InspiredResearch Winter 2022 Issue 21

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

Modelling Antarctic Sea Ice Change

Read more – p12

Wearable authentication in mobile payments – p16

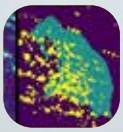

Autonomous detection of disaster events – p20

Automated 3D analysis of the heart – p24

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

CONTENTS

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

Letter from the Head of Department

Welcome to the Winter 2022 Issue of *Inspired Research*. In this issue, I'd like to introduce you to the five new faculty members who have joined us this Autumn.

Professor Christian Coester is an expert in algorithms for decision making under uncertainty. Such algorithms include online algorithms where decisions are made without the knowledge of future inputs and learning-augmented algorithms where algorithm design incorporates predictions of future events. Christian focusses on rigorous, provable results in these frameworks. **Professor Ronald Clark** is an expert on 3D computer vision, computer graphics, and their intersection. He focusses on photorealistic reconstruction and scene representation and understanding. His research has applications in robotics and in the physical sciences. **Professor Giuseppe De Giacomo** is an expert in knowledge representation and automated reasoning, a sub-field of Artificial Intelligence. He is currently leading the ERC Advanced Grant project WhiteMech, which aims to develop the science and tools enabling fully analysable self-programming mechanisms such as manufacturing devices and intelligent robots. **Professor David Parker** is an expert in quantitative verification. His research develops formal techniques for verifying quantitative properties of systems such as safety, reliability, and performance. David led the development of PRISM, which is the most widely-used software tool for verification of probabilistic systems. **Professor Nobuko Yoshida** is the new Christopher Strachey Professor of Computing. Nobuko's research focusses on the foundations of concurrent and distributed systems. Her particular specialism is research on session types which is a type-based approach to the verification of message-passing programs. Nobuko develops theories, semantic concepts, and supporting programming languages and software to support her analysis of concurrent systems. She currently holds an EPSRC Established Career Fellowship. We are delighted to welcome Christian, Ronald, Giuseppe, David, and Nobuko to the department.

In this issue, you can read about the exciting ongoing research of the department. One highlight is new work led by Professors Niki Trigoni and Andrew Markham using self-supervised learning to help autonomous vehicles to achieve safer and more reliable navigation capability, especially under adverse weather conditions and driving situations without GPS. The department congratulates Professor Trigoni on her recent election as a Fellow of the Royal Academy of Engineering. Another recent highlight is work by researchers Sebastian Köhler, Richard Baker and colleagues. This work, which was featured in *New Scientist,* demonstrates a new cyberattack that uses radio waves to interfere with image-recognition systems.

I'd like to close by congratulating our Intercollege Programming Team, which has recently taken a bronze medal in the ICPC World Finals. This team consists of (now graduated) students Stefan Constantin-Buliga, Lukas Michel, and Costin-Andrei Oncescu. It is coached by Professor Bartek Klin and sponsored by OxFORD Asset Management. I also congratulate the whole department for topping the *Times Higher Education* world ranking for Computer Science for the 5th consecutive year.

Enjoy the winter issue.

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

Departmental Programming Team takes Bronze Medal in ICPC World Finals

On November 10 2022, in Dhaka, Bangladesh, a team comprising Stefan Constantin-Buliga, Lukas Michel and Costin-Andrei Oncescu, represented the University of Oxford in the 45th World Finals of the International Collegiate Programming Contest (ICPC). These were the delayed finals of the 2020/21 season. Because of this delay all our team members have already graduated, and we are very happy that they agreed to practice hard and represent Oxford one last time.

The team performed very well. Solving 8 out of 12 problems got them 11th place (out of 132 teams from all around the world, selected in the 2020/21 season of regional contests), and a bronze medal. They were very close to solving a ninth problem (which would have given them 6th place and a silver medal). The finals were won by a very strong team from MIT, the only team to solve 11 problems.

Our competitive programming teams are proudly supported by OxFORD Asset Management, a technology driven investment firm.

The team is pictured below, from the left: Lukas, Costin and Stefan.

NEWS

Hello from Sophie Gibbons, our new Head of Administration and Finance

I am delighted to have joined the Department of Computer Science as the Head of Administration and Finance (HAF), thank you all for welcoming me so warmly since I started in September and while I still settle into the department.

Immediately before joining the department, I spent six years as the HAF at the University of Oxford, Department of Social Policy and Intervention. Prior to joining the University, I served 13 years as the Business Administration Manager heading up the Children's Social Care and Youth Justice Service for Oxfordshire County Council.

As the HAF, I work alongside the Head of Department (HoD) and the Mathematical, Physical and Life Sciences (MPLS) divisional teams to provide strategic leadership in relation to governance, policy, planning, growth and financial sustainability.

I lead the department's team of Operational Managers and Professional Services Staff (PSS) in the following areas:

- Human Resources
- Recruitment and retention
- Budgets and financial control
- Research facilitation, pre and post-award grant management
- Health & Safety, business continuity planning, and security
- UG and PG academic administration
- IT
- Communications and outreach
- CDTs
- Capital planning

Although, along with other departments in the Division, we are navigating our way through some challenges including the rising cost of utilities, there are exciting examples of growth and innovation. I have joined the Senior User Group meetings for the new Informatics Building. This will bring together Computer Science and Information Engineering located in the Radcliffe Observatory Quarter and drive forward our vision to be the leading Informatics centre in the world. The planning stage is exciting and intricate work, and I look forward to sharing progress with you.

Another positive move for the department is to increase the number of Associate Professorships with Tutorial Fellowships as part of the recruitment strategy approved by MPLS Division. This represents an excellent potential for growth in the department and diversification of our research portfolio and continuation of our excellent standards of teaching.

These areas and more make me feel very privileged to be part of such a dynamic and successful department, working alongside exceptional academics, research, and professional services colleagues.

Jennifer Watson: New Learning Technologist

I am Jennifer Watson the Department of Computer Science's new learning technologist. I am part of the IT support team within Computer Science and my work is complementary to the work carried out by the rest of the CS IT support team.

This role is new for the department but it is also new to me, as I previously worked as a Secondary School teacher. I spent 15 years as a Maths teacher and 3 years as a Computer Science teacher. During this time, I realised that the thing I was most passionate about was finding out how to use technology in ways that enhanced my teaching and supported the learning of my students.

I am interested in lots of different uses of technology and aim to support the academics by running training/demonstration/induction sessions, workshops, 1:1 support and just being there to try and answer their questions. I will be making help videos, help guides and other documentation. I cannot promise to have all of the answers but I am there to spend time investigating a solution, working collaboratively with others in the University.

Alongside the normal day-to-day support with existing technology in the department, I have been spending much of my time working on Learning Management Systems (such as Moodle and Canvas), Video recording and editing methods (including Panopto, Lightworks and RapidMooc), 3D imagery (AR, VR, SketchFab, ThingLink) and all sorts of useful websites which offer tools such as polls, quizzes, notes, forums and whiteboards.

I will be working with academics and staff within the Department of Computer Science, and I will be talking to students to get their feedback on whether the techniques we are using adequately support them in their learning. I will also be talking to other departments to see what lessons we can learn from their experiences. I like to also keep an eye on what is going on outside of the University in terms of new technologies and initiatives.

The most important thing you need for this job is the confidence to talk to anyone and the enthusiasm to get stuck into anything. Come and speak to me about your use of technology in teaching or learning - I would love to hear about it.

You can contact me at jennifer.watson@cs.ox.ac.uk

Oxford tops *Times Higher Education* world ranking for Computer Science for fifth consecutive year

The University of Oxford has topped the *Times Higher Education* world ranking for Computer Science for fifth consecutive year, in their newly released subject tables.

The Computer Science subject table uses the same performance indicators as the *Times Higher Education* World University Rankings 2022, but the methodology has been recalibrated to suit the discipline. This year's table includes 974 universities, up from 892 last year. Read more about the results here: https://bit.ly/3VWie6n

The University of Oxford as an institution also topped the *Times Higher Education*'s World University Rankings for the seventh consecutive year. https://bit.ly/3uMNDMA

Department team wins UN PET Lab Hackathon

In November a Department of Computer Science team took part in the United Nations Privacy Hackathon and placed 1st out of 196 teams (325 participants) worldwide! Our team consisted of Manuj Mishra, Vishal Ramesh, and Konrad Kollnig (pictured below, right). Manuj and Vishal are each pursuing the MSc in Advanced Computer Science while Konrad is a 4th year CS DPhil student with a focus on privacy.

The competition included university teams from Harvard, Zurich, and Toronto as well as teams from industry, start-ups, and national agencies including the UK's own Office for National Statistics, who came 3rd. The competition ran over 4 days from the 8th to the 11th November and had physical locations in Yogyakarta, Boston, and Zurich. The Oxford team took part remotely.

The challenge involved extracting valuable insights relating to refugees in Kenya during the COVID-19 crisis – such as the probability of their receiving government aid or the probability of them wanting to return to their home country. To do this the team had to make use of sensitive information collected in surveys by the UNHCR (United Nations High Commissioner for Refugees).

The competition was similar to other data science challenges with the exception that participants had to make use of state-of-the-art PETs (privacy-enhancing technologies) like differential privacy, secure enclaves and synthetic data to maintain the privacy of individuals.

As part of their prize, the team will be presenting their work at the 2023 UN Big Data Conference and will be writing a blog post with the organisers. Team members also won an iPad each.

Sponsors of the competition included UN PET Lab, United Nations, UNHCR, Swiss Federal Statistical Office, IBM Research, OpenDP, OpenMined, Microsoft, and Amazon Web Services (AWS).

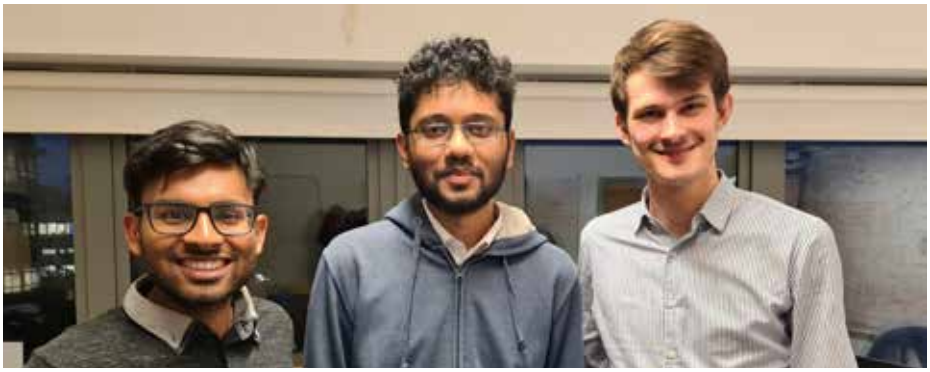

MPLS Equality, Diversity and Inclusion Fellows

The MPLS (Mathematical, Physical and Life Sciences) Division strives to create an inclusive culture where different perspectives and experiences are encouraged, and each person is respected and treated fairly, so that individuals of all backgrounds can fully contribute their talents to our scientific community.

