InspiredResearch Summer 2022 Issue 20

Research News from the Department of Computer Science, University of Oxford

Charging... At What Cost?

Read more – p18

University of Oxford

Department of Computer Science

Inspired Research

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

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

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

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

InspiredResearch
Three New Centres for Doctoral Training Come to Oxford
InspiredResearch
...vid-19 Research ...d the ...artment ...Computer ...nce
InspiredResearch
Oxford ranked first in world
Computer Science

Editorial board

Contributors

Yannis Assael
Richard Baker
Reuben Binns
Matthew Comb
Emma Dunlop
Edith Elkind
Jeremy Gibbons
Keri Grieman
Sebastan Köhler
Ivan Martinovic
Tom Melham
Leto L Riebel
Oliver Sampson
David Sheldon
Anjuli Shere
Martin Strohmeier
Carolyn Ten Holter
Jun Zhao

CONTENTS

Please Note: Photographs in the newsletter are used for illustrative purposes and may have been taken before COVID-19 restrictions came into force.

Letter from the Head of Department

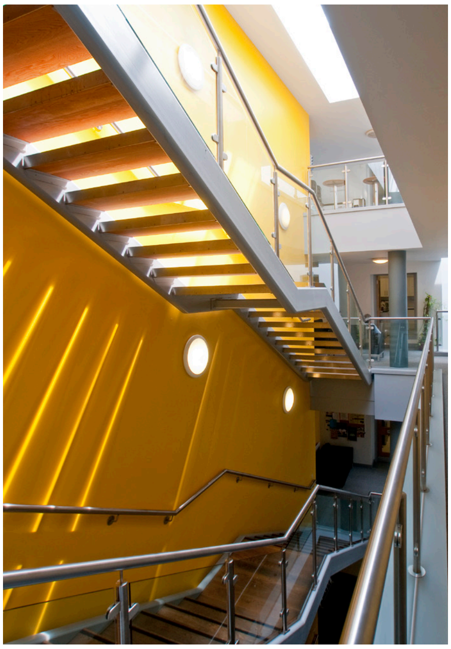

Welcome to the Summer 2022 edition of *Inspired Research*. The recent period has seen more progress with our recruitment plans and an excellent set of results in REF, as well as the many individual achievements and distinctions which you will find covered in this issue.

In January, Michael Bronstein officially started as DeepMind Professor of AI. He is pioneering next-generation AI techniques in his work on geometric and graph machine learning, which has many important real-world applications, and we are delighted to welcome him formally to the Department. Further recruitment activity is ongoing, and we look forward to welcoming more new faculty next Michaelmas term.

In May, the efforts of the whole department were recognised in our outstanding REF results. Overall, **81% of our research activity was rated 4* or 'world-leading'**, the highest possible score, and the rest was rated 3* or "internationally excellent". This is **the second-highest percentage of 4* research activity** (after Imperial College London) out of all 90 UK institutions making submissions in the REF Computer Science and Informatics unit of assessment.
Full results are at https://www.ref.ac.uk/

Overall research quality is based on assessment of research outputs, environment, and impact. The percentage assessed as 4* is the key metric that measures the quality of submitted research activity in each area.

Oxford Computer Science scored strongly in all three categories:

- **82.6% of all submitted research outputs were rated as 4* or 'world-leading'** – the second highest score awarded for outputs among the 90 submitting institutions. Only four submissions had >70% of research outputs rated 4*, and only two >80%.

- **Our research environment was awarded a perfect 100% 4* score**, one of only seven UK research institutions to achieve this rating in the CS unit of assessment.

- **Two thirds of our research impact activity was evaluated as 4*,** or **'outstanding in terms of reach and significance'**.

The results continue an upward trajectory from REF 2014, when we ranked respectively 4th by overall 4* activity, 3rd by 4* outputs, and 8th by 4* environment. They are a great reflection of the outstanding academic work we do here, and of the many and varied impacts that our research goes on to have across diverse sectors and beneficiaries. Thanks are due to the whole department, to our impact study authors (Professors Georg Gottlob, Ian Horrocks, Ivan Martinovic, Oege de Moor, Kasper Rasmussen, and Niki Trigoni), and to the departmental and divisional REF teams and other staff who contributed directly to preparing the submission.

Professor Leslie Ann Goldberg
Head of the Department of Computer Science
June 2022

Group Design Projects

As part of all undergraduate degrees in Computer Science at Oxford, students in their second year are assembled into teams of 5-6 who choose from projects devised by industry affiliates. Student groups are tasked to devote time over Hilary term to create software solutions to the problems posed, working as a team and liaising with their industrial supervisor ('the client') to ensure their solution meets the brief. The projects are a great opportunity for industry affiliates to work alongside students on an extended project, and often result in further opportunities for students to intern or work with their industrial sponsor in the future.

This year we were pleased to work with Amazon Web Services, Arribada Initiative, Apex:E3, Bloomberg, Earth Trust, Ideas Atlas, Micro:bit Educational Foundation, Microsoft and TradeTeq. Project briefs varied from those that were very much open to interpretation such as devising a way to find technical talent, to more specific goals such as the visualisation and management of satellite tracking tags.

There was a host of excellent presentations and live demonstrations of final software solutions however, the judging panel consisting of Oxford academics, Peter Minary and Joe Pitt-Francis alongside Michael Agius from G-Research and Alice Walker from Microsoft, considered two groups to be stand-out and worthy recipients of prizes.

Congratulations to Team 5 (Thomas Aston, Ciprian Florea, Emilia Folta, Godwyn Lai, Tereza Opera and Cezar Trisca-Vicol) for their project 'Real world, real time social media sentiment analysis' working with Amazon Web Services who were awarded the G-Research prize for developing a very accessible and interactive interface and giving an excellent real-time demonstration during their presentation.

Congratulations also to Team 13 (Hugo Berg, Bruno Edwards, Eloise Holland, Tobias Loader, Marcelina Marjankowska and Liam Sawyer) for their project 'Finding Technical Talent' working with Microsoft ,who were awarded the Department of Computer Science prize for giving an excellent presentation and developing innovative games based on graph problems.

The event was capped off with drinks and canapés at Keble College, and it was great to see students and industry representatives discussing their work and future opportunities together. We are always very proud of the attention to detail and quality of the final product of these design projects.

If you would like to know more about affiliation with the department and participating in next year's group design projects, please email industry@cs.ox.ac.uk. We would be happy to discuss participation and provide a full guidance document with key dates.

Science Together showcase

What happens when eight Oxfordshire community groups are given access to the world-class skills, knowledge and resources of the Oxford University and Oxford Brookes researchers? Science Together is a brand new, grass-roots programme that harnesses the power of community-led collaborative research projects to overcome challenges and seize opportunities for people who live and work in Oxfordshire.

To showcase some of the new partnerships and projects being developed, Science Together hosted 'Explore Science Together' on Tuesday 7 June 2022, an interactive day of free workshops and activities at the Oxford University Museum of Natural History.

The Science Together projects included one led by researchers from our department. Leys CDI – Leys Community Development Initiative wanted an app to help the young people they work with better connect with the services available to them. However, rather than developing this independently they have enabled the young people to lead the development – to prioritise collaborative research and user-centred design. Through workshops with Oxford's Department of Computer Science, the young people of Blackbird Leys and Greater Leys estates have co-developed and built the app for their peers.

Senior researcher Jun Zhao commented, 'We are experienced researchers working with user participants to design future technologies. However, the app club has been the first truly user-led process where we let the club students decide what they would like and how they would like to build it. It has been the most rewarding process, even though it may have taken much longer. It was wonderful for us to see how students were able to take the ownership of their project, the leadership of the design process, and create an end prototype that is most meaningful for themselves.'

Read more: bit.ly/3wMxDfd

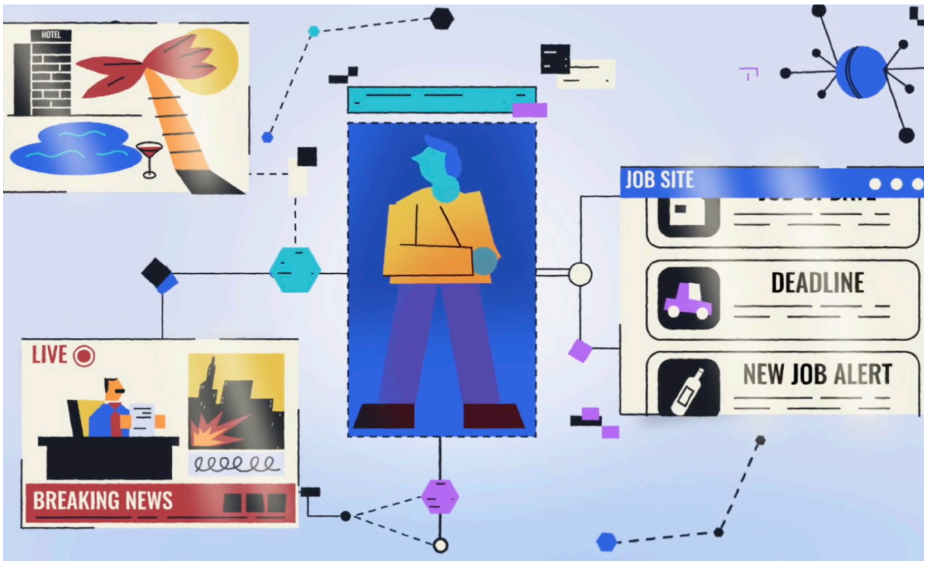

The ReEnTrust project is delighted to announce the launch of a new animated video, which aims to unpack some of the mystery in the algorithms we so often encounter, but may find it difficult to get to grips with. The video is titled 'Algorithms and Us', and draws attention to the pervasive and not-always-positive effects of algorithms on daily life. bit.ly/3ypYKh8

Dan Olteanu and Jakub Závodný win test-of-time award for 2012 paper

The International Conference on Database Theory (ICDT) presented the 2022 Test of Time Award to Dan Olteanu and Jakub Závodný. The award recognises a paper presented 10 years prior at the ICDT conference that has best met the 'test of time' and had the highest impact in terms of research, methodology, conceptual contribution, or transfer to practice over the past decade. The award was presented during the EDBT/ICDT Joint Conference in March 2022.

