InspiredResearch Summer 2020 Issue 16

News from the Department of Computer Science, University of Oxford

Covid-19 Research and the Department of Computer Science

AI Deep Learning:
Restoring Greek inscriptions – p27

Behind closed doors:
How do pilots handle cyber attacks – p20

The 'Digital Twin':
Enabling the vison of precision cardiology – p28

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

Editorial board

Kiri Walden (Editor)
Suvarna Designs (Designer)

Sarah Baldwin
Leanne Carveth
Rosanna Cretney
Emma Dunlop
Stefano Gogioso
David Hobbs
Suzanna Marsh
Elisa Passini
Alex Rogers
Helena Webb
Standa Živný

Contributors

Yannis Assael, Alfonso Bueno-Orovio, Alexis Ciambotti, Michael Clerx, Bob Coecke, Bill Dutton, Edith Elkind, Arnau Erola, Paula Fiddi, Eoghan Flanagan, Alan Garfinkel, Jeremy Gibbons, Georg Gottlob, Manuel Hepfer, Ian Horrocks, Philip Inglesant, Andrew Martin, Ivan Martinovic, Tom Melham, Dominik Peters, Jonathan Prag, Blanca Rodriguez, Matthew Rogers, Yashovardhan Sharma, Matt Smith, Thea Sommerschield, Oana Tifrea-Marciuska, Max Van Kleek, Helena Webb, Carolin Weisser Harris, Michael Wooldridge, Jun Zhao.

CONTENTS


p27

Letter from the Head of Department

This is certainly the most unusual 'Letter from the HoD' I have ever had to write, and I have to say I very much hope I will never have to write another like it. On Thursday 19 March the Department of Computer Science physically closed its doors and, like most of our Oxford University colleagues we are now operating in exile, working from home until we are sure it is safe to return. From first hearing about Covid-19 to having to abandon our offices and set up make-shift arrangements in our homes, our lives, like those of so many people around the world, have been changed in a dizzyingly short period of time. I hope everyone is well and coping as best as can be expected with work,family, and loved ones etc. I realise that for many staff with young families or dealing with other caring duties these are particularly difficult times, and we as a department realise it can't be 'business as normal' in such circumstances.

I would like to again thank everyone for their hard work in transitioning to a remote working mode. I wouldn't describe the transition as 'seamless', but for the most part, it has been remarkably smooth. Whilst the number of weeks ahead is still relatively unknown we want to be sure that our staff and students are well and that no one feels increasingly isolated/alone. I don't want a single member of the department to feel they are being ignored and so staying in touch is now part of our responsibilities as tutors, supervisors, and line managers. Please look after each other, be sensitive and supportive.

Apart from our move from office to home working, the second big challenge has been arranging, at very short notice, the transition from being a tutorial university (that really is our unique selling point, after all) to being an online university. We have worked hard to find a model for assessment that scales to students spread across the globe, with wildly different personal and working circumstances. We have done our very best to develop systems that are fair to all concerned, while maintaining suitable academic standards. Through all of this, our key principles have been:

1. Safety and wellbeing: of all students and staff.
2. Continuity: we aim to teach our students, to assess their work, and to see them graduate successfully and on schedule as far as reasonably possible.
3. Fairness: no student should be disadvantaged by the current circumstances.

I want to take this opportunity to say a big thank you to all members of the Department of Computer Science for calmly getting on with whatever they can at a time of extraordinary disruption. The quick transition to home working has made our task in keeping things running much smoother than might otherwise have been the case – for which we are very grateful.

Finally, one piece of non-Covid-19 news, Louise Botley has been appointed as Head of Administration and Finance and joined the department in January. She was previously Head of Research Facilitation in the Humanities Division. Welcome Louise! I'm so sorry your first few months may not have been quite what you were expecting… it isn't always like this, honestly.

With my heartfelt best wishes to you all: stay safe!

Professor Michael Wooldridge
Head of Department of Computer Science
June 2020

News in brief

DPhil student Bushra Al-Ahmadi spoke as part of Google Mena Digital Days on 'Providing a Safe Cyber Space for Women'. She spoke about the threats women face online and how we can build technologies that protect them. bit.ly/2UzzLTk

The termly Strachey Lecture took place on 19 November 2019 with Dame Wendy Hall (University of Southampton). The talk was titled 'The Windmills of Your Mind: Reflections of a career in Computer Science research'. bit.ly/2VnwCYA

Klaudia Krawieka, a DPhil student in Cyber Security, was shortlisted for the University of Oxford's Vice-Chancellor's Diversity Awards for her exceptional work throughout her studies as an advocate for diversity and women's representation in the department and the wider university.
Read more: bit.ly/3ebbgnS

Head of Department, Professor Michael Wooldridge gave a talk for the BBC radio programme *Naked Scientists* on the limitations of AI. bbc.in/3agPKM0

On 13 March, Head of Department, Professor Michael Wooldridge spoke against the topic 'We are the Architect of our own Demise' at the Oxford Union. Michael was the last of six speakers, with the team winning the debate. bit.ly/34W25nJ

NEWS

Niki (centre) receiving her award.
Photo courtesy of Women in IT Awards

Professor Niki Trigoni wins CTO of the year award at Women in IT Awards London 2020

The winners of the sixth annual Women in IT Awards London 2020 were revealed at a gala ceremony in Grosvenor House Hotel, Park Lane on 29 January 2020. Professor Niki Trigoni won her award in recognition for her role as (founder and) Chief Technical Officer at spinout company Navenio.The awards, organised by Information Age and DiversityQ in partnership with Amazon Web Services, serve to showcase the achievements and innovation of women in technology and identify new role models in a sector where female representation stands at only 19%.

Read more about the awards here: bit.ly/2VPOSZH

Professor Georg Gottlob awarded Royal Society Research Professorship

Professor Georg Gottlob has been awarded a highly prestigious Royal Society Professorship for a long-term research project titled 'RAISON DATA Rule-Based AI Systems for Reasoning on Massive Data'. The overarching goal of this research is to significantly contribute to next-level AI by finding methods, and building systems, that perform hybrid problem-solving tasks automatically, based on symbolic expert knowledge combined with machine learning over massive amounts of available data.

Professor Gottlob said, 'I am very excited to be awarded this Research Professorship. It will allow me to continue my research on efficient rule-based reasoning and study how this method can be combined with machine learning. The Royal Society Professorship allows me to fully concentrate on this research, and to start building a new team. There is no better place for this project than the Department of Computer Science at Oxford. The department is already a top research hub for both machine learning and rule-based knowledge representation and reasoning. The time has now come to study how these two approaches can be combined fruitfully. I hope we will make significant progress over the years ahead.'

Professor Sam Staton awarded ERC Consolidator Grant

Professor Sam Staton has been awarded a prestigious Consolidator Grant, by the European Research Council for a five-year research project 'Better Languages for Statistics (BLaSt)'. This is a project to investigate probabilistic programming languages, which are programming languages for rapidly developing statistical complex models with advanced inference. The research has three ultimate goals: to propose new probabilistic programming languages; better languages for statistics; to devise new general inference methods for probabilistic programs; to build new foundations for probability.

Sam has also recently received a three-year extension to his Royal Society University Research Fellowship titled 'Principled Foundations for Probabilistic Programming'. Amongst other things, this further funding will enable him to continue his research to develop a method called 'synthetic measure theory'. Synthetic measure theory is a way of transporting standard ways of reasoning in probability theory, such as integral notation, to different mathematical universes. This is a foundational method but it has applications to new programming languages.

Sir Tony Hoare named as International Federation for Information Processing Fellow

Emeritus Professor Sir Tony Hoare has been recognised by International Federation for Information Processing (IFIP) as part of its first cohort of fellows. The new members played very significant roles in driving innovation, conducting research and industry development. The global ICT professional body recognised the contributions of 18 high achieving members from across the globe, these included Sir Tony, known for his work developing Hoare logic for verifying program correctness, the formal language communicating sequential processes (CSP) to specify the interactions of concurrent processes, and for being the inspiration for the programming language occam.

Oxford team wins Cyber 9/12 London competition

Cyber 9/12 is a cyber policy competition for students across the globe, who compete to develop national security policy recommendations which tackle a fictional cyber-incident. Students from Oxford University have taken part in the competition before but this is the first year the Oxford team has done so well. Our 'CyberSeals' team reached the finals, and then won the competition.

Team member Matthew Rogers writes, 'we were given a 40 page intelligence brief from multiple sources and a month to write a situation assessment and decision brief, giving a 10 minute oral presentation to UK government experts who acted as judges. For the semi-finals we had to do the same, but this time a 20 page intelligence brief in just 13 hours from 5pm to 9am. In the competition finals we were given 20 minutes alone in a room with a 10 page intelligence brief, and then promptly presented a 10 minute brief, with a 15 minute Q&A on stage to all attendees. Through this competition we evaluate the threats different actions propose to the UK government, technological responses, and the multi-national, inter-agency, and government/industry connections that must be considered when responding to cyber attacks. This is especially when the attribution is unclear and situation only continues to escalate.'

The winning team, pictured above, from left: **Alexis Ciambotti** (MPhil Politics), **Manuel Hepfer** (Centre for Doctoral Training in Cyber Security), **Matthew Rogers** (Centre for Doctoral Training in Cyber Security), **Yashovardhan Sharma** (Centre for Doctoral Training in Cyber Security).

Air traffic surveillance data used in analysis by Bank of England

The Bank of England included in the latest 'Monetary Policy Report' for May 2020 an analysis of data from the OpenSky Network (https://opensky-network.org/), a research project by several universities and government entities with the goal to improve the security, reliability and efficiency of the airspace. Professor Ivan Martinovic, Head of the System Security Lab, a research group at the Department of Computer Science, alongside colleague Martin Strohmeier, Research Officer at the Department of Computer Science, have been working closely with OpenSky Network. The network, which began as a research project that was founded in 2012 by the Department of Computer Science, armasuisse and University of Kaiserslautern, now includes a number of other members from other organisations.

The latest report from the Bank of England, which establishes the economic impact of COVID-19 (using data from the OpenSky Network) displays a sharp decline in the number of aircraft departures from four major airports: Heathrow, London; JFK, New York; Schipol, Amsterdam; Incheon, Seoul. The report states economic activity has fallen sharply around the world as a result of COVID-19 and that, 'Travel has been severely restricted because of the pandemic. The number of commercial flights leaving US and European airports has fallen by around 90% since the start of March'[1]. Data from the OpenSky Network is also referred to as a high-frequency indicator for gross domestic product (GDP) forecasting. The Bank of England use this data to show the rapid decline of aircraft departures over March from UK airports including Birmingham, Gatwick, Heathrow, Luton, Manchester and Stansted. Using this and other sources, the bank staff estimate that monthly GDP will fall by enough in March to pull GDP down by around three percent.

The Bank of England Monetary Policy Report for May 2020 can be viewed in full here: bit.ly/35Z8tej

[1] Bank of England Monetary Policy Report May 2020. 2.3 Current Economic Conditions. P22.

