

NEWS FROM THE DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF OXFORD

## THREE NEW CENTRES FOR DOCTORAL TRAINING COME TO OXFORD

Read more – p5



### HAPPY BIRTHDAY GCSCC:

Celebrating five years of cybersecurity impact – p18



### EVOLUTION AS AN ALGORITHM:

Understanding algorithmic biology – p24



### GENE EDITING:

Improving the safety of CRISPR – p28



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](mailto:editorial@cs.ox.ac.uk)

To subscribe to future issues, e-mail: [editorial@cs.ox.ac.uk](mailto:editorial@cs.ox.ac.uk)

To download previous issues, visit [www.cs.ox.ac.uk/inspiredresearch](http://www.cs.ox.ac.uk/inspiredresearch)



## Editorial board

Suzanna Marsh (Managing Editor)

Kiri Walden (Editor)

Suvarna Designs (Designer)

Sarah Baldwin

Leanne Carveth

Rosanna Cretney

Emma Dunlop

David Hobbs

Stefano Gogioso

Sharon Lloyd

Elisa Passini

Alex Rogers

Helena Webb

Standa Živný

## Contributors

Bushra AlAhmadi, Reuben Binns, Katherine Fletcher, Yarin Gal, Ross Gales, Paul Goldberg, Vojtech Havlicek, Jassim Happa, Laura Jones, Marta Kwiatkowska, Artem Kaznatcheev, Francesca Margara, Peter Minary, Grace Parsons, Blanca Rodriguez, Shirley Sardar, Sean Sirur, Florian Stortz, Max Van Kleek, Irina Voiculescu.

# CONTENTS

p4 Remembering Stephen Cameron

p5 Three new centres for doctoral training

p6–12 News

p13–14 Student focus

## Feature articles

p15 New research in antibiotic resistance

p16 Using evolutionary game theory to identify cell resistance within cancer

p17 Protective: cyber threat intelligence improving cybersecurity incident response

p18–19 Celebrating five years of cybersecurity capacity research impact

p20–21 My home is my castle: rethinking smart home IoT for privacy

p21 Teaching AI new skills with Hanabi

p22 How capturing data in social media posts provides vital information during emergency situations

p23 GDPR: are we there yet? Not quite

p24–25 Algorithmic biology: evolution is an algorithm so let us analyse it like one

p26 Factoring: searching for solutions that are sure to exist but hard to find

p27 Supervised learning with quantum-enhanced feature spaces

p28 CRISPR gene editing: a computational biology challenge

# Letter from the Head of Department

There are many different aspects to being Head of Department and no two days are the same.

I was shocked and deeply saddened to hear of the sudden and unexpected death of Stephen Cameron, one of the longest-serving members of the Department. Stephen was appointed as the first tutorial fellow in Computer Science at Keble College in 1988, and built up the subject from scratch there, nurturing generations of undergraduates and graduates. He contributed greatly to the breadth of our curriculum, enthusing students to get involved in a huge range of projects from robot football, to flying machines, co-ordinating search and rescue teams, and computer animation. We will greatly miss his down-to-earth approach to our subject, and his real concern to see students learn and thrive. There are tributes to Stephen on page 4, and some reflections on his contributions and achievements. I can only add that I continue to miss him, and offer my very best wishes to all his family and friends in their very sad loss.

There have been unexpected pleasures too. It was wonderful to hear the announcement in April that Marta Kwiatkowska had been elected as a Fellow of the Royal Society. It is a great privilege to work in a department with so many talented individuals who are really making their mark and shaping the whole field of Computer Science.

I was also really proud to attend our group project presentations recently and see our 2nd year undergraduates displaying their amazing skills and impressing their industry sponsors. Our students created and demonstrated software that did a huge range of useful things, modelling pollution in London, picking up abnormal X-rays, digesting financial data, and spotting who couldn't be trusted on social media. There were 15 teams this year and each one had achieved something impressive. When we started we had just 6 teams: that's how much we've grown.

Whatever else 2019 brings, our aim as a department will continue to be to deliver world-class education and research.

**Professor Peter Jeavons**  
June 2019



## New Starters

We welcome new members of the department, and some existing colleagues who have changed role. They are (in joining/moving order):

Destiny Chen, Amir Shaikhha, Bei Peng, Cristina Cirstoiu, Elizabeth Wyatt, Konstantinos Meichanetzidis, Nikhil Balaji, Steven Herbert, Ivo Sluganovic, Simon Eberz, Stefano Germano, Xiaoxuan (Chris) Lu, Hector Martinez-Navarro, Agnieszka Zbrezezny, Chaoqing (Bill) Tang, Patricia Esteve-Gonzalez, Amy Hunter, Richard Baker, Matthew Smith, Graham Lee, Sarah Baldwin, Thibault Lestang, Thomas Palmer, Trevor Jennings, Abhishek Dasgupta, Sahar Vahdati, Hugo Paquet, Pierre Pradic, Oliver Sampson.

## News in brief

Sam Staton and Luke Ong are part of a collaborative team that has won a Facebook Research 'Probability and Programming' research award. The award is for the team's research on differentiable probabilistic programming semantics. Probabilistic programming is a method for writing complex statistical models as programmes. Many machine learning methods use gradients, and differentiable probabilistic programming takes advantage of those methods. The team proposed to investigate the foundations of this kind of programming.

Facebook gave unrestricted gifts to ten research teams. Read more about the awards here: [bit.ly/2HG7VhZ](https://bit.ly/2HG7VhZ)

Aidan Gomez has been awarded an Open Philanthropy AI Fellowship. The Open Philanthropy Project awarded approximately \$2,000,000 over five years in PhD fellowship support to eight promising machine learning researchers. These fellows were selected 'for their academic excellence, technical knowledge, careful reasoning, and interest in making the long-term, large-scale impacts of AI a central focus of their research.'

Aidan leads the research group FOR.ai, focusing on providing resources, mentorship, and facilitating collaboration between academia and industry. Aidan's research pursues new methods of scaling individual neural networks towards trillions of parameters, and hundreds of tasks. His work takes a humanist stance on machine learning applications and their risks.

# Remembering Stephen Cameron

The Department of Computer Science is very sad to announce the death of Professor Stephen Cameron. Stephen was one of the longest-serving members of the Department, joining us in 1988 as a University Lecturer in Computer Science and Tutorial Fellow at Keble College. Stephen suffered a major stroke on Friday 5 April and did not recover consciousness. He died peacefully on 15 April. He will be sorely missed.

Many colleagues and friends have contributed to a book of remembrance, and we include some excerpts from that below:

‘From project supervisor, to enthusiastic researcher, to group leader and caring colleague, Stephen has inspired me to persevere in my research field and pursue my ideas even when these were moving away from the mainstream. Over almost 24 years of collaboration he has often encouraged me to take new directions in my career.’

**Dr Irina Voiculescu, Department of Computer Science, Oxford**

‘When I started working with him, 22 years ago, he was much more than my research “boss”. He was full of good ideas for things to do, but also gave me the freedom to explore things for myself. [...] He also taught me quite a bit about folk music and morris. Stephen will be greatly missed.’

**Dr Joe Pitt-Francis, Department of Computer Science, Oxford**

‘Such a lovely guy. When I first ventured into the den of geometry fundamentalists Stephen was always very kind, good-humoured and considerate.’

**Cecil Armstrong, Emeritus Professor, Queen’s University, Belfast**

‘In all situations, whether concerns were academic or administrative, I could trust him to have the right balance of practical suggestions and cheerful encouragement. That voice on the phone, before I was even sure that I belonged in Oxford,

proved to be a tutor, a mentor, and a friend. Stephen’s dedication to his students and their development was immense yet understated, and those of us fortunate enough to have been in his care will, I am sure, remember him with fond admiration.’

**Dr Gwyn Skone, Software Engineer**

‘Stephen had a great sense of humour, something he didn’t hold back in bringing to his work. I remember particularly an undergraduate Computer Graphics practical based on a dancing figure. It had a photo of Stephen’s face wrapped around its head, and a cylinder of tartan as a kilt. It was unmistakably Stephen. Constraining the figure’s movements to something sensible was an optional extension to this task and, as a consequence, some very undignified dance choreographies resulted.’

**Simon Bond, Software Development Lead, Big Data Institute**

‘I consider myself so lucky to have had Stephen as a supervisor. He gave us a lot of space to pursue our research questions but intervened and supported us exactly when we needed it. [...] I wouldn’t be where I am today without him and I think back to my PhD with him very fondly.’

**Dr Julian de Hoog, Research Scientist, IBM Research Australia**



**Stephen Cameron  
1958-2019**

‘[...] quiet but with a good sense of humour, thoughtful, kind, generous with his advice, and modest, yet with an appreciation of good beer and morris dancing. I always thought his good ideas were undervalued.’

**Ralph Martin, Emeritus Professor, University of Cardiff**

‘A friend, a colleague and a renowned Robotics expert, Stephen, you will always be alive in our hearts. You were a fine man with a great heart, the thought of you always brings a smile to our faces, and good memories. Pity that you left so early, as there are so many peated whiskeys we were to enjoy together. We will miss you. Rest in peace.’

**Dr Ashutosh Natraj, CEO at Vidrona, ex-colleague and friend**

Stephen was a keen bell ringer: he was for many years the Senior Member of the Oxford University Society of Change Ringers (OUSCR), and tower captain at Charlton-on-Otmoor. Members of the OUSCR and the Oxford Diocesan Guild of Church Bellringers rang quarter peals and a full peal in Stephen’s honour.



*Members of the Oxford University Society of Change Ringers (OUSCR) after a quarter peal in memory of Stephen Cameron*



# Three New Centres for Doctoral Training come to Oxford

Three Centres of Doctoral Training (CDT) at the Department of Computer Science have received new funding from the Engineering and Physical Sciences Research Council (EPSRC), part of UK Research and Innovation (UKRI), to ‘develop the skills needed for UK prosperity’.

Science and Innovation Minister, Chris Skidmore, said, ‘As we explore new research to boost our economy with an increase of over £7 billion invested in research and development over five years to 2021/22 – the highest increase for over 40 years – we will need skilled people to turn ideas into inventions that can have a positive impact on our daily lives.’

The Department of Computer Science CDTs will be in the areas of health data science, computational biology, and machine learning. The CDT in health data science will be based within the Oxford Big Data Institute (BDI). Over the four-year programme students will undertake a year of training in statistics, machine learning, health data, ethics and responsibility. A three-year project will follow with a leading research group, including placement opportunities with partner organisations such as technology and pharmaceutical companies, research organisations, and healthcare providers. Teaching, mentoring, and research supervision will be provided by leading specialists from the Departments of Computer Science, Statistics, and Engineering Science, the Nuffield Department of Population Health, and the Nuffield Department of Medicine.

Jim Davies, who will head up the new CDT in health data science,

said, ‘thanks to the EPSRC and the BDI, we are able to deliver research training at the heart of the University’s medical campus, producing the data scientists we need to help transform health research and healthcare delivery.’