There are incredibly talented individuals across the Division who are passionate about equality, diversity and inclusion (ED&I) and are dedicated to making positive change within our institution. Gunshi Gupta from the Department of Computer Science is one of 10 ED&I Fellows announced for the 2022-2023 cohort. Gunshi comments,

'I'm Gunshi, a first year DPhil student in CS/Machine Learning from New Delhi, India. As an incoming ED&I fellow I'm excited to share experiences with the cohort, identify common goals and contribute on two fronts: Making STEM a place where BIPOC students and individuals with neuro-divergences can thrive. I hope to apply my experience being a gender and racial minority in different Robotics/CS labs to help inform policy-making for fostering inclusivity in these spaces, especially in a way that doesn't try to always overlook people's cultural backgrounds but instead celebrates and appreciates it.'

The 2022-2023 cohort of MPLS Equality, Diversity and Inclusion (ED&I) Fellows will help advance ED&I work within the Division. If you're interested in getting involved with diversity in MPLS, please contact diversity@mpls.ox.ac.uk.

News in brief

Kasper Rasmussen has been chosen to receive the title of Professor of Information Security in the Recognition of Distinction Exercise for his work on security protocols and embedded systems. Kasper has worked for the Department of Computer Science since 2013. Before coming to the University of Oxford, Kasper worked at the University of California, Irvine and ETH Zurich. His area of research specialises in information security in different contexts specifically centred on protocol design, applied cryptography and security of embedded and cyber-physical systems.

We are thrilled that the amazing research being done here at the Department of Computer Science has been recognised with three MPLS Impact awards and one commendation. Well done to all involved!
Professor Tom Melham won an MPLS Commercial Impact Award for contributions to a commercial formal verification product. **Professor Yarin Gal** and colleagues **Jan Brauner, Soren Mindermann** and **Mrinank Sharma** won an MPLS Social Impact Award for Understanding the effectiveness of interventions against COVID-19, using Bayesian models,
Professor Seth Flaxman won an MPLS Impact Award (Covid-19) for Computational methods underpinning the Global Reference Group on Children Affected by COVID-19, and **Departmental Lecturer Atılım Güneş Baydin** won an MPLS Social Impact Commendation for Artificial Intelligence launched into space to detect floods.

Read more here: https://bit.ly/3rvHbrQ

Oxford partners with the Optiver Foundation to launch new international postgraduate scholarship programme for women in STEM

A new scholarship programme at Oxford will increase the number of women from low- and middle-income countries who take up offers to study science, technology, engineering and mathematics (STEM) subjects at postgraduate level. The programme, which will provide support for 30 taught master's students over a period of five years, has been made possible by a donation from the Optiver Foundation.

The first Optiver Foundation Scholars began their studies at Oxford in October 2023. Scholars may join the Departments of Statistics, Computer Science, Physics, Chemistry and Engineering Science, or the Mathematical Institute, and will receive full financial support for their course fees, as well as a stipend to cover living costs. The donation also enables the provision of tailored induction activities, mentoring and on-course support for the scholars. Thanks to the generosity of the Optiver Foundation, Oxford will also develop and pilot a new way of assessing the socio-economic background of international applicants. This will ensure that the funding provided through the new Optiver Foundation Scholarship Programme is finely targeted at eligible offer-holders who are most in need of financial support. Key findings from this pilot will be used to inform future scholarship programme development at the University.

By targeting support at women from low- to middle-income countries, the Optiver Foundation Scholarship Programme will play an important role in increasing postgraduate diversity in STEM subjects at Oxford, as well as contributing to a more a diverse workforce in the longer term.

Find out more: http://bit.ly/3VfmTQf

DPhil student Arthur Laudrain makes prestigious list of potential future leaders

DPhil Cyber Security student Arthur Laudrain has been named one of 35 promising young minds in the Santander-CIDOB 35 under 35 List. The list recognises the work of 35 people under the age of 35 with the potential to become future leaders in their areas of expertise.

Arthur (who is currently on a placement with the Department of Politics and International Relations) was asked if he could provide a brief statement for the Future Leaders Forum, a collaboration between Banco Santander and CIDOB. The Forum is a platform for the cultivation of ideas among those 35 nominated to the list and an opportunity for them to discuss key issues on the global stage. Arthur spoke at the event on his vision of EU digital strategic autonomy, defining what it means and what the priorities surrounding it should be.

Speaking on his nomination to the list, Arthur said: 'It feels very humbling to be surrounded by such an accomplished, engaged, and diverse crowd. I'm very much looking forward to getting to know everyone better and explore potential collaboration.'

Law and Computer Science course shortlisted for Vice-Chancellor's Education Award

We are delighted that the Law and Computer Science course co-founded by Professor Tom Melham was shortlisted for the biennial Vice-Chancellor's Education Awards. These awards are designed to celebrate and showcase educational activity at Oxford that has a clear and demonstrable impact upon the educational experience of our students.

Law and Computer Science is a unique and ground-breaking course, offered as an option in postgraduate master's degrees in both the Law Faculty and the Department of Computer Science at Oxford. It is also available as a fourth-year option in all three Oxford undergraduate courses in Computer Science.

This multi-disciplinary course stems from world-leading research into the use of AI in the provision of legal services and the mindset, skills and knowledge gaps between lawyers and computer scientists.

It allows students from both disciplines to gain a unique perspective on their own specialist areas as well as learning to communicate and collaborate effectively to respond to the challenges and opportunities of digital tech.

Students are supported in their highly innovative practical work by a group of industry mentors and sponsors, who bring a strong commercial focus as well as sharing invaluable practical insights, data, and technical tools.

Open Days – welcome back to Oxford!

During the pandemic like so many other events our department open days became online events. It has only been this year that it has been possible to invite visitors back into our buildings and there was a real buzz as our student ambassadors and academics hosted three packed days of events for anyone interested in applying to do our undergraduate courses.

DeepMind offers six scholarships for under-represented students

DeepMind, the leading British artificial intelligence company, has renewed its commitment to supporting students at the University of Oxford following the success of its scholarship programme for under-represented groups in AI.

Five more DeepMind graduate scholarships for students wishing to pursue a master's degree in Advanced Computer Science in the Department of Computer Science will be made available for students commencing courses in 2023–24. Three of these scholarships are open to individuals ordinarily resident in the UK who belong to one or more of the following groups:

• identifying as female and/or
• of Black, Bangladeshi or Pakistani (and relevant Mixed) ethnicity and/or
• from more disadvantaged socio-economic backgrounds

Two scholarships will be open to students not ordinarily resident in the UK. These scholarships are open to female students and/or those ordinarily resident in an under-represented region of the world.

A DPhil scholarship will also be funded, for a student undertaking either a DPhil in Engineering Science, a DPhil in Computer Science, a DPhil in Statistics or in the Autonomous Intelligent Machines and Systems EPSRC Centre of Doctoral Training (CDT). This scholarship is open to female students and/or those not ordinarily resident in the UK*, or for UK residents meeting the same criteria listed above for the Masters.

There is no separate scholarship application process for these DeepMind Scholarships. All eligible candidates will be automatically considered when applying to one of the relevant courses ahead of their respective deadlines.

News in brief

Professor David Gavaghan has been appointed in one of three University-level leadership roles created by Oxford University to drive improvements in research culture. The three Academic Leads plan to work with the community and with the sector to create an environment that benefits all Oxford researchers.
Read more at https://bit.ly/3CvIm1A

MSc student Hunar Batra has been awarded the Google Women in Computer Science Generation Scholarship Award. This prestigious award is given to Women in Computer Science for demonstrating outstanding academic achievements and research work in Computer Science. Hunar is one of only four students in the UK to be selected for this award. The Google award website page explains, 'These students have demonstrated a passion for technology, academic excellence, and have proven themselves as exceptional leaders and role models. They will join a community of over 2,800 Google scholarship recipients.'

In a new video, Professor Sam Howison, Head of MPLS Division, talks about the benefits that leaving a legacy can bring (you can watch it here http://bit.ly/3JBKCoe). As we look to the future, we must ensure that our research endeavour and teaching provision is sustainable for generations to come. Leaving a Legacy gift to Computer Sciences will help support our leading research programmes and exceptional students. Whatever the size, and whether for graduate scholarships, academic positions or to support core activities, every gift is greatly appreciated and contributes to our ongoing success. If you would like to know more about leaving a Legacy, contact Caitlin Tebbit at caitlin.tebbit@devoff.ox.ac.uk

Professor James Worrell receives UKRI Frontier Research Grant

Professor James Worrell has been awarded a £1.6m grant for his research project titled BELINDYSYS (Beyond Linear Dynamical Systems). Its goal is to achieve a major advance in the algorithmic theory of fundamental dynamical systems arising in program verification, control theory, and related areas.

Dynamical systems pervade the quantitative sciences: for example, recurrence sequences appear across Computer Science, combinatorics, number theory, economics, and theoretical biology. Characteristically such systems are simple to describe and yet they have a rich algorithmic and mathematical theory. Over the past ten years, James has studied applications of linear dynamical systems to automata theory, loop termination, and hybrid systems. A major step forward in the BELINDYSYS project involves analysing non-linear models, including conditional branching, non-determinism, external control, and polynomial recursivity. If successful, this five-year project will make significant progress on longstanding open problems and will open new lines of research at the boundary of Computer Science and mathematics.

BELINDYSYS was submitted to the European Research Council (ERC) Advanced Grant call in 2021 and received a successful outcome. However, because the UK has not yet associated to Horizon Europe, funding for this grant has been facilitated by UKRI and it will be known as a UKRI Frontier Research grant.

New Fellowship for Professor Niki Trigoni

At its AGM on 20 September 2022 the Royal Academy of Engineering elected Professor Niki Trigoni as a Fellow. She is one of 72 leading figures in the field of engineering and technology newly appointed to the Fellowship.

The new Fellows have all made exceptional contributions to their own sector, pioneering new innovations, leading progress in business or academia, providing high-level advice to government, or promoting wider understanding of engineering and technology. New Fellows were formally admitted to the Academy in London on 8 November, when each Fellow signed the roll book.