The 2012 paper that earned the award is 'Factorised Representations of Query Results: Size Bounds and Readability.' The paper was written while Dan was a professor and Jakub a DPhil with the Department of Computer Science.

The paper introduces the fundamental concept of factorised databases. Such factorised databases give a proper balance between succinctness of the data representation and efficiency of subsequent processing. The paper presents several contributions in terms of mathematical characterizations, algorithmic techniques, and computational complexity. The paper initiated a new line of research and had a significant impact on theoretical and systems research and development in the field of data management, in particular on graph database systems, representation and enumeration of answers to database queries and of their provenance, low computational complexity for answering quantitative queries, and training machine learning models over large databases.

This work also laid the foundations for Dan's 2015 ERC consolidator grant and his 2014 Google Faculty Award.

Together with his former Oxford research team, Dan previously won a Best Paper Award at the ICDT for their 2019 work on the first dynamic algorithm for counting triangles in graphs in worst-case optimal time.

Congratulations to the AAAI2022 award winners

The Association for the Advancement of Artificial Intelligence (AAAI) 2022 Awards were announced during the opening ceremony of the 36th AAAI Conference on Artificial Intelligence (AAAI2022). The list of AAAI award-winning papers includes two papers with Oxford University Department of Computer Science authors, Jiarui Gan and Professor Alessandro Abate.

Jiarui was co-author of *Bayesian persuasion in sequential decision-making* (recognised as AAAI-22 outstanding paper runner-up) and Alessandro was co-author of AAAI-22 distinguished paper *Sampling-based robust control of autonomous systems with non-Gaussian noise*.

This year the AAAI conference received a record 9,251 submissions, of which 9,020 were reviewed. Based on a thorough and rigorous review process 1,349 papers were accepted. This yields an overall acceptance rate of 15%.

Alan Turing Gift

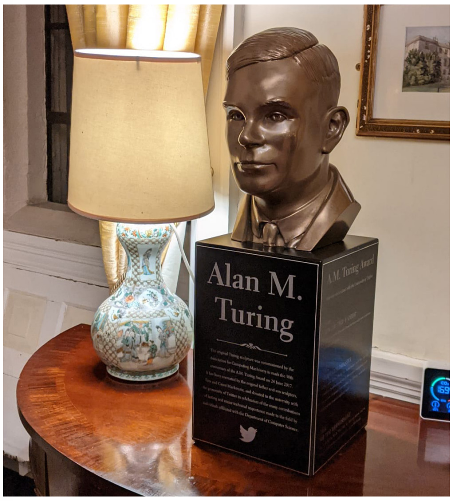

On Friday 18th February we were delighted to be presented with a kind gift from Twitter, a sculptural bust of the eminent Computer Scientist Alan Turing. It was unveiled by Head of Department Professor Leslie Ann Goldberg with department members and the sculptors present.

NEWS

News in brief

Professor Andrew Martin is part of a research team working on a secure networking by design project, with Nquiring Minds as the lead partner. This project is one of 10 which have secured £7.9 million in funding. These projects will enrich and expand the Digital Security by Design software ecosystem. Development of this ecosystem will ensure security benefits of the new, more secure processor architecture developed in the prototype Morello hardware board and can aid to underpin the broader market adoption once the technology is commercially available.

The projects, based at institutions and companies across the UK, will usher in a new age of digital security by designing software and hardware from the bottom up to be more resistant to attacks.

Senior researcher Jun Zhao and doctoral student Claudine Tinsman both made submissions to the government's call for evidence for the new Online Safety Bill. The intention of the government is that new online safety laws will make the internet a safer place for everyone in the UK, especially children, while making sure that everyone can enjoy their right to freedom of expression online.

Jun and Claudine submitted their recommendations based on the research they have been involved with in the Human Centred Computing group here at the Department of Computer Science. Their submissions have been published on the government's website.

Read Jun's submission here:
bit.ly/3bco1Ci
and Claudine's here:
bit.ly/3n772Uw

Professor Standa Zivny has been awarded a five-year €2million ERC Consolidator Grant NAASP

Professor Standa Zivny has been awarded a five-year €2M ERC Consolidator Grant NAASP: New Approaches to Approximability of Satisfiable Problems. He is one of 313 scientists given awards in a €632 million EU investment to boost cutting-edge research.

The NAASP project aims to make advances towards our understanding of approxiability of perfectly satisfiable instances of constraint satisfaction problems (CSPs), such as the approximate graph colouring problem, whose complexity status is open since 1970s.

ERC (European Research Council) Consolidator Grants are highly competitive and provide long-term funding for mid-career researchers to pursue ground-breaking, high-risk projects.

NAASP is the second ERC Grant awarded to Standa; from 2017 to 2022 he held the ERC Starting Grant PowAlgDO (Power of Algorithms in Discrete Optimisation).

Yuhang Song Awarded JP Morgan PhD Fellowship

Yuhang Song is one of 6 students worldwide (and the only one based at a UK University) to be recognised with a JP Morgan PhD Fellowship. Yuhang is a fourth-year DPhil) student in Computer Science and Nuffield Clinical Neurosciences at the University of Oxford, working with Professor Thomas Lukasiewicz and Prof. Rafal Bogacz at Intelligent Systems group and Models of Brain Decision Networks group.

Yuhang's research focuses on deciphering and extracting the learning principles of biological neural systems, so as to reverse-engineer them as algorithms or even specialised hardware. Such a route of research would, on the one hand, bring us one step closer to true artificial intelligence that facilitates our daily life, and, on the other hand, improve our understanding of the most sophisticated part of our body, the brain, so that diseases related to learning, and broadly, to neural systems, can be better understood and treated.

The J.P. Morgan AI Research Awards empower the best research thinkers across AI today – helping them to achieve their goals tomorrow.

Manuela Veloso, PhD, Head of AI Research, JPMorgan Chase & Co. comments, 'Our goal is to recognise and enable the next generation of leading AI researchers. We want to create an environment where researchers can inspire change and make a lasting impact in our communities and across our industry.'

Professor Niki Trigoni wins CTO of the year at the Women in IT Awards

Professor Niki Trigoni won CTO of the year last week at the prestigious Women in IT Awards 2022. Niki demonstrated to the judges how she has put technology at the heart of spin-out company Navenio. Having founded Navenio, Niki's work has gone a long way in reimagining the healthcare sector, particularly during the pandemic.

The Women in IT Summit and Awards series celebrates outstanding contributions to the technology industry. Now in its eighth year, the prestigious awards have recognised and celebrated over 1,000 women, allies, and organisations across the UK for their outstanding contribution to the technology industry.

This year, Niki was part of an exceptionally strong category, having been shortlisted alongside Sharon Moore of IBM, Avril Chester of RIBA, Georgiana Owens of Liberis, Hazel Olivier, Formerly of TradeCrediTech, and Joana Paivaof of intelligent Lab on Fiber Ltd (iLoF).

Success on the Forbes 30 under 30

Congratulations to Xenia Miscouridou, who has made the 7th Forbes Europe '30 under 30' list for Science and Healthcare, and to Tatiana Botskina who has made the Technology list.

Xenia Miscouridou gained her DPhil in the Department of Statistics and is now a Postdoctoral Researcher and Junior Research Fellow in the Department of Computer Science. She aims to use machine learning models for social good to influence public policy. She helped characterise a new Covid variant of concern and its impact in Brazil, published in *Science*, which was used by policymakers such as WHO.

Tatiana Botskina is a DPhil candidate in natural language processing in the Department of Computer Science. She is the cofounder and CTO at Deriskly, an AI-powered software for dispute prevention. The mission of Deriskly is to prevent disputes before they escalate and avoid litigation by detecting early signs of high risk claims. The company is currently in the process of raising its seed round. The Mathematical, Physical and Life Sciences Division has done very well this year, with 4 researchers being recognised by Forbes. The Forbes annual 30 Under 30 Europe List celebrates the latest class of young entrepreneurs, disruptors and rising stars from across Europe. The 2022 list includes 300 changemakers across ten categories, all under 30 years old, who are inspiring change and driving innovation in their respective fields across business, society and culture.

Return to the Hay Festival

2022 saw the return of the famous Hay Festival as a face-to-face rather than online event. The department has a long association with the festival and we are delighted that this year we had Associate Professor Reuben Binns speaking about his Enigma Machine project (more information about this on page 24) and Professor Mike Wooldridge talking about 'Life lessons from game theory'. Mike's talk was so popular he had to be moved to a bigger venue, attracting one of the largest audiences of the festival. Reuben was interviewed for BBC Click for their 'live from Hay' show. Doctoral student Klaudia Krawiecka lead an OxWoCS team who ran a very popular Robotics workshop for families.

News in brief

Congratulations to Andrew Markham and Phil Blunsom, who were awarded the title of 'Professor of Computer Science' in the 2020-2021 Recognition of Distinction Exercise. (The results of the 2020-2021 exercise were delayed by Covid and have only recently been announced.)

Alumna Anne-Marie Imafidon temporarily became a TV star on Channel 4's Countdown, when she did maternity leave cover for Rachel Riley. A farewell Tweet said 'You've been an absolute pleasure, joy & an inspiration'.

The Department of Computer Science is taking part in a new Graduate Scholarship Scheme for Ukraine Refugees being offered at Oxford. Thanks to the generous support of the University and its colleges, as well as donors and funding partners, up to 20 graduate scholarships for one-year full-time postgraduate courses (master's) will be awarded for the 2022-23 academic year across the University. The scholarships are open to candidates who meet the eligibility criteria and have been selected for a place on one of the eligible master's courses. These include two courses in our department, the MSc in Advanced Computer Science and MSc in Mathematics and the Foundations of Computer Science. Each scholarship will cover course fees in full and will provide £7,500 to cover living expenses for the duration of the course.
Read more about the scheme here: bit.ly/3lg7P4P

NEWS

Royal Society cautions against online censorship of scientific misinformation

Governments and social media platforms should not rely on content removal for combatting harmful scientific misinformation online, according to a report from the Royal Society.

The *Online Information Environment* report, created by a working group of leading researchers, including Oxford University computing, internet and media experts, recommends wide-ranging measures to build resilience to misinformation and a healthy online information environment.