The computational biology CDT will focus on Sustainable Approaches to Biomedical Science: Responsible and Reproducible Research (SABS: R<sup>3</sup>). This is an innovative open collaboration between the University of Oxford and 22 partner industrial organisations. The programme aims to train first-rate biomedical scientists equipped with the skills needed to transform their research through the creation of innovative, reusable computational tools and solutions for cutting-edge biomedical research problems. This four-year programme of research and training has strong industrial links, with each student having both academic and industrial supervision.

David Gavaghan, who will head the computational biology CDT comments, ‘The demand for biomedical scientists who can lead the development of the transformative, well-engineered software needed to maintain the world-leading position of the UK pharmaceutical and biomedical industrial sectors is growing exponentially. Despite this, biomedical scientists are rarely trained to develop the well-engineered software tools that could transform the way research is conducted, and hence accelerate global scientific discovery. The EPSRC CDT in Sustainable Approaches to Biomedical Science: Responsible and Reproducible

Research will address this need. It will be developed by a consortium of internationally leading computational biomedical researchers drawn from the University of Oxford and 24 industrial and institutional partners. The doctoral research projects undertaken each year will be industrially-derived and sponsored, and will have joint industrial and academic supervision in the areas of physiological modelling, medical imaging, and computational approaches to structural biology. Our aim is to equip 75 future research leaders with the skills needed to transform the research culture in computational biomedical science.’

The remaining CDT is for research in Autonomous Intelligent Machines and Systems (AIMS). This new centre reflects a national requirement to increase the number of people with advanced skills in machine learning. AIMS incorporates not only machine learning, but also robotics and vision, cyber-physical systems, verification and control. Within the CDT there will be Special Interest Groups (SIGs) on sustainable urban development, extreme and challenging environments, and smart health.

Lead Principal Investigator (PI) of the new AIMS CDT Niki Trigoni says, ‘Our continued vision is to train cohorts with theoretical, practical and systems skills in autonomous systems (comprising AI, machine learning, robotics, verification and sensor systems) and a deep understanding of the cross-disciplinary requirements of these domains.’

## News in brief

We are delighted that Catherine Vlasov, a Department of Computer Science student, was a finalist in the Worshipful Company of Information Technologists (WCIT) Outstanding Information Technology Student Awards 2019. The awards were created to recognise outstanding undergraduate and postgraduate IT students within the UK. Criteria for the award included academic excellence, overcoming adversity, entrepreneurial skills, and contribution to charity or community. Each of the five finalists, accompanied by an academic lead from their sponsoring institution, was invited to the WCIT's 103rd business lunch in February, held in the beautiful Wax Chandlers' Hall in the City of London.

The Department of Computer Science hosted a children's party on 26 January 2019. Members of the department brought along their children for some fun activities including party games as well as arts and crafts, all on the theme of Computer Science.



Adrian Duncan's tortoise Shelly, representing Worcester College, won the Corpus Christi tortoise race on 2 June.

## Schwarzman donation will help Oxford academics to work together on the ethics of AI

It has been announced that a £150,000,000 donation from Stephen Schwarzman will be used to establish the Stephen A. Schwarzman Centre for the Humanities, but will also be used to create a new institute of AI Ethics at Oxford. This will enable interdisciplinary collaboration that will help academics across the university, both in the sciences and humanities, to work more effectively in this exciting area of research. Mr Schwarzman, who was educated at Yale and Harvard, is a New York businessman and philanthropist. He has previously supported leading educational and cultural institutions such as Yale, MIT, The New York Public Library and the Schwarzman Scholars Program at Tsinghua University.

## Bernardo Cuenca Grau awarded funding for three-year project: OASIS

The new EPSRC funded project OASIS (Ontology Reasoning over Frequently-changing Streaming Data) will aim to lay the foundations for a new generation of ontology-based data management systems (ODMSs) capable of ingesting and processing rapidly changing data in real time.

ODMSs play a key role in many sectors of our economy – oil and gas companies continuously monitor sensor readings to detect equipment malfunction and predict maintenance needs; network providers analyse flow data to identify traffic anomalies and denial-of-service attacks; and Internet of Things (IoT) applications require real-time analysis of data stemming from multiple types of device. Despite this, state-of-the-art ODMSs are not well-suited for applications operating in real-time analysis of rapidly changing data. This is the problem that Bernardo Cuenca Grau, together with co-investigators Boris Motik and Ian Horrocks, will be addressing.

## Paper wins Alberto O. Mendelzon Test-of-Time Award

Two Department academics, Georg Gottlob and Thomas Lukasiewicz, are co-authors on a paper which has won the Alberto O. Mendelzon Test-of-Time Award.

The paper is 'A General Datalog-Based Framework for Tractable Query Answering over Ontologies' (co-written with Andrea Cali from Birkbeck, University of London). The ACM PODS Alberto O. Mendelzon Test-of-Time Award was presented for the first time in 2008. It is awarded every year to a paper or a small number of papers published in the PODS (Symposium on Principles of Database Systems) proceedings ten years prior, which had the most impact in terms of research, methodology, or transfer to practice over the intervening decade. Each year, the winner(s) of the award receive plaques and the sum of \$1,000 (divided equally among the winners, if more than one). The award is funded by a generous gift from IBM.

The committee handling the award for 2019 was comprised of Victor Vianu, Jianwen Su, and Dirk Van Gucht.

## Server room cooling now 70% more efficient

The University of Oxford's Environmental Sustainability team have transformed the energy efficiency of the Wolfson Building's server room in a collaborative project with its users, the Department of Computer Science, and the engineers of the Building Services team.

The building's Facilities Manager, Joe Atherton, approached the team with concerns about how much power the building's server room was using – it is cooled all day, every day by three air conditioning units, each with two fans. These were several years old and used an outdated design – they were driven by belts and operated at a fixed speed. Replacing them with modern fans, improving airflow through the system, installing temperature sensors and giving the fans the ability to vary their speed depending on how much cooling is needed, reduced the system's energy

consumption by a total of 70%. By reaching lower temperatures more efficiently, the project should also help the computer equipment in the server room to work better and last longer. It is expected to save 429,000 kWh of energy each year, 131 tonnes of carbon dioxide emissions. In monetary terms that is £43,000 a year – enough for the project to pay for itself in less than seven months. Still greater savings may be possible once the system is integrated with the building's building management system, which will provide even more control over the fans.

Sharon Lloyd, Departmental Administrator, commented: 'Computer Science has had a programme of reducing its carbon footprint over the last 5 years and Joe had implemented with the sustainability group and estates

a number of measures to reduce power consumption from sensor lights to timers on water heaters. Our biggest issue had been the consumption of power in our data centre and so Joe worked with the sustainability group on an ambitious programme to replace the ageing cooling plant whilst keeping the data centre running, with staggering results. Congratulations to Joe for achieving such a fantastic result and reducing our energy bills by such a huge amount.'



## Jennie Charlton wins Most Improved Apprentice Award



Jennie Charlton, PA to the Administrators/Office Manager at the department has won the 2019 Most Improved Apprentice Award at the 2019 University Apprenticeships Expo.

Jennie, who is doing the Operations/Departmental Manager: Level 5 Standard Apprenticeship says 'It's a great scheme which enables existing staff to gain a qualification whilst working, and in a discipline related to their work, rather than the Jewellery

& Silversmithing I studied at university! It has really boosted my confidence in many areas that I really struggled with previously, and that has been helped by having a very supportive manager, as well as the different staff and tutors I have been involved with.'

At the same ceremony former department apprentice Bradley Boodell was awarded a certificate for completing his Finance Assistant Apprenticeship.

## Quantum Computing Seminar in Parliament

On 19 March Professor Marina Jirotko was one of three leading Oxford academics invited to participate in a seminar held at the Houses of Parliament, exploring the opportunities that the quantum computer revolution could bring to the UK. The discussion explored how close the UK is to creating the world's first scalable quantum computer, a

challenge which the University of Oxford is spearheading.

Quantum computing has the potential to forge not just revolutionary technological change, but a whole new industry for the UK. Participants were invited to discover why continued support for the quantum technologies programme, including the decision to invest £35 million in a new National Quantum Computing Centre, is key for the UK to retain

its scientific, economic and security edge.

Marina Jirotko is Professor of Human-Centred Computing and leads the Responsible Research and Innovation (RRI) programme of work within NQIT (Networked Quantum Information Technologies Hub, part of the UK National Quantum Technology Programme).

Read more: [nqit.ox.ac.uk/](http://nqit.ox.ac.uk/)



## Hay Festival Robotics workshop

Three members of OxWoCS (Oxford Women in Computer Science) ran an interactive workshop at the Hay Festival in May. There were two sessions: one aimed at children aged between 10 and 11 years old, another for children aged between

13 and 14 years old. Klaudia Krawiecka, Paula Fiddi and Bushra AlAhmadi taught participants how to code by playing with Sphero SPRK+ robots, programming robots to perform dance moves and navigate a maze.

## UNIQ summer school

In April we welcomed 25 students, studying Computer Science as part of the UNIQ scheme. During their stay UNIQ students are taught by Oxford academics, attend lectures and labs, and use departmental facilities – giving them a taste of life at the University of Oxford. The UNIQ Computer Science course is a taste of the mathematical, scientific and engineering principles underlying every kind of computing system. Throughout the week students are introduced to theories and concepts in lectures, backed up and demonstrated through practical sessions in the department labs. UNIQ is also a great opportunity to meet like-minded people and gain deeper understanding of a subject that interests them.

Read all about UNIQ here: [uniq.ox.ac.uk](http://uniq.ox.ac.uk)

## NASA Collaboration

In January Yarin Gal visited the Arecibo Observatory in Puerto Rico as part of ongoing research with NASA's Frontier Development Lab on the use of machine learning for 3D shape modelling of asteroids from radar data.



## Klaudia Krawiecka awarded Google Women Techmakers Scholarship

Computer Science DPhil student Klaudia Krawiecka has been awarded a 2019 Google Women Techmakers Scholarship. She will receive a 7,000 EUR scholarship, and attend the Google Scholarship Retreat in London. Through the Women Techmakers Scholars Program - formerly the Google Anita Borg Memorial Scholarship Program - Google is furthering Dr. Anita Borg's vision of creating gender

equality in the field of computer science by encouraging women to excel in computing and technology and become active leaders and role models in the field.

Klaudia is a doctoral student at the Centre for Doctoral Training in Cyber Security at the University of Oxford and the recipient of (ISC)<sup>2</sup> Women's Cyber Security Scholarship. In 2018, she was elected President of Oxford

Women in Computer Science Society (OxWoCS) at the University of Oxford. Apart from leading and coordinating technical workshops and hackathons (e.g. OxfordHack 2018), her responsibilities include organising events to promote Computer Science among female students. She is also involved in many outreach activities such as InspireHer! workshops for mothers and daughters. Klaudia's current research focuses on authenticating the Internet of Things devices using out-of-band channels such as visual channels (e.g. Augmented Reality).

## Blanca Rodriguez recognised with two Fellowships

Professor Blanca Rodriguez has been elected Fellow of the European Alliance for Medical and Biological Engineering & Science (EAMBES), and Fellow of the European Society of Cardiology (ESC).