Sir Jim McDonald FREng FRSE, President of the Royal Academy of Engineering, said: 'I am delighted to welcome such an array of enormously talented people to the Fellowship of the Royal Academy of Engineering. From industry and enterprise to education and government – both national and international – these are some of our most pioneering and distinguished engineers and technologists.'

Professor Blanca Rodriguez appointed to Scientific Advisory Committee

On 11 October 2022 the European Commission Joint Research Centre formally appointed Professor Blanca Rodriguez as one of nine experts as the new members of the European Union Reference Laboratory for Alternatives to Animal Testing (EURL ECVAM) Scientific Advisory Committee (ESAC). Renewed every five years, ESAC is composed of external scientists who are appointed on the basis of their scientific expertise and who act independently in the public interest. Blanca comments, 'I am very excited by my appointment to the Scientific Advisory Committee for the European Union Reference Laboratory for Alternatives to Animal Testing. I will contribute with my expertise on computational/in silico methodologies in biomedicine.'

Vodafone and the University of Oxford tackle bias bots

A research collaboration framework has been initiated and established by Jindong Hou at Vodafone and Professor Thomas Lukasiewicz, Oliver Sampson, Associate Professor Ani Calinescu and Doctoral Student Tatiana Botskina at the Department of Computer Science, University of Oxford in the last three years. Thanks to this initiative, Vodafone is working with the University of Oxford to embed ethical AI throughout the customer journey, to ensure fair, transparent, explainable, and responsible AI.

Jindong Hou, Lead Research & Innovation Architect at Vodafone, explained: 'In the age of big data, a company that harnesses the power of AI will have a competitive edge. However, as AI grows in usage and impact across geographies and industries, we have the responsibility to consider how it affects our customers, our employees, and wider society.

'We are looking at the ability of deep learning models to understand hidden patterns and dependencies. These will allow us to better analyse and gain insights on complex business processes, such as customer journeys, to detect and remove human bias,' continued Jindong.

A proposal written by affiliates of the University of Oxford reinforces this view. It describes the benefits of combining attention-based machine learning models (which break down a complex task such as a customer searching for, and buying, a product) with causal inference (working out the cause of a decision from the data available). Getting to grips with this will help Vodafone further understand what frustrates customers and what delights them.

JPMorgan Chase Faculty Awards 2022

Two members of the Department of Computer Science have received awards from JPMorgan Chase (JPMC) to advance their research. Professor Alessandro Abate has received funding to conduct research into 'Learning, Understanding, and Planning in Repeated Games with Partial Knowledge and Information.'

The JPMC-funded project, to which Professor Michael Wooldridge will act as an expert adviser, will enable Alessandro – along with his research assistant Licio Romao, and his DPhil student Will Lee - to investigate a new approach to understanding how autonomous agents repeatedly interact with each other, whether collaboratively or competitively, in the presence of partly known and memory-dependent (stateful) environments that affect them, and that are also possibly affected by their actions. The project builds on an ongoing research initiative in this area by Alessandro's group.

Professor Paul Goldberg's JPMC award has been made for his research entitled 'Price Discovery via Decentralised Networks of Trading Agents.'

The project studies the behaviour of markets composed of many trading agents who buy and sell goods from each other. One basic functionality expected of these markets is the ability to identify 'fair' prices. In their research, Paul and his researcher, Edwin Lock, will address the following issue: what are the obstacles to finding fair prices, and to what extent do they cause the performance of a system to degrade?

News in brief

Doctoral student Ge Tiffany Wang received an Honourable Mention at the Doctoral Researcher Awards for her work on *Informing age-appropriate AI*. The Doctoral Researcher Awards is a UK wide competition for junior academics.

Ralph Abboud, a final-year DPhil student has been awarded the 2022 G-Research PhD Prize in Maths and Data Science, in recognition of his doctoral research in Machine Learning. The prize is awarded to Oxford or Cambridge University students (in a quantitative discipline) in their ultimate or penultimate PhD years. Up to £10,000 is given to the best PhD draft dissertation entered into the competition.

Department Teaching Awards have been given to Professor Bartek Klin and Mate Szabo. Department Teaching Commendations have been given to Professors Jon Barrett, Bernardo Cuenca Grau, Leslie Goldberg, Joe Pitt-Francis and Michael Wooldridge. They are all awarded on the basis of excellent student feedback. Despite the 2021-2022 academic year being disrupted by the effects of the pandemic, the department was still able to deliver teaching of excellent quality. This is thanks to the dedication and genuine love for teaching demonstrated by our distinguished faculty members, as well as the enthusiasm and dedication of our students and researchers. This set of teaching awards celebrates these contributions.

OxWoCS

Oxford Women in Computer Science (OxWoCS) was founded in 2013 with the aim of supporting, promoting, and empowering women and non-binary people in computer science to make an impact in their field. OxWoCS has since expanded, with members across multiple departments in Oxford. The committee consists of a mix of undergraduates, postgraduates, and postdoctoral researchers, from a range of backgrounds. Together they bring knowledge on deep learning for imaging, medical statistics, computational social science, the application of AI to tackle human trafficking, and AI applications in space.

Throughout the academic year, OxWoCS runs a variety events, including the distinguished Seminar Series lectures, social, and industry events. Each year they connect with the Cambridge women@CL society to run the annual Oxbridge Women in Computer Science Conference. Recently, the conference was held at the Mathematical Institute in Oxford and showcased research posters from students in the community, a panel of alumni on kickstarting your career, and a workshop held by Microsoft.

As well as running events for those in Computer Science, OxWoCS also run outreach events for those who want to get into the field. The OxWoCS Challenge Club (run in collaboration with the outreach team at the Department of Computer Science) is a digital programme for Year 12 girls at UK state schools consisting of monthly sessions going through mathematical and Computer Science related problem sheets. GirlsWhoML (started by members of OxWoCS) is another outreach programme to encourage those who identify as female and non-binary into the field of AI and ML. GirlsWhoML holds regular hands-on workshops at universities around the UK, all run by volunteers and completely free to students. OxWoCS likes to keep in touch with alumni from Computer Science or related departments at Oxford. In Summer 2023 they will hold an Alumni Dinner in London, a chance for alumni to reconnect and network. Please visit tinyurl.com/OxWoCSAlumniDinner to register your interest for the alumni dinner and to receive more details closer to the event (women and non-binary people only please).

If you're interested in learning more about OxWoCS, check out their website at oxwocs.com.

ReadySetCode – A Computer Science/ Christ Church Outreach initiative

ReadySetCode is a new sustained engagement programme of three online sessions, beginning in January 2023 and ending in March 2023 with an in-person graduation day at Christ Church, Oxford. The initiative is for 11 (15–16 year olds, Scottish S5 or NI Y12) students of black and mixed heritage who identify as female, from state-funded schools in the UK.

The online sessions, led by members of the Department of Computer Science and Christ Church, will offer students in-depth investigations of topic areas outside of the school curriculum, combining theory with interactive exploration. The sessions will introduce students to Oxford University teaching methods and will involve small-group work.

The final day of the programme will be an in-person graduation day held in Oxford. Sessions designed to support students in their progression to university – for example, in supporting A Level choice – will also form part of the programme alongside the academic elements. There will also be a careers panel talk with Computer Science alumnae during the programme and a chance to meet current students.

For more information about the initiative, visit: bit.ly/3ECjW75

Alumni Profile

James Pavur – An expert in the USA's Directorate for Digital Service within the Office of the Secretary of Defense.

What course did you study at Oxford University and when?
I did a DPhil in Computer Science (2017-2021) as part of the EPSRC Centre for Doctoral Training in Cyber Security. My supervisor was Professor Ivan Martinovic and I was funded courtesy of a Rhodes Scholarship. My DPhil thesis was about the cyber security of space systems, with particular focus on satellite communications and spaceflight coordination.

What was your background before that?
My background was a bit unusual for a CS DPhil actually! My undergrad was at Georgetown University's School of Foreign Service in Washington DC. I received a Bachelors of Science in Foreign Service, majoring in 'Science, Technology and International Affairs.' I've always been interested in the intersection of technology and policy.

What attracted you to studying Computer Science as a subject?
I'm a hacker at heart. Even as a kid, I loved taking toys apart to see how they worked and then putting them back together again. As the world around us becomes ever more digital, studying Computer Science presented an opportunity to peel back the hood and understand some of the magic that makes modernity. It's also just really fun to break systems and then fix them to be even more secure.

What aspects of the course did you particularly enjoy?
For me, it was the community around Oxford. Both within the lab and the broader University I was surrounded by so many brilliant people who managed to simultaneously be humble and challenge me. Beyond Oxford, I had the opportunity to do research work with international collaborators, travel the world to present at conferences, and feel like I was part of a collective effort much bigger than myself. The day-to-day of a DPhil is often solitary, but the highlights for me were all communal.

What did you do when you left Oxford?
I moved back to Washington, DC and joined the Defense Digital Service. Our mission is to help the Department of Defense tackle really hard technical problems and respond to urgent needs. My work runs the gamut from software engineering, to cyber security, to policy guidance; but the one constant is the sense that I'm empowered to build technologies that improve lives and make the world a safer place.

How has the course helped you in your current career?
During the DPhil, I learned a lot about how to be an effective communicator to diverse audiences. Being able to explain complex technical problems to stakeholders with different backgrounds is a superpower in the public sector and the DPhil really helped me develop that skillset.

What advice would you give to current students on applying their knowledge in the workplace, when they leave university?
In my experience so far, it's less about what you learn and more about the skills you develop on the journey to get there. Sometimes the topics I specialised on in my DPhil feel relevant to work outside of academia, but most of the time it's more the technical and soft skills I cultivated in the process of researching those topics that come in handy. As you leave the university and head to the workplace, thinking about what sort of work you've enjoyed doing – rather than simply what you feel most knowledgeable about – is a good heuristic for finding a career fit.

What would the student you have thought about what you are currently doing – would you have been surprised, proud, amazed?
At the start of my undergraduate degree, my goal was to become a cyber security lawyer, so I've definitely strayed a bit from that trajectory. That said, it's also not like I've run off and joined the circus – so I'll go with 'mildly surprised.' Even when I was a student, I felt that there was an acute need for people with deep technical backgrounds to shape policy around cyber security and technology. While I'm not there yet, hopefully I'm on a trajectory to make student me proud someday!