Professor Sir Nigel Shadbolt, Professor of Computing Science, one of the working group, maintains, 'The internet has been one humanity's greatest innovations. The knowledge and information it supports and disseminates is amongst our greatest resources.'

But, he says, 'We face a torrent of misinformation on topics great and small. The report reviews the challenges of misinformation and what steps we can take to deal with them. It does not call for content removal as a panacea, rather it recommends a range of measures that governments, tech platforms and academic institutions can implement – recommendations that build resilience to misinformation and promote a healthy online environment.'

Professor Michael Bronstein, Oxford Deep Mind Professor of Artificial Intelligence and working group member, points out, 'Members of the public often lack the tools to tell authoritative sources from fictitious ones and tend to regard science as the absolute "truth" rather than a constantly evolving picture, and consequently fall victim both to honest mistakes and misreading of scientific results as well as intentional manipulation.'

The working group's report recommends a range of measures for policy makers, online platforms and others to understand and limit misinformation's harms, including:

- Supporting media plurality and independent fact-checking.
- Monitoring and mitigating evolving sources of scientific misinformation online.
- Investing in lifelong information literacy.

Read more here:
bit.ly/3FuXmMb

The open standard on the Ethical Black Box for robots presented to experts and policy makers

One of the main non-technological challenges for autonomous robots interacting with human beings is how to win public trust and acceptance.

A possible solution is coming from the EPSRC funded project RoboTIPS (www.robotips.co.uk), in which researchers are developing a device called Ethical Black Box (EBB), which is the robot equivalent of an aircraft flight data recorder. It is called 'ethical' because – according to the team - it would be irresponsible to deploy robots without one. Indeed, the purpose of the EBB is to allow accident/incident investigations, by ensuring accountability and hence improving the overall safety of robots by learning from errors.

On February 9 2022 the All-Party Parliamentary Group on Data Analytics (APGDA) in partnership with the RoboTIPS project held a roundtable in which an open standard on the EBB was presented to robotics experts and policy makers. The standard is open, namely publicly available, to invite feedback and suggestions for improvements.

The roundtable was chaired by APGDA Co-Chair Lord Holmes of Richmond and the speakers included: Lord Tim Clement-Jones, Co-Chair of the All-Party Parliamentary Group on Artificial Intelligence, Professor Marina Jirotka, Professor of Human Centred Computing at the University of Oxford and PI of the RoboTIPS project and Professor Alan Winfield, Professor of Robot Ethics at the University of the West of England, Bristol and co-PI of the RoboTIPS project.

The EBB open standard is available on: arxiv.org/abs/2205.06564

Varsity Chess win

The Oxford - Cambridge Varsity Chess match is the longest standing traditional chess event in the world. This year it was played for the 140th time at the Royal Automobile Club in London. Even though Cambridge has a slight overall lead, Oxford has done well to level the score with 4 wins in the last 5 years.

This year three members of the Department of Computer Science were playing on the top three boards for Oxford: Tom O'Gorman (first year, Computer Science and Philosophy) who is the team's highest rated player; Victor Vasiesiu (third year, Computer Science) the Oxford University Chess Club president; and Filip Mihov (fourth year, Computer Science) the Varsity team captain.

The match was incredibly close and for a long time it seemed like it could end either with a Cambridge victory, or a level 4-4 score. However, in the final minutes a complete turnaround happened and Oxford managed to secure a win 4.5 - 3.5. We're up for an interesting match next year!

CDT in Cyber Security Research Showcase 2022

On the 21 April 2022 the CDT in Cyber Security held its Annual Research Showcase at the H B Allen Centre, Keble College, Oxford. Here current students and CDT alumni were able to present their research and the event was well attended by members of the wider Cyber Security community. The day consisted of a series of talks given by current CDT students and also CDT alumni covering a wide variety of topics which included, Safer (Cyber) Spaces: Reconfiguring Digital Security as Care (Julia Slupska), Watchauth: Authentication and Intent Recognition in Mobile Payments using a Smartwatch (Jack Sturgess) and Brokenwire: Wireless Disruption of CCS Electric Vehicle Charging (Sebastian Köhler).

Posters were displayed enabling the students to discuss their work with participants and also the event provided a valuable opportunity for networking.

Mental Health First Aiders

The mental health and wellbeing of our staff and students remains a top priority for the MPLS Division and Department of Computer Science.

Mental Health First Aiders are a point of contact if you, or someone you are concerned about, are experiencing a mental health issue or emotional distress, or simply need someone to talk to.

They are not therapists or psychiatrists but they can give you initial support, listen, and signpost you to appropriate help if required.

The Department of Computer Science has five qualified Mental Health First Aiders:
Jen Lockie, Eva Nagyfejeo,
Jordan Summers, Janet Sadler,
and Jo Francis

DeepMind donation creates 12 new research internships

DeepMind, Britain's leading AI company, is boosting the University's flagship access programme, UNIQ+, by creating 12 new research internships hosted by research groups across the University's Departments of Computer Science, Engineering, Statistics and the Mathematical Institute.

Launched in 2019, UNIQ+ provides talented individuals from underrepresented and disadvantaged groups with a real day-to-day experience of postgraduate research at Oxford.

Over the programme length of 10 weeks, UNIQ+ DeepMind interns will also have the opportunity to improve their research skills, receive support and mentoring from Oxford research scientists and doctoral students, and find out more about postgraduate study and careers in AI.

DeepMind is a world leader in AI research and its application for positive impact and the donation is part of its initiative to broaden science participation including support for graduate scholarships in the Department of Computer Science.

Further information:
bit.ly/3O5bSgG

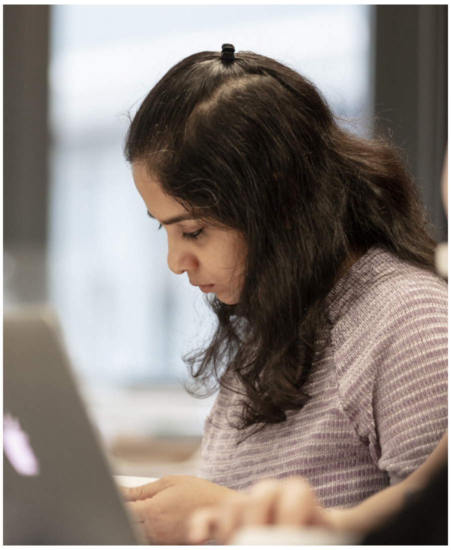

Capgemini and Oxford work together on new research

Capgemini has launched a new research project with the University of Oxford. The innovative joint research project will focus on underlying factors determining safety and trust in autonomous cars.

David Jackson, Chief Technology Officer for Product & Systems Engineering at Capgemini Engineering, will work with a team of researchers supervised by Marina Jirotka, Professor of Human Centred Computing and lead for Responsible Research and Innovation in the Department of Computer Science at the University of Oxford, on a project titled, *Being safe, feeling safe: designing, measuring and evaluating underlying factors determining safety and trust in autonomous vehicles*.

'The impacts of novel technologies on societies and individuals can sometimes get lost in the excitement of new tools and innovations. We believe this project with Capgemini gives us a great opportunity to really examine how we can keep human needs and interests at the forefront of research and development,' said Professor Marina Jirotka.

Capgemini is a global leader in partnering with companies to transform and manage their business by harnessing the power of technology. The Group is guided everyday by its purpose of unleashing human energy through technology for an inclusive and sustainable future. It is a responsible and diverse organization of over 340,000 team members in more than 50 countries.

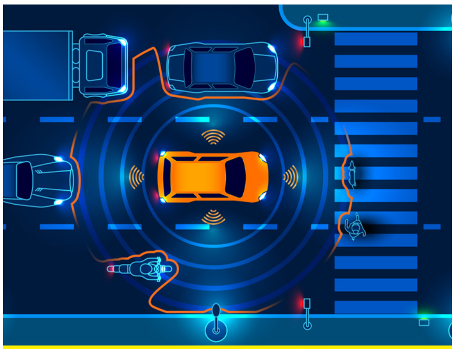

Here to support research

In the Department of Computer Science we are fortunate to have a team of seven professional staff dedicated to supporting research funding. Their aim is to enable researchers and academics to pursue their research goals, foster research collaboration and increase the department's research income.

The pre-award team (led by Jennifer Lockie, with Emma Dunlop, Thomas Palmer and Oliver Sampson) work closely with prospective applicants to discuss their research proposals and determine which funding opportunities would be most appropriate for their career stage. They have a wide variety of responsibilities: they source funding opportunities in the form of calls, competitions, fellowships, industrial funding and grants, and disseminate relevant opportunities. They also offer tailored support to applicants on how to construct a successful funding proposal and coordinate mock panels for candidates who have been called to interview. The pre-award team also produce budgets for research and as well as give PIs detailed guidance based on the funder's criteria and their professional experience of the sector. In consultation with colleagues across the University, they advise on research contracts, assess risk, and give advice on Intellectual Property (IP) and potential conflicts of interest. The team monitors the department's performance in open access ('Act on Acceptance') to ensure compliance with funder open access policies (including for REF), tracks information relating to the impact of the department's research activities, and advises on related queries. They also administer the department's Research Ethics Committee, review ethics applications and provide bespoke advice.

Oliver Sampson, the Industry Liaison Administrator develops links between industry and the department, and funnels industry approaches and specific interests into a variety of engagement routes, from undergraduate projects to DPhil funding. Oliver works alongside the Development Office and the Industry & Business Partnership team to provide wider insight into potential funding opportunities and collaborations.

Once applicants have been notified of funding, both research and non-research, the post-award team (led by Ian Watts, with Vernon Jenner and Gabrielle Alexis) work closely with the principal investigators, project managers and administrators, with regards to all the financial aspects of their projects.

This includes advising on sponsor terms and conditions, advising on the eligibility of costs to reduce the likelihood of costs being rejected by a sponsor/auditor, financial forecasting to enable the PI to make informed decisions regarding available budgets, project extensions, and budget amendments. The post-award team are also involved in the recruitment process of externally funded staff.

Working closely with Research Accounts, who are based in the Central Finance Division, the team are also responsible for all project audits and reporting of costs to sponsors.

