


THE KNOB ATTACK ON BLUETOOTH

Read more – p24



ONLINE RADICALISATION:

Detecting extremist content on Twitter – p17



BAYESIAN DEEP LEARNING:

The work of the Frontier Development Lab – p26



DATA SECURITY:

Educating users with smartwatch games – p28



DEPARTMENT OF
**COMPUTER
SCIENCE**

Inspired Research

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

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

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

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



Editorial board

Kiri Walden (Editor)

Suvarna Designs (Designer)

Sarah Baldwin

Leanne Carveth

Rosanna Cretney

Emma Dunlop

David Hobbs

Stefano Gogioso

Elisa Passini

Helena Webb

Standa Živný

Contributors

Sara Dutta, Paula Fiddi, Yarin Gal, Stuart Gillespie, Václav Janeček, Klaudia Krawiecka, Nic Lane, Peter Marinov, Suzanna Marsh, Hector Martinez-Navarro, Mariam Nouh, Alina Petrova, Elisa Passini, Kasper Rasmussen, Blanca Rodriguez, Oliver Samson, Carolyn Ten Holter, James Parr, Helena Webb, Shimon Whiteson, Meredydd Williams, Michael Wooldridge, Jun Zhao, Maaïke Zwart.

CONTENTS

P4-12 News

P13 Farewell to...

P14-16 Student focus

Feature articles

P17 Detecting extremist content online

P18 New standards to protect children online

P19 Engaging with your children: recipes for bringing up the 'AI natives'

P20-21 New research creates computer simulations of heart attacks

P22 Virtual Assay and experience with industry

P23 Training autonomous vehicles using real-life human behaviour

P24-25 The knob attack on Bluetooth

P26-27 Bayesian Deep Learning for all humankind

P28 About time: changing security behaviour through smartwatch games



Letter from the Head of Department

This is my first 'letter' since my return from sabbatical, and it's clear that 2019 has been an exciting year for the department.

We continue to garner awards and recognition across different areas of Computer Science research. Professor Marta Kwiatkowska, Professor Yarin Gal, and alumnus Daniel Nichol have all won major individual awards, while various teams have made an impact at international conferences and in publications of note. It is no fluke that for the second year running the University of Oxford has been ranked first in the world for Computer Science in The *Times Higher Education* 2020 World University Rankings.

2019 has been a period of change for our community, with many long-standing members of the department leaving us. But it has also been a time of positive development. We now have an interdisciplinary Law and Computer Science course which is run jointly with the Law Faculty. We are also happy to welcome Oliver Sampson in his role as Industry Liaison Administrator.

We continue in our endeavour to widen access to our courses. DeepMind have generously funded four more graduate scholarships, for students identifying as female, BME or from households with traditionally low progression to higher education, who want to pursue a full-time master's degree in the Department of Computer Science. Professor David Gavaghan and his team ran a Uniq+ a summer school designed to give participants the chance to experience what postgraduate study at Oxford is like. The summer school is for people who might find postgraduate study difficult due to socioeconomic or financial circumstances.



We are very lucky to attract many talented applicants each year and our department is so popular that during the Oxford University undergraduate open day, for the first time we had to use a marquee in the University Parks to accommodate all our visitors.

Once again this newsletter is full of interesting items about our department and the work we do. I am looking forward to 2020, which promises to be another busy year for our ever-growing community.

Professor Michael Wooldridge

December 2019

DPhil student, Adrian 'AJ' Duncan, sadly passed away in October. Adrian worked in the cyber security theme, and was a student at Worcester College. He had been unwell for some time. He leaves a wife and two daughters and our thoughts are with them, and his many friends and colleagues here at Oxford, who will miss him greatly.



New Starters

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

Tracy Marshall,
Nikki Theofanopoulou,
Mital Kinderkhedra,
Yongchao Huang,
Abhishek Dasgupta,
Hugo Paquet,
Alexandra Pendleton,
Jennifer Lockie,
Gabrielle Alexis,
Sahar Vahdati,
Michaela Walsh,
Stefano Germano,
Qiyi Tang,
David Hobbs,
Patrick Hohenecker,
Robert Newnham,
Bo Li,
Georgios Birmpas,
Sarah Retz-Jones,
Josephine Francis,
Hadrien Pouget,
Lei Wang,
Ahmad Atamli,
Jenny Wang,
Stefano Gogioso,
Anita Hancox,
Alina Petrova,
David Tena Cucala,
Luca Laurenti,
Mirco Giacobbe,
Emanuele La Malfa,
Leema Chapman,
Pan Hu,
Eike Neumann.

News in brief

Professor Michael Wooldridge was part of a prestigious line up of guests for the 'Tech Tonic Stage' at the FT Weekend festival. He was a member of a panel which discussed the question, 'Is Our Future Going To Be One Long Gap Year?' The panel also featured Calum Chace of the Economic Singularity Club and Kathryn Parsons, founder of Decoded, with Madhumita Murgia the FT's European technology correspondent as moderator.