Alumni groups

The Responsible Technology alumni network includes Oxford alumni from various disciplines interested in responsible technology with the aim of breaking down disciplinary silos and connecting alumni across the world. We engage in a variety of activities from networking events to workshops and campaigns to raise awareness about relevant issues and we welcome all interested alumni, current students and other people affiliated with the University to join us here: https://bit.ly/3Thl3xm

The University of Oxford is a global leader in Artificial Intelligence (AI), producing cutting-edge research and supporting innovation in applied AI. We are launching an official **University of Oxford alumni network for artificial intelligence**, providing a community for students, researchers, academics and alumni from all backgrounds to engage with AI from a wide range of perspectives. Find out more here: https://bit.ly/3eMnfy3

Using Random Forests to Model 20th Century Antarctic Sea Ice Change

By Doctoral student Alfie Brazier

During the summer, I ventured to Hobart, Tasmania to work at the Institute for Marine and Antarctic Studies (IMAS). I have long wanted to use my Computer Science knowledge to aid the fight against climate change and this was a prime opportunity to do just that. I worked under Will Hobbs in the sea ice team as part of the Australian Antarctic Program; our goal was to model 20th century Antarctic sea ice change using chemical data taken from Antarctic ice cores.

The first satellites capable of detecting Antarctic sea ice to a sufficiently useful level of detail launched in the 1970s. These satellites use passive microwave radiometers, which detect radiation thermally emitted from the Earth's surface, and are able to determine the concentration of sea ice in a given area. That complete records only exist from 1979 onwards is problematic, as this time period is not long enough for us to draw solid conclusions regarding the impact of anthropogenic forcings on sea ice levels. Therefore, a 20th century reconstruction of such records would be useful in understanding how humans may have affected the extent of annual sea ice over approximately the last 125 years.

Current reconstructions of sea ice records rely on linear regressions using the chemical markers found in ice cores. On expeditions to Antarctica, teams of geoscientists dig trenches and drill down into the ice to extract large cores, which are either analysed in situ or kept frozen until they can be chemically analysed in the lab. The depth within the core from which a sample is taken correlates to the year in which said sample was formed, allowing us to construct a time series for each chemical marker. It is these time series, in conjunction with the satellite records, which are used to construct a linear regression for 20th century sea ice.

However, due to decadal shifts in wind patterns that drag snow over sea ice and then deposit it at the ice core site, many chemical markers have a non-linear correlation to sea ice extent, and so it can be challenging to reconstruct wider sea ice records out of historic proxies if we constrict ourselves to only linear regressions. My models showed that this difficulty may be overcome by combining those proxies using a non-linear approach - in this case a random forest.

The regressions for sea ice are three-dimensional: longitude section, season of year, and extent. I built several random forest models, each of which grouped varying sections of these dimensions together to try and find the sweet spot in which the ratio of predictors to predicted values allowed for most accurate prediction, whilst also providing insight into how chemical markers correlated to different areas of sea ice throughout the seasons.

A random forest model combines multiple decision trees to form a 'forest' which outputs a prediction. Each of these decision trees is essentially a flowchart, with questions regarding the levels of different chemicals in the sample, or season of the year, leading to branches with different predictions for the sea ice extent. The 'forest' contains lots of decision trees, each asking slightly different questions and getting slightly different predictions. At the end of the process, each tree casts its vote and the forest chooses the prediction which garnered the most votes.

Upon finishing my time at IMAS, I am proud to say my team remarked that my model outputted significantly better results than anything in the current literature and represented a step change in the field of sea ice prediction. I hope to publish my research next year and pursue further work in bringing together the fields of computer and climate science.

Oxford University Researchers Develop New AI to Enable Autonomous Vehicles to Adapt to Challenging Weather Conditions

Researchers at Oxford University's Department of Computer Science, in collaboration with colleagues from Bogazici University, Turkey, have developed a novel artificial intelligence (AI) system to enable autonomous vehicles (AVs) achieve safer and more reliable navigation capability, especially under adverse weather conditions and GPS-denied driving scenarios. The results have been published in *Nature Machine Intelligence.*

Yasin Almalioglu, who completed the research as part of his DPhil in the Department of Computer Science, said: 'The difficulty for AVs to achieve precise positioning under challenging adverse weather is a major reason why these have been limited to relatively small-scale trials up to now. For instance, weather such as rain, fog, or snow may cause an AV to detect itself in the wrong lane before a turn, or to stop too late at an intersection because of imprecise positioning.'

To overcome this problem, Yasin and his colleagues developed a novel, self-supervised deep learning model for ego-motion estimation, a crucial component of an AV's driving system that estimates the car's moving position relative to objects observed from the car itself. The model brought together richly-detailed information from visual sensors (which can be disrupted by adverse conditions) with data from weather-immune sources (such as radar), so that the benefits of each can be used under different weather conditions.

The model was trained using several publicly available AV datasets which included data from multiple sensors such as cameras, lidar, and radar under diverse settings, including variable light/darkness levels and precipitation. These were used to generate algorithms to reconstruct scene geometry and calculate the car's position from novel data. Under various test situations, the researchers demonstrated that the model showed robust all-weather performance, including conditions of rain, fog, and snow, as well as day and night. The high level of location accuracy provided by the self-supervised deep learning framework GeometRy-Aware Multi-Modal Egomotion estimation (GRAMME) enables AVs to reliably understand their environment and make safer decisions. Researchers demonstrate that the complementary and redundant perception that AVs gain from multiple sensors improves the reliability of vehicles in challenging situations, especially in adverse weather conditions. Besides, the self-supervised aspect of the model enables AI systems deployed on AVs to learn localisation from orders of magnitude more data, which is important to quickly recognise and understand new driving conditions.

The main potential of the AI model is to act as a complementary software solution for positioning to improve the performance of AVs, especially in challenging conditions, as demonstrated by the study. The model improves the safety and reliability of AVs by providing precise location information as a basis for numerous core software components such as perception, planning and control.

The team anticipate that this work will bring AVs one step closer to safe and smooth all-weather autonomous driving, and ultimately a broader use within societies.

Professor Niki Trigoni, who co-supervised the study, said: 'The precise positioning capability provides a basis for numerous core functionalities of AVs such as motion planning, prediction, situational awareness, and collision avoidance. This study provides an exciting complementary solution for the AV software stack to achieve this capability.'

Professor Andrew Markham, also a co-supervisor for the study, added: 'Estimating the precise location of AVs is a critical milestone to achieving reliable autonomous driving under challenging conditions. This study effectively exploits the complementary aspects of different sensors to help AVs navigate in difficult daily scenarios.'

The full paper, Deep learning-based robust positioning for all-weather autonomous driving, is published in 'Nature Machine Intelligence': bit.ly/3BczvzO

Developing a Responsible Sustainability Framework for the Digital Economy

By Lucas Somavilla – Research Associate, Responsible Technology Institute (RTI)

The world is rapidly warming due to emissions produced by anthropogenic activity. Action to address climate change through understanding and reducing emissions is urgently required across all sectors, including the digital economy.

In this context, the PARIS-DE project (EPSRC) brings together researchers from Lancaster University, University of Oxford, King's College London, University of Sussex, and Small World Consulting to investigate and co-design an actionable framework for sustainable and responsible innovation in the digital economy.

The expansion and utilisation of digital technologies in the economy result from digital innovations driving new exponential and ubiquitous connectivity via data from devices, applications, and digital platforms. Innovations such as cloud computing, machine learning, natural language processing, and distributed ledger technologies require a significant amount of energy consumed intensively during the whole innovation cycle and use, producing considerable carbon emissions and environmental impacts. It is estimated that carbon emissions from the ICT sector are currently between 2.1 and 3.9% of global carbon emissions, and it is expected these figures will continue to increase.

Digital innovations play an important role in helping us understand our environment and the impact we have on it. For example, sensors in nature and our homes provide significant amounts of data to be processed through ever-improving computers and AI systems. However, deep uncertainties abound in the measurement and management of carbon emissions resulting from digital innovations and the tensions between short-term versus long-term solutions to enduring social, economic, and environmental crises.

As the project progresses, we have learned that digital innovation has transformative impacts on the economy and broader society, posing both benefits and risks. For example, the increasing computational costs, the energy and material consumption and the overall environmental impacts of hardware use and disposal. There is a strong need to develop evidence-based approaches to account for carbon in the digital economy; to ensure responsible research and innovation (RRI) is embedded in the development of digital tools and that planetary boundaries are considered in the design and deployment of digital innovations.

By engaging in multidisciplinary research, the project interacts with Computer Science, human-centred design, anthropology, philosophy, ethics, environmental science and policy. We work with various case studies focusing on artificial intelligence, autonomous systems, natural language processing and cloud computing to co-design an evidence-based, sustainable, and responsible innovation framework for digital innovation.

Findings from empirical research show competing narratives regarding digital technologies' benefits and risks to sustainability. For example, positive narratives leverage digital tools' power to better understand, monitor and act on climate change. In contrast, other narratives point to the risks of rising computational costs, supply chain governance challenges, and ethical problems linked to production and consumption.
A peer review publication in the journal *Sustainability* explores these findings from different stakeholders' practices and perspectives. Due to the competing values present within sustainability research and public policy, the demand to increase awareness of how value-laden assumptions affect decision-making is of paramount importance as it pervades all epistemic and expert practices, including in sustainability, Computer Science research and digital innovation policy.

It is essential that narratives are evaluated and balanced by those innovating in the digital sector and capture the complexities of digital technology's potential impacts. It is crucial to help these innovations with evidence on reducing the carbon emissions of the products and systems they create – ensuring the digital sector can contribute positively toward more sustainable futures.

The project is currently developing a wider characterisation of systemic issues at the intersection of digital innovation and policy. For example, in a workshop with stakeholders, we explored the problem of efficiencies in digital products and processes that you would think would help reduce emissions, but which in fact lead to increases in demand, infrastructure and production that counteract and sometimes surpass those efficiency gains, producing rebound effects and other unknown consequences. A research output has been accepted for publication in the journal *Patterns*.

We are conducting expert interviews to explore salient and legitimate ways to build a robust digital sustainability framework by learning from across industry, research, and policy. In addition, a combination of empirical case studies, experimentation with narrative computing, and expert participation will lead to co-design workshops with stakeholders and produce valuable information and digital tools for those working in the innovation of digital technologies to understand, anticipate, plan and act to enable and embed responsible sustainability governance at the heart of the digital economy.

Understanding the Young Digital Mind : a Four-Year Challenge

The digital world can be a challenging space for young people online, with tools and technology intended to support youth having their own unforeseen consequences when used irresponsibly. The Medical Research Council has recently funded a four-year, interdisciplinary and cross-institutional project to investigate how the use of different psychological and technological interventions can support young people's mental wellbeing, titled Adolescence, Mental Health and the Developing Mind. There are a range of themes across this project : from soft robotics, to stimuli investigation for people living with neurodivergences such as ADHD and autism, to exploring the societal and unintended consequences of social media or digital mental health technologies' use online. The Responsible Technology Institute (RTI) within the Department of Computer Science is focussing on the latter theme: to examine the unintended consequences of the use and appropriation of mental health technologies for young people.