NEWS

Law and Computer Science: Final project presentations

What do the EU Emissions Trading Scheme, Green Energy Trading Microgrid, shareholder proxy voting, and smart tickets have in common? They have been scrutinised by our interdisciplinary teams of students as part of the inaugural Law and Computer Science course led by Professor Rebecca Williams (Faculty of Law) and Professor Tom Melham (Department of Computer Science).

Four teams of students demonstrated that collaboration between the Oxford Faculty of Law graduates and their colleagues from the Oxford Department of Computer Science can swiftly turn into a number of fascinating entrepreneurial ideas as well as impressive practical implementation of those ideas. During the last five months, the students have developed their projects from scratch, building on their domain expertise and on new insights gained from the series of seminars, each of which was jointly delivered by faculty members of both departments.

On Thursday, 12 March, the student teams of six presented their projects and pilot products concerning the four themes mentioned above to a diverse professional audience. A funding pitch-like environment got the students to think about some difficult legal, technical, business and other questions from representatives of Allen & Overy, Slaughter and May, Clifford Chance, Innovate UK, Oxford Sciences Innovation, and Barclays Eagle Labs. The entire event was conducted online.

The course and practical lab sessions for our students are closely linked with a large-scale research programme titled *Unlocking the Potential of Artificial Intelligence for English Law*[1]. If you want to learn more about the course and the programme, or are interested in sponsoring opportunities, please write to the project administrator: admin.ai@law.ox.ac.uk

[1] https://bit.ly/2yCxFL8

Clare Lyle named as Open Phil AI Fellow

Clare Lyle, a DPhil student at the department, has been named as one of the recipients of the prestigious Open Phil AI Fellowship 2020. The Open Phil AI Fellowship is awarding $2,300,000 (approximately £1,800,000) over five years in PhD fellowship support to 10 promising machine learning researchers around the world. Clare is the only recipient of this award who is studying at a university in the UK.

Clare is pursuing her DPhil in Computer Science here as a Rhodes Scholar, advised by Professor Yarin Gal and Professor Marta Kwiatkowska. She is interested in developing theoretical tools to better understand the generalisation properties of modern ML methods, and in creating principled algorithms based on these insights.

The students chosen to be 2020 fellows were selected from more than 380 applicants for their academic excellence, technical knowledge, careful reasoning, and interest in making the long-term, large-scale impacts of AI a central focus of their research.

Read more here: https://bit.ly/2yO0btU

MIT Technology Review names Professor Yarin Gal on their Innovators Under 35 Europe 2019 list

Since 1999, *MIT Technology Review* has recognised young innovators and talented entrepreneurs from different countries who are developing new technologies to help solve the problems that affect our society. These bright young people, aged between 24 and 34 years old, have become Innovators Under 35 Europe 2019 after being selected by a judging committee of 116 specialists, from a pool of more than 1,000 candidates. The objective is to reward innovation – the development of new technologies or the creative application of existing ones – to solve current problems, ingenuity and advances on issues of global concern. Professor Yarin Gal has been included in the 'pioneers' category, because according to the citation 'AI still makes mistakes, but his tools can alert us when a system is about to go wrong.'

The winners offer a broad picture of how technology is evolving, as their projects cover different industries: Biotechnology & Medicine, Nanotechnology & Materials, Software, Artificial Intelligence & Robotics, Computer & Electronics Hardware, Internet & Web, Transportation and Telecommunications. The category with the most laureates this year has been that of Biotechnology & Medicine; 19 of this year's winners work in such fields, while Artificial Intelligence and Robotics has come in as the second category with 6 winners developing projects related to those topics. More that 75% of this year's laureates have already founded a start-up in the aim of transferring their value propositions to the market and generate a positive impact on society as a whole.

DPhil graduate Dominik Peters receives distinguished dissertation award

Dominik Peters has been announced as the recipient of the 2019 Victor Lesser Distinguished Dissertation Award. This award is sponsored by the International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS). It is awarded for dissertations written as part of a PhD defended in 2019, which show originality, depth, impact and written quality, supported by quality publications.

Dominik's award citation reads, 'We award the prize to Dominik for his ground breaking thesis making key contributions to MAS. His work on using SAT solvers was able to improve the bounds for a classic impossibility result in social choice, while generating computer-generated proofs that are comprehensible to humans. He designed new strategy-proof mechanisms for participatory budgeting, which has the potential to transform how budgets are allocated in thousands of cities in the real world. Dominik's work has put forward a number of new techniques that are already being used by other researchers in the field.'

Dominik's DPhil work titled 'Fair Division of the Commons' was supported by an Engineering and Physical Sciences Research Council studentship, and his supervisor was Professor Edith Elkind.

News in brief

The Department of Computer Science has created an online tour of the department, a great new resource, giving people who would be unable to visit in person the opportunity to see inside the department. The video is available to view here: bit.ly/3bDhTy4

Shangtong Zhang, Wendelin Boehmer, and Professor Shimon Whiteson are contributors to a paper which has won the Best Paper Award in the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020). The paper titled 'Deep Residual Reinforcement Learning' revisits residual algorithms in both model-free and model-based reinforcement learning settings. A copy of the paper can be found here: bit.ly/3alWCaN

Eman Alashwali and Professor Andrew Martin have contributed to a paper which won the Best Paper Award in the 15th International Conference in Security and Privacy of Communication Networks (SecureComm19). The paper titled 'Towards Forward Secure Internet Traffic' was co-written with Pawel Szalachowski from Singapore University of Technology and Design.

Antoine Amarilli (Télécom Paris) and the department's Dr Ismail Ilkan Ceylan won the Best Paper Award at the 23rd International Conference on Database Theory (ICDT 2020), for their paper titled 'A Dichotomy for Homomorphism-Closed Queries on Probabilistic Graphs', which can be accessed in the open proceedings: bit.ly/2ZjQAWj

NEWS

Navenio announces £9m in Series A funding for its efficiency-boosting location technology

Navenio, a Department of Computer Science spinout company co-founded by Professor Niki Trigoni, has announced an impressive £9m in Series A funding for its infrastructure-free indoor location solutions, which can be used to power a range of apps and platforms in sectors including healthcare.

Of particular interest during the Covid-19 pandemic is the potential use of Navenio's Artificial Intelligence (AI) led 'intelligent workforce solution' in a hospital setting, where it can assign tasks to healthcare teams based on their location. This will help hospitals to prioritise workload in real-time.

The secure system uses smartphones, with no need for investment in expensive new infrastructure. The benefits of Navenio in healthcare include faster cleaning of infectious areas in hospitals, enhanced visibility of vulnerable patients being moved between hospital departments, and better communication between porters and staff, to ensure that equipment is transported to where it is needed as soon as possible.

The funds will be used to increase the scope of Navenio's offering in the UK, as well build on existing collaborations in the US and Asia.

Philea Chim, senior vice president at QBN Capital, said, 'We're delighted to lead this funding and be closely involved with Navenio's mission to help frontline teams. Navenio's technology has the potential to transform the way that organisations deal with tasks across the board for the better.'

Tim Weil, CEO of Navenio, added, 'Our technology is all about giving hospitals the oversight needed to allocate frontline teams the tasks they are best suited to carry out, in a safe and secure manner – this is critical both for right now, and for the huge patient flow challenge that will be faced as hospitals re-open for business as usual.'

Read more about Navenio here: www.navenio.com

Collaboration between Samsung Research UK and the Department of Computer Science

by Eoghan Flanagan / Samsung Research UK

In order to develop knowledge representation techniques to advance the next generation of AI solutions Samsung Research UK (SRUK) have announced a collaboration with Professor Ian Horrocks of the Department of Computer Science.

One of the key challenges this collaboration will address is integration of logical AI and statistical machine learning. Ian's team at Oxford has been researching and developing advanced techniques to integrate knowledge from different sources and produce a broader and more powerful ontology representation. SRUK will use its experience with machine learning natural language processing (NLP) to improve the current state-of-the-art in this field. NLP techniques could be used to better understand that while the strings 'salt' and 'sodium chloride' are not alike, they occur interchangeably in text, and that they likely refer to the same substance in the real world. This semantic matching could then be used to better integrate knowledge bases that might use inconsistent labels for the same ingredient.

By integrating techniques from the logical-reasoning AI field and the machine-learning NLP field, SRUK and Oxford's collaboration looks to bring the best from both worlds to create powerful knowledge representation that will drive the next generation of Samsung products and services.

Waymo acquires Latent Logic to accelerate progress towards safe, driverless vehicles

Waymo has acquired Latent Logic, an Oxford-based technology company, spun out of Oxford University's Department of Computer Science. Waymo originated as the Google's self-driving car project then became a stand-alone company in December 2016.

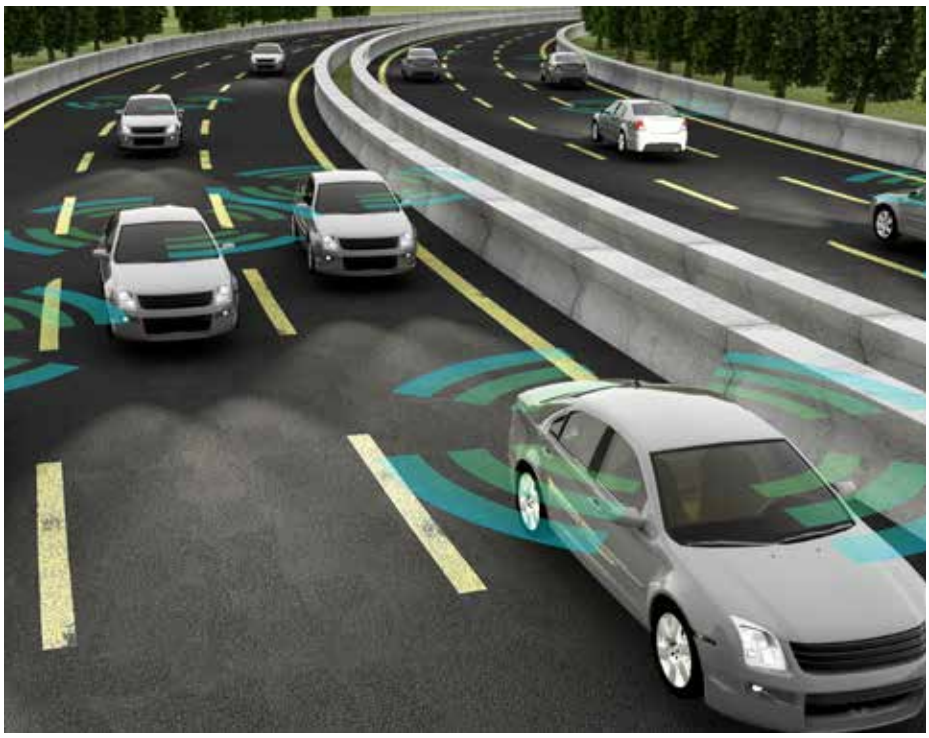

Latent Logic uses a form of machine learning, called imitation learning, to develop realistic simulations of the behaviour of motorists, cyclists, and pedestrians to help autonomous vehicles coexist and interact safely with humans in the real world, at scale.

The company's two founders (Professor Shimon Whiteson and João Messias) and CEO (Kirsty Lloyd-Jukes) join Waymo, along with key members of the engineering and technical team. The team will remain in Oxford and join Waymo's Research efforts.