The EAMBES represents the academic aspects and interests of biomedical engineering in Europe. The Division of Fellows is formed by individuals who have distinguished themselves by identifiable contributions or accomplishments in medical and biological engineering and science. Blanca delivered a lecture in Brussels in March for the EAMBES fellowship.

Fellow of the ESC (FESC) is a title of honour that is awarded to healthcare professionals who have made a significant contribution to cardiology. There are more than 4,300 Fellows of the ESC, many of whom are actively involved in ESC activities, from congress and research to education and advocacy. The FESC title is recognised throughout the cardiology community as a symbol of excellence. Find out more: [eambes.org](http://eambes.org) and [escardio.org](http://escardio.org)

## Marta Kwiatkowska elected as a Fellow of the Royal Society

Marta Kwiatkowska is one of five eminent Oxford scientists who have been elected as Fellows of the Royal Society for their exceptional contributions to science.

The Royal Society is a self-governing Fellowship made up of the most eminent scientists, engineers and technologists from the UK and the Commonwealth. Fellows and Foreign Members are elected for life through a peer review process on the basis of excellence in science. Each candidate is considered on his or her own merits and can be proposed from any sector of the scientific community. Every effort is made to encourage nominations of female candidates and candidates from the emerging disciplines.

There are approximately 1,600 Fellows and Foreign Members, including around 80 Nobel Laureates. Each year up to 52 Fellows and up to 10 Foreign Members are elected from a group of around 700 candidates who are proposed by the existing Fellowship.

New Fellows are formally admitted to the Society at the Admissions Day ceremony in July, when they sign the Charter Book and the Obligation of the Fellows of the Royal Society.



## Dave Gavaghan appointed as University Advocate for Research Staff Development

Professor Dave Gavaghan has become the University of Oxford's first University Advocate for Research Staff Development.

He says of his new role, 'research staff comprise some of the brightest minds in their generations drawn to Oxford from across the globe to conduct cutting-edge research across the University's four academic divisions. The importance of the University's research staff in maintaining its position as one of the world's leading research institutions cannot be overstated.'

If the University is to continue to attract the best researchers to Oxford, particularly in the face of current external challenges, then it



needs to provide an environment within which each member of research staff can determine and successfully move on to the next stage in their career, whatever that next stage might be. In short, the University needs to provide a world-leading researcher development environment to support its research staff.

My appointment to the role of University Advocate for Research Staff Development will enable me to continue to take a leading role in the development of this environment, and in ensuring that the University continues to be the destination of choice for researchers wishing to move on successfully to the next stage of their careers.'



## Marina Jirotká awarded an Established Career Digital Economy Fellowship by EPSRC

Professor Marina Jirotká has been awarded a five year Established Career Digital Economy Fellowship by the EPSRC.

The fellowship project, titled RoboTIPS: Developing Responsible Robots for the Digital Economy, aims to develop a new approach to responsible innovation (RI) in the context of trustworthy and secure technologies. The principal challenge is to embed responsible innovation into technology developers' practices and to create positive cases of RI in action.

The project is in collaboration with Alan Winfield, Professor of Robot Ethics at the Bristol Robotics Lab (BRL). Marina and Alan have been working together for some time to develop the concept of the 'ethical

black box'. This advocates that robots should be equipped with a 'black box', equivalent to the flight data recorders used in aviation.

The black box continuously records sensor and relevant internal status data and can be extended in scope to also capture the AI decision-making process and environmental factors occurring before an adverse incident. This data can provide crucial evidence following accidents or adverse incidents involving robots. The information provided can help us to understand why a robot behaved in the way it did and then make recommendations for changes to prevent similar incidents or limit the potential damage caused. The ethical black box thus provides a pathway to greater safety and public trust in robots.

Marina will bring together a team of experienced researchers and companies developing social robots to work on the project. The team will include Dr Helena Webb, who works closely with Marina as part of the Human Centred Computing theme at Oxford.

Throughout the five years, Marina and her team will develop: a grounded approach to RI for technology in the digital economy; practical methods for developers of autonomous systems to design accountable technologies; new modes of accountability and new concepts of liability suitable for collectives of both human and robot agent; and a first concrete example of a digital economy technology developed through a responsible innovation approach.

## The 6th Oxbridge Women in Computer Science Conference

The Oxbridge Women in Computer Science conference took place on Saturday 9 March in the Mathematical Institute with over 70 delegates in attendance, more than 90% of whom were female.

This was the sixth iteration of the Oxbridge Women in Computer Science conference, and for the first time, the conference was opened to delegates outside of the Universities of Oxford and Cambridge. We welcomed people from a number of other institutions including University College London and the University of

Warwick. We also worked closely with our new industrial sponsors and we very much look forward to working with these partners again in the future.

We had two female keynote speakers, Dame Stephanie Shirley and Professor Rebecca Eynon (Oxford Internet Institute) who spoke on 'The story of a woman in computing' and 'Coding as a "pipeline to prosperity"?' respectively. Dame Stephanie Shirley sold signed copies of her memoirs with all proceeds going to the UK charity Autistica.



© César Manivet

In addition to our keynote speakers we had eight student presentations, a small poster session and two workshops conducted by one of our sponsors, Google.

The scientific programme rounded off with a panel, lasting 75 minutes, discussing the topic 'what kinds of gender-bias exist in Computer Science workplaces?'

## Marta Kwiatkowska awarded ERC Advanced Grant FUN2MODEL

Marta Kwiatkowska has been awarded a five year 2.4m euro ERC Advanced Grant FUN2MODEL: From FUNction-based TO MModel-based automated probabilistic reasoning for DEep Learning. She is one of 222 scientists given awards in a €540 million EU investment to boost cutting-edge research.

The FUN2MODEL project aims to make advances towards provably robust ‘strong’ Artificial Intelligence. In contrast to ‘narrow’ AI perception tasks realised by deep learning, which are limited to learning data associations, and sometimes referred to as function-based, ‘strong’ AI aims to match human intelligence and requires model-based reasoning about causality and ‘what if’ scenarios, incorporation

of cognitive aspects such as beliefs and goals, and probabilistic reasoning frameworks that combine logic with statistical machine learning.

The objectives of FUN2MODEL are to develop novel probabilistic verification and synthesis techniques to guarantee safety, robustness and fairness for complex decisions based on machine learning, formulate a comprehensive, compositional game-based modelling framework for reasoning about systems of autonomous agents and their interactions, and evaluate the techniques on a variety of case studies. Marta comments, ‘I am amazed by the award of the grant. It will allow me to focus on safety and robustness for AI, which are very important topics in view of the

widespread deployment of such systems.’

ERC (European Research Council) Advanced Grants are highly competitive and provide long-term funding for exceptional leaders to pursue ground-breaking, high-risk projects. Carlos Moedas, European Commissioner for Research, Science and Innovation, said, ‘The ERC Advanced Grants back outstanding researchers throughout Europe. Their pioneering work has the potential to make a difference in people’s everyday life and deliver solutions to some of our most urgent challenges. The ERC gives these bright minds the possibility to follow their most creative ideas and to play a decisive role in the advancement of all domains of knowledge.’

FUN2MODEL is the second Advanced Grant awarded to Marta; from 2010 to 2016 she held the Advanced Grant VERIWARE. [fun2model.org](http://fun2model.org)

## New book ‘The Theory and Practice of Social Machines’ published

A new book has been published which features research from the SOCIAM project, in which the Department of Computer Science was a partner.

The term ‘Social Machines’ was introduced by Tim Berners-Lee in 1999. They comprise networks of people and devices at scale, their overall behaviour emerging from the interaction of humans and our computing infrastructure. Social Machines harness the power of the crowd, enabling anyone to potentially contribute – to generate content, cooperate on tasks, exchange information, or simply interact for work or pleasure.

Professor Sir Nigel Shadbolt, member of the Department of Computer Science, was the

Principal Investigator on the project and is a co-author of the book. He says, ‘social machines such as the global encyclopaedia Wikipedia, systems to carry out citizen science such as Galaxy Zoo or support humanitarian relief such as Ushahidi, augment our intelligence and extend our human abilities. They promise to help us address the challenges we face in the 21st century. Our book describes how these systems emerge, how they can be improved, new types designed and what opportunities and challenges they present us with’.

The book, co-authored with Kieron O’Hara, David De Roure and Wendy Hall is the fullest and most comprehensive discussion of Social Machines written to date. It appears



© Brendan Lea

as part of the Lecture Notes in Social Networks series.

The SOCIAM project ran from 2012-18 and was funded by EPSRC. The academic partners were the University of Oxford, University of Southampton, and University of Edinburgh.

Read more about the research project here: [sociam.org](http://sociam.org)

## J.P. Morgan awards

J.P. Morgan is establishing and supporting partnerships in research and innovation with universities, and the Department of Computer Science has recently received three awards from them as part of their aim to advance cutting-edge AI research to solve real-world problems.

Two Faculty Awards were granted under the scheme's 'Explainability and Interpretability' theme; the first to Professor Michael Wooldridge (PI), Doyne Farmer and Ani Calinescu to develop a data-driven stress testing methodology for monitoring profitability and risk for investment strategies; and another, to Shimon Whiteson (PI) to develop new multi-agent learning algorithms.

The PhD Fellowship was granted to Oana-Maria Camburu, one of only 14 recipients worldwide of this award. Oana's research interests lie in developing interpretable neural network models that can learn from human-provided guidance at train-time and can provide natural language explanations of their predicted decisions at test-time. Oana attended a fellowship reception event at the J.P. Morgan headquarters in New York in May 2019 which gave attendees opportunities to present their research and network with other future leaders in AI.

## Alan Turing Institute latest

The Alan Turing Institute, headquartered in the British Library, London, was created as the national institute for data science in 2015. The University of Oxford was one of the five founding universities. The Department of Computer Science currently has nine Turing Fellows and two Turing Research Fellows.

In April, Professor Michael Wooldridge was appointed as Programme Co-Director for Artificial Intelligence at the institute. He will support the development and delivery of key strategic themes, spanning research, translation/knowledge exchange, training and thought leadership.

'I am delighted to be joining the Turing at such an exciting time,' said Professor Wooldridge. 'I will be working with the Turing team to build up the institute's portfolio in AI – not just in my own area, multi-agent systems, but AI more generally. We have a really exciting portfolio of activities in the pipeline, which will confirm the UK's place at the heart of the 21st century AI revolution.'

Adrian Smith, CEO of The Alan Turing Institute, said, 'AI continues to be a thriving area of research for the Turing, focused on understanding how we can realise its positive potential for society and the economy. Michael's appointment will add to our strategic leadership and world-class talent in AI research.'

## Accelerating the Impact of CS Research

It is recognised that there is often a long lead time between research discoveries and their impact in the 'real world'. Impact Acceleration Awards (IAAs) are designed to speed up or amplify the impact arising from research. Oxford is one of 33 UK institutions since 2012 to have been awarded an EPSRC IAA grant to support activities which will reduce this lead time, and help accelerate the impact from past, current or future research that falls within EPSRC's remit.

Oxford's IAA projects support three main types of small grant: **partnerships** - to increase impact through engagement between Oxford researchers and end-users of the research who are outside academia; **impact workshops** – to support impact-focussed workshops and meetings; **technology grants** – to develop new technologies to the point where they are suitable for commercial exploitation, or to enable knowledge transfer end-

users in non-commercial contexts (e.g. the NHS).