Responsible Research and Innovation (RRI) is positioned to consider the intended and unintended consequences of conducting research. It is considered early in the research and development process, often to influence how research is conducted or technology is designed. When adopting RRI approaches, we think about how, for example, the use of technology itself may engender a cycle of negative experiences, even when the app or device is designed to promote better mental resilience. Using the example of social media: if we develop tools for young people that are designed to help them manage the impact of social media, we must question whether the use of such tools might do harm in itself.

Anticipatory governance is therefore often discussed in relation to RRI to anticipate how we can employ oversight to decrease the risk of impact on society in the future. We consider the societal effects of both research within and through technology, for example, online 'trolling', or so-called 'suicide forums'. As research specialists on this topic, we want to ensure that we make the greatest possible impact attacking these issues at their intersection with government policy. Already this has been a priority for those working on this project with some contributions from the project members to round table discussions with the UK Government on mental health policy.

At the core of this project is the ability to develop interdisciplinary research via cross-cutting themes with members of the team covering a variety of different areas, including medicine and education. With such a varied set of academic avenues of investigation, from soft robotics, to interactive cards, assessment tools for clinicians and the digitisation of formerly manual psychological tools eg for Cognitive Behavioural Therapy (CBT), the team are working on bringing them in line with the digital age.

Also at its core is the contribution of RRI to this work existing at a metalevel to the research, with RTI researchers embedding themselves and consulting other researchers on executing their own research protocols. This in itself is a valuable end, as it imbues the project with self-awareness and self-reflection, but also instils the aforementioned anticipatory considerations within everyone's daily work.

The outcomes of this project are open ended and the RTI is looking forward to working with young people, as we develop our own pathways to research on this project. In the coming months, we will explore young peoples' perceptions of mental health technologies and what it means to design new mental health technologies with them, responsibly. We will also be expanding our contributions to the project and young people's mental health more generally by creating exciting opportunities to embed RRI within mental health research.

More information about this project and others is available on the Responsible Technology Institute website: www.rti.ox.ac.uk

Wearable Authentication in Mobile Payments using a Smartwatch

By Doctoral Student Jack Sturgess, Researchers Simon Eberz and Ivo Sluganovic, and Professor Ivan Martinovic

The popularity of cashless payment systems continues to grow. Early payment cards used magnetic strips to facilitate transactions; then NFC (near-field communication) technology unlocked contactless payments. This quickly became entwined with the rise of the smartphone: mobile payment systems (also known as tap-and-pay systems), such as Google Pay, enabled the user to store payment cards digitally in a virtual wallet so that they can make payments over NFC using the NFC module of the smartphone. Some systems, such as WeChat Pay and Yoyo, also explored the visual channel and enabled payments to be made via QR code.

In recent years, smartwatches have evolved rapidly. Modern smartwatches typically have an NFC module and, by sharing the virtual wallet of a paired smartphone, they too can enable the user to make payments over NFC. However, unlike smartphones, they lack fingerprint readers and cameras for facial recognition. In order for the system to remain contactless (ie without requiring the wearer to enter a PIN on a point-of-sale terminal), new means of authentication are needed. This is especially pressing in Europe where, in 2020, the EU overhauled its banking regulations with the Updated Payment Services Directive (PSD2), mandating the use of multi-factor authentication in payment transactions.

For a better user experience, users prefer not to be pestered for the sake of security and favour so-called *implicit* factors for authentication, ie those that require no additional effort from the user because they are inherent in the task being performed, such as keystroke dynamics when typing or gait when walking.

The development of smartwatches has been driven with continuous health monitoring in mind, such as step-counters, sleep-trackers, and heart rate monitors.

Figure 1: Our test terminals.

This means that smartwatches offer and support high resolution, always-on sensing by design using inertial sensors, such as accelerometers and gyroscopes, that measure wrist motion.

In our recent work, entitled WatchAuth, we made use of these sensors to provide an implicit authentication factor when making payments. We showed that the tap gesture, performed as the user taps a smartwatch on a terminal to make a payment, can be used as a biometric to authenticate that user. When the smartwatch initiates an NFC connection with the terminal, we take the last few seconds of inertial sensor data (which is collected continuously) to represent the tap gesture, and we found that this data contains features that are sufficiently unique to each user that it can be used as an authentication factor.

We constructed a user study by placing some common terminals at different (and broadly representative) angles. We had participants wear a smartwatch and make payments on these terminals to collect their inertial sensor data. We then trained random forest classifiers on this data in a leave-one-out manner to create a terminal-agnostic authentication model (which can work on any terminal, regardless of its position).

Payments systems are becoming optimised towards convenience, eg Apple Pay Express Mode allows payments to be made with Apple Pay at busy transport barriers without the user needing to authenticate at all, just by hovering the device over the terminal. In the near future smartphones will be able to accept NFC payments as well as making them, meaning that there will be many new potential terminals on the street. These optimisations cause uncertainty for payment providers as to whether a user intended to make a given payment or if it was the result of an accidental swipe or a skimming attack (where an attacker presses a terminal against a card or device to trigger a payment).

As a solution, we showed that the tap gesture can also be used for intent recognition. That is, we showed that the series of movements required to fulfil a mobile payment, as measured by the inertial sensors on the smartwatch, is sufficiently obscure and deliberate that it can be identified. This means that, whenever an NFC payment is initiated, the smartwatch can check the last few seconds of inertial sensor data to see if a tap gesture was performed and if not, reject the payment.

To prove this, we had our participants wear smartwatches outside the lab to collect a large dataset of activity data. We identified three activities where the user might be the

continued on next page ▶

from previous page ▶

victim of a skimming attack, namely when commuting on a bus or train, when walking along a busy street, or when in a shop. We then trained random forest classifiers to distinguish between the features of this data and those of a tap gesture.

Our intent recognition model can be used together with our authentication model, such that a single tap gesture can simultaneously and implicitly both authenticate the user and recognise intent-to-pay. The authentication model requires a training phase, where it gets to know the user. The intent recognition model, on the other hand, does not require any training, meaning it could even be deployed as a stand alone module to provide a layer of security for an anonymous user.

We found that wrist rotation is a key discriminator between tap gestures and other movements. If a terminal is placed in a comfortable position then it improves results for authentication, but slightly worsens results for intent recognition and vice versa. Awkwardly placed terminals elicit movements that are only associated with tap gestures, helping to differentiate them.

Figure 2: Our smartwatch apps.

More details can be found in our paper, 'WatchAuth: User Authentication and Intent Recognition in Mobile Payments using a Smartwatch' (available on arXiv and presented at EuroS&P'22), and in a forthcoming journal extension.

Predicting SARS-CoV-2 Mutational Landscape

By Hunar Batra, MSc Advanced Computer Science student, under the supervision of Professor Peter Minary, Computational Biology Group.

Protein-protein interactions between the SARS-CoV-2 spike protein, human receptors, and antibodies are important factors in the virus virulence and ability to escape the human immune system. With the proliferation of SARS-CoV-2 pandemic globally since December 2019, numerous variants have been emerging on a regular basis containing distinct transmission, infection rates, fitness levels, risks and impact over evasion of antibody neutralisation.

The ability of RNA-based coronaviruses to mutate and the possibility of emergence of mutations with higher fitness rate, calls for the need to leverage SARS-CoV-2 proteomic data for anticipating viral features and future alterations to considerably improve disease control, prevention and drug development. Early discovery of high-risk mutations is critical towards undertaking data-informed therapeutic design decisions and effective pandemic management.

In a recent research project by the Computational Biology group, the question of deciphering evolutionary mutations in SARS-CoV-2 spike protein was explored with the introduction of a Machine Learning-based-model MuFormer. As evolution amongst protein structures is mostly neutral and the majority of mutations usually occur within protein sequences, MuFormer leverages both the proteins sequential and geometric space to learn from the encoded evolutionary representations to design mutational sequences in an iterative gradient-based fixed backbone design process.

The proposed model consists of an inverted implementation of AlphaFold2[1] as a structure prediction oracle used for inverse folding, injected with frozen sequence embeddings from a pre-trained protein language model as an inductive bias. At each design step, MuFormer maximises the likelihood of amino acids appearing at each position in a sequence, resulting in mutation of amino acids at positions with low sequential or structural likelihood. Without any information about the target sequence, MuFormer exploits the phylogenetic information from the sequential space to mutate the protein sequence at each design iteration to fit the backbone atoms configuration with high confidence.

The generated mutational sequences have been validated with historical SARS-CoV-2 data from GISAID[2], which exemplified the ability of MuFormer to capture phylogenetic and evolutionary properties for generating mutations. The model was able to mutate Alpha variant's sequence into Delta and Omicron variants, showcasing the ability of MuFormer to learn evolutionary landscape, with no additional training. MuFormer outperformed vanilla AlphaFold2 by DeepMind for the in-vitro mutagenesis sequence generation task.

While most of the work until now has focussed on evaluating emerging SARS-CoV-2 variants for their fitness levels, MuFormer marks the first model capable of predicting protein sequence mutations directly as well as flagging generated mutations with high fitness rate. The mutational sequence generation capability of MuFormer highlights the ability of transformer based models to explore the representational language of biology which could assist in controlling spread of diseases by predicting mutations with higher infectivity and fitness in advance.

[1]AlphaFold is an AI system developed by DeepMind that predicts a protein's 3D structure from its amino acid sequence.
[2]The GISAID Initiative promotes sharing of data from all influenza viruses and COVID-19

Signal Injection Attacks against CCD Image Sensors

By Doctoral student Sebastian Köhler, Researcher Richard Baker, and Professor Ivan Martinovic

The importance of cameras for vision-based intelligent systems, such as autonomous vehicles, is undeniable. The captured video frames are often used as part of the decision-making process, making their integrity crucial for the correct behavior of the system. However, due to the nature of analog sensors, it is not easy for a sensor to verify whether the captured information has been manipulated. In other words, a sensor cannot distinguish between a naturally occurring physical signal and an artificially generated one. We discovered that this issue also applies to charge-coupled device (CCD) image sensors, as used in professional and scientific applications, such as ground and space astronomy, microscopy, industrial automation, military surveillance, and defence systems.