The Latent Logic team's expertise in reinforcement learning and imitation learning help further accelerate Waymo's progress in areas from simulation to behaviour prediction and planning. In leveraging cutting edge machine learning techniques, Waymo will be able to scale more easily to complex and diverse environments.

This acquisition also signals the creation of Waymo's first European engineering hub, in Oxford, which is at the cutting-edge of global machine learning and robotics research.

Shimon comments, 'in joining Waymo, we are taking a big leap towards realizing our ambition of safe, self-driving vehicles. In just two years, we have made significant progress in using imitation learning to simulate real human behaviours on the road. I'm excited by what we can now achieve in combining this expertise with the talent, resources and progress Waymo have already made in self-driving technology. We're proud to have started in Oxford and to now be part of Waymo's first European engineering hub in the UK, surrounded by such amazing talent and entrepreneurship in machine learning, robotics, and autonomous vehicles.'

Donation from Mastercard to benefit biometric authentication and advanced technology research

Mastercard has donated £150K which will be used by Professor Ivan Martinovic and his Systems Security Lab in ongoing research relating to various uses of biometric authentication and AR technology.

The focus of the latest research project will be on the security and performance analysis of various biometrics authentication methods, especially voice-based biometrics, eye-movement biometrics, and authentication in mixed-reality environments.

Ivan comments, 'we are very thankful to Mastercard for their continuous support of our research. These are important projects as biometrics are at the centre of the tensions between seamless user experience and underlying security guarantees. This donation will be used to better understand the interaction between biometric authentication, AI methods, and the overall system performance under realistic adversaries'.

Inauguration of HKU–Oxford partnership in quantum information and computation

Initiated by the Quantum Information and Computation Initiative (QICI) of the Computer Science Department, Hong Kong University (HKU) Faculty of Engineering, the quantum groups at the Computer Science Departments of the HKU and the University of Oxford recently signed a Memorandum of Understanding (MoU). The MoU consolidates a long-standing collaboration between the Oxford Quantum Group and the Quantum Information, Foundations, and Technologies Group at HKU to establish a 'HKU-Oxford Joint Lab for Quantum Information and Computation'. In addition, the collaboration will include the exchange of mutual research visits, joint participation in grant applications, and joint supervision of PhD students in order to facilitate the exchange of talents between the two partner groups and the training of young researchers in the growing area of quantum information and computation.

Furthermore, the HKU and Oxford quantum groups have recently been awarded a grant from the John Templeton Foundation, which provides them with funding to hire joint postdoctoral fellows. With the support of the grant, a workshop on the intersection between quantum information and quantum gravity was held in January 2020 in Hong Kong.

An MoU signing ceremony was held on 15 January 2020 at HKU. It was attended by HKU Acting-Pro-Vice-Chancellor (Research) Professor Alfonso Ngan, HKU Dean of Engineering Professor Christopher Chao, Joint Head of the Quantum Group at the University of Oxford Professor Bob Coecke, HKU Head of Department of Computer Science Professor Lam Tak-Wah, Director of HKU Quantum Information and Computation Initiative (QICI) Professor Giulio Chiribella and members of the HKU and Oxford groups. Professor Alfonso Ngan and Professor Bob Coecke introduced the background and motivation of the HKU-Oxford partnership at the ceremony. Professor Ngan also expressed the strong support of HKU to this collaboration.

Professor Christopher Chao, Dean of HKU Engineering, said, 'The establishment of the MoU with the University of Oxford is an important step for the research and development of quantum information and computation within the Faculty of Engineering of HKU. We have high expectations for the collaboration with the Oxford Quantum Group, and for the establishment of a Joint Lab for Quantum Information and Computation. It will push forward the boundaries of our knowledge of quantum information, and will lead to potentially disruptive new technologies.'

Under the MoU, members of the Joint Lab will have opportunities to spend extended visits and conduct research on the broad spectrum of topics covered by the quantum groups at the two universities. Hlér Kristjánsson, PhD Student at the University of Oxford, said, 'The HKU-Oxford collaboration provides a unique opportunity for students at both universities to be immersed in the vibrant academic community and traditions in the University of Oxford, whilst also experiencing the innovative research and cosmopolitan city life in Hong Kong.'

University of Oxford signs Memorandum of Understanding with Waseda University

To support exchanges of graduate students and staff and to collaborate on research, a Memorandum of Understanding has been signed between the Departments of Computer Science, Mathematics, and Physics at Oxford and Waseda University in Tokyo. The new Memorandum of Understanding was established following the visit of Professor Hironori Kasahara (senior executive vice president of Waseda) in November to give a lecture on green computing; University of Oxford Professors Jeremy Gibbons (Computer Science) and Peter Braam (Physics) made the return visit to Waseda held in January 2020 to set up the memorandum. Waseda is one of the top private universities in Japan, with particular strengths in robotics and green computing, and this agreement will provide new opportunities for working together, particularly in machine learning and programming languages.

Oxford Computer Scientists and Amazon Web Services create a test-bed for cloud-based research

The Department of Computer Science is playing a key part in a new strategic collaboration between Oxford University and Amazon Web Services (AWS). The collaboration will focus on building a portfolio of new research projects relating to AI, robotics, cyber-physical systems, human-centred computing, and support to the University's new 'Lighthouse' Doctoral Scholarships.

This new university-industry collaboration, supported by a £7 million gift from AWS to the Mathematical, Physical and Life Sciences Division, will accelerate advances in AI and Data Science across the entire research portfolio of the University.

Max Peterson, VP International Sales Worldwide Public Sector, AWS, said, 'We are excited about this collaboration with the University of Oxford. With AWS, the University will be able to accelerate time-to-science as multiple, large experiments can be conducted in parallel with greater ease and in less time. And by driving cost down, researchers can dramatically increase the scale of computational experimentation. The collaboration demonstrates how academia can use the cloud to deliver excellent science with greater speed, flexibility and security, compared to using on premises data centres. Through our donation we will also support a new generation of researchers accessing cloud-native tools and technology for research through the University's "Lighthouse" Doctoral Scholarships programme.'

Oxford University has a vibrant, large and growing programme in Data Science, AI and Robotics research and development. To support and inspire the research students and staff, access to fast-moving, state-of-the-art large-scale computing resources is critical.

AWS, one of the leading global cloud providers, offers secure, reliable, scalable, low-cost cloud infrastructure that underpins millions of customers around the world and this collaboration will significantly accelerate the cloud-based research conducted at the University.

Professor Marina Jirotka, Professor of Human-Centred Computing, said, 'This gift will open up many amazing opportunities. My team is excited to expand our work on the ethical black box for autonomous robots, and lay the foundations for an Institute of Responsible Technology. We look forward to extending the impact of our responsible innovation work into industry.'

Professor Niki Trigoni leads the Cyber Physical Systems group, which focuses on infrastructure-free indoor positioning systems, human robot interaction and autonomous collaboration systems for workplaces ranging from hospitals, to office environments and warehouses. The group has set up an international centre-to-centre collaboration, across the UK, the US and Australia.

Professor Trigoni said, 'We are focusing on building state of the art situation awareness algorithms for blue light emergency services that combine variety of sensing modalities - inertial, radio, visual, mmWave and thermal. This additional cloud infrastructure will enable efficient, often health/life-critical solutions for human-machine collaborations in the workplace and beyond.'

Visiting Academic Profile

Alan Garfinkel : is currently the Newton-Abraham Visiting Professor and a member of the Computational Cardiovascular Science Group, part of the British Heart Foundation Centre of Research Excellence at Oxford. Here he describes what he has been involved with at Oxford.

'I arrived on 1 October 2019 as the Newton-Abraham Visiting Professor, based in the Department of Computer Science, and working in the Computational Cardiovascular Science Group and will have visiting professor status until the end of August 2020. My home institution is the University of California, Los Angeles (UCLA), where I'm Professor of Medicine in the Division of Cardiology. As a Professor of Medicine I have found a happy home in Oxford's Department of Computer Science, which says something about the department; it is exceptionally strong in using computational science to address real-world problems in the sciences.

Many critical problems in modern biology and medicine require massively parallel supercomputing. For example, at UCLA I studied the processes that account for the development of branching in the lung. Having been inspired by the biological work of Alan Turing, we postulated that the fundamental features of lung branching could be accounted for by a 5-variable Partial Differential Equation, modelling the reaction and diffusion of key chemical morphogens.

At Oxford, I'm delighted to be a part of the Computational Cardiovascular Science Group, using similar techniques in the simulation of cardiac electrical conduction and arrhythmias, as well as cardiac contraction and its pathologies, such as heart failure. At the same time I've also been extremely fortunate in being appointed to a Professorial Fellowship at Lincoln College. The camaraderie, intellectual fellowship and stimulating conversations have led to a number of productive collaborations.

During Hilary Term, I gave a DPhil course called Dynamical Systems Modelling. The course introduced differential equation modelling and computer simulation as tools to understand the dynamics of biological systems at all space and time scales, with an intended audience of physiology students with low levels of maths and computational skills. The Department of Computer Science served as the ideal platform to introduce these students to the use of simulation as a research tool, and I was pleasantly surprised to find that 54 DPhil students ended up taking the course, coming from a variety of other departments.

Since COVID-19 hit I have been involved in creating a Computer Science webpage to help teachers explain the role of simulation in understanding the spread of an epidemic through a population. The page features an interactive differential equation model of populations of susceptible, exposed, infected, medically symptomatic, and recovered individuals, with sliders enabling the student to change parameters such as the transmissibility of the disease, social distancing, and to see the effect immediately. You can take a look at the teaching resources here: bit.ly/2Vxihsk

While Covid-19 has taken my work in a new direction, I have enjoyed my time here so far and continue to enjoy collaborating with Oxford University colleagues.

MPLS Enterprise Fellow: Jiaoyan Chen writes about his role

The Mathematics and Physical Life Sciences (MPLS) Enterprise Fellows represent and liaise with MPLS research staff to provide an evidence base for the enterprise training provision and support of fellow researchers.

Since October 2019 I have been a part of a group of five Enterprise Fellows including Dr Martine Abboud (Department of Chemistry), Melanie Ghoul (Department of Zoology), Weixin Song (Department of Materials), and Tianyou Xu (Department of Engineering). We organise a monthly meeting in The University Club, in order to share views and opinions with each other and the wider research staff community. The meetings are open to all university staff and students, as well as entrepreneurs from a variety of different disciplines, who are invited to share with us their own experience of creating or working with companies.

As part of our work we also discuss new programs or activities which could support university staff, especially for postdoctoral researchers and students. For example, we are planning to invite senior academics from different departments and Oxford colleges to listen to current postdoctoral researchers and students talk about their motivation (and apprehension) regarding entrepreneurship.

Throughout the COVID-19 pandemic, the MPLS Enterprise Fellows have continued to meet online to shape a year-long enterprise development programme tailored for research staff. We have also been taking part in online meetings with the participants of Ideas 2 Impact at the Saïd Business School, a programme to enable postdoctoral researchers and postgraduate students to build enterprise skills. We use these meetings as an opportunity to discuss and explore where technical ideas and business ideas meet, and will be rolling this same opportunity out to current research staff within MPLS at the University of Oxford.

