of Integrated Network-based Cellular Signatures (LINCS) project, The Cancer Genome Atlas (TCGA), or national healthcare programs, such as NHS-supported UK Biobank, provide scientists with tens of thousands of high-quality samples. This biomedical data flood should be handled using special approaches. Even high-quality biomedical data is highly heterogeneous and its integrative analysis is difficult. Computational biology methods are routinely used in various fields of healthcare and are incorporated in discovery pipelines of pharmaceutical companies. Machine learning techniques are among the leading methods of in silico analysis.

Machine learning provides general purpose tools for data analysis without a prior assumption on the functional dependencies of the data. This approach is especially useful for the analysis of biological data that usually contains patterns of information without defined relationship between those patterns. Raw biological data can also be sparse and high-dimensional, which complicates statistical analysis. Classical machine learning approaches, such as unsupervised k-means clustering and principal component analysis or supervised methods like support vector machines and random forests, are usually used to analyse different types of biological data. However, these shallow machine learning methods are limited in terms of modelling and predictive power due to the inability to learn high-level representations of data.

These challenges could be addressed by modification of raw features which already incorporate a certain level of complexity. Thus, computational pharmacology and computational transcriptomics facilitate the design of molecular descriptors and the annotation of gene networks and signalling pathways, respectively. Yet, the raw data analysis has the potential to be more efficient and methods with deep architecture, that are effective in extracting high-level abstractions from raw heterogeneous data, hold a great promise in biomedicine.

An increasing volume of data, coupled with the availability of more powerful and cheaper computational resources, is leading to the rapid development of the entire field of deep learning that goes far beyond text, voice and image recognition problems. The number of literature sources on the applications of deep neural networks in biomedicine is growing rapidly. These approaches include different types of biomedical data analysis, versatile online omics tools, such as the epigenetic framework 'DeepSEA' or aging biomarker 'Aging.AI', and first proofs of concepts in drug design and drug discovery. [Omics in this context is an informal way of referring to any field of study in biology which ends 'omics'].

Deep learning is a state of the art method for many tasks, such as machine vision and image analysis, and it has been successfully translated into the field of biomedicine. Architectures of convolutional neural networks enable superior performance and, more importantly, transfer learning.

Convolutional neural networks could be pre-trained on large datasets, such as ImageNet, and then applied on more specific biomedical datasets with a limited number of samples. Feed-forward deep neural networks are well equipped to handle sparse omics data with non-linear dependencies. Because of their high generalization ability, deep neural networks can be used for the binding and interpretation of data profiled using different microarray and sequencing platforms or cell lines.

One of the most promising areas of deep learning algorithm research are generative adversarial networks, which allow for sampling from the learned latent distribution and generation of synthetic data. Modern drug discovery process is a blind search for promising molecules that are called leads (which can potentially become drugs). New molecular structures with the desired properties can be generated with the use of the deep adversarial autoencoder, trained on the possible 2D structure representations called molecular fingerprints.

While deep architectures can be significantly more efficient than shallow machine learning models, deep neural networks are hard to train and regularise. Also, deep learning models, like many other machine learning methods, lack transparency, which complicates quality control and implementation of results, and that is why they are often referred to as black boxes. One of the prospective areas in deep learning is the development of tools for feature selection and feature extraction from hidden layers.

British Heart Foundation Fellowship for Alfonso Bueno-Orovio

Congratulations to Alfonso Bueno-Orovio who has been awarded a British Heart Foundation Intermediate Basic Science Research Fellowship, with funding for 5 years.

The fellowship will be based in our department and co-sponsored by the Radcliffe Department of Medicine.

Alfonso has been awarded the fellowship for research into Hypertrophic Cardiomyopathy (HCM). This is a complex hereditary heart disease and the leading cause of sudden cardiac death in children and young adults, including young athletes. An emblematic example of the risk of death in HCM was the cardiac arrest experienced by Fabrice Muamba in March 2012 in the televised football match between Bolton and Tottenham Hotspur: he was clinically dead for more than an hour, and his heart needed 15 electrical shocks to be restarted, saving his life. Whereas many HCM patients remain asymptomatic, sudden cardiac death is also in many instances the first manifestation of this disease.

Hypertrophic Cardiomyopathy affects both the muscle and the electrical function of the heart. However, little is yet known about how these two factors combine to increase the risk of deadly episodes in the disease, as this cannot be easily studied in living patients. In this context, computational models of human HCM electrical function and anatomy will be constructed based on a unique dataset at different levels of human physiology. Simulation studies will be conducted to investigate how the various components of the disease play a role in establishing arrhythmias in human HCM hearts. These new technologies will help in a better diagnosis of patients at risk of cardiac death, and in the identification of novel therapies for the disease.