In recent years, various attacks against camera-based systems that compromise their integrity have been demonstrated. Since image sensors are optical sensors, the most obvious attack vector is the injection of light. However, injecting light in a controlled way is almost unfeasible and only partially possible for CMOS image sensors that implement an electronic rolling shutter mechanism that reads the captured image information row by row, rather than all at once. CCD image sensors always implement a global shutter inherent to their design. This means fine-grained signal injection attacks using light are not possible. Moreover, a light-based attack requires line of sight between the adversary and the target camera. Finally, attacks that leverage optical emission tend to be suspicious and easily detected by simple mechanisms – for example, if the frame is suddenly over or under-exposed, an alarm is triggered.

In this article, we present an approach that overcomes these limitations and allows the injection of fine-grained perturbations using intentional electromagnetic interference (IEMI).

Signal Injection Attack

We hypothesise that, due to their architecture, CCD image sensors are susceptible to intentional electromagnetic interference, making them vulnerable to post-transducer signal injection attacks. Normally, a sensor should only react to one specific physical stimulus to which it is intended to respond. In the case of an image sensor, the stimulus is light. Incident light causes the generation of electronic charges that can be measured and quantised. Yet, the image sensor itself cannot determine whether the signal charge was generated by the photodiode array during the integration period due to the incident light, or resulted from electromagnetic interference that coupled onto the circuit. A malicious actor could leverage this fact and emit electromagnetic waves at the resonant frequency of the target CCD image sensor to induce a voltage and subsequently alter the captured image information.

Evaluation

We validated our hypothesis on two different CCD cameras, namely a DFM 25G445-ML and a 420TVL CCTV board camera. To prevent interference from other signal sources, we placed the target camera and the antenna in a closed RF shielded box. This also prevented uncontrolled radiation of the attack signal, which could otherwise interfere with legitimate communication channels in the tested frequency spectrum or couple onto other equipment. Nevertheless, to rule out the possibility that the attack signal is induced directly into the cabling and not into the CCD image sensors themselves, we also performed the attack with the cameras switched off. Finally, placing the camera in the EMI shielded box also prevented the generation of legitimate signal charge, making it easier to detect whether the attack was successful or not.

Method

To find the most effective carrier frequency, we captured video frames while running a frequency sweep with a step size of 1 MHz from 50 to 5000 MHz. We then calculated the Structural Similarity Index Measure (SSIM) between the collected frames. More specifically, for every carrier frequency f_c, we collected ten frames, three *legitimate* frames during normal operation, and seven *malicious* frames while emitting a sine wave with frequency f = 1 kHz modulated onto the carrier wave.

continued on next page ▶

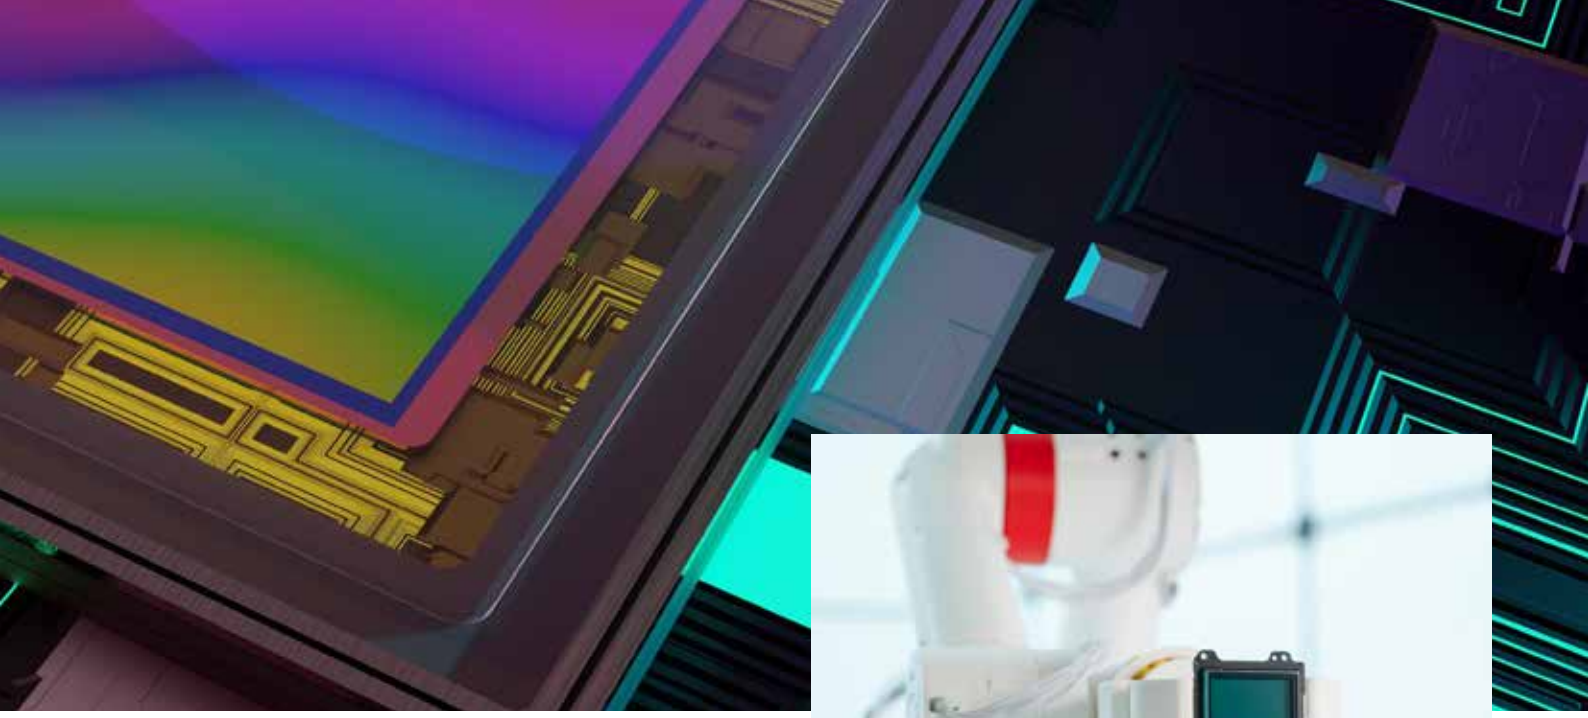

from previous page ▶

For this experiment, the cameras were placed around 3cm away from the transmitting antenna and the output power of the software-defined radio was set to the maximum (20.1 dBm). The experiments were conducted inside a shielded and completely dark box, leading to almost entirely black frames during normal operation. Hence, the SSIM between consecutive legitimate frames was high, meaning they were almost identical. In contrast, for a successful attack, the SSIM between legitimate frames and malicious ones should be as low as possible. The most effective carrier frequency was selected based on the smallest SSIM value. In other words, the frequency that caused the smallest SSIM values induced the most significant perturbations. On the other hand, an ineffective carrier frequency did not induce any signal charge and led to high SSIM values similar to those measured between legitimate frames.

Results

The results of the frequency sweep revealed that the most effective carrier frequency was 190 MHz for the DFM 25G445-ML and 341 MHz for the analog CCD. Interestingly, while the analog CCD camera was only affected at around 341 MHz, the DFM 25G445-ML was vulnerable at various frequencies. The findings indicate that a malicious signal modulated onto a sinusoidal carrier wave at the appropriate frequency is highly likely to couple successfully onto the CCD image sensor. As a result, an adversary would pick 190 MHz or 341 MHz, depending on the target camera.

Use Case: Barcode Scanning

To illustrate an end-to-end attack, we evaluated the attack in the scenario of automated barcode scanning, as used in manufacturing or logistics. We considered an attacker that seeks to remotely disrupt the performance of the barcode scanning, thereby either inhibiting the efficient flow of tracked items or corrupting the inventory management of the facility. As automated CCD barcode scanners often handle hundreds of barcodes per second, even a short attack can quickly impact a large number of items. Scanning a barcode relies on the color contrast between bright and dark bars.

We found that injecting random noise into a CCD image sensor can break this contrast and substantially reduce the reliability of the scanning system. For the lowest selected exposure of 20,000µs and no additional amplification of the signal charge, the captured frames were slightly underexposed. This led even during normal operation to a detection rate of only 50%. As such, it is not surprising that the injected noise reduced the detection rate even further. However, increasing the exposure time and the gain improved the performance under normal operation significantly, leading to a consistent detection rate above 99%. At the same time, the attack effectiveness diminished with increasing exposure time, and contrary to our expectations, for higher gains. This observation can be explained by the increasing contrast between the white background and the black bars of the barcode. Nevertheless, under optimal settings, for instance, for exposure time = 20000µs and a gain of 8.7, the attack caused the detection rate to drop to 1%.

Conclusion

The results of our evaluation show that CCD image sensors are susceptible to intentional electromagnetic interference. We have proven our hypothesis by successfully conducting signal injection attacks against two different CCD image sensors, showing that IEMI can be used to manipulate the captured frames down to the granularity of a single pixel. To demonstrate the impact of such an attack, we evaluated the consequences in the context of automatic barcode scanning, showing that the attack can reduce the detection rate to 1%. Although CCD image sensors are not as widespread nowadays, we argue that signal injection attacks are a real threat to applications relying on input from cameras equipped with CCD image sensors.

AI Onboard of Satellites for Autonomous Detection of Disaster Events

By Doctoral Student
Vit Růžička

Timely disaster detection onboard a satellite is important for the prioritisation of relevant data to downstream, and for the scheduling of following observations. A new AI system called RaVAEn allows for the detection of a variety of disaster events onboard low-powered satellites, allowing an unprecedented amount of autonomy for satellite constellations.

Remote sensing and Earth observation is on the verge of its own unique revolution. Multiple factors come into play here, but the synchronisation of several technical advances make AI-powered intelligent decision-making onboard possible but also needed. For now, most satellites work as cameras that simply observe the Earth and send terabytes of image data daily to the ground – that is, if they are scheduled to do so, if someone has paid for their capture, or if they are one of the few missions that provide free data for the whole community of researchers, such as the Sentinel-2 satellite from the European Space Agency (ESA).

With better sensors and larger quantities of satellites in orbit, the quantity of data will grow uncontrollably. Relying on systems that capture imagery only on-demand would, on the other hand, miss a lot of interesting, unexpected events – such as disasters. Thankfully, this need is matched with new opportunities that come with the capabilities of hardware on even small satellites. We are entering the age where it is possible to run Artificial Intelligence models on these devices, which will allow us to make choices about what to do with the data.

Several experimental satellite missions have been launched to date, serving as demonstrators that onboard AI processing is possible and beneficial. The PhiSat-1 satellite from the ESA is, for example, running - among other applications - a cloud detection model to select and downlink only cloud-free images to the ground. Another example is the WorldFloods system, deployed on D-Orbit satellites, which can detect flooded areas from space.