Read more: bit.ly/2KGy41R

OxWoCS Round-up for 2019-20


© John Cairns

Thus far it has been a great year for the Oxford Women in Computer Science Society (OxWoCS). Before the lockdown we were fortunate to have a number of sponsor's events, insightful talks as part of our Distinguished Speakers Series and collaborations with other societies. Thanks to our generous industry sponsors, Google, QuantumBlack, Zilliqa, Bloomberg, GitHub, Marshall Wace, Oliver Wyman, Optiver, and the Department of Computer Science, without whom we would not be able to continue to support and promote women in Computer Science.

Workshops with Zilliqa Blockchain
Starting on 31 October 2019 we took part in a series of workshops organised by one of our gold sponsors, Zilliqa, to introduce students to general concepts of blockchain, smart contracts, security issues, and their use cases. They even had opportunities to code and deploy their own smart contracts! The workshops culminated with a Demo Day on 21 November, 2019.

Women in STEM Networking event
On 25 January 2020 we collaborated on the annual Women in STEM Networking Event (formerly called the Welcome Event) with other Women in STEM societies including Mirzakhani, Oxford Women in Physics, Women in Engineering and OxFEST. The event took part in the Oxford Town Hall and created an opportunity for attendees to interact with inspiring women from industry partners (including local spinouts).

OxWoCS Distinguished Speaker Seminar Series
Throughout January and February 2020 we also had a number of speakers as part of the OxWoCS Distinguished Speaker Seminar Series run jointly with the Department of Computer Science, our speakers included:

- Dr Julia Neidhardt (TU Wein, Austria) on 24 January 2020
- Professor Mihaela van der Schaar (University of Cambridge) on 29 January 2020
- Dr Aishwarya Agrawal (Google DeepMind) on 4 February 2020
- Dr Irina Higgins (Google DeepMind) on 25 February 2020.

We were delighted to have such a varied number of women speaking about their research.

Introduction to Python
On the weekend of 25 January 2020 and 1 February 2020, we also had the opportunity to collaborate with the Oxford Coding Society (CodeSoc), Oxford Foundry and OxFEST in a series of Introduction to Python workshops. These were very well attended with participants drawn from across the University of Oxford.

Sponsor-led activities
On 13 November 2019 we held technical interview preparation workshops with Google. This was an opportunity to walk through coding challenges with their Software Engineers, learning some tips and tricks along the way.

On 29 November 2019 we had an office visit and Machine Learning workshops led by Quantum Black (by McKinsey & Company). We were fortunate enough to be able to network with some of the Data Scientists and Software Engineers after an afternoon of technical presentations and workshops on Algorithmic Fairness, Machine Learning, and Software Engineering.

This year we visited the Bloomberg offices in London on 30 January 2020. We were given a tour of their spectacular building and office space and students got to learn about their work and had the rare opportunity to shadow some of their engineers.

We also visited the Marshall Wace office in London on 27 February 2020 for a day of technical workshops and lunch with the tech team, an eye-opener event into the world of hedge funds!

Oxford Emerging Technologies Party
We ended Hilary term at the Oxford Emerging Technologies Party on 28 February 2020 with other Computer Science societies including Blockchain Society, CodeSoc, VR/AR Society, Fintech and Law Society, OxRam Society and AI Society. It was an evening of Computer Science trivia, pizza and drinks.

International Women's Day
On 8 March OxWoCS, Oxford Women in Business and OxFEST organised an International Women's Day brunch as an opportunity to socialise and network.

We have had largely positive feedback from those who have taken part in our events and attendees have learned about new opportunities available to them in industry and academia. We are particularly grateful for the joint effort of OxWoCS committee members in coordinating and organising these events; we hope to continue activities in 2020-2021!

Alumni Profile

Oana Tifrea-Marciuska : now works at Bloomberg part of the Graph Analytics team. She read for her DPhil between 2012 and 2016 and is one of the co-founders of the Oxford Women in Computer Science Society.

What course did you study here and when?

I read for my Oxford DPhil between 2012 and 2016 researching Personalised Search for the Social Semantic Web under the supervision of Professor Thomas Lukasiewicz. Personalised Search aims to adapt search results based on the tastes, interests, and user needs.

What was your background before that?

I am originally from Romania, where I completed my BSc in Computer Science at Alexandru Ioan Cuza University. It was during this time that I also took part in the Erasmus programme and spent some time working at the Lille University of Science and Technology in France. After the completion of my degree I enrolled in the European Master's Program in Computational Logic, studying one year at the Technical University of Vienna in Austria and one year at the Free University of Bolzano in Italy.

What attracted you to studying Computer Science as a subject?

I started programming at the age of 14, while in high school in Romania. I was fascinated that you were able to make computers do things for you. At the time my dream was to make computers write my homework so that I didn't have to do it! I was always curious about how I would go about making a program do things for me and as I enjoy programming and Maths I decided to study Computer Science.

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

My DPhil research offered me the freedom to explore and learn a range of new skills. In my first year, I took courses, authored research proposals and also started to conduct my own research. I also had the opportunity to present my work to people from academia and industry at AI conferences at venues in amazing parts of the world.

As a co-founder of a startup during my studies I received funding from the University of Oxford Innovation Fund to help create the company. I took courses on building your own business and so pitched ideas at various events. I also used some of the data in my DPhil research.

Another great opportunity I had was to be one of the co-founders of the Oxford Women in Computer Science Society (OxWoCS), which supports and promotes women in Computer Science. I learned a lot on how to set up a society, how to get industry and internal support, set up interesting events, and I met amazing role models. The society is still running today and I can see how beautifully the society has grown and matured, I feel very grateful to have been involved in the beginning of such a wonderful project.

What did you do when you left Oxford?

I had funding from the EPSRC Doctoral Prize and an Alan Turing Institute spin-off research project, which enabled me to undertake postdoctoral research at the Alan Turing Institute on semantic parsing. After a year, I decided to go into industry and took a research engineer position with Bloomberg's AI Group in London. I am now part of the Graph Analytics team, which develops and maintains the Bloomberg Knowledge Graph and its value-add analytics. The team collaborates with academics by publishing papers, providing research funding, participating in conferences and program committees, hosting interns, and peer-reviewing scientific articles.

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

My studies at Oxford helped me be an independent researcher. During your DPhil, you're not just learning about your research topic, you're also learning core skills that apply to jobs both in and out of academia, such as: written and spoken communication, project management and organisation, leadership, critical thinking, collaboration, analysis and problem solving, research and information management, as well as self-management. What is most dear to my heart is that I continued to work on a semantic web-related topic in my current job.

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

Make sure you can show an employer that you have been building your CV, including doing research, publishing, and developing your professional skills. Try to find a job that you enjoy and you will be very happy with every day. During my DPhil studies, I tried to explore many things and this allowed me to understand better what I would want from my job in the future.

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

I would have been grateful of how well things turned out for me as during my time as a student, I sometimes had a feeling of uncertainty about my future. I am so grateful that my DPhil at Oxford gave me so many skills as well as confidence that I am able to use in my current job at Bloomberg.

New module focuses on teaching Ethics and Responsible Innovation

By Helena Webb and Max Van Kleek

© Ross Gales

Are some facial recognition systems racist? How might AI systems, machine learning and data-driven systems affect society? How much information is it acceptable for Amazon and Google to collect about us? How can we stop the spread of fake news on social media without taking away rights to freedom of speech? What does a virtuous and responsible Computer Scientist look like?

These are some of the questions we challenged our first year Computer Science students to tackle on a new module this year: Ethics and Responsible Innovation. The Department of Computer Science had not previously run a dedicated ethics module for undergraduates but the past few years have seen a growing enthusiasm amongst students and staff to shine a spotlight on questions of ethics and responsibility as they relate to current issues in computing. We were motivated to run this course by the increased focus on fairness, trust, accountability and transparency in Computer Science as a discipline as well as awareness that the increased prevalence and influence of computer based innovations in our lives has been accompanied by significant concerns that these innovations are not always safe and do not always act in our best interests. In addition, as students who study with us go on to work in roles in which they design, develop and implement innovations that shape the lives of people across the world, it is crucial that as part of their studies they are encouraged to reflect on debates around ethics and to consider their own values and responsibilities with regard to innovation.

The module was taken by all first year undergraduates and was led by us. It began with a series of lectures that introduced students to ethics as an academic discipline, normative theories, and the overall relevance of ethics to Computer Science. We then focused on contemporary ethical issues in computing including those relating to AI and data-driven algorithms, such as algorithmic bias, societal feedback loops, algorithms as 'weapons of math destruction', as well as controversies surrounding digital surveillance. We then covered the pros and cons of regulation and codes of ethics in computing, and the use of the responsible innovation framework (anticipate, reflect, engage, and act) to develop innovations that meet societal needs and expectations.

© Ross Gales

One key learning objective for the module was to encourage students to discuss and debate ethical and responsibility issues in depth in order to develop their own understanding of the complex and multi-sided nature of these issues. For this reason, the lectures were followed by seminar sessions in which the students worked on discussion-based tasks in small groups. These tasks required them to consider how they would take a responsible approach to the development of an AI system, and how they would seek to prevent the spread of harmful content on social media whilst also recognising the value of freedom of expression. These seminar sessions highlighted that whilst there is rarely a single 'correct' answer to ethical dilemmas in contemporary computing, talking them through and listening to different perspectives provides a pathway through which we can identify ways to address different concerns and take our own responsibilities seriously.

The module was a big success. The students engaged with the topics and tasks with intellectual curiosity and enthusiasm. They showed an awareness of contemporary controversies in Computer Science and a keenness to apply academic understandings of ethics and responsible innovation to them. They also enjoyed debating these issues and came up with some highly creative and thoughtful responses to the tasks they were given.

After the module we asked the students to prepare a short piece of reflective writing to tell us about their experiences of the course. We were very pleased that their responses indicated how much they recognised the relevance of ethics to Computer Science and how seriously they regarded their responsibilities as students and professionals of the future. Here are some examples:

'All the things I have learnt are useful for me as a Computer Scientist, from a user and innovator point of view… I expect in my studies to see a lot of things that may or may not be used in some debatable purposes. What I can do in such situations is to analyse it as we did during the course.'

'As Computer Scientists we should remember that technology we develop should be for the good of society; only by engaging in constructive dialogue with all stakeholders in our technology can we determine what this means in practice.'

'Ethics in Computer Science is very special, as it requires both a high level of social consciousness and a high level of expertise in Computer Science... Hence, it is our responsibility as Computer Scientists to be active in discussing ethical issues, as we have the necessary expertise to make a real change.'

GCSCC welcomes newest regional partner in South Africa and moves to online training

In light of the Covid-19 pandemic the University of Oxford and thousands of other educational institutions have moved to online teaching. The Global Cyber Security Capacity Centre (GCSCC) – normally not as involved in teaching activities as many others in the department – found itself involuntarily being one of the first movers in this regard as part of the centre's contingency planning. One of the first actions was to change travel plans for a visit to South Africa and Uganda to train participants on the deployment of its Cybersecurity Capacity Maturity Model for Nations (CMM).