Recent award recipients in the Department of Computer Science include:

- Professor Alex Rogers for his project 'Open Acoustic Devices' which bring low-cost smart acoustic sensors to conservation professionals.
- Professor Michael Benedikt for his project applying modern declarative data integration and query planning techniques to precision medicine applications, tailoring medical treatment to individual patients.
- Professor Ivan Martinovic for his project 'Augmented Reality Security Assistant'. IAA funding will help real-world evaluation and commercialisation of the system, making it usable in everyday settings on a range of mobile devices.

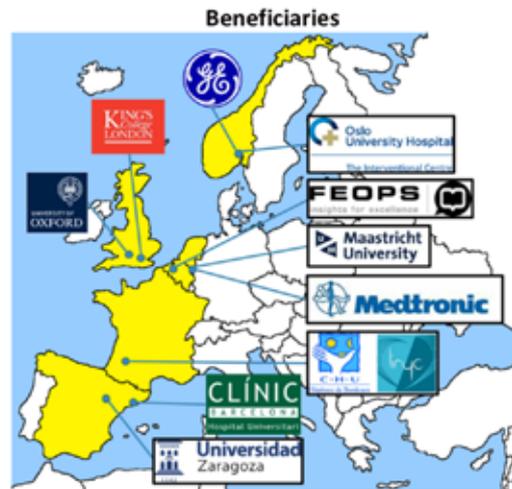
# What is it like to be a Marie Skłodowska-Curie fellow?

My name is Francesca Margara, I am a DPhil student in Computer Science, Keble College, and I am a Marie Skłodowska-Curie fellow within the Personalised In-silico Cardiology (PIC) European Innovative Training Network (ITN). PIC is training a cohort of 15 PhD students, hosted across Europe, working towards the vision of a 'Personalised In-silico Cardiology' where the management of cardiac conditions is informed by computational modelling and simulation. My individual project is focused on investigating the mechanisms explaining the cardiac safety and efficacy of pharmacological therapy through the in-silico investigation of the electromechanical response of the human ventricles in non-diseased versus diseased conditions.

In the past year, this inspiring and challenging programme has given me the opportunity to work on my project here in Oxford within the Computational Cardiovascular Science team under the supervision of Professor Blanca Rodriguez and Dr Alfonso Bueno-Orovio, and be at the same time exposed to a network of leading experts from four universities, three medical equipment companies, and three hospital trusts in Europe.

Networking and high-skilled training are two pillars of our ITN, including the opportunity of spending clinical and industrial secondment periods in our partner institutions.

This amazing research journey has already taken me to London, Oslo, Beerse, Bordeaux, and Canfranc



[Above] Participating countries in the PIC project

for summer schools, workshops, and secondments, with many more travels yet to come. I'm grateful and honoured to be part of such a network of excellence and look forward to the next two years of research.

This project has received funding from the European Union's Horizon 2020 Research and Innovation programme under the Marie Skłodowska-Curie grant agreement No 764738.



PIC students and some principal investigators in Oslo (Norway) during the first summer school

# Student Profile

**Sapna Patel:** student on our part-time professional Masters programme in software and systems security.



## What led you to choose to study software and systems security?

I wanted a change of career to work in cyber security and chose a masters degree to obtain knowledge of the area. I have worked in IT for many years and love technology but I did not have any prior cyber security experience. I was looking for a masters course with GCHQ accreditation. When following the GCHQ module pathway I have been trying to attend the more technical modules, in contrast to my most recent career roles which had become less technical in nature. Although we have the option to take software engineering modules, all my selected 10 modules have been security modules, as having many years software engineering experience, I wanted to gain new knowledge and avoid modules covering ground where I have previous experience/knowledge.

## What was your academic background before doing this course?

After focusing on science subjects during A levels, I chose a BSc in Economics at the London School of Economics for my first degree. Immediately after my BSc, I completed an MSc in Computer Science at University College London.

## How will your studies contribute to your professional career?

My IT career comprised of software development, software development management, project and programme management. The course has exposed me to an IT sector that I intend to work in and having completed

just under three quarters of the course in 18 months, I am optimistic it will take my career in a new direction soon.

## How have you found being a part-time student? How have you fitted studying into your working and home life?

The assignments require significant time and dedication to be invested in order to do well in them. I have been working through the modules quickly with some overlap between them. I tend to allocate blocks of time and prioritise assignments to complete them.

## What has been your most useful experience on the course so far?

I really like catching up with other students that I have met from previous lectures. I really wish we could do more socialising. I have enjoyed all the modules, but my favourites so far have been SIM (incident management) and MSS (mobile and IOT security): both surprised me. MSS had a large amount of content and was hard work but I learnt a lot from it. SIM surprised me as it didn't only focus on management perspectives, but was technical in nature too, if delved into in all aspects. I suspect that my last 3 modules I have yet to do will also be interesting: Design for security, Malware, and Forensics.

I have enjoyed returning to University and studying once again so much, and knew I had the right mind set to study this intense course. I enjoy the social aspects of the course and expanding my knowledge.

## Mental health initiatives

The department's commitment to raising awareness and support for mental health and fostering wellbeing has continued into 2019.

For 'Time to Talk' day in February, students and staff pedalled their way the equivalent distance from the department to the Alan Turing Institute in London – and back again – 135km on an exercise bike over six hours!

We held one of our almost-famous bake sales with healthy food options, raising nearly £300 for Mind, the Mental Health Charity.

Through regular communications about our Employee Assistance Programme, Validium, we provide an additional layer of support for staff and DPhil students 24/7, 365 days per year, by way of specialist telephone advice and/or counselling services across a range of issues we can face as part of our lives.

The department organised a range of activities for Mental Health Week, aimed at promoting positive mental health and wellbeing for everyone,

including exercise (yoga/walking), massage, healthy eating, and learning knitting together.

We now have over 30 staff who are trained as Mental Health Champions.



# New research in antibiotic resistance

Professor Peter Jeavons and recent DPhil graduate Daniel Nichol are co-authors of a Nature Communications paper which details a new approach for helping overcome antibiotic resistance. This approach combines mathematical modelling and biological experiments to determine the probability that second-line antibiotics will prove effective in treating an infection that has already gained resistance to an initial antibiotic treatment in a patient.

Antibiotic resistance represents a growing health crisis that necessitates the immediate discovery of novel treatment strategies. Developing entirely new drugs is difficult, time consuming and expensive, so it is important to try and identify ways in which existing drugs can be used more effectively. The identification of trade-offs in evolution forms the basis of an emerging strategy for combating drug resistance; prescribing sequences of drugs wherein the evolution of resistance to the first induces susceptibility to the next. Where this occurs, the first drug is said to induce collateral sensitivity in the second.

Experimental studies have attempted to identify drug pairs or sequences which exhibit this collateral sensitivity. However, drug pairs that are identified as collaterally sensitive will not necessarily induce collateral sensitivity each time they are prescribed, because evolution is driven by randomly arising mutations. Through mathematical modelling, the research team quantified the extent to which the existence of multiple evolutionary paths to resistance can lead to variable collateral sensitivities.

These model predictions were verified using high-throughput experimental evolution of *E. coli* exposed to the widely-used antibiotic cefotaxime. Exposing 60 independently derived cefotaxime-resistant strains to a panel of follow-up drugs, the team found that the sensitivity profiles were indeed highly variable.

In the case of the antibiotic ticarcillin, another commonly prescribed antibiotic, the researchers found that the efficacy can increase by as much as three times when following cefotaxime. However, such collateral

sensitivity was found to be rare and cross-resistance to be much more common, arising in 89% of the samples, and suggesting ticarcillin could prove a poor follow up drug.

‘These results clearly demonstrate the need to quantify the likelihoods of collateral sensitivity as a first step towards making informed decisions about antibiotic sequencing in the clinic’ said Daniel Nichol, the study’s lead author. ‘Unfortunately, the number of possible drug sequences is astronomical and direct testing of each sequence is likely intractable. We anticipate that improved theoretical models may be the only way to make the evolutionary predictions necessary to evaluate drug sequencing.’

The research team comprises Daniel Nichol, Peter Jeavons, Joseph Rutter, Christopher Bryant, Andrea M. Hujer, Sai Lek, Mark D. Adams, Alexander R. A. Anderson, Robert A. Bonomo and Jacob G. Scott

Read the paper here:  
[go.nature.com/2H98922](http://go.nature.com/2H98922)

## Diffblue: spin-out

Diffblue, a spin-out from the University of Oxford’s Department of Computer Science, was founded in 2016 by Professor Daniel Kroening as a way to use Artificial Intelligence to fundamentally improve how code is developed.

The company raised \$22m in Series A funding in June 2017, led by Goldman Sachs Principal Strategic Investments and Oxford Sciences Innovations (OSI). Diffblue’s first product, Diffblue Cover, intelligently writes unit tests, a time-consuming but

important process usually taken on by software development teams. This spring, Diffblue launched the Diffblue Playground, a new free tool that allows developers to experience the power of AI for themselves by generating unit tests for their imported code in seconds. This tool can help students, hobbyists, and professionals better understand their code and learn testing best practices.

‘This is our way to connect with the next generation of developers,’ Daniel Kroening explains. ‘Since we founded Diffblue, we’ve wanted to get AI for Code into the hands of as many developers as possible. With Diffblue Playground, we’re excited to show the world the power of

AI for code and the boundless potential of using next-generation technology to improve the way in which we test and develop software. We know that testing remains a real and growing pain point for developers and we anticipate significant demand from developers around the world who can use this for free.’

To try the Diffblue Playground yourself, visit:  
[playground.diffblue.com](http://playground.diffblue.com)



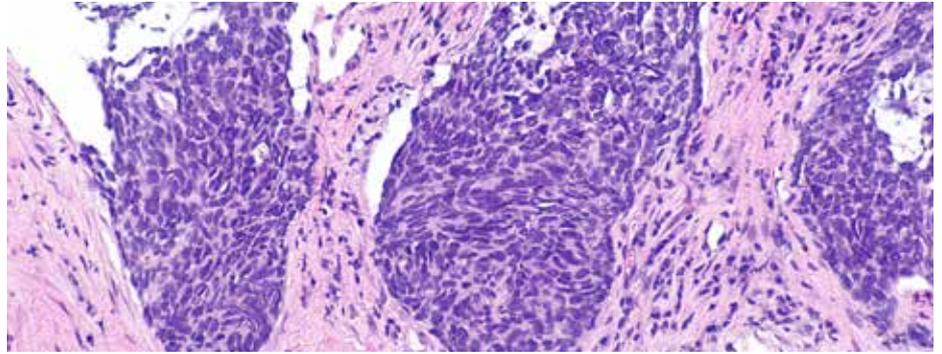
# The only way is ethics

As the frontiers of Computer Science keep expanding, we are finding that more and more research in the department involves (or has the potential to impact) human participants. To make sure that this is properly supervised, we have set up the first Departmental Research Ethics Committee in the Mathematical, Physical and Life Sciences Division.