UnBias uses Youth Juries in research

By Helena Webb, Marina Jirotka and Menisha Patel from the UnBias team.

Young people today are digital natives who grow up online. It is crucial that research seeks to support their safe and productive engagement with the digital world.

The ongoing 'UnBias' project is using an innovative approach called youth juries, developed by our project partners at the University of Nottingham, to empower young people to reflect on their digital rights in relation to the operation of algorithms on online platforms. These youth juries are interactive sessions designed to stimulate discussion and co-produce scenarios to illustrate young peoples' concerns. They allow us to understand youth views on the ways that algorithms filter and personalise online content. They also explore how young people feel these algorithms shape their online experience and what trade-offs they feel are 'fair' between ease of use and freedom of choice.

Further youth juries will co-develop educational materials that will encourage young people to develop a healthy and engaged perspective on online information provision. For more information see unbias.wp.horizon.ac.uk

Research Highlights: EEG Privacy Attacks

By Associate Professor Ivan Martinovic and Doctoral Student Ivo Sluganovic.

Brain-Computer Interfaces (BCIs) enable non-muscular communication between a user and a computer by measuring the brain's activities. In the last decades, BCIs have been primarily applied in the medical domain with the goal to increase the quality of life of patients with severe neuromuscular disorders. Additionally, as particular patterns of brain waves have been found to differentiate neurocognitive states of both disabled and healthy users, EEG has been used for wide variety of research applications. These include neurofeedback therapy in attention deficit hyperactivity disorder (ADHD), epilepsy and sleep disorders monitoring, studying underlying processes of skilled performance in sports, or even monitoring the mental workload of air-traffic control operators.

Besides medical and research applications, EEG-based BCI devices are becoming increasingly popular in the entertainment and gaming industries. In the recent years, several companies (Emotiv Systems, NeuroSky, etc) have started offering low-cost EEG-based BCI devices and software development kits that allow building third-party apps for their application markets. Currently available applications range from accessibility tools (such as a mind controlled keyboards) and hands-free arcade games, to so-called serious games, ie, games with a purpose other than pure entertainment, such as attention and memory training

continued on next page ▶

from previous page ▶

games. The increasing importance of BCI technologies have been recently emphasised by CEO's such as Facebook's Mark Zuckerberg and Tesla's Elon Musk, who have both in the last two months announced research projects and plans to develop brain-computer interfaces and integrate them in their future products.

Given that this technology provides information on our cognitive processing and allows making inferences about our intentions, conscious and unconscious interests, or emotional responses, our research focuses on understanding if EEG-based BCIs can be turned against its users to reveal their private information.

In our initial work on the topic of privacy attacks on brain-computer interfaces, we emphasise that such devices may make the raw EEG signal available to potentially untrusted third-party applications. If such an application is malicious, it could in turn abuse the BCI device to infer private information about a victim, such as her/his preferred bank, area-of-living, or recognition of a familiar face.

The general idea of this attack is similar to a polygraph, where the interrogated person's physiological reactions are used to reason about his/her knowledge. As we show in our research, the attacker that controls a malicious BCI-enabled app can manipulate the visual stimuli presented on screen and analyse the corresponding event-related responses in the EEG signal to learn information about the user that he would otherwise disclose. While the experimental results show that an attacker can not learn private information with perfect confidence, the guessing entropy of the private information is decreased on average by approximately 15% - 40% compared to random guessing attacks, indicating that event-related potentials indeed leak private information.

While such attacks could potentially be very damaging, their fundamental limitation is that they rely on repeatedly showing a similar consciously perceived (supraliminal) visual stimuli, which makes them detectable by users. Our current work focuses on researching if the threat of previously proposed privacy attacks can be further increased by preventing their conscious detection. We propose a subliminal attack that infers private information by probing the victim at a level below his/her cognitive perception. Similar to subliminal advertising, the key idea is to show the visual stimuli within the screen content that the user expects to see, but for a duration that is too short for conscious perception (several milliseconds), yet still sufficient to result in detectable activation of certain parts of user's brain.

To show feasibility of such attacks, we implement small snippets of visual stimuli within a few frames of the video, for a duration not exceeding 13.3 milliseconds, and run a proof-of-concept study with 27 subjects who are wearing EEG-based BCI devices. We use machine learning techniques to analyse users' subliminal reactions to the embedded visual stimuli and conclude that for 18 out of the 27 subjects, our classifier was able to correctly infer whether the user recognises a specific face or not, thus reducing the attacker's guessing entropy by 20.8%.