RoaRQ: Robust and Reliable Quantum Computing

This programme, funded by a £3m grant from the Engineering and Physical Sciences Research Council, will establish a vibrant and cross-disciplinary community of researchers in quantum computing and Computer Science, who will collaborate to address the global challenge of delivering quantum computing that is robust, reliable, and trustworthy. With substantial recent progress internationally in building ever larger quantum computers, verifying that they do indeed perform the tasks they were designed for has become a central unsolved problem in the field.

From complex software articulated in high-level languages down to the silicon chips made in foundries, 60 years of computer science and engineering has defined and refined a tower of abstractions that constitute the solid foundations of today's classical computer systems. Challenges to reliability and correctness have been faced—and overcome—at many levels in the stack, and there is a wealth of insight and expertise in the diverse community of computer science researchers who work across it. Verification and testing are done at each level, with clearly defined protocols and acceptance criteria. Decades of classical computing systems research has worked out the architectures, languages and translations that bring it all together to make reliable digital systems.

Achieving reliable quantum computation faces unique challenges—not least the fragility of quantum systems due to their interactions with their environment and the fact that the state of the system during a computation cannot be measured to confirm its correctness. The very feature that makes quantum computation powerful, the exponential size of the space of states in the number of qubits, makes it hard to emulate and hence assess behaviour.

This programme will bring quantum computation research into close contact with the scientific tools, methods and (especially) mindsets of the computer science research community—across a broad spread of the key classical computing stacks. Together, they will define the beginnings of a general framework and advance specific solutions for robust and reliable quantum computation, at key layers across the principal quantum computing stacks needed to achieve trustworthy quantum computing systems.

Over the first year, the programme directors will invite engagement from across the UK's scientific community to co-create a portfolio of funded, cross-disciplinary projects that address this ambitious goal. A series of scoping workshops will be convened to propose and discuss technical directions and to facilitate the formation of project investigator teams. Projects selected for funding will commence from April 2023.

Tom Melham, Professor of Computer Science, University of Oxford

'This innovative programme, funded by the EPSRC, will create an entirely new scientific community in the UK aimed at making trustworthy quantum computing a reality. Our ambition is to seed innovation in the design of reliable quantum computing systems as far reaching as the revolution in VLSI chip design of the late 1970s and 1980s.'

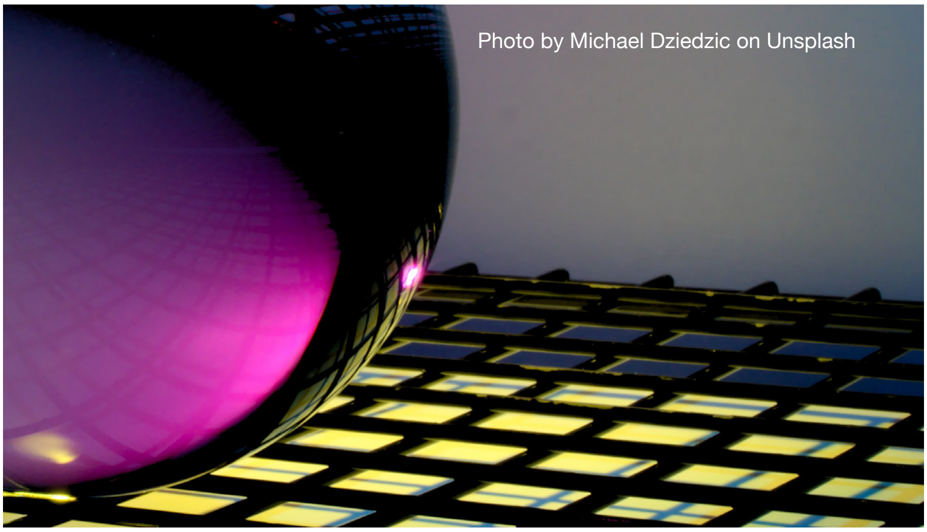

Photo by Michael Dziedzic on Unsplash

Simon Benjamin, Professor of Quantum Technologies, University of Oxford

'It's an incredibly exciting time for quantum computing, when we need people to come together from diverse backgrounds so that these machines achieve their potential as enabling tools for everyone—not just people with doctorates in quantum physics! This project is an important step in making that happen.'

Dan Browne, Professor of Physics, University College London

'I'm excited to be taking part in such an innovative research programme. Quantum computing can learn a huge amount from the know-how in the established computer science community. I am looking forward to sharing ideas with this community and building new collaborations.'

Paul Kelly, Professor of Software Technology, Imperial College London

"This is an unusual and exciting opportunity to reach out to, establish, expand and seed the network of UK computer systems and software researchers to exploit the capabilities of quantum computing—and to bridge the gap to deliver quantum-accelerated applications to realise new computational capability across diverse application domains."

Noah Linden, Professor of Theoretical Physics, University of Bristol

'At its most ambitious, our programme—with its focus on reliability and robustness—could lead to a completely new view of the quantum computing stack, with implications for hardware and software at every level.'

Many vaccine passports have security flaws

By Matthew Comb, doctoral student in the Department of Computer Science.

COVID vaccination passports have proved extremely divisive during the coronavirus pandemic, due to issues relating to civil liberties or their potential to discriminate against the more vaccine-hesitant groups within society.

But as many governments around the world push forward with their implementation in an attempt to curb the spread of COVID-19, the security of our data has become a major cause for concern.

Many COVID passes work by producing a QR code or 2D barcode for each user, which can be scanned as proof of vaccination. The barcodes used in some of these passports are not that secure because they are not generated with encrypted data. However, they can be made secure if national governments, international organisations and global tech companies work together to make the most of the exciting possibilities this technology presents.

Embedded within the barcode is a verifiable credential that proves vaccination status, and a number of personal details depending on the barcode's format. These are likely to include the user's full name and date of birth. To ensure authenticity and prevent fraud, the barcode also contains a unique digital signature, which is generated based on its contents.

International certificate of
Vaccination
Status : Immune
This is to certify that the holder of this certificate been successfully vaccinated or prophylaxis against :
Coronavirus Covid-19

A number of vaccine passport programs have already come under fire for a lack of security, including those in New York and Quebec, which have been criticised for allowing people to obtain other people's barcodes by entering their details. To mitigate some concerns, the EU has established its own open standard for vaccine passports – the EU Digital COVID Certificate (EUDCC). It has been adopted by the 27 EU states and 18 other countries.

However, this hasn't addressed the fact that the contents of the certificate are not encrypted, so anyone with access to the barcode (and the necessary skills) can decode it and retrieve the personal information contained within. This applies to COVID passports in the EU, Canada, UK, California and New Zealand. There are only slight differences in how the data is encoded – but in all these cases it is not encrypted.

To encrypt the COVID certificate's contents, there must be what's known as an encryption key associated with the certificate and the owner's digital identity. Currently, most COVID barcodes do not encrypt their contents due to the lack of digital identity infrastructure as well as the requirement to operate offline. This puts a user's personal information at risk.

There is also another problem with the current COVID certificates. They are signed by the issuer (for example the NHS) using a region- or country-specific key, or code. If someone should attain the key, they could create a false certificate. The authorities would have to respond to the fraudulent COVID passports by revoking the compromised key, which would mean that all pre-existing COVID certificates would become invalid.

Up until recently, digital identity management for a computer user has consisted of a simple username and password credential. It's a system that has worked, in the main, for more than 60 years. But the current explosion in online content, cybersecurity challenges and privacy concerns are driving the need for a user to have more control of their own digital identity.

Our identity is essentially made up of millions of small truths about ourselves. Verifiable credentials in a barcode could enable us to share just a single truth rather than our whole identity, to suit the particular situation if the data is adequately encrypted.

To its credit, the COVID certificate does just that. It is a simple proof of an individual truth, in theory enabling you to demonstrate you have been vaccinated without giving any other details away. The fact that the certificate is not entirely secure indicates the absence of a more robust digital identity infrastructure.

continued on next page ►

from previous page ▶

The absence of this piece of the digital identity puzzle must be rectified at some point in the future. Until then, the current COVID passports could be open to abuse.

The personal information involved in the vaccination certificate is not particularly sensitive at face value, because it is often easily found in other places such as a driver's license, school records or passport. But in the future, when this technology is more widespread, we will probably be using similar certificates which contain verifiable credentials in pretty much every aspect of our lives – such as to access a building or services, or to approve purchases (both instore and online).

This has positive and negative consequences for users. On the plus side, we will only need to provide the minimum amount of personal information in a very user-friendly way. For example, we will be able to sign into websites without even entering a name.

But if we present non-secure barcodes in many places, each containing small single truth about ourselves, then eventually these can potentially be combined together and the identity of the individual to whom they relate may be compromised.

This is how many cybercriminals currently work, combining data from different sources of information, which allow a person's digital identity to be constructed over time. This could lead to an increased risk of identity theft, and potentially be used as a basis for a variety of cybercrimes.

However for all these concerns about digital passports, we should remember that if it can be made secure on an international scale, this kind of digital identity technology has significant potential upside for citizens – and not just for vaccination certificates.

Originally published in *The Conversation*: bit.ly/3xNbyMD

Research exploring the 'hidden world' of proteins attracts prestigious grant

A project to harness three groundbreaking technologies developed by Oxford researchers has been awarded £5.5m by the BBSRC (Biotechnology and Biological Sciences Research Council).

Due to begin this year and continue for five years, the project will involve a multidisciplinary team led by Justin Benesch, a professor in the Department of Chemistry who will work colleagues from the Department of Chemistry (Professors Dirk Aarts, Hagan Bayley, and Philipp Kukura with Yujia Quing) and Computer Science (Professor Yarin Gal) and Zoology (Professor Kevin Foster). They will collaborate with researchers at the University of Liverpool's Centre for Proteome Research and the Wellcome Sanger Institute.

The team aims to develop and apply a novel approach for identifying proteins and their common modifications. These modifications are difficult to detect with existing technology – meaning they remain largely hidden. The new approach will help scientists understand how proteins function in health and disease, enabling improved understanding of microbial life, helping better combat infection and antimicrobial resistance.

The project was awarded a Strategic Longer and Larger grant by the BBSRC. The programme is designed to support frontier research that will address significant fundamental bioscience questions and improve our understanding of the fundamental 'rules of life'.