Sir Nigel Shadbolt was named one of *Business Insider's* UK Tech 100: The 100 most influential people shaping British technology in 2019.

Read more: bit.ly/2PkKaRQ

We were delighted to host Dame Wendy Hall (University of Southampton) whom as part of the Department of Computer Science Strachey Lecture series gave a talk this term entitled 'The Windmill of Your Mind: Reflections on career in Computer Science research' in the Mathematical Institute, University of Oxford.

Members of the department contributed to the IF-Oxford Science and Ideas Festival, presenting two workshops. 'Managing your Digital Self' was a visual exploration of the personal data routinely generated in daily life, through smart devices at home, digital mobile devices used by millions of families, as well as various portable devices which surround us in everyday life. In the 'Making Faces' workshop Niki Trigoni talked to members of the public about the latest in 'multimodal recognition', which uses facial recognition technology in combination with voice and walking gait analysis to help identify individuals.

Sir Nigel Shadbolt judge of Royal Society book prize

Professor of Computer Science Sir Nigel Shadbolt chaired a panel of five judges to choose the 2019 Royal Society Insight Investment Science Book Prize. The prize celebrates the very best in popular science writing from around the world.

Sir Nigel said, 'It was a privilege to be chairing the judging panel for the Royal Society Science Book Prize 2019. The prize celebrates great writing about science and the scientists who make it happen. There has never been a more exciting time for popular and accessible science writing. There is a huge appetite for books across all topics of science and writers who can

explain and inspire, entertain and inform. Despite the distemper of the times, and claims we live in a post-truth age, there is a real hunger for books that explain our scientific understanding of world around us. There is huge interest in the triumphs and tribulations of science and the human stories behind the scientific endeavour.'

The winner of the 2019 Prize was announced at an awards ceremony at The Royal Society on Monday 23rd September, hosted by Professor Brian Cox. Caroline Criado Perez was announced as the winner, for her book *Invisible Women: Exposing Data Bias in a World Designed for Men*. Caroline, a writer, broadcaster and feminist campaigner received a cheque for £25,000, (and £2,500 was awarded to each of the other five shortlisted authors).



Mixing music with maths

The Barbican launched 'Life Rewired', an arts and learning season, which ran throughout 2019 exploring what it means to be human in the face of technological and scientific forces that are dizzying in their speed, scale and complexity. As part of this season Sir Simon Rattle opened the London Symphony Orchestra's 2019/20 season on 14 Sep 2019 with the world premiere of *Antisphere*, a new work by composer Emily Howard.

Emily is an Oxford graduate in Mathematics and Computer Science

who went on to take a PhD in Composition from the University of Manchester. The Barbican commission completes a triptych of orchestral works, following *Torus (Concerto for Orchestra)*, which won the orchestral category of the 2017 British Composer Awards, and *sphere*, which was premiered by the Bamberg Symphony in the same year. Emily is the Director of PRISM, the Centre for Practice & Research in Science & Music at the Royal Northern College of Music in Manchester, which brings together creative collaborations between the sciences and music including cutting-edge research involving AI-assisted composition.

European Research Council Starting Grant awarded to Nic Lane

Nic Lane has been awarded a prestigious ERC Starting Grant for his project 'REDIAL: Re-thinking Efficiency in Deep Learning under Accelerators and commodity and processors'.

Over a short period of time, breakthroughs from the field of deep learning have transformed how computers perform a wide-variety of tasks such as recognising a face, driving a car or translating between languages.

Not only has deep learning become an everyday tool, it is also the most promising direction for tackling a number of still open problems in machine learning and artificial intelligence. However, routine deep learning activities (such as training a model) exert severe resource demands (eg memory, compute,

energy) that are currently slowing the advancement of the field, and limiting full global participation in this research to only the largest of companies.

The goal of REDIAL is to solve core technical challenges that span the areas of machine learning and system research which collectively can enable a radical jump in the efficiency of deep learning. It aims to address both the challenge of high training costs and time, as well as the barrier to deploying models on constrained devices (like wearables, sensors) that currently require new efficiency techniques to be invented each time a deep learning innovation occurs.

To accomplish this REDIAL takes two complementary approaches. First, it seeks to build a theoretical understanding of current approaches

to deep learning efficiency, a desperately needed step given current over reliance on empirical observations.

Second, it aims to develop new architectures and methods for training and inference that tackle core efficiency bottlenecks, such as: dependencies preventing parallelisation and excessive on-chip data movement; while also opening new opportunities including the greater adoption of analog processing within accelerators. REDIAL aims to change the way the world trains its models, and deploys them to constrained devices, by producing a series of new deep architectures and algorithms with properties that promote high efficiency that can serve as a foundation for new machine learning innovation.