But let's start from the beginning: In January this year the Global Constellation of Regional Cybersecurity Capacity Research Centres, which has been established by the GCSCC over the last years, welcomed its newest partner, the Cybersecurity Capacity Centre for Southern Africa in South Africa. C3SA ('Ceesa') is the second regional partner since the Oceania Cyber Security Centre (OCSC), in Melbourne, Australia, became partner in 2017 with funding from the Government of Victoria. C3SA is based at the University of Cape Town (UCT) and is a consortium between the GCSCC, UCT's Department of Information Systems, Research ICT Africa, a digital policy think tank based in Cape Town, and the Norwegian Institute of International Affairs (NUPI). The funding for this two-year project comes from the Ministry of Foreign Affairs of Norway. The three centres of the Global Constellation, GCSCC, OCSC and C3SA, aim to drive regionally informed cybersecurity capacity research and lead the deployment of the GCSCC's CMM, in their respective regions, and sustain the CMM's impact from the past six years.

'Collaboration on cybersecurity is crucial to all of us,' said Sadie Creese, Professor of Cybersecurity at the Department of Computer Science, and founding Director of the GCSCC. 'Our work with colleagues at the new C3SA is aimed at enhancing cybersecurity capacity-building in the Southern African region. It will not only contribute to building cyber-resilience for the countries involved, but also to further the academic excellence in cybersecurity across South Africa and the wider region. We welcome this addition to the growing constellation of partnerships around the world.'

Cybersecurity capacity plays a crucial role in Africa, even more so during a global pandemic when more and more activities move online and users are more exposed to the risks in cyberspace. According to Enrico Calandro, Co-Director of C3SA and Senior Research Associate at Research ICT Africa, 'cyber threats and risks are particularly challenging for African countries and nations affected by structural inequalities, conflicts and fragility. Having, generally, developing digital and physical infrastructures, weak institutional arrangements and governing mechanisms compounded by limited resources, the digitalisation of these countries is often characterised not only by low levels of connectivity, digital inequality, and limited opportunities, but also by insufficient security which places societies and individuals at a high risk of further marginalisation.'

The establishment of a research centre focusing on the region, C3SA provides 'a platform for governments, intergovernmental organisations, the private sector, and civil society organisations to leverage cyber capacity research in their efforts towards building more equal,

continued on next page ▶

from previous page ▶

inclusive, safe, and resilient digital economies and societies. C3SA will achieve its objectives by conducting regionally-focused research on cybersecurity capacity by deploying the CMM that provides countries with a baseline for capacity-building and resource allocation, and by developing and implementing locally-informed educational programme', says Enrico.

Professor Wallace Chigona, Co-director of C3SA, adds, that 'C3SA will achieve its objectives by taking a multi-faceted approach to cybersecurity capacity maturity advancement. This approach includes: Train-the-trainer activities to prepare our local researchers to deploy the CMM; the completion of national-level assessments to better understand what constitutes national cybersecurity capacity from an African perspective; capacity-building initiatives in the form of post-graduate training; and analysis and dissemination of findings to ensure that lessons learnt can have an impact on cybersecurity policy-making'.

The CMM facilitates the assessment of the maturity of a country's cybersecurity capacity across five dimensions. The deployment of the CMM is a quite intensive process which is led by the GCSCC or its partners and was described in the Winter 2018 edition of *Inspired Research* (issue 13). It includes research and facilitation skills, but also requires a very good understanding of cybersecurity as a multidisciplinary subject which relates to political, social, economic, technical and also psychological issues. Therefore to ensure C3SA fulfils is objectives and as part of the 'onboarding' activity for the GCSCC, there are plans to train researchers who come from different disciplines in the deployment of the CMM. Each CMM consists of desktop research, but the core part of the assessment process involves 10 focus group discussions over three days for which the research team travels to the country on invitation by the government. The researchers speak to no fewer than 200 stakeholders from the government, private sector, critical infrastructure, civil society and others. Afterwards the researchers draft a report including recommendations for each of the dimensions which is submitted to the respective host government. Therefore, detailed training is needed to produce high quality reports.

The CMM Deployment training is normally done in two parts: in-person sessions and 'learning on the job' during a CMM review. This was the intended approach for the new colleagues in Cape Town and a team from Oxford had been booked to travel to South Africa in March to attend the launch event of C3SA and to conduct training for the researchers affiliated to the Centre, including two research fellows and two PhD students, as well the outreach staff and other members of UCT who are affiliated to the project.

The week afterwards the teams from Oxford and Cape Town would have travelled to Kampala to conduct the first joint CMM on invitation of the Ugandan government. However, both South Africa and Uganda were amongst the first countries to impose travel restrictions for incoming visitors due to the COVID-19 pandemic and therefore the launch event had to be postponed indefinitely. However, both partners decided nevertheless to proceed with the planned train-the-trainer activities and jointly decided to move the training online.

Over March and April in 20 online sessions the GCSCC team faciliated the C3SA researchers to understand the GCSCC and the Constellation's purpose and aims, and the context of this global project. Through presentations by the researchers and the members of the GCSCC's Technical Board they took a deep dive into the CMM, its rationale and structure and how it enables countries to have a benchmark for cybersecurity capacity building and investment. Through interactive and home exercises they learned in theory how to deploy the CMM, conduct focus group discussions, and to plan, organise, and conduct a CMM review – from the agreement with the host government, to data collection and report drafting, to the publication of the report by the government. The training also included an introduction to the global data set that the GCSCC has built with the input of over 80 CMM assessments since 2015, and the qualitative research arising from it. The training was conducted using Microsoft Teams, which turned out to be a stable system considering the quality of the Internet connection.

Self-study was also part of the training. The participants looked into the GCSCC's resources including the CMM document itself but also the CMM Deployment Manual, a CMM Structured Field Coding Tool, various example reports, academic publications along with online resources such as the Cybil Portal, a global knowledge resource on cyber capacity building that the GCSCC has developed with international partners. Colleagues of the Oxford Cyber Security centre supported the training and shared lessons they have learnt from their work with the GCSCC and the CMM.

Overall, this part of the training worked well considering that everyone involved had been pushed into remote format on quite short notice. This online format also had the nice side effect that it allowed our new colleagues to 'virtually' meet everyone who is involved in the GCSCC – and it also was an opportunity for the GCSCC to get to know the C3SA which is in South Africa, Malawi and Zimbabwe.

The joint country visit, which is the second (and most crucial) part of the training, needed to be adapted and required a new approach to the usual CMM methodology. With the Ugandan government, who are going through the CMM process for the second time, this has been a challenge that GCSCC and C3SA has taken in consultation with the Technical Board of the GSCCC. Between May-July 2020, the researchers from Cape Town will conduct detailed desk research about the Cybersecurity capacity of Uganda based on the 2015 CMM report, followed by online expert interviews and focus groups with key Ugandan stakeholders. The outcome will be a preliminary CMM report that GCSCC and C3SA will submit to the government later this year which will also complete the CMM training for C3SA.

Despite the challenges this training imposed on the team, it was a good learning experience and when the whole process has been completed and evaluated, this online programme may inform future CMM training activities for partners and stakeholders.

Showcasing Responsible Innovation in Quantum computing

By Philip Inglesant

In the Summer 2018 edition of *Inspired Research* we asked, 'will quantum computing make the world better?' This was anchored in our Responsible Innovation research as part of the Networked Quantum Communications Information Technologies Hub (NQIT) centred in Oxford and part of the UK National Quantum Technologies Programme. NQIT incorporated Responsible Innovation from the outset with work led by Professor Marina Jirotka from the Department of Computer Science.

The first five-year phase of this programme came to an end in November last year, but its work continues through the new Quantum Computing and Simulation (QCS) Hub, also centred in Oxford, but with an extended set of 17 universities and over 25 industry partners. The transition to the new phase, as quantum computing starts to move from the laboratory to real-world adoption, is an opportune time to look back at the Responsible Innovation work which was a core part of NQIT, and forward to the QCS future.

In October 2019, we held a successful day-long showcase event at Imperial College in London to mark the culmination of our Responsible Innovation work in the NQIT Hub. There were 25 participants from research, industry, government and public policy that took part in a lively discussion led by six invited speakers: Professor Tom Rodden, Chief Scientific Adviser for the Department for Digital, Culture, Media and Sport; Professor Ian Walmsley, Director of the NQIT Hub and now Provost of Imperial College; Robin Williams, Professor of Science Technology and Innovation Studies at the University of Edinburgh; Elham Kashefi, Personal Chair in Quantum Computing at the University of Edinburgh, researcher at the Centre National de la Recherche Scientifique in Paris and Associate Director of applications development in NQIT; Jonathan Legh-Smith, Head of Scientific Affairs at BT; and Jack Tindale, Policy Manager, Design & Innovation at Policy Connect, a not-for-profit, cross-party think tank and coordinator of the All-Party Parliamentary Group on Data Analytics.

The discussions centred on issues such as the ethical challenges that will be wrought by these new, potentially very powerful technologies, and the specific opportunities, and concerns they may bring. Will they engender social changes following the pattern of existing computing applications such as the internet, mobile, and big data, or will they bring something quite new? Who will have access to the technology, and who will decide? How can we ensure that the benefits of technology are widely shared by society? Responsible Innovation – sometimes known as Responsible Research and Innovation – tries to look forward and address such questions. Even if we cannot yet know all the answers, we can anticipate and prepare for the future, whatever it might hold.

Until recently, it has seemed that practical quantum computing is always a few years into the future, but now there is increasing interest in what is known as Noisy Intermediate-scale Quantum Computing (NISQ), looking at what might be possible with existing quantum computers, (even if 'universal', fault-tolerant quantum computing is still some years away), and in quantum simulation for applications such as investigating new materials. These interests, together with the recent announcement by Google of 'quantum supremacy' and the availability of publicly available functional prototype quantum computers, suggest that there is an urgent need to think about the social, economic, and ethical issues which may arrive sooner than expected.

Professor Ian Walmsley quoted the well-known statement that quantum computing is as different from existing 'classical' computing as a conventional computer is from an abacus. It is not a simple case of quantum computing being, in general, much faster than classical but rather, there are certain specific problems which are completely impossible for any conceivable conventional computer, that become possible using quantum methods.

Quantum computing is expected to work alongside the constellation of existing computing, each one doing the tasks for which it is best suited. Even so, there is still much uncertainty about when quantum computing will become widespread, and the usefulness and impact of NISQ computing and quantum simulation. It is also not yet clear which form of quantum computing will turn out to be the winning combination. As Professor Elham Kashefi explained at the showcase event, there are many possible ways to make qubits – the fundamental building block of quantum computing – and many viable architectures. She drew the analogy, in computing terms, of early ENIAC computers (the

continued on next page ▶

from previous page ▶

first electronic general-purpose digital computer) or ARPANet (the technical foundation of the internet) that were not yet in use beyond a small group of pioneers, but had plenty of promise and went on to become the technology we know today.