We have had more than 60 applications in our first 18 months, covering a huge range from conventional surveys, usability studies and interviews, to more complex projects looking at questions like how parents and children understand privacy risks of tablet computers, or how online radicalisation spreads.

Since starting this process, we've learned there is much more to be done – and that there's space for Oxford to take the lead in developing organisational best practice. We are currently developing a project to put the research ethics application forms online and improve the user experience for applicants; we are also helping drive efforts across the University to put the right safeguards in place for emerging research fields like software/hardware vulnerability and data mining/pattern extraction, where ethics frameworks are still being developed. In June, we will be co-leading a workshop with colleagues in Cambridge, to help develop better understanding of ethics issues across the community of National Cyber Security Centre-recognised Academic Centres of Excellence in Cyber Security Research.

For more information on the Computer Science Departmental Research Ethics Committee (including how to apply for ethics review), please see [cs.ox.ac.uk/research-ethics/](http://cs.ox.ac.uk/research-ethics/)



## Using evolutionary game theory to identify cell resistance within cancer

Cancer is a complex adaptive system and the interaction of various cancer cell types can be described as an evolutionary game. Tumours are heterogeneous, evolving ecosystems, comprising subpopulations of neoplastic cells that follow distinct strategies for survival and propagation. We can see the interaction of these strategies as games. The success of a strategy employed by any single cell type is dependent on the distribution of other strategies, and on various components of the tumour microenvironment, such as cancer-associated fibroblasts (CAFs).

Such games and their transformations through treatment have been largely taken as a theoretical postulate in mathematical oncology. Now, the Department of Computer Science's Artem Kaznatcheev and his colleagues have developed a method for directly measuring these games, and show how the games are transformed by treatment and the presence of non-cancerous cells.

In a paper published in *Nature Ecology and Evolution*, Artem and his co-authors describe research in which they develop a game assay to measure effective evolutionary games in co-cultures of non-small cell lung cancer cells that are sensitive and resistant to the drug alectinib. They show that treatment with drugs or the elimination of fibroblasts transforms

the type of game from 'leader' to 'deadlock'. By using evolutionary game theory, this research helps to identify why some tumour cells are resistant to treatment, and quantifies how tumours that have acquired resistance are often more aggressive than before they were treated, even in the absence of drugs.

Artem's game assay measures a linear gain (of fitness) function in cell co-cultures. This gain function allows one to represent the interdependence between parental and resistant cells as a matrix game. Experimentally cataloguing these games provides an empirical grounding for existing theoretical work in mathematical oncology that considers treatment (or other environmental differences) as changes between qualitatively different game regimes.

In this framework, treatment has the goal not to directly target cells in the tumour, but instead to perturb the parameters of the game they are playing to allow evolution to steer the tumour towards a more desirable result. This gives hope for therapies that don't treat just the player (ie directly targeting cancer cells) but treat the game (ie targeting the interactions between cells).

Read the research paper here: [go.nature.com/2V759ao](https://go.nature.com/2V759ao)

# PROTECTIVE: cyber threat intelligence improving cyber security incident response

Oxford Computer Science researchers are leading a pan-European pilot to assess PROTECTIVE: a Cyber Threat Intelligence (CTI) sharing tool, and its ability to improve incident response in Computer Security Incident Response Teams (CSIRTs).

CTI is information about the latest cyber threats, threat actors, vulnerabilities, malware and indicators of compromise. This insight can help CSIRTs – security teams that protect organisations and people working for those organisations – better combat threats and attackers in incident response. CSIRTs have traditionally focused on protecting their own constituency, or sharing CTI over closed communication channels. Wider CTI-sharing across CSIRTs is becoming increasingly important because of its ability to enable security teams to work in concert against malevolent attackers. CTI sharing can enable organisations to improve automation in incident response while also considering security in context and to learn from the attacks suffered by other organisations.

PROTECTIVE is a Horizon 2020 project that aims to provide CSIRTs and Managed Security Service Providers (MSSPs) with improved security alert management and new CTI sharing capabilities through novel correlation, prioritisation, asset management, computational trust, meta alerts (summaries of similar intelligence in order to optimise delivery) as well as information-sharing compliance run-time monitoring. The project consists of ten European partners, coordinated by the Athlone Institute of Technology. Dr Jassim Happa has, over the last two years, led Oxford's contribution to the project with the requirements gathering and analysis for the PROTECTIVE tool, and designed a multi-CSIRT pilot

that is now in effect, with hundreds of thousands of CTI events being shared between CSIRTs and MSSPs across the EU. The aim of the piloting activities is to showcase and validate the PROTECTIVE tool in a variety of CSIRTs who, combined, help protect several hundred thousand data subjects across the EU. The pilot examines the tool's usability and functional capabilities in production environments. After the delivery of the pilot design, the European Commission project-assessment panel requested the design to be published as a white paper, so other H2020 projects can learn from the pilot design on how to design and conduct a multi-national pilot. The white paper is set to be published following the completion of the pilot later this summer.

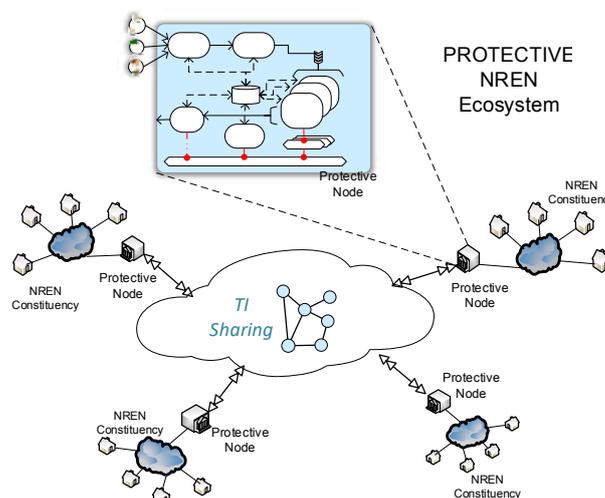
As part of the project, Oxford has also led the development of an information-sharing compliance monitoring component that provides additional assurances of data protection in CTI sharing. The tool acts as a run-time compliance checker and misuse-prevention tool should any CTI contain sensitive data that is not of legitimate interest



to CSIRTs in incident response. All outgoing CTI is automatically scanned for personal data and sensitive organisational information. The component enables analysts to specify their own rules on top of a library of template rules. To create the template rules, the Oxford team ran a series of interviews at CSIRTs across Europe and determined the requirements. This was also supported through a workshop that the Oxford team led in which CSIRT participants and researchers collaborated to solve a series of challenging information-sharing use-cases. The compliance-checking component is highly configurable and adaptable, so that it can be tailored to the specific information-sharing needs of any constituency through a straightforward configuration language. This ensures the technical implementation of policy is kept as direct as possible, facilitating review and auditing. The statistics gathered during the pilot will enable the project to further refine and develop the ruleset underpinning the component.

The PROTECTIVE tool and the information-sharing compliance component have been open sourced and can be downloaded from [bit.ly/2WXTjAJ](https://bit.ly/2WXTjAJ), with additional publications and other media at: [bit.ly/30DPBPg](https://bit.ly/30DPBPg)

by Jassim Happa



# Celebrating five years of cybersecurity capacity research impact

In February 2019, the Global Cyber Security Capacity Centre (GCSCC) held its annual conference at the Oxford Martin School, commemorating the end of its fifth operational year. The 2019 conference provided a chance for the GCSCC, its key sponsors and partners, and the global cybersecurity capacity building community to reflect on the research impact success and future challenges of cybersecurity capacity. This occasion was also marked by the launch of the Centre's Global Impact Publication exploring knowledge and policy contributions to date.

Looking back on the journey that led to this moment, in 2013 the GCSCC was created with the mission to research what constitutes cybersecurity capacity for a nation and the nature of effective capacity-building practice. This guided a research programme that led to the development of an evidence-based Cybersecurity Capacity Maturity Model for Nations (CMM). The CMM addresses these knowledge requirements, the output of which can be used to underpin the strategic investment decisions, and in doing so helps to accelerate and optimise cybersecurity capacity-building around the world.

In order to fulfil its research mission the Centre has achieved actionable and impactful knowledge contributions through strategic partnerships and the socialisation of the CMM across the public, private and civil society sectors.

Some highlights of the centre's impact on the global cybersecurity capacity community to date include:

- the completion of more than 100 CMM deployments in over 70 countries, working with national governments in all regions of the world;
- the development of the Cybersecurity Capacity Portal (<http://bit.ly/2HfXUYm>), a publicly-available online platform to disseminate knowledge generated and share best practice approaches;
- integration of GCSCC research and knowledge into sector resources such as the World Bank's Combatting Cybercrime Toolkit;
- the use of the structure of the CMM to underpin its cybersecurity capacity building programme;
- the development or revision of national cybersecurity strategies through the adoption of CMM report recommendations in countries such as the Republic of North Macedonia, Lithuania and Sierra Leone.

Building on the work of the CMM, the Centre is also in the process of developing a complementary holistic and robust model for understanding the harm experienced by nations as a result of a lack of capacity.



The Cyber Harm Framework (CHF) expands the existing CMM with a methodological underpinning, backed up by a data collection system, for relating cybersecurity capacity indicators to the areas in which harm might be reduced. The results aim to enable prioritisation of capacity investments towards harm reduction.

Reflecting on this journey and looking to the future, it is clear that this original mission of the Centre has never been more important. The world's economies continue to develop with ever-increasing dependence on technology. If we do not ensure that cybersecurity capacity exists across the entirety of cyberspace, we will inevitably develop cyber ghettos, places where



*continued on next page* ►



from previous page ►

harm is prevalent and where attacks can be successfully deployed, and also from where they can be easily launched. Ultimately, a lack of progress on cybersecurity capacity could result in harm to the prosperity and the well-being of those economies and nations so dependent on cyberspace – increasingly the vast majority of humanity.

With our partners in the global cybersecurity capacity building community, the Centre plans to tackle these challenges head on. In the next couple of years, the GCSCC is looking forward to further developing its regional partnerships. With a successful collaboration already in place with the Oceania Cyber Security Centre (OCSC) in Melbourne, Australia, 2019-20 will see the continuation

of the OCSC CMM reviews in the Pacific region, as well as the start of similar partnerships in Southern Africa and South East Asia. The regional partnerships will be pivotal to the uptake and sustainable use of the CMM and CHF around the world, to the expansion of global cybersecurity networks that support this activity, and to the multiplication of opportunities for increased collaboration in capacity-building activity internationally.

The GCSCC looks forward to advancing its research and engagement programme to promote and embed the use of the models as opportunities arise, strengthening synergies and fostering further collaboration and knowledge-exchange with more countries around the world.

GCSCC Global Impact Publication:  
[bit.ly/2WyZifg](http://bit.ly/2WyZifg)



## New role for Reuben Binns

The Information Commissioner's Office (the UK's independent regulator for data protection) has appointed Dr Reuben Binns as its first Postdoctoral



Research Fellow in AI. Since 2015, Reuben has been a postdoctoral researcher in the department's Human Centred Computing theme, and continues to work on the PETRAS project led by Professors Sir Nigel Shadbolt and Max Van Kleek.