Daniel Nichol wins Reinhart Heinrich Award

Recent Computer Science graduate Daniel Nichol has won the 2019 Reinhart Heinrich Award for his doctoral thesis '*Understanding Drug Resistance through Computational Models of the Genotype-Phenotype Mapping*'. The award is made by the European Society for Mathematical and Theoretical Biology (ESMTB) and will be formally given in a ceremony in Heidelberg in 2020.

The committee who chose to make this award commented, 'Daniel Nichol's thesis is likely to become a landmark in an issue of critical significance. It studies and proposes strategies to deal with drug resistance in the treatment of cancers and infectious diseases taking into account the genotype-phenotype mapping. New mathematical models are used to better understand the underlying mapping mechanisms, integrating genetics, environmental signals and stochasticity. The thesis gave rise to three papers already published in

highly prestigious journals and to a few more papers under publication.'

Dan's supervisors were Professor Peter Jeavons (Oxford, Computer Science) and Professor Alexander R. A. Anderson (Moffitt, Integrated Mathematical Oncology).



Samsung invests in spinout company Oxford Semantic Technologies

Samsung Ventures has led an investment round that raised over £3 million for Oxford Semantic Technologies Limited, an Oxford University spinout company founded by pioneers in the

fast-growing field of semantic technology. The investment round was also supported by the University of Oxford and Oxford Sciences Innovation PLC.

Spun out of the Department of Computer Science in 2017, Oxford Semantics is developing and commercialising RDFox, the high-performance knowledge graph and semantic reasoning engine developed by Professors Boris

Motik, Ian Horrocks and Bernardo Cuenca Grau. The company's technology turns an organisation's data into machine-readable knowledge, combining diverse data sources and then enabling users to query them with powerful reasoning.

Samsung's investment will enable Oxford Semantics to bring its powerful data integration and reasoning engine to more clients.

News in brief

David Rogers is an external lecturer who has been awarded MBE for services to cyber security. David taught a module on Mobile Systems Security for our part-time professional MSc in Software & Systems Security and is himself a graduate of our MSc programme.

Andrew Markham has featured as an Oxford Sparks scientist in the spotlight. More details on how Andrew became a Computer Scientist and Engineer here: oxfordsparks.ox.ac.uk/content/andrew

In July St Edmund Hall hosted a solo exhibition entitled Oxford: An AI-art View' showcasing artificial-intelligence created digital paintings and video art pieces. Developed by Mayur Saxena (2018, MSc Computer Science), the AI-agent algorithmically captures the essence of iconic Oxford landmarks from images on social media then reproduces these as Impressionistic digital artworks.

Katie Dicks has been awarded a certificate for her participation on the OBEA (Oxfordshire Business Education Alliance). The group work with students from four local state comprehensive schools on a project. The project has involved with working with members of the department of Computer Science to work together on solving a problem.

Professor David Gavaghan and project team ran UNIQ + a summer school designed to give the chance to experience what postgraduate study is like at the University of Oxford, for talented individuals who would find continuing into postgraduate study a challenge because of socioeconomic or financial circumstances. Read more: ox.ac.uk/uniqplus

Oxford ranked first in world for Computer Science

For the second year running the University of Oxford has been ranked first in the world for Computer Science in The Times Higher Education 2020 World University Rankings. This is ahead of California Institute of Technology,

University of Cambridge and Stanford University.

The ranking is based on criteria measuring teaching, research, industry income, international outlook and citations, which are combined to provide a comparison of universities worldwide. Since 2017 Oxford has also annually been ranked the top university in the world according to the same publication.

Oxford brings together experts and students from Computer Science and Law

adapted from an article by Václav Janeček

On Tuesday 15th October, the University of Oxford launched its new interdisciplinary course, Law and Computer Science.

This course, jointly offered by the Department of Computer Science (co-convenor Professor Tom Melham) and the Law Faculty (co-convenor Professor Rebecca Williams), will introduce students from both backgrounds to the terrain at the boundaries of their two disciplines. The overarching theme is understanding Law and Computer Science at their interface.



Such interdisciplinary understanding requires both lawyers and computer scientists to develop an appreciation of the way in which they typically approach problems, with very different analytic tools. A key pedagogical strategy for the course is to teach Law and Computer Science students together, and in particular for them to collaborate on a group-work practical exercise. This will accelerate both groups' acculturation to each other's analytic perspectives through learning from each other as well as from faculty.

The opening lecture as well as the lively follow-on discussion was led by Professor Richard Susskind, a OBE former Oxford student, who himself has combined the two disciplines since late 1980s and is now considered one the leading figures in the field.

The research behind this course is part of an ongoing Oxford project titled 'Unlocking the Potential of Artificial Intelligence for English Law' which is run by researchers in the Oxford departments and faculties of Law, Economics, Computer Science, Education and the Saïd Business School, and led by Professor John Armour of the Faculty of Law.