Expansion of AI at the Alan Turing Institute

The Alan Turing Institute has announced a significant expansion of its artificial intelligence (AI) work. This new overarching work strand will be headed up by Professor Michael Wooldridge, who has been a member of the Turing community since 2018. He currently holds a UKRI Turing AI World-Leading Researcher Fellowship and will continue to work at Oxford's Department of Computer Science whilst undertaking this new role.

Under Michael's leadership, the Turing's AI work will focus on a range of foundational areas, including fundamental AI science that complements and adds value to the existing work of the UK AI community.

Michael said: 'The UK has always led the way in understanding the foundations of AI, and I am delighted to have the opportunity to build on these strengths in the UK's national AI centre.'

How the Internet of Things poses a threat to journalists

Anjuli Shere, an analyst, writer and cybersecurity researcher discusses the IoT devices journalists might encounter at work and at home – and how these devices can threaten their work and wellbeing. Anjuli is a fellow at Harvard Kennedy School's Shorenstein Center on Media, Politics and Public Policy, while also pursuing a doctorate in Cyber Security at the University of Oxford.

In its 2021 World Press Freedom Index, which ranks countries and regions according to the level of freedom afforded to journalists, Reporters Without Borders noted that independent journalism is partially or totally stymied in 73% of the 180 countries ranked. While the press has a tendency to shy away from self-reflective coverage, there has been recent acknowledgement of the many journalists targeted by threats such as surveillance, censorship and harassment: Maria Ressa (a current Shorenstein fellow and co-founder of the Philippine news site Rappler) and Dmitry Andreyevich Muratov (editor-in-chief of the Russian newspaper Novaya Gazeta) won the Nobel Peace Prize 'for their efforts to safeguard freedom of expression, which is a precondition for democracy and lasting peace.' The Norwegian Nobel Committee noted that both winners 'are representatives of all journalists who stand up for this ideal in a world in which democracy and freedom of the press face increasingly adverse conditions.'

Some of the dangers faced by journalists are overt — the physical attacks on press when covering protests in 2020, for example — while others, like national security laws encroaching on source protections, are more insidious. Both kinds of threats can be facilitated by the so-called 'consumer Internet of Things' (the IoT): networked devices that are growing in prevalence and ability, ranging from smart cars to fitness trackers. The general risks associated with such systems have been reported on by technology and security journalists (for example, here, here, and here). Similarly, more specific examples of journalists targeted through their smartphones are scattered throughout the media and raised as issues in journalist-focused materials.

One prominent example of a high-profile smartphone-related threat is that of the NSO Group's Pegasus spyware. Clients of the Israeli technology firm, including various national government officials, reportedly identified many journalists as surveillance targets, in countries including Canada and Mexico.

In comparison, there is limited awareness of the implications of other devices, specifically for journalists and their sources, with mentions of the dangers of the IoT notably absent from safety guides for journalists. That this makes the IoT effectively an 'unknown unknown' is particularly concerning, given the ubiquity of such technologies, which can be found in homes, offices, shops — even on the street. Furthermore, they are often designed to blend into their environments, subtly replacing older versions with less intrusive functionalities — an example being the rise of the smart doorbell. Like spyware, these devices can be co-opted to monitor messages, location information and daily actions. Unlike spyware, the IoT can also facilitate cyber-physical threats.

This article outlines how journalists can begin to think about the various environments they pass through, which IoT devices they might encounter on their travels in each place, and how these devices may pose a risk to their work and wellbeing. It draws on my DPhil research, which began with a pilot study assessing the extent to which journalists recognise and understand IoT threats (spoiler: not well). My work then involved mapping IoT threats to journalists by environment (information that will be shared here, as an awareness aid). So far, my research has involved interviewing over 70 cyber security experts and journalists in the US, UK, Australia and Taiwan, and the initial findings have been presented to both computer science and public policy audiences.

continued on next page ▶

from previous page ▶

The pilot study results indicated that abstract concerns regarding technological threats are causing some journalists to move back to analogue methods of information gathering, communication and storage — like using pen and paper rather than voice recorders, and choosing physical dead drops over online ones. Cybersecurity expert recommendations spanned both immediate and long-term mitigation methods, including practical individual actions that are technical and socio-political in nature. However, all proposed individual mitigation methods are likely to be short-term solutions, as 76.5% of the 34 cybersecurity experts who participated in the study answered that within the next five years it will not be possible for the public to opt-out of interaction with the IoT.

Four categories of IoT threats

Bearing in mind the most likely journalistic workflow, my research has divided the common environments in which to consider IoT threats into four categories: (1) private homes, (2) public spaces, (3) workplace and (4) wearable devices. There is overlap between the categories — for example, many journalists' homes are also their workplaces — especially amid the pandemic and budgetary cutbacks that are closing physical newsrooms. Still, this method of categorization should enable journalists to get an initial idea of the scale of the issue, and to cross-reference relevant categories, as needed. Each of the four sections has been further subdivided by function of device, to make it easier for journalists to spot these likely poorly secured devices as they hide in plain sight.

In private homes, there are three kinds of IoT devices: those used primarily for leisure, for security, and for household management/utilities. Here, journalists should consider threats such as:

Leisure: Internet-connected children's toys are easy tools for espionage, as demonstrated by the 'My Friend Cayla' doll, the spiritual descendent of the Furby. Cayla was banned from Germany because of the ease with which hackers could access her microphone to listen in on private conversations, which could be a goldmine for discovering potential passwords (for example, children's and pet's names).

Security: A smart doorbell that someone uses to check on home deliveries from the office can be an easy way for an attacker to livestream footage of the surrounding area, including neighbouring properties, providing useful pattern-of-life information for residents in and around the owner's home – even those who do not themselves own a smart doorbell.

Household management/utilities: Voice assistants, popular for their ability to order shopping and cue mood-setting music without users needing to lift a finger, have been known to 'wake' and start listening even without 'hearing' their name, as well as to send out snippets of recordings to people on their owners' contacts lists, which could compromise confidential calls with sources or editors regarding unpublished stories.

In public spaces, there are three sub-environments where networked devices of different kinds may be found: transportation, indoor public areas, and outdoor public areas. In each, journalists should consider threats such as:

Transportation: Smart cars' GPS systems could be hacked to track the vehicle and the brakes could be hijacked to cause a crash.

Indoor public areas: Smart alarms, which are controlled remotely through smartphones and other wireless devices, can be subject to flaws and hacks that could trap people in buildings or keep people out, which could inhibit journalistic work.

Outdoor public areas: Drones can be hijacked to surveil those below, and are now commercially available in forms small and quiet enough to fly by surreptitiously — threatening even clandestine outdoor meetings with sources in areas otherwise devoid of closed-circuit TV.

In workplaces, there are three kinds of IoT devices: those used primarily for meeting/waiting room entertainment, for security, and for utilities. Here, journalists should consider threats such as:

Meeting/waiting room entertainment: Smart televisions can come with cameras, microphones, and access to online accounts that link to credit cards. These devices could easily be hacked to show the array of people meeting with newsroom executives, even before partnerships have been publicly announced, perhaps endangering the already-limited sources of funding afforded to the media.

Security: Remote-access closed-circuit TV systems could be hacked to allow continuous video-monitoring of employees at a news organization.

Utilities: Internet-connected printers could be an entry point to a network, or they could even log both the content and metadata associated with a document, enabling it to be reprinted by unauthorised users.

In all environments, journalists should consider threats stemming from wearable devices, such as:

Smartwatches and fitness trackers can perform many of the functions of a smartphone, and have other intimate purposes. If hacked, they can divulge a journalist's vital signs and tracking information, leading to the publicization of confidential locations through apps — akin to recent military debacles relating to drinking (UnTappd) and running (Strava).

What now?

Building on my pilot study, I am developing a risk assessment framework of strategic and tactical countermeasures that journalists and news organizations can use to protect themselves from these emerging technological threats.

How to investigate when a robot causes an accident – and why it's important that we do

By Research Associate Keri Grieman

Robots are featuring more and more in our daily lives. They can be incredibly useful (bionic limbs, robotic lawnmowers, or robots which deliver meals to people in quarantine), or merely entertaining (robotic dogs, dancing toys, and acrobatic drones). Imagination is perhaps the only limit to what robots will be able to do in the future.

What happens, though, when robots don't do what we want them to – or do it in a way that causes harm? For example, what happens if a bionic arm is involved in a driving accident?

Robot accidents are becoming a concern for two reasons. First, the increase in the number of robots will naturally see a rise in the number of accidents they're involved in. Second, we're getting better at building more complex robots. When a robot is more complex, it's more difficult to understand why something went wrong.

Most robots run on various forms of artificial intelligence (AI). AIs are capable of making human-like decisions (though they may make objectively good or bad ones). These decisions can be any number of things, from identifying an object to interpreting speech.

AIs are trained to make these decisions for the robot based on information from vast datasets. The AIs are then tested for accuracy (how well they do what we want them to) before they're set the task. AIs can be designed in different ways. As an example, consider the robot vacuum. It could be designed so that whenever it bumps off a surface it redirects in a random direction. Conversely, it could be designed to map out its surroundings to find obstacles, cover all surface areas, and return to its charging base. While the first vacuum is taking in input from its sensors, the second is tracking that input into an internal mapping system. In both cases, the AI is taking in information and making a decision around it.

The more complex things a robot is capable of, the more types of information it has to interpret. It also may be assessing multiple sources of one type of data, such as, in the case of aural data, a live voice, a radio, and the wind.

As robots become more complex and are able to act on a variety of information, it becomes even more important to determine which information the robot acted on, particularly when harm is caused.

Accidents happen

As with any product, things can and do go wrong with robots. Sometimes this is an internal issue, such as the robot not recognising a voice command. Sometimes it's external – the robot's sensor was damaged. And sometimes it can be both, such as the robot not being designed to work on carpets and 'tripping' Robot accident investigations must look at all potential causes.

While it may be inconvenient if the robot is damaged when something goes wrong, we are far more concerned when the robot causes harm to, or fails to mitigate harm to, a person. For example, if a bionic arm fails to grasp a hot beverage, knocking it onto the owner; or if a care robot fails to register a distress call when the frail user has fallen.

Why is robot accident investigation different to that of human accidents? Notably, robots don't have motives. We want to know why a robot made the decision it did based on the particular set of inputs that it had.