These proof-of-concept missions demonstrate the usefulness of AI onboard, and hint at exciting future developments in this area.

Our team of researchers participating in the Frontier Development Lab 2021 (Vit Růžička and Daniele De Martini from Oxford University and six researchers from other Universities) – has recently published a paper proposing an AI system called RaVAEn for the task of unsupervised change detection of extreme disaster events, which we demonstrate could run on a small EO satellite with limited processing power.

Simply put, change detection compares a series of images and tries to point to locations within these images where there is a change. With real-world satellite imagery, this change could, for example, be to the colour of a river when it is flooded, in comparison with its usual state. Similarly, a burn area will look different

continued on next page ▶

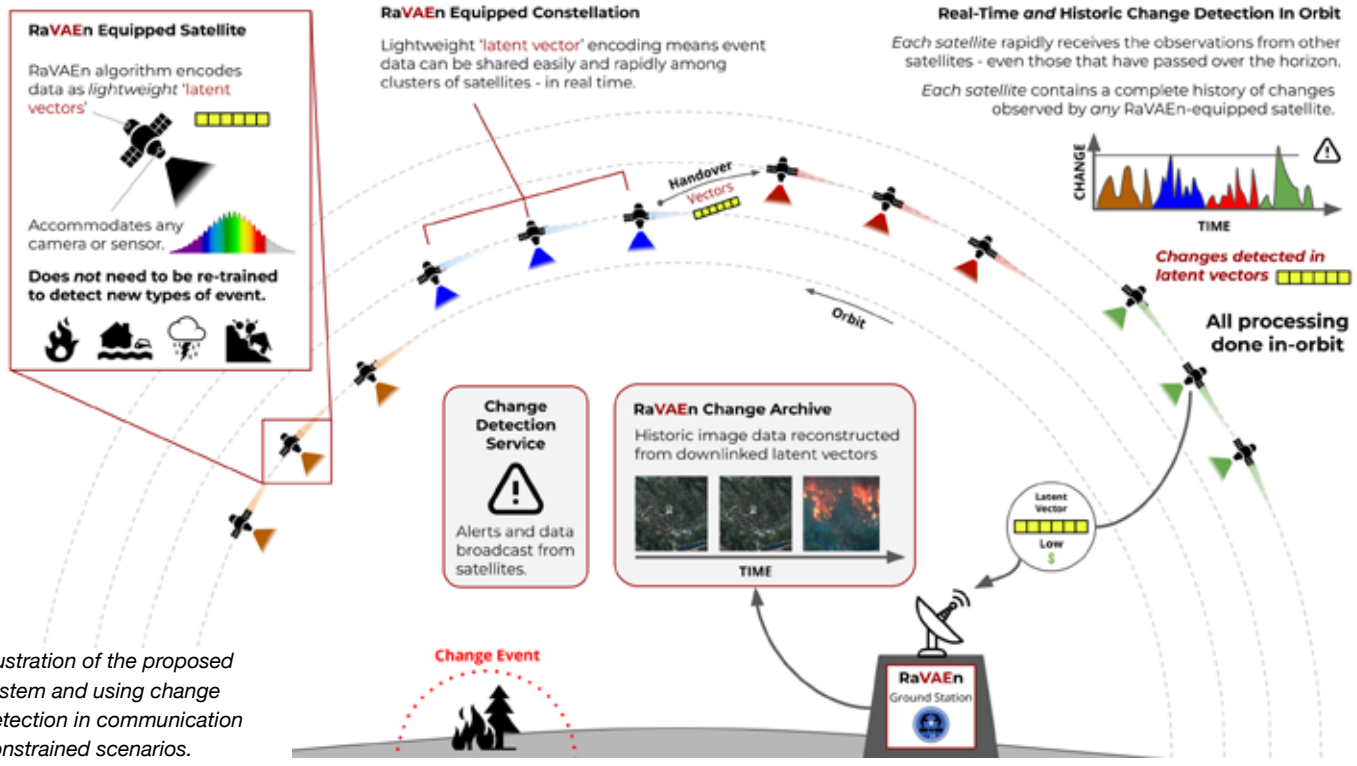

Illustration of the proposed system and using change detection in communication constrained scenarios.

from previous page ▶

when compared to the previous satellite capture of that location before a wildfire. In some cases, it is important to also measure the magnitude of this change, because small changes may be something we want to ignore, such as noise from the sensors or natural variations of the data.

A specific property of our work is leveraging unsupervised learning methods, which in general means that there is no requirement for data annotation. This is useful in cases such as remote sensing, where there is a near never-ending stream of data, which couldn't be feasibly annotated and checked by human experts. Another benefit of this approach is that by using unsupervised models, we remain sensor-agnostic, robust to near-sensor noisy data, and finally, that we can detect any type of event.

More concretely, we use a variational auto-encoder (VAE) model, pre-trained on a representative dataset of satellite images from Sentinel 2, using the L1C level of data (level of processing before removing the effects of atmosphere, also known as the Top-Of-Atmosphere product). Later we evaluate this model on a dataset of four types of disaster events and we show that we can reliably detect any of these as changes despite not training with these disasters in the original dataset.

Unsupervised approaches typically learn features useful for understanding the data, usually on some auxiliary tasks. Our example is using the auto-encoder model, which learns to reduce the dimension of the observed representative data into a bottleneck description, from which it then reconstructs the original data. Without having any labels, the model needs to understand enough about the data that it can reconstruct it, in the meantime finding an efficient representation. In our approach we compare the representations of images in a time series to select areas with a high amount of change.

To evaluate the feasibility of deployment of our model on-board of satellites, we test it on a Xilinx Pynq board, which is a representative piece of hardware that simulates the low compute capabilities of a small CubeSat satellite (these are miniaturised satellites based on standardised size, each unit corresponding spatially to a cube of 10cm). At present, we reach very fast speeds without loss of accuracy (25 km2 area in 2 seconds) when deploying our model on the CPU of this board alone. Excitingly, we could also deploy the model on the FPGA component (which is an integrated circuit capable of reconfiguration based on a program) to achieve even faster processing speeds.

In our work, we are reporting detection capabilities that outperform simple non-machine learning baselines. Our proposed system also allows for ingestion of longer time series of data - for example, we demonstrate better performance when allowing larger memory of images. Namely, comparing three previous passes with the last observation gives significantly better performance than when comparing only with the last remembered pass. Naturally, this would require larger storage for remembering earlier passes – however, using the learned feature encoders, we benefit from having to store only the feature representations occupying 60x less storage than if we used the raw images. This would also be beneficial in scenarios where we need to communicate these representations throughout a larger system of satellites in a constellation. A satellite flying over an area with detected disaster event could provide an update on the evolution of the event, or provide more robust detection in constellations with mixed sensing capabilities with a tip-and-cue regime.

Furthermore, as our approach is sensor agnostic: the proposed method could be developed and deployed for any satellite without any need for costly manual annotation, outside of having a representative sample of images to re-train the auto-encoder models.

Finally, if deployed, we see our system as the first system capable of the detection of general changes, and as a logical next step following the previous experimental missions of PhiSat-1 and WorldFloods, which focused on a concrete single class detection. Timely and accurate detection of disasters could be a reality thanks to the deployment of our model. The next phase of our project is to work towards putting our system onboard a real CubeSat.

With this work, we are also one step closer to the futuristic vision of sending out a probe into deep space which would be capable of learning its own representations from raw observations, meanwhile sending anomalous detections to the home planet while allowing human curation. Cataloguing resources on other planetary bodies, early detection that would allow fast follow-up measurements with other instruments – these are just a few of the interesting future directions that occupy similar research space.

Read more at: https://rdcu.be/cYaTc

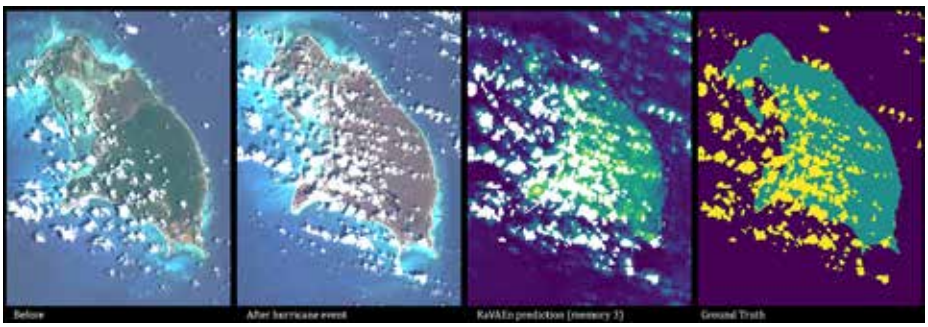

Example of using our method to outline changes after a hurricane event. The green colour in the third image corresponds to areas of large change (loss of vegetation in this case) and is confirmed by the manually annotated ground truth in the fourth image (this one is only used for evaluation). Interestingly, small clouds and cloud shadows are also detected; in a real-world example these can be separated from the detections.

Six ways the Internet of Things poses security threats to journalists

Cyber security researcher Anjuli Shere writes a follow-on from her article in our last issue.

There is plenty of research showing that many journalists have insufficient support, inadequate training and incalculable numbers of adversaries looking to cause digital harm. Most journalist cybersecurity guidance focuses on legacy devices — laptops, tablets and phones. While these threats are by no means over (spyware, for example, is still very much a concern), it is important to acknowledge and address the invasion of newer networked technologies all around us, such as Amazon Alexa devices and smart light bulbs.

In a previous article in *Inspired Research*, I wrote about the multiplying numbers of consumer Internet of Things (IoT) devices in private and public spaces and the threat that they pose to journalists' security. This article further categorises threats to journalists from the IoT, pairing example threat-types in each category with descriptions of potential consequences. The information presented here is based on a forthcoming paper in Springer's Proceedings of the International Conference on Cybersecurity, Situational Awareness and Social Media. Rather than providing an exhaustive or overly-technical list of potential threats, this system represents an initial step toward illustrating new and upcoming threats. It is designed to appeal to a narrative-driven audience, such as the media, to help them navigate the uncertainty that shrouds IoT threats, such as surveillance.