At the Information Commissioner's Office (ICO), Reuben will be investigating a framework for auditing algorithms and will conduct further in-depth research activities. The ICO has responsibility for upholding data privacy and information rights in the public interest. Artificial Intelligence, and its impact on data protection, is one of the ICO's current priorities.

Reuben will head a team from the ICO's Technology Policy and Innovation Directorate to develop a solid methodology to audit AI applications, to ensure measures are in place to prevent data protection risks arising from them. This work will also inform future guidance for organisations to support the use of AI within the law. Risk areas covered include: fairness, transparency, accuracy, the role of human oversight, security, and impacts on rights and freedoms of data subjects.

The team will be sharing their progress through a series of blogs, which can be found at [ai-auditingframework.blogspot.com](http://ai-auditingframework.blogspot.com)



# My home is my castle: rethinking smart home IoT for privacy

As connected computational sensing devices become miniaturised and embedded throughout our everyday surroundings how will these devices change the ways we live, work, and play? While there are huge advantages that such devices will undoubtedly bring to people's health, safety, and well-being, they also could introduce other side effects, including unprecedented risks to people's privacy and autonomy.

In this context, should privacy be preserved as a fundamental human right, or should it be considered an obsolete ideal? If we take the position that privacy is fundamental, how can it be preserved in a society where so many stakeholders (including those creating devices) have so much to gain from data collection and exploitation? This problem is particularly challenging as digital devices are granted physical access to the most private and intimate settings of people's lives – our own homes, our beds, and even in our bodies. How can privacy be preserved, and ethical uses of data be ensured, in such settings?

Such have been among the fundamental questions addressed by Oxford researchers as part of the UK's PETRAS (Privacy, Ethics,

Trust, Reliability, Acceptability, and Security) of the Internet of Things project, a large, five year, multi-institution UKRI research network that featured UCL, Oxford, Warwick, Edinburgh, Southampton, Surrey, Cardiff, and Bristol at its core. PETRAS looked broadly at each of its constituent core topics in its name as interrelated streams, with dozens of sub-projects spread across one or more streams.

Professors Max Van Kleek and Sir Nigel Shadbolt led a project called *Respectful Things in Private Spaces* (ReTIPS) associated with the privacy and ethics streams. Its focus, as its name suggests, was focusing on the design needs and for Internet of Things (IoT) for the most intimate settings of people's lives: the smart home and its various intimate sub-environments, including the bedroom or bathroom, or the most intimate locations of all, on and within the human body.

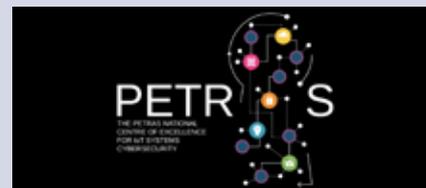
Many kinds of devices already exist on the market for both sets of settings: from voice assistants such as Amazon's Echo Dot, to smart scales, 'smart' beds, baby monitors, wearable or even implantable health monitors and other sensors. Thus far, these devices have been deployed

with minimal privacy controls or data legibility. This means that it is difficult, if not impossible, for end users to actively control how and when such devices, once activated, are capturing and disseminating data to various entities, nor how, further, such entities might store, and later access and use them.

A number of high-profile data breaches have shown that even device manufacturers are, themselves, having difficulty securing and ensuring the safe and intended uses of sensitive data captured by such devices, yet, little recourse has yet been seen. ReTIPS thus focused on potential ways to improve privacy for intimate IoT. One outcome was a system called the Privacy X-Ray that used a combination of multiple analysis methods to infer data disclosure practices by devices and apps, for the ultimate purpose of making such practices discernible and understandable to end-users.

Beyond PETRAS and ReTIPS, which ends in August 2019, Max will be Oxford Co-Investigator (jointly with Professor David De Roure, in Engineering), for a continuation project called PETRAS 2.

# UK launches National Centre of Excellence for IoT Systems Cybersecurity



As part of the UK Government's wider initiative to become a world leader in tackling cyber threats, a National Centre has been launched to provide a step-change in research associated with the Internet of Things (IoT).

The PETRAS (Privacy, Ethics, Trust, Reliability, Acceptability, and Security) National Centre of Excellence for IoT Systems Cybersecurity will provide national capability enabling the UK to become a world-leader in IoT and associated systems security. Its research focus will be on the opportunities and threats that arise when edge computing is deployed more widely in the economy and society, when Artificial Intelligence and Machine Learning technologies move from being centralised systems to being run at the periphery of the internet and local IoT networks.

Professor David De Roure (Professor of e-Research, Department of Engineering Science) leads the University of Oxford's participation in the consortium. David is joined by human-computer interaction and privacy researcher Professor Max

Van Kleek (Department of Computer Science) as PETRAS co-investigator, further developing a successful partnership which also includes Cyber Security Oxford. PETRAS 2 focuses on 'edge computing', a systems concept that has been described as a next major stage in computing beyond cloud computing, in which advanced computation, sensing, and data analysis are integrated into systems embedded within the lived physical environment, instead of sending such data in bulk back to data centres. This may end up, in the long term, being better for users by keeping sensitive data close to the source.

Max comments, 'It is both timely and important to see PETRAS continuing work on aspects of privacy, security, and resilience with a focus on computing at the periphery (a.k.a. 'the edge'), due to the increased capabilities being embedded within the world – including within smarter infrastructure and transport systems, in the home, and the workplace. Associated with the tremendous challenges with keeping such systems secure and managed, edge computing also has the potential to bring this computation within users' reach and control, as it becomes

situated within the physical and social contexts in which they live and work.'

Funded by UK Research and Innovation through the Engineering and Physical Sciences Research Council (EPSRC) as part of the Securing Digital Technologies at the Periphery (SDTaP) programme, this second phase of PETRAS will strengthen an already established and successful platform which, since 2016, has coordinated and convened 11 universities and 110 industrial and government user partners in cross-disciplinary collaboration. PETRAS has created a dynamic shared research agenda that addresses social and physical science challenges in equal measure and has worked across a broad range of technology readiness levels.

PETRAS will manage a series of open academic funding calls throughout its lifetime, the first of which will be announced in summer 2019.

Twitter: @PETRASiot



## Teaching AI new skills with Hanabi

The Department of Computer Science's Jakob Foerster is co-author on a new paper which argues that the card game Hanabi is the next great challenge for our AI agents.

Jakob worked as part of a team with researchers at DeepMind and Google Brain on 'The Hanabi Challenge: A New Frontier for AI Research'. Hanabi requires 'theory of mind' – which is about

understanding the mental states of others – just the sort of skill that humans learn as children but is far more difficult to teach AI.

Speaking to Digital Trends, Jakob comments, 'as a researcher, I have been fascinated by how AI agents can learn to communicate and cooperate with each other and ultimately also humans. Hanabi presents a unique opportunity for a grand challenge in this area, since it requires the players to reason over the intent, beliefs, and point of view of other players.'



## How capturing data in social media posts provides vital information during emergency situations

In October 2018, category 5 Hurricane Michael made landfall in Florida as one of the most intense hurricanes on record. Michael would cause over \$25 billion in damage and impact thousands of lives. As affected cities struggled to recover from the damage caused by the 160mph winds brought by Michael, intense rainfall flooded streets and buildings, hindering the efforts of relief organisations and confining many residents to their homes.

The changing landscape of a city after a disaster of this magnitude presents a key challenge to emergency response organisations. Running operations in areas where information is uncertain and existing maps have become unreliable requires significant resources dedicated to updating intelligence: observing which roads are now flooded, or which buildings have collapsed, is critical when determining a relief strategy. Typically, this information is collected from aerial observations, satellite imagery and survey teams which report back to their headquarters. However, there is an emerging source of data: civilians with

smartphones and internet-enabled devices are increasingly posting live updates to social media platforms. From photos of rising water levels, to calls coordinating their own relief efforts, these posts contain live information covering a wide geographic area which can supplement the information collected by emergency response groups.

By considering smartphone-equipped civilians in an area as a range of 'sensors', first-hand information can be gathered in real-time. For example, following Hurricane Michael, photos of a collapsed bridge might appear on social media long before official reports, allowing for the redirection of relief vehicles before the obstruction caused delays.

The primary challenge when trying to use this data is handling volume: large amounts of data are posted to social media each hour, and isolating posts created by eye-witnesses in an affected area is challenging. During times of crisis, response organisations simply don't have the resources to filter social media posts for useful information.

Research by Ross Gales in the Human-Centred Computing group at the Department of Computer Science looks at ways in which AI can handle this filtering automatically. By observing a Twitter user's unique characteristics and their social networks, an estimation can be made on their relevance to the disaster event. By focusing on relevant user accounts, content which describes the environment in an affected area may be extracted from the larger social media streams and added to existing intelligence platforms.

This project is currently being trialled by several emergency response organisations, allowing them to discover first-person perspectives of crisis events in real-time. The system allows responders to see through the 'eyes' of the geographically distributed population and identify the people and locations that need help the most. Pending the success of these trials, similar approaches will be attempted on other social media platforms and in other scenarios, such as event detection or political forecasting.

# GDPR: are we there yet? Not quite

A paper published by Sean Sirur, Jason R.C. Nurse and Helena Webb at a workshop at the 25th ACM Conference on Computer and Communications Security (CCS) highlights the struggle faced by small and medium organisations when complying with GDPR.

The EU General Data Protection Regulation (GDPR), enforced from May 2018, aimed to revolutionise how organisations handle the personal data of EU citizens. Its scope was unprecedented, threatening escalated potential penalties and prosecution across the globe for non-compliance. We investigated the challenges faced by organisations when implementing GDPR in the period shortly before the deadline, through a series of interviews. The participants ranged from entrepreneurs to national security advisors from multinational companies, each discussing the GDPR document, the compliance implementation process, and compliance verification.

While some organisations were satisfied with their compliance attempts, the results also showed that predictions were accurate concerning the difficulties faced by SMEs (small and medium sized enterprises) in compliance and the large budgets expended on GDPR. Organisations of different capabilities in terms of size and resources had a clear disparity between their attempts at compliance. Broadly, smaller companies struggled to keep up with the regulation unless already focused on security or data protection.

Larger organisations felt that GDPR compliance was achievable before the deadline but smaller organisations felt their compliance levels to be rudimentary. Participants extolled the benefits of state-issued guidance, training and education, record keeping, and (to some extent) automated monitoring. On the other hand, common problems faced when complying included deciphering the qualitative expectations of GDPR, verifying compliance without preceding case law for context, and

mapping organisational data flow. While the last of these issues is an archetypal data protection challenge, the first two have emerged in recent years. The metrics used for measuring GDPR compliance rested on the assumption that it will be assessed realistically, with participants agreeing that it was for the courts to draw the line between GDPR being pragmatic and beneficial, or bureaucratic and stifling.

Existing software engineering paradigms appear capable of handling GDPR. While several respondents were well-versed in the technical aspects of compliance, feedback on this area was still limited. While GDPR could largely be integrated into existing software engineering practices, this was not always possible in short time-scales or with limited resources (again, primarily affecting smaller organisations). This could be remedied by larger companies acting as compliant data-service providers (e.g. cloud, database services) to SMEs.