One vision of quantum computing, in whatever form it takes, holds out the hope of it becoming pervasive, underpinning many application areas. This would be comparable with earlier technologies such as lasers – themselves reliant on quantum physics – which, originally developed for quite specific purposes, have become ubiquitous in mundane as well as highly advanced applications. In fact, lasers are essential components in quantum computing. However, as Robin Williams cautioned, the rhetoric which surrounds ambitious visions of innovation trajectories is often not fulfilled: the path from laboratory to adoption is not linear and often does not follow the predicted course. Indeed, there have been claims of 'supremacy' around earlier technologies, which have not always led to the expected global dominance.

Professor Tom Rodden argued that in order for quantum computing to become widespread, it will require not only technical feasibility but also a sustainable and diffuse, open ecosystem to support it; there will be implementation and manufacturing problems to overcome, as well as the need to train a skilled workforce of quantum programmers and engineers. As with other new technologies, this innovation may be disruptive, painful for some even, if beneficial in the long term. Small companies, so often at the forefront of innovation, are likely to play a key role, therefore standardisation and protection for intellectual property will be integral to their continued operation.

Turning to potential application areas and their likely implications, Jack Tindale, from Policy Connect, discussed the implications if quantum computing were to be applied to big data analytics. This raises the importance from the policy and governance perspective of maintaining public trust. Companies have been analysing data in various ways for decades, but with increasingly pervasive and large-scale data collection, concerns around privacy and the control of data have taken a much higher public profile, especially since the adoption of GDPR.

These points raise wider questions of ethics and Responsible Innovation. For Jonathan Legh-Smith, human rights provide an ethical basis for these issues. Ethics has been widely rehearsed in the context of Artificial Intelligence, but a recourse to ethics should not be a mask for what are clearly bad practices; he gave the example of facial recognition trialled in London's Kings Cross without the consent or prior knowledge of the unwitting public. How can we proceed with Responsible Innovation in the face of these known and unknown challenges?

Several presenters suggested the need for a framework such as a 'licence to operate' for companies exploiting data analysis to provide services to the public – a suggestion which was also a key recommendation in the recent Trust, Transparency and Tech report (bit.ly/2TYSSH2) produced by Policy Connect and the All-Party Parliamentary Group on Data Analytics, for which Professor Marina Jirotka acted as an adviser. Although not specific to quantum computing, these principles could also be applied here: the need for consistent and practicable governance; the need to engage with all stakeholders in dialogue and co-creation; and to find the right balance to protect civil liberties valued by society.

Ultimately, Responsible Innovation encompasses ethics and regulation, but provides a broader framework as it dynamically responds to social and technical changes 'upstream' as well as 'downstream' in innovation, anticipating, reflecting, and engaging widely. The showcase brought these messages to researchers in quantum computing to understand, perhaps for the first time, the capacity of Responsible Innovation to prepare for some of the more negative as well as positive impacts for society of powerful new technologies.

Quantum physics is noted as particularly difficult for non-specialists to grasp. Richard Feynman is quoted as saying, 'if you think you understand quantum mechanics, you don't'. The complex interplay of science, technology, and innovation requires the insights of experts in human centred computing, social scientists, policy experts, figures representing industry and innovation, as well as a groundwork of highly advanced scientific research.

Responsible Innovation brings together and brokers dialogue and shared understandings between different stakeholders. It aims to close the distance between the most advanced research and those charged with leading society's response by creating a more agile pipeline between researchers and policy makers. It aims to open up dialogue and look at multiple perspectives through engagement with diverse publics. It recognises that powerful technologies can give disproportionate power to those who control them, and can deny power to those who are unable to access them.

In summary, while Responsible Innovation cannot predict the future, it can anticipate and create the capacity to respond to uncertainties in the future. This requires strategic, coordinated, cross-government leadership, learning from other models, and thinking about Responsible Innovation from the outset or, as Professor Tom Rodden put it, concluding his message to the Showcase, 'in a decade's time, will a future Chief Scientific Adviser be advising a minister on the need to introduce mandatory regulation and enforcement of the quantum technology sector?'

A nine-minute video of highlights from the showcase is available here: bit.ly/38IO12c

Behind closed doors: how do pilots handle cyber attacks?

by Matt Smith

Aviation is often seen as having a 'gold-standard' level of safety. According to UK Government figures between 2009 and 2018, air travel saw 0.1 fatalities per billion passenger kilometres, against 1.3 for car travel – over 100 times safer than arguably our most common form of travel.[1] A critical component in achieving this is well-defined cockpit procedure; pilots spend hundreds of hours learning how to deal with situations, both routine and emergency. When licensed, they continue to be regularly reassessed, with special attention paid to their fault and emergency handling skills.

One area in which the aviation industry has been slow to adapt is security. As a result of development taking many years, systems coming into use today are likely to have been designed without realistic threat models. Consequently, security researchers have identified a range of potential attacks on important avionic systems.

Despite many of these attacks having the potential to erode the safety of an aircraft, the research usually treats the aircraft as a unit. In doing this, it overlooks the humans in the loop – the flight crew – and how the attack might affect these pilots. With the flight crew being ultimately in charge of flying the aircraft, we wanted to understand how they might deal with such scenarios. Pilots rely heavily on their systems and are taught to trust them; when the systems do fail, they typically fail safely and in consistent ways. Systems under attack might not behave so consistently, leading to the question: can existing pilot training just mitigate attacks anyway?

Systems and Attacks

To test this, a team consisting of Matt Smith, Martin Strohmeier, Jonathan Harman (independent), Vincent Lenders (armasuisse) and Professor Ivan Martinovic were awarded a Departmental Impact Fund grant to construct a small flight simulator. Using commercially available simulator software, we implemented the likely effects of attacks and invited 30 Airbus A320 pilots to fly scenarios subject to malicious interference.

Figure 1: Flight simulator setup.

Typically, we consider the case where an attacker can transmit arbitrary signals, on the same frequencies used by avionic communications, using inexpensive equipment. With some expertise and a budget of around £10,000 (though usually much less), these signals can imitate real avionic communications. Since authentication is rare in these systems, both parties have few tools to check if a signal is legitimate or malicious. This allows an attacker to – in theory – attack aircraft as they fly overhead.

Based on existing research, we simulated attacks on three key systems, none of which have inbuilt security mechanisms. All systems are deployed and used on all commercial airliners, and are:

- Instrument Landing System (ILS), which enables precision, safe landings in most weather conditions,
- Traffic Collision Avoidance System (TCAS), which provides 'last resort' instructions if aircraft are on a collision course and,
- Ground Proximity Warning System (GPWS), which is a high priority alerting system, triggering if pilots fly too close to the ground.

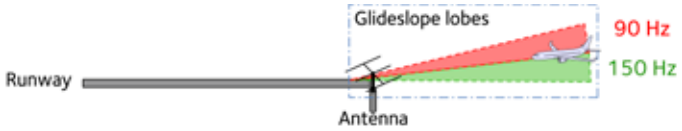

Figure 2: Instrument Landing System (ILS) glideslope diagram. The aircraft follows the path where the 90 Hz and 150 Hz lobes are of equal strength.

For each case, we recreated the cockpit effects of a wireless interference attack. Pilots flew scenarios unaware of which system might be affected and were asked to respond as they would in a real aircraft. In this article, we focus on the results of the Instrument Landing System (ILS) attack, but full details of all attacks are available in our Network and Distributed Systems Security (NDSS) 2020 conference paper.[2]

What Happened?

Our ILS scenario saw attackers transmitting a 'false' glideslope – a vertical landing path used for aircraft guidance. Glideslope signals, whilst complex, can be replicated to spoof the real glideslope. Following a false glideslope until landing caused pilots to overshoot the runway.

continued on next page ▶

[1] UK Government. (2020, 04 17). Passenger casualty rates for different modes of travel (RAS53). Retrieved from GOV.UK: bit.ly/3asXCdj

from previous page ▶

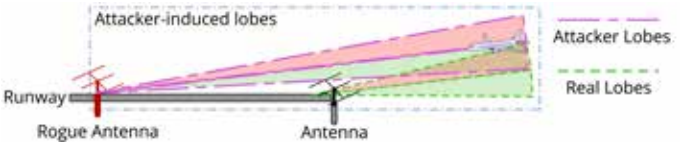

Figure 3: An under-attack glideslope. The attacker transmits a false glideslope from the far end of the runway, leading to the aircraft being too high on approach.

With our scenarios in place, we set the pilots off!

As expected, all pilots landed safely, but a wide variety of responses occurred on the way. Most pilots realised they were too high on the approach path by about a minute until landing, roughly 1,000 ft in altitude and 1 mile horizontally out from the runway. This is likely down to the fact that the attack only introduced a small change in altitude between the correct position and the 'under attack' position.

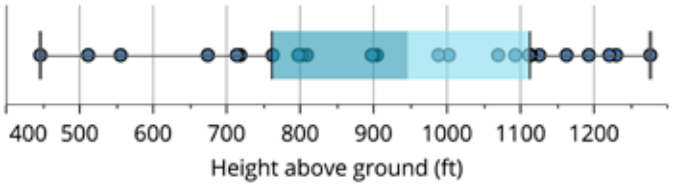

Figure 4: Box plot of altitudes for each participant choosing to abort the approach, at the point of beginning the missed approach procedure.

At this stage, 26 participants chose to abort the landing and attempt a second approach. This is a 'default' option, giving the pilots time to think, attempt to diagnose the issue and prepare for the next approach. However, this burns extra fuel and incurs delay – not to mention that aircraft are only fuelled to have around two missed approaches. If multiple aircraft are affected by this, longer backlogs could occur and aircraft may end up diverting to other airports, which becomes extremely costly.

On the subsequent approach, all the 26 pilots had lost faith in the glideslope so opted for a range of different approach methods and eventually went on to land safely. No single approach method saw more than a third of the remaining participants using it. This demonstrates that even when the pilots diagnosed the issue, there was no established response – a very different situation to typical aviation procedure. The gap creates a 'grey area' between the end of typical fault handling steps and the eventual choice of action they take. Aviation tries to eliminate these areas as they can lead to unsafe situations and security vulnerabilities are highlighting new grey areas.

Of the four participants who chose to land on their first approach, they had to take rapid action in order to correct the landing. This led to the aircraft descending close to its maximum acceptable rate for this point in flight – safe, but not ideal. Although responses did vary, one pilot reaction was consistent – when pilots noticed the glideslope was acting unusually, they stopped following it. This was a trend observed across the study; attacks led to pilots being forced away from using certain systems, ceding some control to attackers.

Beyond our scenario, our participants highlighted complicating factors which might amplify the effect of such an ILS attack. One of the most prominent concerns was poor weather. Rain, clouds and fog remove many visual references – even the runway itself sometimes – until the aircraft is close to the ground. Participants felt that under such conditions, it could be hard to spot the attack until late on, making the situation much less safe.

How do we fix it?

It is not easy to secure all the systems we have looked at in this study; many are reliant on legacy or bespoke technologies which are difficult to modify. Simply adopting cryptographic approaches is unlikely to be a realistic option in the short-term. Furthermore, there are challenges surrounding the use of security controls in the safety-critical aviation context that the industry is yet to address.