While not perfect, our experimental results show that the attacks on privacy of EEG-based BCI devices could indeed be feasible, and that some of the attacks could even be performed subliminally.

On the other hand, considering that this is a truly interdisciplinary topic that sparks the interest of many individuals, and thus has a large potential of being misinterpreted by the media, it is important to put the existing results into perspective. The results do not mean that your thoughts and brainwaves can undetectably be read by adversaries whenever you use a BCI device. In order to rule out negative results due to environmental factors, the discussed experiments were performed under strictly controlled conditions. Furthermore, the subliminal probing was done using a high-end medical-grade BCI device, which achieves significantly better measurement than the current user-grade devices.

However, as a first indication of new threats, this research is a crucial important step towards starting the discussion about security and privacy of brain computer interfaces – a topic which is only gaining in importance as more and more BCI devices become commonplace in our lives.

Modern laws of robotics:
What would it mean for AI to act responsibly?

Niel de Beaudrap ponders the question, 75 years after writer Isaac Asimov first proposed his Three Laws of Robotics.

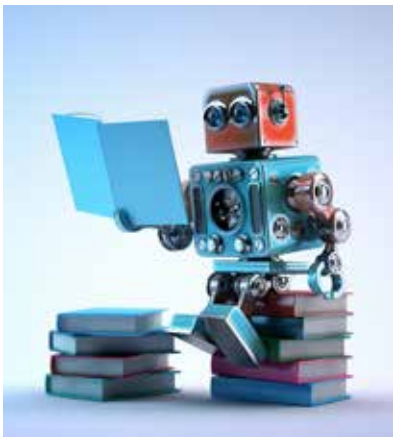

Isaac Asimov was a brilliant scholar and author who influenced millions. But we should remember that his famous Laws of Robotics were written in the United States before the Civil Rights movement, not as a technical specification, but to tell interesting stories, often about robots being wrong-footed by social constraints. His three 'laws' as introduced in the 1942 short story *Runaround* were:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

Asimov was not alone in using Artificial Intelligence (AI) as social commentary. Alan Turing's article *Computing Machinery and Intelligence* can itself be read as a study of how or why we regard someone as a moral equal – or at least a peer – as much as it does machine intelligence. So, we should consider what we hope to achieve with new 'Laws of Robotics', for AI as we understand it at the current time.

When we think of autonomous AI, we often think of autonomous military drones and self-driving cars. But possible applications of Artificial Intelligence which do not require a localised 'body' must also be considered, such as algorithms which may in the future control access to government services, direct economic policy, or monitor human behaviour. In all cases, the concern is one of our right to life and self-determination – specifically as it might be affected by inscrutable, unaccountable, and powerful entities.

In short, the issue is one of trust, and becomes important before autonomous AI makes an appearance in society. Any job which can be done by a sufficiently versatile machine, is a job which no longer requires a human to do – a job which no longer supports a human's welfare, nor needs to pass a human's judgment. The problem becomes compounded once fully autonomous AI begins to play a role in more complex tasks, as the impact and complexity of work which is performed without oversight by members of society increases.

Is the problem that inherently moral humans might be replaced and displaced by entirely amoral AI? On the contrary: some of what we describe as 'distrust in AI' is really distrust in those who might use AI as a tool for their own interests, disrupting the lives of many in order to benefit a few. Such concerns are nothing new: there is a long tradition of distrust in those who own the means of production, or occupy seats of power. The new risk is that AI may allow powerful people to rely less on the approval of society and care less about its welfare. It would be naive to suppose that AI will be deployed for the good of humanity by default.

The question of the livelihood of displaced human workers requires a thorough reconsideration of how we order society. But as for the moral impact of work, a possible answer is to replace human ethical oversight with *robotic* ethical oversight. If we can design AI with flexibility of behaviour, we can go one step further to design it to make ethical considerations in how it behaves. One social safeguard could then be to require that no work be done without the oversight of a moral agent – human or otherwise.

How does one make a moral machine, to design it to act ethically? We cannot simply 'program in' ethics: we cannot even agree amongst ourselves what 'ethical' is, or how best to teach morals to children. But society depends on us developing and teaching ethics to our children, and to detect warning signs in the behaviour of others. That is achieved through social engagement – tutelage, discussion, debate, and shared experience.

So, perhaps autonomous AI should be designed to engage in the franchise of humanity, as described by Turing.

This suggests two guiding principles for new 'Laws of Robotics':

1. AIs must be engaged members of society; and
2. AIs may not damage the social welfare or, through inaction, allow the social welfare to come to harm.

This is no small order: it is a design feature implemented only imperfectly in humans. But perhaps this is a place to start…