GDPR was universally regarded as a move in the right direction by these organisations, providing a much needed incentive to step back and reassess data protection practices. Individuals versed in

data protection appreciated the freedom of interpretation given by the regulation's qualitative passages. Some lauded the regulation as an example of improving pragmatism and realism in comparison to its predecessors. However, this same freedom risked alienating smaller companies which were not always equipped to readily interpret imprecise statements. Indeed, this is where external support was key. Governmental advice played a key role in clarification and organisations felt it was invaluable though rather belated. Industry involvement received mixed reactions, focusing on consultancy, training, bespoke data protection services and cyber insurance. Participants generally felt academic research to be somewhat impenetrable, highlighting a key issue to be addressed.

One year has passed since the publication of the study: necessary future endeavours would likely focus on understanding the key properties of successful compliance as confirmed by the proceeding case law. This would provide insight both into how smaller organisations can undertake mature data protection practices despite resource restrictions as well as how to improve the introduction of new regulations as a whole.



# Algorithmic biology: evolution is an algorithm so let us analyse it like one

By Artem Kaznatcheev

Everybody recognises that Charles Darwin is the first evolutionary biologist in the modern sense of the field. But I think that he was also the first algorithmic biologist, or at least he anticipated the modern movement of algorithmic biology. In particular, he saw the importance of asymptotic analysis and the role of algorithms in nature.

However before Darwin, let us read Malthus, who wrote in his *Essay on the Principle of Population*, 'Population, when unchecked, increases in a geometrical ratio. Subsistence increases only in an arithmetical ratio. A slight acquaintance with numbers will show the immensity of the first power in comparison of the second.'

Computer scientists will recognise this as an asymptotic argument. Although Malthus goes on to estimate the exact geometric ratio involved, he also recognises that the specifics are not essential to his argument, and that given any geometric factor greater than one, and any arithmetic factor, eventually the geometric growth would surpass the arithmetic. Malthus was using the fact that exponentials are little- $\omega$  of linear and then building a theory on top of this asymptotic separation.

Darwin saw the importance of this observation. He also recognised that the essential aspect of it was the asymptotic separation. It did not matter which particular resources implemented the limiting factor. More importantly, it did not even matter what specific sub-exponential function those resources scale with — just that it was sub-exponential. Darwin abstracted this principle as the

basis for his struggle for existence. Thus, he provided an abstract cause for natural selection.

Today, when we make appeals to evolution, we usually place natural selection as primary. For example, when we describe evolutionary medicine, we use terms like 'using natural selection to achieve therapeutic goals'. In other words, we use a natural process to achieve our artificial goals. Similar to how one might use the flow of the river to turn a mill.

But this was not the direction from which Darwin approached evolution. Instead, he started with domestication before moving on to variation in nature, laying out the struggle for existence, and only then — in chapter 4 of *On the Origin of Species* — did he finally define natural selection. It is only after this, late in chapter 4 that he refers to domestication as 'artificial' selection.

Structurally, Darwin's argument proceeds from looking at the selection algorithms used by humans and then abstracting it to focus only on the algorithm and not the agent carrying out the algorithm. He realised that the breeder's role as selector can be replaced by another actor: the struggle for existence. He sees the importance of the algorithm of selection and that it can be implemented in many ways. That evolution is realisable in multiple ways. It is only after we've accepted Darwin's explanation that we proceeded to reify natural selection and redefine or explain artificial selection in reference to it.

In other words, Darwin's approach was to project onto nature the human algorithm corresponding

to the actions of breeders. Like a computer scientist or mathematician today, he was using his understanding of human procedures to look at nature.

But Darwin did not mathematically formalise his insights in the way that theoretical computer scientists might today. This was probably because he did not see himself as a mathematician and even wrote in his autobiography that,

'I have deeply regretted that I did not proceed far enough at least to understand something of the great leading principles of mathematics, for [people] thus endowed seem to have an extra sense.'

So science had to wait until around the 1930s for the mathematical theory of evolutionary biology to take shape in the form of the modern synthesis. The modern synthesis developed at the same time as Computer Science. From the beginning there was interaction between Computer Science and biology — just look at Alan Turing's highly influential work on morphogenesis, or his unpublished work on neural networks and artificial life. But the connection has largely focused on tools and practical concerns: on computational biology.

There is an alternative to computational biology in theoretical Computer Science and the algorithmic lens. The algorithmic lens is not about computers or computer programmes. In the same way that astronomy is not about telescopes and that thermodynamics is not about steam engines. Rather, the algorithmic lens highlights the fact that our

*continued on next page* ►

from previous page ►

theories, models and hypotheses are algorithms in their own right. Thus, we can use the conceptual tools built by theoretical computer scientists for analysing and designing algorithms to evaluate and refine our scientific theories, models, and hypotheses.

Evolution is such an algorithm. And we can analyse it using the tools of theoretical Computer Science. This is algorithmic biology.

Whereas computational biology is a practical branch of biology, algorithmic biology is a theoretical branch of biology. Algorithmic biology is a suite of mathematical techniques taken from theoretical Computer Science and applied to the conceptual objects of evolutionary biology. A paper in computational biology might feature simulations, data crunching, and computer programmes as central characters. A paper in algorithmic biology is more likely to include theorems, lemmas, proofs and conceptual analysis.

Recently, I wrote a paper in Genetics using theoretical Computer Science to understand evolution. I showed that computational complexity can act as an ultimate constraint on evolution and thus enable open-ended evolution.

Ever since Sewall Wright's contribution to the modern synthesis, evolution is often conceptualised as a population climbing up a fitness landscape until it reaches a local fitness peak. Recent experiments have shown that these landscapes can have a rich combinatorial structure. But theoretical Computer Science

shows us that this combinatorial structure can make hill-climbing difficult, effectively transforming the mountain into a maze. On these hard fitness landscapes, no evolutionary dynamic can find the equilibrium of a local fitness peak. At least not in polynomial time.

This perpetual maladaptive disequilibrium can result in a slow power-law decay in selection strength that is consistent with ongoing long-term evolution experiments in bacteria. Such effectively unbounded growth in fitness can help us understand how evolution can continue to innovate. This is especially useful when thinking about the evolution of costly learning or the maintenance of cooperation in populations of complex organisms like humans.

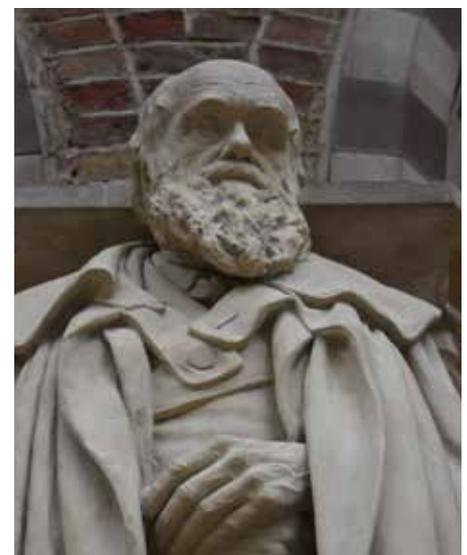
To establish results about particular evolutionary dynamics, I use techniques from the analysis of the simplex algorithm to show that strong-selection weak-mutation (SSWM) dynamics will take exponentially longer to find a fitness peak. I am able to show the more general results that abstract over-arbitrary evolutionary dynamics (not just SSWM) by relating gene-interacting networks to the weighted 2SAT-problem or more general valued constraint satisfaction problems (VCSPs). This gives us asymptotic separations and provides an abstract cause for open-ended evolution.

Now, Computer Science undergraduate Alexandru Strimbu is translating my asymptotic results into concrete simulations of evolution on hard fitness landscapes. And Peter Jeavons, David Cohen (Royal Holloway) and I are using the connection between

biological fitness landscapes and VCSPs to find the boundary that separates easy from hard fitness landscapes.

We are not the only ones looking at evolution through the algorithmic lens. Gregory Chaitin (Rio de Janeiro) and colleagues have tried to connect algorithmic information theory and evolution to 'prove Darwin' through metabiology. Leslie Valiant (Harvard), Varun Kanade and colleagues have used computational learning theory to look at evolution as a special kind of machine learning. And Adi Livnat (Haifa), Christos Papadimitriou (Columbia) and colleagues see sex as an algorithm.

As a community, we are increasingly identifying theoretical Computer Science as foundational to our understanding of the natural world. And algorithmic biology is the part of Computer Science that helps us see the foundations of life and evolution. I think that Darwin would appreciate the insights this 'extra sense' is providing.



# Factoring: searching for solutions that are sure to exist but hard to find

By Paul Goldberg

The computational problem known as factoring is stated: given a number, compute its prime factorisation. This problem is of huge importance in cryptography since some of our most cherished and widely-used cryptographic protocols rely on the belief that it is a genuinely hard problem to solve, computationally.

Factoring has a special feature: for every number, the problem has a solution (the prime factorisation of the number), and if you are told that solution, it's easy to check that it's correct.

Factoring is one of a select group of problems that are hard to solve but always have solutions. Some of these problems arise from certain mathematical theorems having appealing names such as the 'Ham Sandwich Theorem' and the 'Hairy Ball Theorem', discussed below.

Another such problem comes from game theory, namely the problem of computing a Nash equilibrium. Nash's celebrated theorem shows that any game has such a solution, but the proof is non-constructive as it does not show how to compute one, at least not with a computationally efficient algorithm. The 'gold standard' for intrinsic computational difficulty is NP-hardness, but there are reasons to believe it is not applicable to problems that have solutions that are guaranteed to exist.

The Ham Sandwich Theorem, informally stated, says that a ham sandwich (a 3D structure of bread, ham, and lettuce) can be cut in half with a plane such that all three ingredients are equally split. It dates back to a 1938 paper by Steinhaus (for details, see its Wikipedia page [bit.ly/2vVoNf0](http://bit.ly/2vVoNf0)). The computational

challenge is to find such a plane (known to exist by the Theorem), given a description of how the three ingredients are distributed in 3D space.

More generally, in  $d$ -dimensional space, for any arrangement of  $d$  ingredients there exists a hyperplane that evenly splits all  $d$  ingredients. See Figure 1 [below] for a 2-dimensional example.

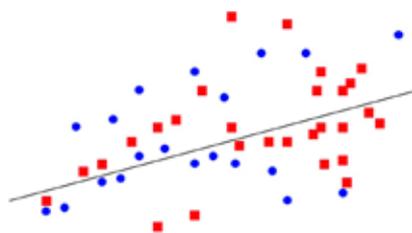


fig 1 'Ham sandwich cut'

The technical statement of the Hairy Ball Theorem is that every continuous vector field on an even-dimensional sphere has a fixed point. It has various informal statements, such as 'however you try to comb a hairy ball, there will be a point where the direction of the hairs is undefined', or 'every cow has a cowlick'.