As a global leader in education, Oxford University is committed to sharing knowledge and insights related to the course and will release multiple open access materials as the course progresses.

DeepMind renews commitment to under-represented students at Oxford

DeepMind, the leading British artificial intelligence company, has renewed and extended its commitment to supporting students at the University of Oxford following the success of its inaugural scholarships for under-represented groups, announced in 2018.

Four more DeepMind graduate scholarships for students wishing to pursue a master's degree in the Department of Computer Science will be made available for students commencing courses in 2020–21. These are open to individuals ordinarily resident in the UK who belong to one or more of the following groups: identifying as female, BME or from households with traditionally low progression to higher education.

Three new DPhil scholarships, with a preference for students of machine learning, in the Department of Engineering Science will also be launched thanks to a further gift from DeepMind. These scholarships are

open to applicants identifying as female who are ordinarily resident in the UK.

Vice-Chancellor, Professor Louise Richardson, says: 'We are very grateful for DeepMind's continuing support of underrepresented groups at Oxford. This supports our strategy to increase the diversity of students and the diversity of perspectives in the University.'

Dr Demis Hassabis FRS, co-founder and CEO of DeepMind, says: 'AI has incredible transformative potential, but if we're going to develop a technology that benefits everyone, it is critical that the people working on it are representative of the wider world. That's why the next generation of researchers, from all backgrounds, should feel able to access the best educational opportunities. We're proud to be working with leading universities, such as the University of Oxford, to offer talented students from under-represented backgrounds a chance to reach their full potential through scholarships and other programmes, collectively broadening access to science and research.'

Professor Michael Wooldridge, Head of the Department of Computer

Science, says: 'We are delighted that DeepMind has renewed their scholarship programme for under-represented groups. Our goal is to deliver the best possible education to the best students, irrespective of their background. This scheme is a huge boost for us in that mission.'

Professor Ronald Roy, Head of the Department of Engineering Science, adds: 'The Department is very much appreciative of DeepMind's commitment to supporting postgraduate education in information engineering. Our goal to engage more women in this rapidly emerging field will be significantly advanced by the company's generosity.'

DeepMind was established in London in 2010 and is now a world leader in artificial intelligence research and its application for positive impact. The DeepMind Scholarships are part of a wider initiative by the company to broaden participation in science and support wider research in the UK.

More details of the DeepMind Scholarships will be available soon on <https://deepmind.com/blog>.

To read more please visit: <https://bit.ly/33aFGAY>



Yarin Gal announced as Turing AI Fellow

Yarin Gal is one of five new Turing AI Fellows announced by The Alan Turing Institute. The Turing AI Fellows, who have been appointed for five years, are drawn from a wide variety of disciplines and backgrounds and will be tackling research challenges ranging from sustainable aviation to AI for discovery in data intensive astrophysics.

Yarin will work on democratising safe and robust AI. While already in use in industry and academia, major obstacles still stand in the way of deploying deep learning AI safely and responsibly. Yarin proposes to tackle these problems by building community challenges derived from real-world applications of AI in industry. With the community competing on these public

challenges, new safe and robust AI tools will be developed for responsible use in industry.

Adrian Smith, Institute Director and Chief Executive of The Alan Turing Institute, said: 'We look forward to what this talented group of Turing AI Fellows will bring to our vibrant research community and we welcome their contributions to our growing Institute. There is vast potential for their diverse work to be transformative in both the foundations and applications of AI and I am confident they will push the boundaries of what these new technologies can do for the good of society.'

To read more: <https://bit.ly/34Sxwyw>

Marta Kwiatkowska named as BCS Lovelace Medal Winner 2019

Professor Marta Kwiatkowska has been named as the recipient of the 2019 BCS Lovelace Medal, the top award in computing in the UK, awarded by BCS, The Chartered Institute for IT.

The award is presented annually to individuals who, in the opinion of the BCS Academy Awards Committee, have made a significant contribution to the advancement of Information Systems.

Marta has been recognised for her major contributions to probabilistic and quantitative verification. She has made significant contributions across the breadth of theory, applications and software tools. Her research and her software system PRISM have made a huge impact on Computer Science in the UK and worldwide. Since 2001 she has led the development of the highly influential probabilistic model checker PRISM.

Marta explains: 'Computing infrastructure has become indispensable in our society, with examples ranging from online banking, to intelligent vehicles and electronic medical equipment. Software faults in such systems can have disastrous consequences. My research is concerned with

developing modelling and automated verification techniques that can guarantee the stable, safe, secure, timely, reliable and resource-efficient operation of computing systems.'

PRISM is widely used for research and teaching across the globe; it is highly cited, and has been downloaded over 70,000 times and used in over 500 publications to date. It has been shown to be effective for modelling a wide range of applications including communications protocols, molecular and cell biology, and more recently, adaptive and Artificial Intelligence systems. It has proved its usefulness in many distinct fields, including distributed and cloud computing, wireless networks, security, robotics, quantum computing, game theory, biology and nanotechnology.