Instead, we need to look at other ways to more immediately help flight crew. In the first instance, this will require the development of procedures to handle attacks. This could even involve switching systems off, but importantly would be done as a controlled step rather than a last resort to minimise distraction.

One approach, based on our work, would be to provide attack-focussed simulator scenarios for pilots to train with. This would allow them to test these new scenarios and be prepared for attacks. Care must be taken to avoid negatively training pilots into distrusting their instruments, but this should be possible considering the extent of failure training they already do.

Longer term, aircraft systems need to be designed with security in mind. This would avoid leaving the pilots to be a last line of defence and instead let them focus on flying the aircraft. Until then, our work indicates they handle the attacks well – though without better security, they might be flying a few extra laps of the approach leg!

[2] bit.ly/2S47IAC

FBI follows Oxford academic's guide to beat the Zoom-bombers

It's not every day that the FBI takes a close interest in an Oxford academic's blog[1]. But, then again, a million television programmes suggest maybe it is. Whatever the case, Professor Bill Dutton's recent blog about the conferencing platform, Zoom, preceded a lively international debate – and interest from the Feds in his native US.

Ahead of the curve, Bill, who was a founding director of Oxford's Internet Institute (OII), wrote a blog with colleague, Arnau Erola, Research Fellow at the Department of Computer Science, highlighting the potential problems and issues around video-conferencing: Zoom-bombing the Future of Education.

Since the beginning of the Covid-19 crisis, the highly-accessible provider, Zoom had gone from geek-obscurity to household name and in a matter of days the number of its users increased from 10 to 200 million, as people around the world tried to overcome social distancing to connect with colleagues, friends and students.

But, as Bill, (who is also a fellow of the Global Cyber Security Capacity Centre, based at the Department of Computer Science) observed, problems were not far behind, 'It took about a week.'

With the increased use of Zoom and other platforms, came issues. In his blog, he wrote, 'One particular challenge that has risen in prominence is efforts of malicious users to sabotage classrooms and discussions, such as by what has been called Zoom-bombing (Zoombombing). Some have defined it as ''gate-crashing tactics during public video conference calls" that often entail the flooding of Zoom calls with disturbing images.'

And, since the 26 March publication, the topic has hit the headlines in the UK and around the world, as 'bombers' have interrupted and disturbed video-conferencing events from Jewish religious services to corporate meetings, often with malicious intent. News stories have appeared, voicing concerns about the security and protection of video-conferencing. Blame has been levelled against key providers. But, many of the issues are actually related to the users moderating the conference, rather than the software, says Bill.

He maintains, there are numerous ways in which meetings and events can be safeguarded from malicious intent. He says, 'There has been exaggerated coverage of the problems. It's not usually a problem with the software. Many of these issues can be addressed by the moderator.'

Bill adds, 'It's the job of the moderator to set up the meeting...Zoom is incredible, it's brilliant technology.'

A main issue is the overnight success with new users, who are unfamiliar with the technology and failing to use the safety and security measures that are available. Bill maintains, 'Part of the problem is that Covid-19 moved so many people online so quickly. Teachers and people with no background are using [this technology] because it is so simple. But it made them vulnerable to malicious intent [because they did not take the security measures that were available].'

But, says Bill, there are 'all sorts of settings' that could be used by the moderator [or organiser of the meeting], from using passwords, only accepting participants with known email addresses or invitations and control of visuals and audio. These can prevent unauthorised persons gaining access to private meetings and provide the safe space that organisers and legitimate participants seek. As well as commending guidance and training for novice staff, Bill made six key recommendations in his blog, aimed at ensuring safety:

- Authentication – limit the connection to specific users;
- Authorisation – restrict the technical facilities of participants, so they can't disrupt or show offensive material;
- Monitoring – although a laborious process, participants should be reviewed to prevent gate-crashers;
- Moderation – participants' activities can also be reviewed – particularly useful in an educational context;
- Policies – each institution using such technology needs to have set-down policies of acceptable and unacceptable behaviour;

continued on next page ▶

from previous page ▶

- Procedures – anyone breaching the rules could lose authorisation or be dealt with using laid-down procedures.

Not long after the blog was published, the FBI launched its own recommendations – which are very similar to those from Bill's team.

'They pretty much aligned with our recommendations,' he says, clearly amused.

With so much official and corporate business being conducted online, it is evidently a high-priority for the Bureau, which recently threatened 'zoombombers' with the possibility of jail sentences.

Bill's blog, meanwhile, saw particular problems for educators, he wrote, 'It is clear that zoom-bombing has become an issue for schools and universities, threatening to undermine the vitality of their teaching and relationships with faculty, students, and alumni of their institutions.'

Reflecting on the need for security, he says, 'It undermines the whole culture of education, which should be open and accommodating.'

In usual times, lectures can be open to any student. But these are not usual times and when setting up a class remotely, it will be necessary for the moderator to take steps which they would not otherwise consider. He recognises that this will not come naturally to all. Bill says, 'Universities need to provide the resources for IT staff to brief academics and do some hand-holding.'

This blog has been made available via the Oxford University Science Blog.

[1] The blog was written by Professor Dutton and Arnau Erola and was based on their discussions with Louise Axon, Mary Bispham, Patricia Esteve-Gonzalez, and Marcel Stolz

Picking probable parameters with PINTS

by Michael Clerx

Simulations with differential equations are used to model almost any dynamic system that has been written about in the physical and biological sciences. These systems encompass all scales: from the sub-molecular electrochemical processes in fuel cells; the intra- and inter-cellular interactions in cancer and heart disease; the population dynamics of infected individuals in epidemics; to the orbits of stars around black holes.

A common problem in applying these models is inference-finding parameter values that make model simulations match real-world observations. This might be finding a single best parameter set using optimisation but, because real-world data is noisy, estimating a distribution of parameters often gives a more complete and informative result. If you already have some idea of what the likely parameter values could be, Bayesian inference lets you incorporate this knowledge to find an even sturdier estimate.

PINTS, which stands for 'Probabilistic Inference on Noisy Time-Series', is a Python package developed in the department (in collaboration with researchers at Imperial College London and University of Cambridge) which provides users with inference methods particularly suited to time-series problems.

PINTS started out as a simple attempt to reduce duplicate effort but now contains a zoo of inference methods ready to be deployed against a large slice of time-series problems. These methods include powerful optimisers, likelihood-free approximate sampling routines and fast exact approximate sampling methods. We also include a hoard of toy problems that we hope encourages researchers to develop and benchmark their methods within PINTS. Despite being in an alpha-development stage, PINTS has already been used in publications in computational cardiology and electrochemistry.

If you're interested in parameter estimation, Bayesian inference, or derivative-free optimisation methods, please feel free to join in! PINTS is being developed in an open spirit, by a cross-disciplinary team of applied researchers, but using a high-degree of unit testing, documentation, and software best-practices. We also occasionally have actual pints.

Documentation and examples are given on github.com/pints-team/pints, but feel free to email us too at pints@maillist.ox.ac.uk

How to keep children safe and connected during home schooling

By Jun Zhao

Since Monday 23 March 2020, most schools in the UK have embarked on school closure in response to the COVID-19 outbreak in Europe. For most parents, aside from juggling between the daunting prospect of home schooling and their work, they are also overwhelmed by their children's need to stay connected with friends.

This is particularly a challenge for parents of primary school age children who have not been extensively exposed to independent online social communications until the COVID-19 crisis. Under this stressful and rapidly developing situation, many parents handed over to their children some general purpose video conference applications or social media platforms as a quick and easy solution. This article highlights several things that parents should look out for when facilitating their young children's first online chat or social communication, to keep their children safe, and enjoy a happy and rewarding home schooling time.

What are the risks?

Online social communication of teenagers has been more widely studied than primary school age children. Research has shown that online social communication provides a wonderful mechanism for teenagers to keep in touch with their friends during the increasingly dynamic modern family life, although they can also cause stressful or even damaging consequences to teenagers' mental health, wellbeing or safety[1].

In the UK primary school age children (5-11 years old) are reported to spend most of their online time on gaming or watching videos[2]. However, the emerging home schooling situation raises a new need for them to stay in touch with their friends to retain a social connection in a self-isolating or social-distancing context. Therefore, we see a surge of children of all ages embark on various online communication platforms without a full awareness of associated risks.

Many parents and children are probably aware of the risks of being approached by strangers, who are not their immediate friends from school or family members. Many platforms (such as Houseparty or Zoom) do not necessarily stop strangers from joining a video chat unless properly configured. Parents should ensure that any discovery function of such platforms (such as finding friends nearby) are disabled, particularly for young children. If you decide to give your child an email address or an old smartphone so that they can chat with their friends, be sure that they do not give out their contact details to strangers or receive contact from people they do not know.

Even when children are amongst their friends, they could still become emotionally overwhelmed as everyone may get easily excited, having not seen their friends for weeks and then 'meeting' them in a completely new and unfamiliar way. A lot of the anxiety that teenagers have been reported to experience online is caused by feeling being left out[3] or having their secrets exposed in public. This can also happen to young children and they would find it even harder to cope with or comprehend such behaviour[4]. Make sure that you talk regularly to your children about the importance of their own and other people's privacy (such as secrets), using respectful and appropriate language, and letting them know that they can also come over to talk to you if they encounter trouble. Online interactions should have lots of synergies to their social interactions at schools. Examples that schools have used to show children how to be kind can also help them understand that it is equally important to act responsibly and respectfully on the digital space.

Given the rapid rising need, parents are seeking solutions through general purpose tele-communication tools (like video conference tools used for the workplace) or social communication tools. These tools could be associated with various security or privacy risks. Some tools like Houseparty do not come with

continued on next page ▶

from previous page ▶

end-to-end encryption, which means that your children's conversations are not necessarily private. Other tools, like WhatsApp or Messenger, are not expected to be used by children under the age of 13 and parents need to take up the legal responsibility.

What should parents do to keep your children safe:

1. Think about the age restrictions: Unlike the age rating associated with apps or films, which indicates the appropriateness of the content, most social media platforms impose a minimum age restriction for accessing their services. This age restriction may vary in different countries. In the UK, you need to be 13 years old to have your own YouTube or WhatsApp account. Parents should think carefully when providing this consent on their children's behalf.

2. Respect your child's privacy: For children as young as 4, they have already developed a sense of personal privacy or secrecy[5]. Parents should tread carefully when trying to protect children from online risks and try not to break their mutual trust. For example, just because you know your child's password for their email, you should not look into their conversation with their friends without their presence or permission.

3. Talk about being respectful: Talk to your children about how to communicate respectfully to others, and when and from whom they should look for help if they find anything unpleasant. A private social network (eg setting 'Houseparty' to private mode) that prevents children from being contacted by strangers is a good start. However, parents should talk to their children about what it means to be a good friend online and how to respect your friends' privacy.

4. Make careful choice of parental control mechanisms: There are lots of parental control tools on the market that parents can install on children's devices. They offer a range of capabilities, including restricting the content or sites that your child can access, monitoring all their activities online, or sending alerts to parents for any inappropriate content or contact etc. Although these tools seem to offer peace of mind for parents, they may be pushed back or overcome by children[6]. If such tools are to be introduced at home, parents should talk to their children about why these tools are used. Parents and children should treat the tools as offering an opportunity to co-develop online risk coping skills. Every alert you received can be a great learning opportunity for you and your child, instead of an initiative for punishment.