In the example of the bionic arm, was it a miscommunication between the user and the hand? Did the robot confuse multiple signals? Lock unexpectedly? In the example of the person falling over, could the robot not "hear" the call for help over a loud fan? Or did it have trouble interpreting the user's speech?

The black box

Robot accident investigation has a key benefit over human accident investigation: there's potential for a built-in witness. Commercial aeroplanes have a similar witness: the black box, built to withstand plane crashes and provide information as to why the crash happened. This information is incredibly valuable not only in understanding incidents, but in preventing them from happening again.

As part of RoboTIPS, a project which focuses on responsible innovation for social robots (robots that interact with people), we have created what we call the ethical black box: an internal record of the robot's inputs

continued on next page ▶

from previous page ▶

and corresponding actions. The ethical black box is designed for each type of robot it inhabits and is built to record all information that the robot acts on. This can be voice, visual, or even brainwave activity.

We are testing the ethical black box on a variety of robots in both laboratory and simulated accident conditions. The aim is that the ethical black box will become standard in robots of all makes and applications.

While data recorded by the ethical black box still needs to be interpreted in the case of an accident, having this data in the first instance is crucial in allowing us to investigate.

The investigation process offers the chance to ensure that the same errors don't happen twice. The ethical black box is a way not only to build better robots, but to innovate responsibly in an exciting and dynamic field.

First published in *The Conversation*

Responsible Technology Institute welcomes new Autonomous Vehicle project

The Responsible Technology Institute, led by Professor Marina Jirotka, is delighted to add another new project to its roster of research focused on responsible innovation in novel technologies.

RAILS (Responsible AI for Long-Term Autonomous Systems) builds on the work of the RoAD (Responsible AV Data) project. Both projects are part of the UK's Trustworthy Autonomous Systems Hub, which was set up by the UKRI Strategic Priorities fund to enable the development of socially beneficial autonomous systems that are both trustworthy in principle and trusted in practice by individuals, society and government. Society is seeing enormous growth in the development and implementation of autonomous systems, which can offer significant benefits to citizens, communities, and businesses. The potential for improvements in societal wellbeing is substantial. However, this positive potential is balanced by a similar potential for societal harm through contingent effects such as the environmental footprint of autonomous systems, systemic disadvantage for some socio-economic groups, and entrenchment of digital divides. The rollout of autonomous systems must therefore be addressed with responsibilities to society in mind. This must include engaging in dialogue with society and with those affected, trying to anticipate challenges before they occur, and responding to them.

To that end, these two projects have focused on the deployment of self-driving cars and other autonomous systems, as these are likely to have significant societal impacts in many respects.

RoAD looked at the challenge of data-collection by AVs, in the legal, ethical and social context of using the data for accident-investigation purposes. The project also examined civil society responses to the prospect of video-recording being carried out by AVs, as this may be a tool used in accident investigation. Our research found that not only are there gaps in the identification of vulnerable road users (for example equestrians seem to be rarely considered in the deployment of AVs), but there may be gaps between what is required by the public to enable them to trust that AVs are safe, and current standards in manufacturing and governance.

RAILS builds on this work by looking at the problem of change in autonomous systems. Autonomous systems are not designed to be deployed in conditions of perfect stasis, as they are unlikely to encounter such conditions in real-world environments. Systems that are designed to interact with humans or to operate independently will inevitably encounter unfamiliar situations for which they have not been trained - this is likely to have effects both on the system and on those engaging with it. Changing situations and environments will also affect 'learning' systems over the longer term that are in essence designed to be trained by their environment. Not only that, but these changes in deployed systems and in their operating conditions are also likely to take place against a shifting contextual background of societal alteration (eg other technologies, 'black swan' events, or simply the day-to-day operation of communities). The effects of such change, on the systems themselves, on the environments within which they are operating, and on the humans with which they engage, must be considered as part of a responsible innovation approach.

In particular, RAILS will explore how the notion of responsibility is affected by open-ended dynamic environments - situations that change over time, and lifelong-learning systems - ie systems that are designed to adapt themselves to their circumstances and 'learn' over time.

RAILS will look at social and legal contexts, as well as technical requirements, in order to assess whether and how these systems can be designed, developed, and operated in such a way that they are responsible, accountable, and trustworthy.

For more information about this and other Responsible Technology Institute projects, please visit our website at rti.ox.ac.uk

Brokenwire: Wireless Disruption of Combined Charging System Electric Vehicle Charging

Sebastian Köhler, Richard Baker and Professor Ivan Martinovic (University of Oxford) with Martin Strohmeier (armasuisse S+T)

With governments worldwide trying to reach net-zero emissions within in the next few decades, the importance and adoption of e-mobility is growing at a rapid pace.

Besides electric cars, which are a first and important step to achieve the specified goals, we will see the electrification of other sectors and vehicles, such as battlefield vehicles, emergency vehicles, buses, heavy trucks, private boats, public ferries, mining machines, and even small airplanes. At the same time, e-mobility will be contributing towards power grid stability. Bi-directional charging in combination with Vehicle-to-Grid (V2G) communication enables vehicles to act as energy storage to buffer excess energy and feed it back into the grid to meet peak demand.

With increasing reliance on electric vehicles (EVs), the security aspects have become more and more critical. The Systems Security Lab at the University of Oxford, led by Professor Ivan Martinovic, is one of the leading research groups in the area of EV charging security. Together with our partners from armasuisse Science + Technology, we conduct research to find security vulnerabilities, identify their underlying cause and develop countermeasures to prevent potential exploitation. In addition, we are working closely with industry and government bodies, to whom we disclose our findings and recommended mitigation strategies.

Our research has shown that the transition towards fully electric vehicles enables novel and unprecedented attack vectors. We discovered that one of the most widely-adopted direct current rapid charging standards across North America and Europe is vulnerable to wireless attacks. Using software-defined radios, a malicious actor can eavesdrop from afar on the charging communication and cause the charging session to abort. Since the charging cable is unshielded, an unintentional wireless side-channel that enables the radiation of electromagnetic waves exists.

Background

Charging an EV can be done with either Alternating Current (AC) or Direct Current (DC). In contrast to AC, DC charging enables high-power charging, often referred to as rapid charging, with up to 350 kW. Therefore, DC chargers have become the de-facto standard for recharging a vehicle in a short period. Nowadays, four competing DC charging standards exist — Combined Charging System (CCS), CHAdeMO, Tesla's supercharger, and GB/T 20234. For safety and efficiency reasons, all standards rely on communication with the vehicle to exchange vital information, such as maximum voltage, required current, and the State of Charge (SoC). The main difference between these standards is the technology used for the communication. The Combined Charging System is the only standard that implements power-line communication (PLC), whereas all the others use CAN.

Brokenwire Attack

Our most recent work, dubbed Brokenwire, exploits the unshielded charging cable and usage of PLC to cause a denial-of-service. It can be conducted wirelessly from a significant distance and allows individual vehicles to be disrupted stealthily, or even entire fleets to be denied charging en masse. More specifically, electromagnetic waves emitted in the same frequency range in which PLC is operating can couple onto the charging cable and cause the necessary control communication between the vehicle and the charger to fail. The loss of communication triggers an error state, forcing the charging session to abort. While the communication can continue as soon as the attacker stops transmitting, the charging session will not, leaving the car and charger in a failure mode. Once disrupted, it is necessary to manually start an entirely new charging session from scratch. In other words, the user must unplug the charging cable from the vehicle and plug it back into the charging station, wait for a short period, plug the cable into the car again, authenticate the charging via the preferred authentication method, and

continued on next page ▶

from previous page ▶

wait for the charging to start. Requiring only temporary physical proximity enables wardriving and makes Brokenwire a stealthy and hard-to-detect attack that can prevent the start of a charging session or interrupt one at any point during power delivery.

Brokenwire affects DC fast-chargers using the Combined Charging System. This includes all plugs marked as 'CCS' in public-facing chargers, along with any compliant implementation of the ISO 15118 and DIN 70121 standards (with the exception of implementations using solely ISO 15118-8 wireless communication). The other charging technologies mentioned above are not affected. While most home chargers in domestic environments use AC charging and a different communication standard (IEC 61851), which is not vulnerable to Brokenwire, the transition towards Vehicle-2-Grid communication and bi-directional charging requires the shift to higher-level communication as defined in ISO 15118 and used in DC fast-chargers.

Evaluation

We evaluated the attack under controlled laboratory conditions, using a testbed composed of two PLC evaluation boards equipped with the same chips as used in real-world deployments. On the attacker side, we used a LimeSDR , a popular and low-cost software-defined radio, as the transmitter connected to a 1W RF amplifier. We simulated the charging communication by running IPerf, an open-source software project used for network performance testing and, analysed attack success by measuring packet loss. We found that 10 mW was sufficient to cause a 100% packet loss and completely disrupt the communication of the testbed from 10m away.

Following our experiments in the lab, we also examined the effects of the attack in real-world scenarios. We tested seven vehicles and 18 chargers and found all to be vulnerable. Moreover, we observed that none of the vehicles restarted charging automatically after the attack – all had to be manually unplugged and reconnected to start charging again.

The use of electromagnetic waves as an attack vector allows the attacker to operate beyond physical barriers, without physical access and no line-of-sight to the target. To demonstrate the capabilities of Brokenwire, we simulated a multi-storey parking lot by conducting the attack between floors in a limestone building with thick walls and floors, and a ceiling height of about 3.5m. Around 100 mW output power was sufficient to disrupt the communication. While precise ranges and power budgets will vary between environments, we argue that this result amply demonstrates that the attack can be conducted beyond a physical barrier and, given the nature of the test building, in a multi-storey car park.

Conclusion

Given the prominence of CCS as a charging standard, we believe this attack represents a threat to a substantial proportion of the approximately 12M battery EVs owned worldwide. The affected DC chargers are only one, increasingly important, part of the charging infrastructure for personal cars – but are critical for high-usage fleets such as buses, HGVs, and taxis, which depend on frequent, fast recharging. Moreover, CCS is also poised to play a decisive role in the future of the power grid by enabling bi-directional charging, intertwining EVs even further into critical infrastructure.

As part of the responsible disclosure process, we informed industry and government bodies about the vulnerability, and they have independently verified and acknowledged our findings and issued CVE 2022-0878. We are now working closely with the industry on mitigation strategies. To ensure that the vulnerability is not exploited in the wild, we redacted attack details from public materials. However, comprehensive details of the evaluation and evidence of the impact are openly available, and our source code and detailed instructions on how to verify our findings will be publicly available after the embargo period.