The award will be presented to Marta at a ceremony at the Royal Society in London next spring.

Steve Furber, Chair of BCS Academy Awards Committee says: 'Professor Kwiatkowska has made an outstanding contribution to the understanding and advancement of computing and I am delighted that we are awarding her the Lovelace Medal in recognition of her wide range of



achievements. The impact she has made to Computer Science, and her contribution to academic research, makes her a richly deserving recipient of this prestigious award.'

Marta says of being presented with the award: 'I'm incredibly honoured to be receiving the 2019 Lovelace Medal. It's really exciting to have my work recognised via the award and emphasises the importance and potential of research in probabilistic and quantitative verification. Over the years, I have worked in conjunction with many collaborators and am grateful for their contributions, none more so than Dr Gethin Norman and Professor David Parker, who worked on the first implementation of PRISM and have provided great support ever since.'

Nominations for the 2020 award open on Monday 21 October 2019 and close on Friday 24 January 2020.

More information about the award can be found at www.bcs.org.uk



Photo above: (left to right) Deniz De Barros, Edd Salkield, Alexandra Manciu, Șerban Șlincu, Josh Smailes, Sam Boyer.

Oxford Hack 2019

Oxford Hack 2019 ran again this year at the University of Oxford on 16 – 17 November 2019. It is an event designed to encourage student teams to build something brilliant in 24 hours. Among the group were a number of Department of Computer Science teams including some building: a 'distributed multi-agent reinforcement learning' system for web browsers (<https://bit.ly/33bWEOR>); a machine-learning android app that lets you have conversations with yourself (<https://bit.ly/2pFw6YT>); and a satirical General Election twitter bot using Tensorflow (<https://bit.ly/2O8yGjC>). The hackathon had 19 sponsored awards such 'most entertaining hack' and 'most stylish hack' which were given to various teams and their projects.

Further information on the awards and winning hacks available here: <https://bit.ly/2rdeT9D>

100+ Brilliant Women in AI & Ethics



The ORBIT project has created Oxford's first carbon-neutral conference. The sold-out 100+ Brilliant Women in AI & Ethics took place in September and focused on three key themes - the climate crisis threatening animal and human life; the pushback against untrammelled AI; and the need for diversity in STEM disciplines.

The conference included speeches from Baroness Beeban Kidron; Angela McKane; Jacquelyn Kronos; Beena Ammanath and Dame Wendy Hall. A prevalent theme was the competitive advantage to be obtained in AI through greater diversity, and the need for ethical considerations to ensure that the impact of tech on society is more positive. Offering a robust perspective was author Jeanette Winterson, who has spoken out before on her belief that AI will have a disproportionate negative effect on women. The day was enriched by a panel of A-level students who had entered a competition to attend - they shared their thoughts on artificial intelligence as well as receiving their prizes.

The conference was always intended to be carbon-neutral, and every decision was taken with a view to minimising the carbon footprint. The carbon that was created is being offset with tree-restoration projects in the northern hemisphere and rainforest-protection in the south, ensuring that the event was accredited as carbon-neutral by an assessment company.

The breakout workshops, each chaired by an expert in the field, were the highlight of the day. Conference delegates brought their insight to bear on four key topics, ensuring that all those who attended the conference contributed to its outputs.

The incredibly positive response to the event demonstrated the importance and international impact of the work undertaken by ORBIT and the relevance of responsible research and innovation to cutting-edge technologies. ORBIT is creating a report from the data produced on the day that will be released later this year.

Read more: www.orbit-rri.org/

Oxford Computer Science Conference award winners

The 2019 Oxford Computer Science Conference took place on 6 June and was a great success. The conference featured talks and poster sessions, and concluded with an award ceremony and conference dinner.

The winners of the Oxford Computer Science Conference 2019 Awards are:

1st presentation award

Ulrik Lyngs

2nd presentation award

Sean Sirur

1st poster award

Lingyi Yang

2nd poster award

Isaac Dunn

1st abstract award

Ben Fernando

2nd abstract award

Emma Bluemke

AI@Oxford conference

The AI@Oxford conference took place in September and many academics from our department were involved. On the first day of a packed two-day schedule Mike Wooldridge gave the opening talk 'The World of AI', Daniel Kroening spoke about 'The Augmented World', Jun Zhao contributed to a session on 'Impact of Trust in AI', Alex Rogers was a speaker in a session on 'Earth Systems and AI' and Marina Jirotko spoke about 'Responsible Research and Innovation'.

On day two, Sir Nigel Shadbolt gave a talk on Ethics in AI, then Niki Trigoni, Alex Rogers and Andrew Markham all spoke as part of a session on 'Cyber Physical Systems', Thomas Lukasiewicz contributed to the session on 'AI and Legal Reasoning' and the whole conference wrapped up with a 'Future Makers Question Time Special' panel discussion which included Peter Millican (chair) and Mike Wooldridge.