This is a super-busy time for all families and everyone is experiencing a steep learning curve of all things digital. Be mindful that we cannot all be perfect, and this is a new experience for all of us! Therefore we also compiled a list of resources to help make this an enjoyable time for us all.

- Our privacy analysis of remote work tool (bit.ly/3bImhfe)
- Childnet's social media guide (bit.ly/3cH2XPP)
- Common sense media's guide on safe social sites (bit.ly/2RXThso)
- NSPCC's NetAware (bit.ly/2xRAHLA
- Google's Interland (bit.ly/2x4wDqY)
- LSE's blog on Parenting for a Digital Future (bit.ly/3cHb7YA)

This research was funded by IAA KOALA.

[1] bit.ly/34ZtYvc
[2] bit.ly/2VtTEgg
[3] bit.ly/34ZunxI
[4] bit.ly/2S0iS3O
[5] bit.ly/2Ku861s
[6] bit.ly/2yBOW7c

The COVID-19 infodemic: why false and misleading information is spreading online during the current crisis

By Helena Webb

The spread of COVID-19 has been accompanied by a spread of false and misleading information about the virus, its origins and its potential cures. Much of this misinformation has been shared online, particularly over social media. Fake claims about cures and remedies have gained significant traction and we are also seeing a growth in conspiracy theories – for instance connecting the onset of COVID-19 with the development of 5G technology.

The sharing of misinformation or unverified information is a social phenomenon that has occurred across human history. Studies within the sociology of rumour show us that misleading information particularly spreads at times of societal unease and tension. In these circumstances, it can provide certainty and fill knowledge gaps, allowing people to make sense of what is occurring around them. Given that we are currently in the midst of a huge global crisis, it is not surprising that unverified information is spreading as people attempt to understand what is happening and gain some feeling of certainty about the future.

However, although the spread of misinformation and unverified content is not new, the capacity for this content to spread online means that rumour and fake news surrounding COVID-19 may be unprecedented in both scale and scope. The factors contributing to this relate to the affordances of online platforms and the ways in which we use the Internet at the currrent time.

In the first instance, our capacity to share content with numerous others through social media means that information can have a very rapid and broad reach. Where this information is false, this speed of spread means that misleading ideas can take hold long before counter information is shared to correct it. Social media sites also enable like-minded users to gather together online and reinforce each other's viewpoints. Each of these factors is further enabled by the filtering and personalisation algorithms used by social media sites. Fast spreading content is highlighted in users' newsfeeds, helping it to spread further. Users are shown content similar to what they have already 'liked', reinforcing their views rather than challenging them to reconsider.

Further factors relate to our increased reliance on the Internet as a source of news and our 24-hour news cycle. The appetite for continuous and fresh news makes it more likely that content might be reported and then shared as fact before it is fully verified. Most major news organisations (both online and offline) now place emphasis on 'eye witness' accounts from the general public and encourage us to share our views and experiences. Whilst this greatly democratises participation in public discourse, it also enables unverified content to reach a wide audience – and perhaps gain greater credibility by being shared alongside verified content. Finally, those deliberately manufacturing misinformation to be shared online can take advantage of the fact that online users will frequently click on a link without reading its URL and that fake news pages can be designed to resemble those of trusted news sources.

This combination of social and technical factors means we are now experiencing a global 'infodemic' of false and unverified information about COVID-19. Tackling it will require significant action and we are now seeing various steps being taken by social media platforms to label or remove 'fake' content when it is posted. Similarly there are steps that individual users can take too, to critically assess the content they see and think about their personal responsibilities before sharing it online. Previously these platforms have typically argued against deleting content on the grounds of upholding freedom of expression, so this change in stance reflects the seriousness of the current infodemic. It is likely that we will only fully understand the way in which this 'infodemic' developed when we are post-pandemic and able to look back on and analyse what happened during this unusual and extreme situation.

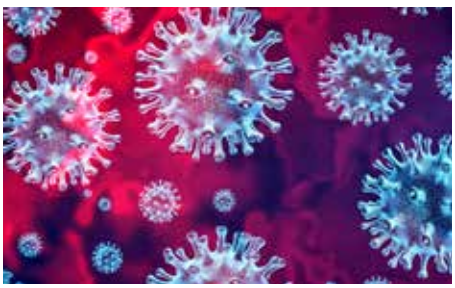

Restoring ancient Greek inscriptions using AI deep learning

Thea Sommerschield and Jonathan Prag (Faculty of Classics, Oxford), alongside Yannis Assael (Department of Computer Science and Google DeepMind) explain their work to restore damaged ancient Greek inscriptions using DeepMind AI.

Historians rely on different sources to reconstruct the thought, society and history of past civilisations. Many of these sources are text-based – whether written on scrolls or carved into stone, the preserved records of the past help shed light on ancient societies. Such sources include inscriptions, texts inscribed on a durable surface (such as stone, pottery or metal). Inscriptions are one of the main direct sources of new evidence from the ancient world, but the majority have suffered damage over the centuries, and parts of the text are illegible or lost (Figure 1). Restoring the missing or damaged text is one of the main undertakings of the discipline of Epigraphy. It is a complex and time-consuming task, but ancient historians can estimate the likelihood of different possible solutions based on context clues in the inscription – such as grammatical and linguistic considerations, layout and shape, textual parallels, and historical context. Although complex, the restoration of these documents is necessary for a deeper understanding of civilisations past.

Figure 1: Damaged inscription: a decree of the Athenian Assembly relating to the management of the Acropolis (dating 485/4 BCE). IG I3 4B. (CC BY-SA 3.0, WikiMedia)

We have been using machine learning trained on these ancient inscribed texts to build a system that can furnish a more complete and systematically ranked list of possible restoration solutions, which we hope will augment historians' understanding of a text.

Pythia

Pythia – which takes its name from the woman who delivered the god Apollo's oracular responses at the Greek sanctuary of Delphi – is the first ancient text restoration model that recovers missing characters from a damaged text input using deep neural networks. Bringing together the disciplines of ancient history and deep learning, this work offers a fully automated aid to the text restoration task, providing ancient historians with multiple textual restorations, as well as the confidence level for each hypothesis.

Pythia takes a sequence of damaged text as input, and is trained to predict character sequences comprising hypothesised restorations of ancient Greek inscriptions (texts written in the Greek alphabet dating between the seventh century BCE and the fifth century CE). The architecture works at both the character- and word-level, thereby effectively handling long-term context information, and dealing efficiently with incomplete word representations (Figure 2). This makes it applicable to all disciplines dealing with ancient texts (philology, papyrology, codicology) and applies to any language (ancient or modern). To train Pythia, the largest digital corpus of ancient Greek inscriptions (PHI Greek Inscriptions) was converted to machine actionable text (called PHI-ML). On PHI-ML, PYTHIA's predictions achieve a 30.1% character error rate, compared to the 57.3% of evaluated human epigraphists. Moreover, in 73.5% of cases the ground-truth sequence was among the Top-20 restoration hypotheses of Pythia, which effectively demonstrates the impact of this assistive method on the field of digital epigraphy, and sets the state-of-the-art in ancient text restoration.

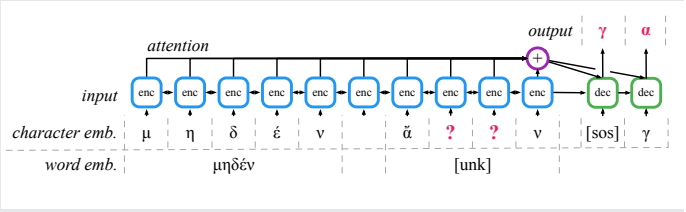

Figure 2: Pythia processing the phrase μηδέν ἄγαν (Mēdèn ágan) 'nothing in excess,' a fabled maxim inscribed on Apollo's temple in Delphi. The letters 'γα' are the characters to be predicted, and are annotated with '?'. Since ἄ??ν is not a complete word, its embedding is treated as unknown ('unk'). The decoder outputs correctly 'γα'.

The combination of machine learning and epigraphy has the potential to transform the study of ancient texts, and widen the scope of the historian's work. For this reason, the Oxford and DeepMind teams collaborated to create an open-sourced online Python notebook, Pythia, and PHI-ML's processing pipeline on GitHub. By so doing, we hope to aid future research and inspire further interdisciplinary work.

Read more about this work on the DeepMind blog here: bit.ly/2X7xqAc.

This article has been made available via the Oxford University Arts Blog.

The ‘Digital Twin’ to enable the vision of precision cardiology

Increasing computer power and novel algorithms could predict future heart health of patients.

In recent times, researchers have increasingly found that the power of computers and artificial intelligence is enabling more accurate diagnosis of a patient’s current heart health and can provide an accurate projection of future heart health, potential treatments and disease prevention.

Now in a paper published in *European Heart Journal*, researchers from Oxford University, working in collaboration with academic, industry, clinical and regulatory partners, and several other international initiatives as the Personalised In-Silico Cardiology consortium, show how linking computer and statistical models can improve clinical decisions relating to the heart.

Alfonso Bueno-Orovio (Oxford University Department of Computer Science) says, ‘the future of medicine is to provide therapies tailored to each patient. I am firmly convinced that the enabling pillar towards this vision is the increasing power of computers and algorithms to learn, reason and build the “Digital Twin” of a patient.’

The team have coined the term the ‘Digital Twin’ to describe this integration of the two models, a computerised version of our heart which represents human physiology and individual data.

‘The Digital Twin will shift treatment selection from being based on the state of the patient today to optimising the state of the patient tomorrow,’ the researchers wrote in the paper. This could mean that a trip to the doctor’s office could be a more digital experience.

‘The translation of the Digital Twin concept, which has been around in the engineering field for several years now, to cardiovascular research is a fundamental step in order to materialise the vision of precision cardiology’ said Francesca Margara, another Oxford member of the research team.

Mechanistic models see researchers applying the laws of physics and maths to simulate how the heart will behave. Statistical models require researchers to look at past data to see how the heart will behave in similar conditions and infer how it will do it over time.

Models can pinpoint the most valuable piece of diagnostic data and can also reliably infer biomarkers that cannot be directly measured or that require invasive procedures.

Pablo Lamata of King’s College London said more information about how the heart is behaving could be retrieved by using these models.

‘We already extract numbers from the medical images and signals, but we can also combine them through a model to infer something that we don’t see in the data, like the stiffness of the heart. We obviously cannot touch a beating heart to know the stiffness, but we can give these models with the rules and laws of the material properties to infer that important piece of diagnostic and prognostic information. The stiffness of the heart becomes another key biomarker that will tell us how the health of the heart is coping with disease.’

The team of researchers believe that the power of computational models in cardiovascular medicine could also provide us with more control over our daily heart health. Much like the popularity of wearable monitoring devices, a Digital Twin of our hearts could inform us about its current health and alert wearers to any risk factors.

Alfonso says, ‘future medical treatments will be tailored not only to current health status and data, but also to restore health by Digital Twin predictions.’

The research predicts we could see the technology in action within the next 5-10 years.