More information can be found at:
brokenwire.fail.

Using computational tools to progress the challenge of regenerating the injured human heart

In the Computational Cardiovascular Science group, at the Department of Computer Science, our goal is to utilise computational power to investigate the response of the human heart to disease and therapy and to augment experimental and clinical investigations. In my project, I am simulating the response of the injured human heart to regenerative therapy, specifically, stem cell therapy.

Despite improvements in therapies and clinical interventions, cardiovascular disease remains the leading cause of death worldwide. Myocardial infarction, commonly known as a heart attack, is often caused by the build-up and rupture of plaques in the coronary arteries, resulting in a lack of oxygen and nutrients supplied to the heart's tissue. The affected tissue becomes necrotic and unable to contract properly, with consequent reduction of cardiac function, which can pave the way for heart failure.

After myocardial infarction, reperfusion therapies can restore blood supply and limit infarct size. However, once damage has occurred, treatment options to reverse it are very limited. Stem cell therapy is being explored to restore cardiac function by replacing damaged tissue with new lab-grown heart muscle cells. Experimental studies and early clinical trials have been investigating optimal stem cell culture, maturation, and delivery conditions. However, despite continuous progress in the field, many challenges remain. Lab-grown heart muscle cells remain largely immature, expressing electrophysiological, structural, and functional properties more similar to foetal than adult human heart muscle cells. Slowed conductivity and spontaneous beating activity are some of the most concerning immature characteristics, as once the cells are introduced into the adult human heart, they could lead to desynchronised propagation of the heart's electrical signal and ultimately hamper coordinated contraction.

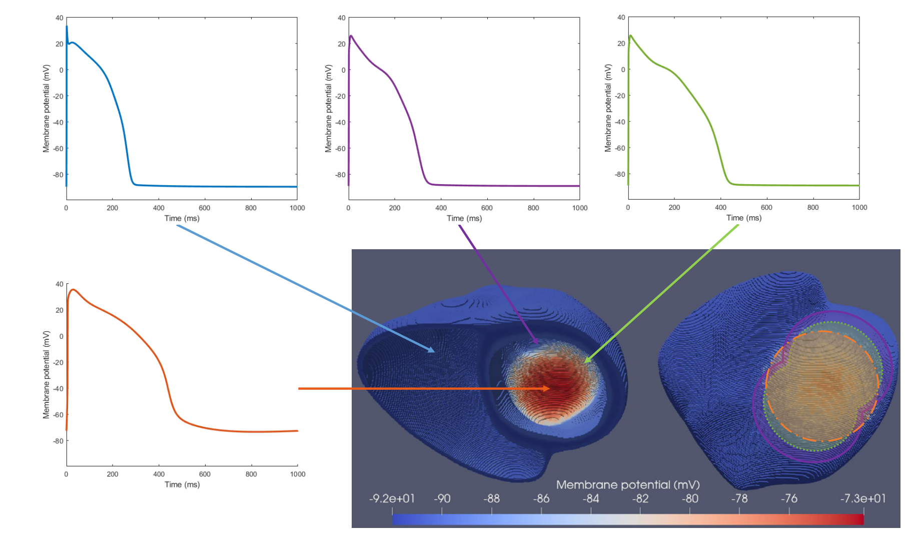

The top row shows action potential traces of the single cell ToR-ORd model (Tomek et al., 2019: bit.ly/3zZIDI6) of human adult ventricular heart muscle cells with ionic current changes following myocardial infarction applied; blue trace (top left): healthy zone, purple trace (top middle): border zone, green trace (top right): central infarct zone. The orange trace (bottom left corner) depicts the Paci2020 model (Paci et al., 2020: bit.ly/3OsMLo7) of ventricular-like human stem cell-derived heart muscle cells. The image in the bottom right corner shows the membrane potential of a 3D human-based biventricular model after myocardial infarction following the delivery of stem cell-derived heart muscle cells into the centre of the infarct. Arrows in the left half, inside view, indicate where each single cell model was applied to, with the infarct located in the left ventricle. The right half shows the epicardial surface, outside view, with border (solid purple outline), central infarct (dotted green outline), and stem-cell delivery zone (dashed orange outline).

Using computational methods, we can build detailed bio-physically accurate and multiscale models, from a single heart muscle cell to the whole human heart. Such models have shown great capabilities, for example in simulating the heart's function under disease conditions and predicting drug effects. Computational modelling and simulation provide a unique tool to control individual parameters and investigate mechanisms on a multiscale level, unfeasible in experimental or clinical settings. They can therefore add to insights from and even guide experimental and clinical studies, whilst being fast, cost-effective, and reducing the need to use animal testing.

The goal of my DPhil project, which is funded by a BBSRC Industrial CASE scholarship in collaboration with AstraZeneca, is to develop a computational modelling and simulation framework to investigate the safety and efficacy of stem cell therapy in the human heart after myocardial infarction. Offering a pivotal perspective on the development and application of disease therapies, the collaboration with AstraZeneca has been crucial in defining and driving my project forward.

During my first year, I focused on investigations at the cellular level, to identify the key model parameters characterising the differences between human stem cell-derived and healthy adult heart muscle cells. From these single cell insights, I then progressed on to introducing the stem cell-derived heart muscle cells into a three-dimensional computational model of the human heart. The 3D model, which consists of two heart chambers (left and right ventricle), also implements information acquired experimentally, such as observed regional differences in cellular properties and patient-specific features. Our 3D

continued on next page ►

By doctoral student,
Leto L Riebel

from previous page ▶

model also includes an infarcted region with specific properties, eg, slowed conductivity and ionic changes, based on experimental and clinical measurements. In total, the 3D model contains up to a few million elements, each represented by a single cell model, linked together through terms of conductivity. Solving the underlying mathematical equations in space and time is time and resource consuming. We are therefore using the fast state-of-the-art parallelised GPU model solver MonoAlg3D (see Sachetto Oliveira et al., 2018: bit.ly/3bjWKOk). Close collaboration with the developer team in Brazil has allowed us to utilise the software's full capabilities and efficiently adapt it to our projects' needs and has sparked new joint research questions.

In the coming months, I will use this computational framework to explore the safety of different stem cell delivery conditions (such as delivery location relative to the infarcted region), patient-specific characteristics (such as infarct size and progression), and ultimately suggest therapeutic targets to increase synchronicity of the heart's electrical signal and reduce the risk of abnormal activity. My project joins many others in the Computational Cardiovascular Science group in continuing to display the promising capabilities of computational tools to investigate disease mechanisms and inform clinical decisions, aimed at optimising personalised treatment of cardiovascular disease.

Madeleine Wyburd takes her work to Parliament

Doctoral student Madeleine Wyburd has been awarded the Bronze award at the STEM for BRITAIN Competition in Engineering. The awards event was held at the House of Commons on 7 March 2022. There was a poster hall, where MPs, researchers, and judges went round to view and discuss the award-winning competition submissions.

Madeleine's poster on research about using AI to analyse the developing fetal brain in ultrasound was judged against dozens of other scientists' research in the only national competition of its kind. Madeleine was shortlisted from hundreds of applicants to visit Parliament.

Madeleine comments, 'I was over the moon to be awarded the bronze engineering prize at STEM for BRITAIN for my work on using AI to analyse the developing fetal brain in ultrasound. STEM for BRITAIN is a fantastic opportunity for early career scientists to present their work to policy makers and researchers across a wide range of fields and it was an honour to be selected to take part in this year's competition.'

Chair of the Engineering judging panel, Professor Mary Ryan FREng, said, 'It is always a thrill to see the sheer variety of high-quality engineering projects and to meet so many great young researchers who want their work to make its mark and who present their work with such skill and enthusiasm.

It has never been more important for us, as engineers, to engage with policy makers and explain the ways in which our work can contribute to the UK's competitiveness and prosperity. The STEM for BRITAIN competition provides a fantastic showcase for the rising stars of engineering in the heart of Westminster.

The Parliamentary and Scientific Committee runs the event in collaboration with the Royal Academy of Engineering, the Royal Society of Chemistry, the Institute of Physics, the Royal Society of Biology, The Physiological Society and the Council for the Mathematical Sciences, with financial support from Dyson, Clay Mathematics Institute, United Kingdom Research and Innovation, Society of Chemical Industry, the Nutrition Society, Institute of Biomedical Science, the Heilbronn Institute for Mathematical Research, the Biochemical Society and IEEE UK & Ireland Section.

Madeleine with fellow winner Jorge Corral Acero (Department of Engineering, Oxford University) and members of the Parliamentary and Scientific Committee. *Copyright John Deehan Photography*

Alumni Profile

Yannis Assael – Staff Research Scientist at Google DeepMind

What course did you study here and when?

In 2013, I was accepted to study for an MSc in Computer Science, a journey that was made possible with a scholarship from Hellenic Scholarship Foundation. I could observe myself changing and improving week by week, while at the same time, making some lifelong friendships. That course changed me forever. The journey continued across classrooms, college common rooms, and libraries, and I was honoured to receive the Tony Hoare Prize for the best overall performance in my year. I continued for a second master's at Imperial College London, and I came back to Oxford for a DPhil.

What was your background before that?

I finished my undergraduate diploma at the University of Macedonia, Greece, on a joint degree between Economics and Computer Science. During my studies, I undertook over 50 freelance projects, from websites to mobile applications of TV stations and universities, and participated in a number of small research projects. These activities enabled me to create an engineering and research skill set, obtain some early management experience, and were possibly among the reasons I was given the opportunity to pursue a degree at Oxford University.

What attracted you to studying Computer Science as a subject?

My drive is creativity and to contribute and expand the greater good. From the very first moment I felt that Computer Science could enable one to solve complex problems, automate processes, and as a result, contribute to society on a larger scale, which was also what later led me to Machine Learning.