Eighteen Oxford Computer Science papers at NeurIPS-2019

Members of the Department of Computer Science have co-authored 18 papers accepted for the forthcoming NeurIPS conference to be held in Vancouver, Canada, in December 2019. NeurIPS is the premier international forum for research on machine learning: this year 6743 papers were submitted, of which only 1428 were accepted (an acceptance rate of just 21%).

For more information on NeurIPS-2019 see: nips.cc/Conferences/2019/

Oxford Women in Computer Science

Oxford Women in Computer Science actively supports and promotes women in the field of Computer Science, running social and academic events to provide networking opportunities and a support network for students. OxWoCS events are inclusive and opened to everyone in the university.

Earlier this year we continued our OxWoCS Distinguished Seminar Series with joint lectures with the Department of Computer Science. Talks included those by Julie McCann (Imperial College), Mary Wootters (Stanford Uni), Jenna Wiens (Uni. Of Michigan) and Olga Sorkine-Hornung (ETH Zurich). We will continue with more tech talks on various areas including but not limited to mobile systems, computer vision, machine learning and artificial intelligence. We have made trips to sponsors HQs, visited GlaxoSmithKline and attended a QuantumBlack sponsored event,



CogX 2019 – The Festival of AI and Emerging Technologies. All those who took part had a great time; it was amazing to see the opportunities within the industry and networking with professionals.

We are kick starting the 2019-20 academic year with our annual welcome event and a tech talk by our sponsor, Google, on ‘The Art of Clean Code.’ These will be followed by events and activities organised by our other sponsors: office visits to QuantumBlack; workshops on algorithm fairness and machine learning by QuantumBlack; blockchain technology by Zilliqa Blockchain; and applying tech skills

to management consulting by Oliver Wyman.

[Pictured above] The 2019-20 OxWoCS Committee Members
Top (left to right): Yuge Shi (Sponsorship Officer), Jasmine Rienecker (Conference Coordinator), Stacy Xingyue Pu (Seminar Series Coordinator), Eleonora Giunchiglia (Conference Coordinator), Marilena Bescuca (Secretary) **Bottom (left to right):** Clara Pavillet (Industry Events Officer), Tajwar Nasir (Outreach Officer), Paula Fiddi (President), Klaudia Krawiecka (Outreach Officer), Lonie Sebahg (IT Officer) **Committee members not in picture:** Catherine Tong (Industry Events Officer), Romy Minko (Researchers Rep), Annazita Barry and Teodora Musatoiu (Undergrad Rep), Asmita Poddar (MSc./ Postgrad Rep), Lucy McEvoy and Eleanor Williams (Social Events Coordinators)

A trip to GHC 2019

The Department of Computer Science and OxWoCS co-funded three scholarships to attend this year’s Grace Hopper Celebration (GHC). GHC is the world’s largest conference of women technologists gathering students, academics and professionals across a vast range of specialties and areas of expertise.

DPhil students Klaudia Krawiecka, Paula Fiddi, and Alina Petrova attended the 2019 GHC in Orlando, Florida. It was the largest GHC event to date, with more than 25,000 attendees, over 250 industry sponsors, and hundreds of talks, seminars, and workshops.

Klaudia said: ‘The key theme this year was “We will. I will.”, and my personal motos are: we will keep supporting one another and empowering one another. I will continue to be an advocate for young children and educate them about technology.’

Alina said: ‘The conference is a perfect demonstration of how truly diverse the area of Computer Science is. It was not only a huge inspiration for me, but also a great opportunity to improve my networking skills, as I am transitioning from being a student to being a postdoc.’

Paula said: ‘The Grace Hopper Conference was filled with such amazing energy from 25,000+ women

in technology across the world. I feel refreshed and ready to face new challenges. The experience was empowering and inspiring and I look forward to building and maintaining connections I made during the event. My quote for this year’s theme is: “I will actively share women’s contributions in tech and create technology to improve lives.”



OxWoCS collaborates with Zilliqa to bring diversity to the forefront of Blockchain

Zilliqa, a high-throughput public blockchain platform, has launched the return of Blockchain A-Z, a series of blockchain education workshops, this time with the University of Oxford. Blockchain A-Z aims to encourage greater gender inclusion from lecture halls and beyond. The workshops have been launched in collaboration with Oxford Women in Computer Science Society (OxWoCS) and were open to all students currently enrolled at the University of Oxford. The workshops ran from 31 October 2019 and culminated with a Demo Day on November 21, 2019.

The interactive workshops aim to equip students with a foundational understanding of blockchain fundamentals from both the

business and technological aspects of implementation. Taking place over the course of four weeks, these workshops aim to provide a well-rounded overview of concepts, including philosophical foundations; technical topics; industry insights; and business advisory. As part of the Demo Day, winning participants were invited to submit their solutions to Zilliqa's US\$5 million Ecosystem Grant Programme, in the hopes of receiving mentorship, technical tutorship, and a guarantee of funding for 0% equity in return.