An associated computational problem is: given a programme that computes a continuous vector field (the direction of the hairs on the ball) find a point (known to exist, by the theorem) where the direction is undefined. See Figure 2 [below] for a 2-dimensional example.



fig 2 'Hairy ball'

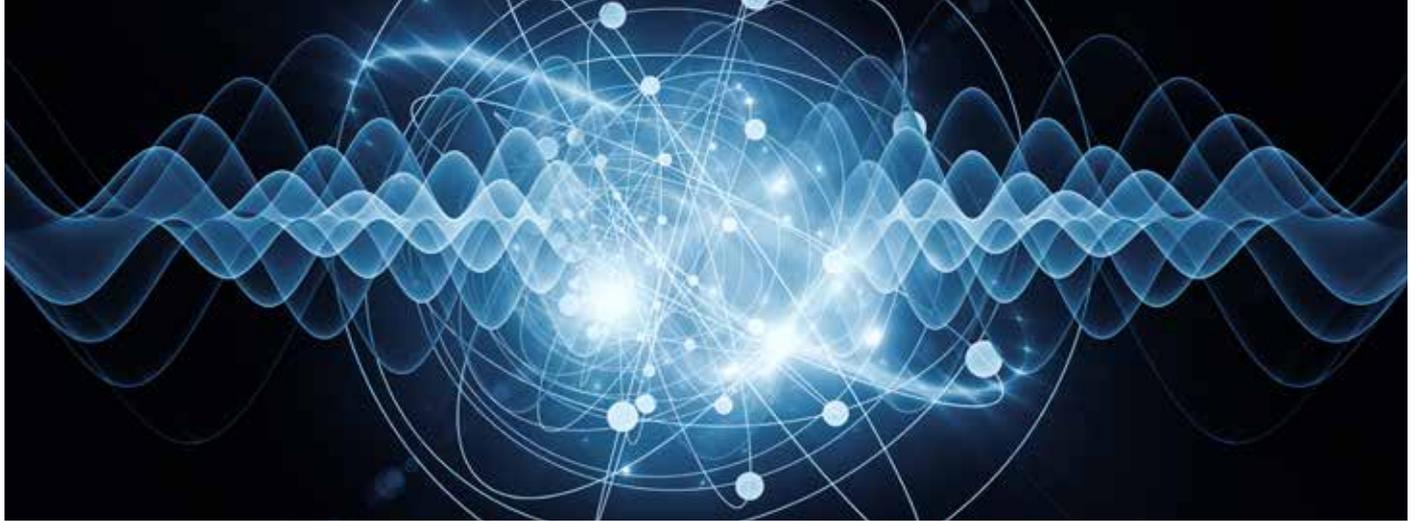
In seeking to understand the complexity of factoring and related problems, researchers have classified them according to the (often, quite simple) mathematical principles that ensure that instances of these problems always have solutions.

Some of these problems look very different but have close connections.

This author's main contributions began about 12 years ago, when (with colleagues then at UC Berkeley) we identified the computational complexity of Nash equilibrium computation. This year (with a colleague now at EPFL, Switzerland) we identified the computational complexity of the Ham Sandwich problem (along with a problem known as necklace splitting, whose computational complexity was first investigated in the 1980s).

What we currently know is that Nash equilibrium computation is equivalent in difficulty to searching for Hairy Ball solutions. The latter problem has been classified in a new paper with a doctoral student, Alex Hollender. Meanwhile, Ham Sandwich solutions and necklace splitting seem, in a strong sense, 'harder'.

It would be of considerable interest to know in more detail how factoring relates to these: one thing we know (due to the work of Emil Jeřábek at the Czech Academy of Sciences) is that in a sense, factoring is no harder than the search for Ham Sandwich cuts. In more detail, if there were an efficient algorithm for the latter, there would be an efficient randomised algorithm for factoring.



# Supervised learning with quantum-enhanced feature spaces

By Vojta Havlicek

Quantum computing studies how computation may be improved by quantum phenomena. The de-facto manifesto of this field was laid out by Feynman in 1981 when he argued that classically hard-to-simulate quantum systems should be better simulated with other quantum systems - the quantum simulators. Following the work of Oxford's own D. Deutsch that formalised quantum computation and subsequent works by Bernstein, Vazirani and Simon on quantum and classical query complexity separations, quantum computing saw a kickstart in 1994 with the invention of Shor's algorithm. This algorithm showed that quantum devices could solve problems exponentially faster than classical computers if factoring integers is classically hard.

Shor's algorithm also sparked interest in experimental implementation of the technology and small-scale NMR quantum computers followed soon. While factoring large integers on quantum computers still seems out of reach over two decades later, giant steps have been made towards the experimental implementation of quantum computing. Several research groups now possess quantum devices on the verge of demonstrating computational tasks with theoretical guarantees of hardness for classical strategies.

While being a clear milestone for quantum computing, these 'quantum supremacy' tasks are not practical and are very hard to certify.

Naturally you would ask 'what should be the next step?' I believe that it is the right time to think about quantum algorithms that could leverage the capabilities of the near-term quantum hardware. During an internship at the IBM Thomas J. Watson Research Center, whose quantum team has a strong focus in this area, I worked on a proof-of-principle implementation of a quantum classifier.

The original motivation was to benchmark the IBM quantum hardware, for which myself and my colleague Kristan Temme (IBM) designed a heuristic classification algorithm with a quantum-classical feedback. I spent several weeks verifying our ansatz numerically and achieved unexpectedly positive results, which motivated us to pursue the proposal further. Inspired by related work on classical tensor network simulations by Stoudenmire and Schwab, we found connections between our algorithm and support vector machines that helped us isolate regimes in which the algorithm resisted straightforward classical simulation. This work subsequently led to an extension of our algorithm by A. Harrow (MIT), that came with

somewhat abstract, yet strong backing for non-classicality of the implemented task.

The team at IBM decided to implement our proposal. After months of experimental work, A. Córcoles, A. Kandala and J. Chow demonstrated one of the first (if not the first) proof-of-principle algorithms for classification that could be both trained and executed on quantum devices. Besides the algorithm, a key ingredient that allowed the experimental implementation was an error mitigation technique developed earlier by Kristan and colleagues that helped the team achieve unexpectedly good results in the experiment.

The work can be seen as a small step towards implementation of pragmatic algorithms on near-term quantum hardware. We need better arguments to see if the ansatz scales, along with better guarantees for its non-classicality that would allow the proposal to become a basis for better and possibly useful algorithms. Thanks to that however, our work seems to have opened promising avenues for further research.

The results have been published in Nature (Vol. 567 Issue 7747, 14 March 2019). The arXiv can be found here: [arxiv.org/abs/1804.11326](https://arxiv.org/abs/1804.11326)

# CRISPR gene editing: a computational biology challenge

There have been recent advances in the development of genome editing techniques based on programmable nuclease proteins, allowing us to make specific changes in genomes of eukaryotic cells. By facilitating the creation of more accurate cellular and animal models of pathological processes, genome editing is enabling us to better understand the relationship between genetics and disease.

At present, the most versatile and widely used gene-editing tool is the CRISPR/Cas9 system, which is composed of a Cas9 nuclease and a short oligonucleotide guide RNA (gRNA). The latter guides the Cas9 nuclease to the targeted DNA sequence, which is subsequently cleaved.

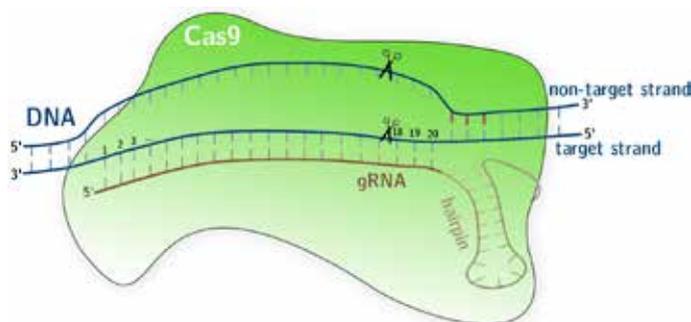


Figure 1. Schematic view of the interaction of the Cas9 nuclease (green) with the target DNA (blue) and the guide RNA (gold). The target DNA strand anneals with the matching region of the guide RNA, guiding the Cas9 protein to the indicated cutting site. Figure adapted from a publication by Alkan et al. in *Genome Biol.* 19, 177

The precise editing of the genetic material is achieved by engineering the gRNA such as to target the genetic region one wishes to edit. Even though DNA cutting is sequence-specific, off-target cutting events can occur due to various factors, including the tolerance of base pair mismatches between the gRNA and genomic DNA sequences. Full understanding of all the factors influencing the outcome of gene-editing experiments and related off-target effects remains to be established and is required for the safe application of the CRISPR/Cas9 technology.

A project pursued by DPhil student Florian Störtz working with Peter Minary uses computational modelling to expand our understanding of said factors in order to aid the design of 'optimal' guide RNAs which minimise the probability of unwanted, possibly dangerous off-target effects. By investigating rational factors (e.g. energetic, structural, mechanistic) affecting guide

RNA binding to target/off-target DNA and subsequent cutting, this research aims to build computational models to facilitate the design of optimal guide RNAs. It also includes the building of a common data storage platform to assist researchers with experiment design, allow algorithm benchmarking and thus bridge the gap between experiment and computation.

Successful execution of the project will contribute to improving safety of gene editing, ie the preservation of normal function of non-targeted genes (and proteins for which they code) in any applications of CRISPR/Cas9 gene editing technology. This could have positive implications for maintaining human health in gene therapies and for developing safe gene modification techniques for industrial biotechnology.

For further information, see [nature.com/articles/s41467-018-04252-2](https://www.nature.com/articles/s41467-018-04252-2) and [youtube.com/watch?v=2pp17E4E-O8](https://www.youtube.com/watch?v=2pp17E4E-O8).

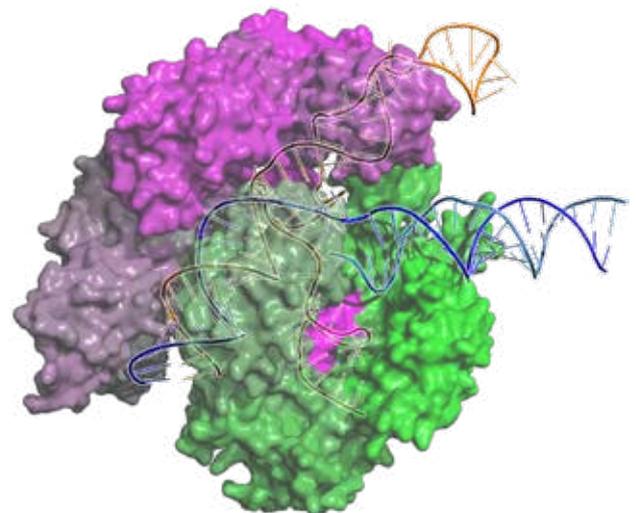


Figure 2. Crystal structure of Cas9 protein with guide RNA (gold) bound to DNA target (blue). The figure was created based on an experimental structure (PDB ID: 4UN3) capturing the CRISPR/Cas9 system in complex with its target DNA. The pink regions on the top are the HNH and RuvC domains, which cleave the target DNA. The rest of the Cas9 domains are coloured in green. The highlighted strand (dark blue) of the target DNA anneals with the guide RNA.