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

During my Master's, one of the leading academics of Machine Learning had just joined the department, Professor. Nando de Freitas. In Oxford, I took my first Machine Learning course, and I will always remember Nando's class and how inspired I felt by his teaching. I knew I had a passion for research, but through his courses, I realised that Machine Learning could help me strengthen my commitment to contributing and expanding to the greater good. It was one of the most beautiful gifts one can ever give - the gift of discovering what you enjoy doing the most. So, after graduation, I came back to continue my journey in Oxford for a DPhil in Machine Learning with Nando.

continued on next page ▶

Department of Computer Science Represented at Nobel Symposium on One Hundred Years of Game Theory

On December 19, 1921, the mathematician Emile Borel published a paper that laid the foundations of game theory. The framework that he proposed became a standard tool of analysis in economics, political science, as well as several other social sciences and even biology. To celebrate the centennial of this paper, the Royal Swedish Academy of Science and the Norwegian Nobel Committee organised a three-day symposium dedicated to future applications and challenges in the field of game theory. The symposium took place in Stockholm in Dec 17-19, 2021. There were 37 speakers, including five winners of the Nobel Memorial Prize in Economics: Robert Aumann, Roger Myerson, Eric Maskin, Paul Milgrom and Al Roth.

Two of the sessions were devoted to recent advances in the field of algorithmic game theory, a fast-growing research area that investigates problems at the interface of game theory and Computer Science. Professor Edith Elkind (Department of Computer Science, University of Oxford) was invited to present her work at one of these sessions. She gave a talk about her work on computing proportional outcomes in multiwinner elections, which built on the theory developed in her recent ERC grant. Other Computer Science topics that were covered during the symposium ranged from the design of blockchain protocols (Tim Roughgarden, Columbia University) to computation of equilibria in non-concave games and their connections to machine learning (Constantinos Daskalakis, MIT) and building agents that can beat word-class poker players (Tuomas Sandholm, CMU).

The tradition of Nobel symposia started in in 1965. They are usually devoted to areas of science where breakthroughs are occurring, or deal with topics of primary cultural or social significance. The game theory symposium was one of the first symposia to be held in a hybrid format, after an interruption caused by the Covid-19 pandemic. While the participants were offered the option of participating online, 20 of them (including four of the Nobel Prize winners) chose to come to Stockholm, and were rewarded by lively scientific exchanges.

from previous page ▶

What did you do when you left Oxford?

During my second year studying for my DPhil, Speech Recognition showed enormous potential but was lacking robustness in the presence of background noise. Humans use cues as mouth movement to filter out sounds. Inspired by this, our research solved Audio-Visual Speech Recognition for the first time, surpassing the performance of professionals. The value of our invention was picked up by investors, so we formed a start-up, LipNet. Soon, our research team joined Google DeepMind. Today, I am a Staff Research Scientist, and I use Machine Learning to address some of the most intriguing challenges in science and humanities.

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

During my studies at Oxford, I had the opportunity to interact with some of the world's leading academics in a plethora of fields and learn from their teaching, research, and ethics. At the same time, with my classmates and colleagues, I learnt how to collaborate as a team and to enjoy working towards a novel common goal. But, most importantly, the dreaming spires of Oxford and my friends made me realise that the quintessential part is not the academic knowledge but the feeling that, after Oxford, you can pretty much do everything.

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

My early research focused on how intelligent agents can learn to communicate and collaborate. Then in my second year, I started working on speech, which later led me to join DeepMind. In my early days at DeepMind, I would visit Oxford quite often. During lunch, I was discussing with a friend of mine doing a DPhil in Ancient History. She was describing to me some of the most difficult challenges in her field. My area of research was far from Ancient History, but immediately we saw the cooperative potential. After a few years of work, we managed to design a model for restoring ancient texts and placing them in their original place and time of writing using Machine Learning. In March, this work was on the cover of *Nature*. I have two messages to share with current students: taking on challenges and adapting quickly to different circumstances helped me pursue my dreams; at the same time being a student at such a reputable university allowed me to reach out to people, and open doors I had never imagined.

What would the student you have thought about what you are currently doing – would you have been surprised, proud, amazed?

This beautiful journey was much beyond my dreams when I was first offered a place at Oxford. I feel sincere gratitude to every day in Oxford for shaping the person I am today and providing me with a solid background able to take on any challenge while enjoying the beauty of collaboration and creativity. Most importantly, as the song says, we all get by with a little help from our friends (and supervisors), and I could not feel more thankful to them.

Professor Richard Bird Obituary

It is with great sadness that we report that Professor Richard Bird passed away on Monday 4 April, after a long illness. Richard was appointed to our Department (then known as Computing Laboratory) in 1983, initially as a non-tutorial fellow at St Cross College before moving to a tutorial fellowship at Lincoln College a few years later. He served as Director of the Computing Laboratory from 1998 to 2003, and retired in 2008.

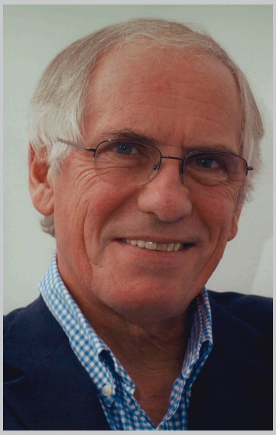

He was hugely influential in functional programming, including in the design of the Haskell language and its predecessors, and in the mathematics of program construction. He was known particularly for the elegance of his work and his writing. We will remember him as a wonderful lecturer and, particularly, public speaker.

Leaving a lasting legacy

The Development Office has launched a new MPLS (Mathematical, Physical and Life Sciences Division, University of Oxford) legacies webpage. It features a film in which the Head of Division Professor Sam Howison talks about the impact of legacies. Watch it here: bit.ly/3JBKCoe

As we look to the future, we must ensure that our research endeavour and teaching provision is sustainable for generations to come. Leaving a Legacy gift to suppor Computer Science will help support our leading research programmes and exceptional students. Whatever the size, and whether for graduate scholarships, academic positions or to support core activities, every gift is greatly appreciated and contributes to our ongoing success.

If you would like to know more about leaving a Legacy to support Computer Science, please contact Caroline Reynolds at caroline.reynolds@devoff.ox.ac.uk

Building an Enigma machine – the challenges

By Reuben Binns, Associate Professor,
Department of Computer Science

During the first lockdown of 2020, I needed a distraction. Something that took me away from the computer screen, and ideally something creative with my hands. I have been interested for some time in the history of cryptography and thought it might be fun to have a go at building an Enigma machine.

The original Enigma machine was invented in the 1920's and was used extensively by Nazi Germany during the second world war. It uses a combination of mechanical rotors and electrical circuits to scramble a secret message into a 'ciphertext' - a series of apparently random letters. When you press a letter on the keyboard, the corresponding ciphertext letter lights up on the lamp board (and vice-versa for decryption). If the ciphertext is intercepted by an adversary, they can't decrypt it unless they know the precise settings of the sender's machine. This means the message can be sent over an insecure channel like radio.

I decided to break this slightly daunting project into smaller chunks. First, I made the basic mechanism which turns the rotors around as the user types each letter, but left out all the electrical components. I used spare pieces of MDF and cut them into gears, ratchets and pawls using a handsaw. The mechanism worked to push the rotors around like the original Enigma, but only just. I realised, I wasn't going to be able to make something precise enough by hand.

So I went back to square one and began designing a new, complete electro-mechanical system. This time, I wanted to make a fully-functional replica of the Enigma, featuring rotors that actually conduct electrical current, along with the keyboard, lamp board, and plug board. For this to work, I needed to get much more precise, so rather than hand-cutting parts, I designed it all using computer-aided design software and sent it off to a workshop to have the parts cut out of thin MDF with a laser cutter. I also took advantage of something Enigma's original manufacturers didn't have: cheap, reliable modern electronic components. While there has been an ongoing global shortage of integrated circuits since the pandemic, these components are so ubiquitous I had no trouble sourcing them.

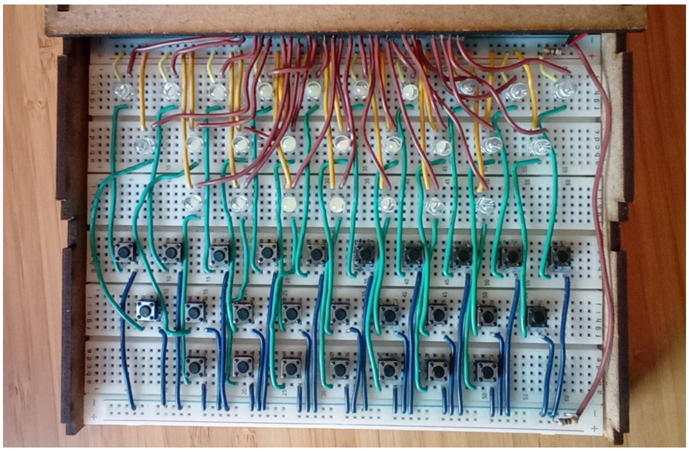

The parts all slot together without any glue or screws, so it can be put together by hand. It does require some patience though; each of the 3 rotors contains 52 contact points which need to be connected by jumper cables and take around an hour to assemble. I also had to re-design the keyboard, lamp board and plugboard several times to overcome the jumbled mess of crossed wires.

This project was a good way to unwind after work, but also prompted me to reflect on my research.

First, I teach a course on the human side of computer security, and one of the things I stress to students is that a system can be secure in theory, but to be secure in practice it needs to be usable. Just as with modern computers, people struggle to remember dozens of complex random passwords, the operators of the Enigma machine didn't always follow the correct protocol; this weakness was exploited by Bletchley Park.

Second, this project gave me a hands-on understanding of the power and perils of electromechanical computing. Modern digital computers represent the world in terms of 1s or 0s, and nothing in between. The Enigma machine requires the position of the rotor to be in exactly 1 of 26 positions, and nothing in between. The difficulty is that most of the physical world is continuous, so things are generally along a sliding scale between states. This makes it hard to model discrete states exactly. We overcame this challenge with modern computers by using transistors, which force electrons to either flow (representing a 1) or not (representing a 0). But when you use a mechanical component to represent a state, like with the Enigma machine, you need very precise engineering. If the rotors turn a little bit too far, and slip out of place relative to each other, the machine will eventually make an error. I learned this lesson the hard way, giving me a newfound respect for the designers of mechanical and electro-mechanical computers, and gratitude for modern transistors that we all depend on today.

I'm now designing the second version of this machine which I hope will be more robust, reliable, and simpler to put together.