OxWoCS was founded in 2013, with the mission of supporting and empowering female students enrolled in programmes. Zilliqa is joined by Google, QuantumBlack, Facebook, G-Research, Semmler

and Oliver Wyman as OxWoCS sponsors.

Paula Fiddi, OxWoCS President said, 'At OxWoCS, our goal is to ensure that female scientists are presented with equal opportunities to engage with various areas of the tech industry. From established fields of research to nascent disciplines ripe for exploration, it's important that women are adequately represented in order to champion diversity both in action and in practice. With Zilliqa's emphasis on nurturing growth from the ground up, this collaboration serves to prepare our students with much-needed support, guidance, and industry connections, to truly excel in the blockchain space.'

Industry Liaison Administrator: Oliver Sampson

I have recently been appointed Industry Liaison Administrator for the Department of Computer Science in order to provide a link and initial point of contact between industry and academia. My background is in the life sciences having studied for a DPhil in Radiobiology at the University of Oxford and prior to that working for GSK and RAFT in analytical chemistry and molecular biology positions respectively. Drawing together my experience in both industrial and academic sectors I hope to use this as a foundation to foster links between the two sectors here in Department of Computer Science, University of Oxford.

As Industry liaison administrator I have the broader interests of the academic community in focus. We are not merely seeking to gain funds to undertake research projects but rather to learn and benefit from what industrial partners can bring to the equation. For example this might be data resources, knowledge or

further opportunities for collaboration. My role also incorporates our students, and I believe better industry interaction will result in a wider understanding of where their career can take them. We currently have industrial representation from finance, automotive, pharmaceutical and telecommunication sectors, each offering a different perspective on how our graduates can contribute to their business. We aim to increase this list in the coming years. This is important for the advancement of the department, the career prospects of our students, and also in attaining the best outcome as part of the Research Excellence Framework review in 2021.

In my new role I am in a position to support nascent projects where there is a mutual benefit to the Department and where industry funding and/or assets will provide synergistic benefit. Typically, time and administration involved in forming these relations is inhibitory, so being able to dedicate effort to this process should provide

greater success, and we are already seeing benefit from a number of projects with industrial links.



As an initial project, the industry relations section of the departmental website has been given a facelift and now contains much more information about the ways industry can partner with us. Please visit www.cs.ox.ac.uk and click on industry relations under the innovation tab. Of course the list of ways to interact with the department is not finite and I am always open to new suggestions on how we can further develop links from inside the department or through our industrial partners.

Please feel free to stop by my office in room 104 in the Wolfson building, Department of Computer Science for a chat or email industry@cs.ox.ac.uk

The Royal Institution Computer Science Masterclasses

Generously sponsored by Morgan Stanley

The Royal Institution Masterclasses are a series of hands-on and interactive extracurricular activities that are designed to encourage, inspire and engage young people in the creativity and practice of Computer Science, helping to

highlight the links between this wide-reaching discipline and other subjects. The series of 6 Saturday morning workshop each lasting 2.5 hours were lead each week by a different expert speaker offering students in-depth investigations of topics in Computer Science, combining theory with interactive exploration working with schools that otherwise do not often engage with

the University of Oxford. Topics run previously have included: computing meets biology, machine learning, sphero robots workshop, and turtle graphics.

This year the Department of Computer Science welcomed back our 24 x Y10 students, who attended in 2018 as Y9s, and also welcomed a new cohort of 22 x Y9 students. Between them the two cohorts represent 24 different state schools with up to two students from each school from within our catchment area of state schools in Oxfordshire, Berkshire, Wiltshire, Gloucestershire and Buckinghamshire. We hope to run the programme for a total of three years per cohort and aim to offer these masterclasses in Computer Science again in 2020.

Further information available: bit.ly/2Jk6TJS



Open day in the Park

The department regularly opens its doors as a part of the university-wide open day(s) during the year. An ideal opportunity for prospective students to explore Oxford, find out more about courses and colleges, as well as speak to tutors and students. We have had growing interest in the last few years in relation to undergraduate Computer Science

courses offered by the Department of Computer Science and as a result the number of people attending our open days has increased significantly. The most recent open day took place on Friday 20 September and to accommodate the increasing numbers the programme was split between the Wolfson Building and a marquee in University Parks. Throughout the day we had just over 500 prospective students and guests visit the Department of Computer Science. In comparison we had only 123 prospective students and guests in September 2013 whilst if we go back 10 years to September 2009 we only had 48 prospective students attending the open day!

The September 2019 open day could not have gone ahead without



the large number of people involved giving talks, demonstrating projects, running social media, and mingling with prospective students. Thank you to all members of the department involved in the open day. We will be taking part in the next university-wide open days on 1 July, 2 July and 18 September 2020.