My goal is to give journalists ways to understand these threats, to easily communicate them to their sources and audiences, and to incorporate the IoT into regular risk assessments. My system includes six categories, comprising 19 IoT-specific threat types that are relevant to members of the media. These categories are:

1. Regulatory gaps
2. Legal threats
3. Profiling threats
4. Tracking threats
5. Data and device modification threats
6. Networked devices threats

One key theme across all six categories is that the commodification of data by the technology industry colors the design of IoT devices such that information leakage is often an intentional feature, rather than a bug. For example, the smartwatch app Strava is intended to facilitate profile sharing and tracking of exercise, such that users can share how much exercise they're doing and where they are when they do it. But this functionality can have unintended consequences; it has also enabled the mapping of secret military bases. Some apps, hosted on IoT devices, only allow users to maximise their functionality if users agree to long and murky terms and conditions — and many devices require constant connectivity to work. The IoT is particularly menacing because even if you opt out of interacting with one device, you can't necessarily escape its friends; it is called an Internet of Things precisely because devices form whole ecosystems. This means that threats can overlap deliberately (with attackers deliberately employing multiple threat categories), or inadvertently (because journalists may be reluctant to report IoT issues due to hostility from law enforcement).

Although these threats can coincide with and compound one another, it is necessary to pull them apart and examine them separately. This can help prevent journalists from becoming overwhelmed and experiencing decision paralysis due to the sheer scale and severity of cyber-threats to their work and wellbeing. This article will therefore detail one threat type from each of the aforementioned categories, to explore the impacts and implications of IoT risk for the media.

Many threats from the IoT stem from inadequate government regulation. While journalists may seek to avoid the devices entirely, they are increasingly pervasive.

One example: There is no legal requirement for IoT designers and manufacturers to secure their technologies, so each small, low-powered device can easily be infected with malicious software (malware) that can be used for various illegal purposes. For example, many poorly secured IoT devices can together be formed into a botnet, which is a network of corrupted IoT devices that can be used to power big, targeted attacks. These attacks can involve attempts to access more secure information while hiding the perpetrator's identity, as well as further malware delivery that can drastically affect services, including publishing news stories.

So what? Botnets could also be used to launch large-scale campaigns to intimidate journalists and amplify disinformation, as reported by Brian Krebs in 2017 when Twitter profiles belonging to him and the investigative journalism outlet ProPublica were suddenly harassed online by thousands of similar accounts.

Legal threats, which are well-documented as the basis of many journalists' fears, refer to ways in which IoT data or actions might be used either in law enforcement investigations or to embroil journalists in lawsuits.
One example: Criminal attacks on news organisations

continued on next page ▶

from previous page ▶

are a one-two punch: first, the media must mitigate the damage of the breach itself and then deal with repercussions that add more legal and financial pressures. The first impact may include being hit by ransomware (malware that holds a news organisation's systems ransom by encrypting their files or locking a news site), as in the case of Portuguese Media Giant Impresa. For the second part, the cost of a cyberattack can be devastating, and, for companies in the United Kingdom, has doubled since last year due to post-incident fines.

So what? These secondary consequences may significantly hinder a news organisation's ability to focus on and sufficiently fund news reporting. Even if they survive the cyberattack, organisations could be hamstrung by regulatory fines, investigation costs and compensation payments, causing staff to be laid off and news stories to be derailed.

One of the creepiest risks associated with the IoT is that of profiling, which means creating a comprehensive outline of a journalist's life and character. This can include recording behaviors, associates (friends, family and colleagues), frequent locations, habits and even health information.

One example: Even when devices seem like they are not listening, they might be eavesdropping. Logs from smart home devices can allow residents' routines to be recognised from their hours of use. Or there are multiple ways in which common IoT devices can be used to recognise typing patterns and reconstruct what was written, even when it was sent through secure channels. So what? A malicious actor could eavesdrop to learn about an investigative story a journalist is reporting — and could threaten to use personal information to blackmail the journalist into not publishing on that topic.

Much of the data generated and interpreted by the IoT enables devices (and their users) to be located by anyone who has access to the location data illegally or legally, which can include foreign, state-sponsored hackers or domestic government actors. Technologically-tracked movement patterns are frequently sold to or shared with third parties, including private companies.
One example: Law enforcement's use of social media analysis may be cross-referenced with data from IoT devices. For example, footage from many camera-equipped doorbells is automatically accessible by local police forces, effectively making the privately-owned doorbells an extension of state surveillance networks.

So what? A fear of identification and reprisals could lead to self censorship, both among journalists gathering information on the ground and their sources. This is especially likely if the devices can provide evidence identifying individuals who are suspected of attending protests at risk of government violence, such as at Black Lives Matter protests.

Journalists rely on their reputation as a source of accurate information; any single discredited individual can have ramifications for public trust. Alteration of any data, from published information to account details, via IoT devices can not only undermine the credibility of a journalist but also potentially endanger them and their sources.

One example: Certain IoT devices — including some smart fridges — can access users' social media and email accounts.

So what? These poorly secured devices can be hacked to plant stories and communications that are falsely attributed to journalists. This can seriously undermine a journalist's credibility and job security.

As intelligence agencies have publicly noted, the interconnectedness of the IoT means that hacking into one device can allow an adversary to compromise an entire network or take down a website. In fact, this happened to *The Guardian*'s service provider via a Distributed Denial of Service attack in 2016.

One example: A journalist may depend on a given device (or network that includes an insecure IoT device) for their work, such as a camera-equipped drone for a photo- or video journalist.

So what? If a bad actor deliberately makes the device inaccessible to the intended user, a journalist could be left vulnerable to extortion — such as demanding a ransom in exchange for access to the device. Internet-access denial tactics are also used to reinforce aggressive treatment of journalists. Attacks unexpectedly limiting device functionalities can trigger detrimental physical, psychological and financial effects — for example, if executed against someone's automobile.

I am currently designing a multi-piece toolkit, which will help members of the press determine the specific IoT threats that are most relevant to them. The toolkit also will help journalists determine the best countermeasures for their circumstances so they can continue their work safely. The toolkit also highlights protections and mitigations at the organisational and industry levels, which I argue are highly necessary. The threats detailed in this and my previous article compound each other to magnify personal and professional consequences for targets. Media organisations should learn about IoT threats and work together to combat them, as no individual can remove these threats alone. Legacy institutions, civil society groups, freelancers — every media stakeholder must make it a priority to share information on these threats and incorporate the IoT into security routines and risk assessments.

I'm keen to connect with others who are interested in emerging technological and legal risks to journalist safety. You can contact me at anjuli.shere@new.ox.ac.uk or via Twitter at @AnjuliRKShere.

Saving Lives after Myocardial Infarction: Automated 3D Analysis of the Heart

By Jorge Corral Acero and Professor Alfonso Bueno-Orovio

Myocardial Infarction, commonly known as 'heart attack', occurs due to an inadequate blood supply through the coronary artery, causing myocardial cell death and, consequently, damage to the heart muscle. This is the deadliest cardiovascular disease accounting for a staggering total of 15% of mortality rates worldwide. Luckily, thanks to medical advances, nowadays most people survive an infarction (infarct). However, these survivors become likely to suffer from other cardiac complications as result of the injury caused by the infarct. Identifying which patients are at risk is crucial for treatment and the subsequent 1-year survival likelihood, but remains challenging. Knowing and predicting these risks is the gap that we have addressed with this research.

The infarct causes an injury that impairs the ability of the heart to contract. How the heart changes as a result of this injury has proven crucial for recovery, but it is still unknown how to use this information in full for diagnosis. The two most used criteria to predict risk of death after an infarct are ejection fraction, or how much blood can leave the heart, and the heart volume in contraction, known as end-systolic volume. In clinical practice, both are manually measured from 2D magnetic resonance images of the heart. This is time consuming for doctors (who could otherwise be using their skills in patient care), prone to errors and inconsistencies in the three-dimensional shape of the heart, and limited to global information – ie global volumes.

In collaboration with King's College London and the cardiologist team from University Medicine Göttingen, we have developed a new method to automate the analysis and to reconstruct the actual 3D shape and contraction of the heart, overcoming the aforementioned limitations. This was attained by combining artificial intelligence with mathematical models of the shape of the heart. The accurate 3D reconstruction does not only allow us to obtain the conventional metrics to predict risk in a precise and automated manner, but opens up a 3D exploration of the shape of the heart and its pumping function. The underlying hypothesis is that, while there are many ways our heart can contract to push the same amount of blood and reach the same volume in contraction, there may be a specific contraction pattern that is especially bad for your future health after an infarct. 3D analysis enables us to capture local details.

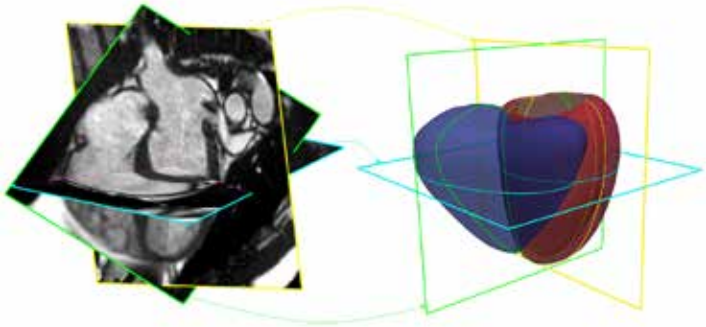

From 2D Magnetic Resonance standard scans to accurate reconstruction of the 3D shape and contraction of the heart.

Indeed, thanks to the proposed approach for first time to the best of our knowledge, we have been able to identify specific 3D patterns that are linked to the risks after the infarct, such as lateral- and basal-specific contraction impairments on the left ventricle of the heart. This allows us to propose new metrics, based on these 3D patterns, that are more informative than the conventional ones and that improve the prediction of risk after infarction by a large margin.

What does this mean in clinical practice? The patients will know two days after the infarct, and thanks to a conventional MRI scan and our advanced analysis tool, the risks of suffering a major adverse cardiovascular event. As such, the patients will be able to take the right medical therapy and lifestyle choices to address that risk. This means that the 10 million-plus patients that suffer from a heart attack each year worldwide could be better identified and provided with the best treatment. This, in turn, prevents secondary outcomes such as reinfarction or death, while avoiding unnecessary procedures like invasive surgery.

In essence, our work improves risk management after infarction, saves time for clinicians that can be reinvested in patient care, and sheds light on how our heart responds after an infarct episode, advancing our understanding of heart function to enable a move towards more preventative medicine. This could potentially alleviate pressure on health care systems – important, given the socioeconomic burden of this disease and the current socio-political framework with tightening financial constraints and aging population. The study ultimately demonstrates the potential of 3D computational models to improve heart disease management and to contribute to personalised and preventive cardiology.

The methods and results of this study were demonstrated in a large clinical cohort of more than 1,000 subjects and have been made publicly available in our *JACC Cardiovascular Imaging* article. This work has been awarded the engineering bronze medal at the 2022 STEM for Britain awards and 1st prize in the 2022 Conference on Clinical Translation of Medical Image Computing & Computer Assisted Intervention.
Read more: bit.ly/3FXNrAo