For further information on open days or other outreach events please see: cs.ox.ac.uk/opendays



Farewell to...

The Department of Computer Science is a complex community of people. Some are only here for a short time on their way up the career ladder, but a few have been an integral part of the team for years, even decades. This year we are saying a fond farewell to several people who have helped to shape this dynamic department as it has grown and evolved throughout the years.

Sharon Backer-Holst

Sharon's official job title was 'Department Administrator / Research Facilitator / Project Manager' but even that complicated title doesn't really get close to describing what she did, and the huge contribution she made to the department over the years.



Sharon worked tirelessly to find this ever-growing community more space, using imagination and determination to somehow create new space in our existing buildings, and acquiring new buildings when she could.

In her role she made sure that the academic and support staff worked together, helping us maintain a sense of community as we become a far bigger department than the one she originally joined.

Our Head of Department Mike Wooldridge comments, 'It would be hard to overstate the role that Sharon Backer-Holst (formerly Lloyd) has played in the growth of the department over the past decade. The mission to grow and broaden the department would simply have been impossible without such a dedicated, focussed, and determined administrator and financial manager. Oxford is a complex university, and making change happen here is not easy. Sharon transformed the finances of the department, and has fought for us endlessly to get people, space, and resources. I could not begin to count the number of times I have relied on her utterly. Everybody working in the department today is in her debt.'

Sharon has always had a keen interest in the research undertaken by our academics, and especially in the various spin outs that have been

launched, so it is fitting that we are 'losing' her to department spin out company Navenio.

There is not a person here who doesn't know Sharon, and we will all miss her hugely.

Julie Sheppard

When we celebrated our 60th anniversary in 2017 it was time to look back over the history of our department, and one person was able to tell us all there was to know. Graduate Studies Administrator Julie Sheppard joined the department in 1984, and retired this summer as our longest serving member of support staff.



During her career Julie held a number of different jobs, most of them involving student support. As a result many students, past and present have very fond memories of her.

Former student Joe Loughry commented that, 'Julie Sheppard was the lighthouse I navigated by throughout my DPhil. I owe her.'

Laura Jones

Laura was an HR Manager who really cared about the community she was part of. Apart from the considerable job of overseeing the constant comings and goings of staff in a busy academic department, she also introduced many initiatives to look after our physical and mental wellbeing.



Laura was also a key member of the Athena Swan team, coordinating our award submissions and making sure that we continue to strive towards better opportunities for women in our department.

Maureen York

Another stalwart of the department, after six years as Centre Administrator for the Cyber Security CDT and over 25 years working in the department, Maureen York has departed for an exciting new life in Spain. In the time she was with us Maureen was particularly involved in the creation of several Centres of Doctoral Training (CDT's). This included the most recent project she worked on, helping Jim Davies set up the new CDT in Health Data Science. Maureen was great at bringing these complex projects together and we will miss both her negotiating skills and unique purple hair!



David Hobbs, Centre Administrator for the CDT in Cyber Security comments, 'Maureen's calm and professional attitude, effervescent personality and deeply caring perspective on life in addition to a love of everything purple, will be keenly missed. We have no doubt that Maureen will make a massive positive impact helping others with her new business.'

Kelly Ryan

Finally, Kelly Ryan left earlier in the year. Kelly was our Senior Research Facilitator, and instrumental in the success we have had (and continue to have) in applying for and being awarded research funding. She became the 'Oxford-Turing University Liaison Manager' before leaving us to work full time at The Alan Turing Institute in London.



Lego lunch

We have launched a regular Lego lunch hour. Every Wednesday in the Atrium a selection of kits and loose bricks are available so that everyone in the department can spend some time out, doing something relaxing and fun together. The idea of adults building with Lego to relax is not a new one, much has been written about how effective it can be to help take the stress out of everyday life. We welcome all members of the department, and hope to grow the Lego collection as more people get involved.



Sport at Computer Science

There tends to be a popular idea of what a Computer Scientists might be like, and ‘sporty’ is not the first description that comes to mind. But in this department we have (and have had) students involved in a wide variety of University sports, from competition-level rowing to...Quidditch!

Peter Marinov [pictured below] is a Doctoral student working within the computational biology research theme. He is involved in a number of different sports, and says, ‘I’m a tennis player, volleyball player and windsurfer.’



I love tennis as it requires me to concentrate and be present in the moment, as well as to think on my feet. I play for the Uni 4th men’s team and we won varsity last year.

I play volleyball for more or less the same reason but also really love the team goal component, where the team has to play well for us to win a match. I used to play for the Blues and now play for the 2nd Uni team.

I like windsurfing because its relaxing and frees up my mind. I like how unpredictable nature can be and the challenge to ride it.’

Maike Zwart [pictured right with her dance partner] is a Doctoral student working on research in quantum computing. She is involved in an acrobatic dance group.

She says, ‘I’m part of Flight Club, an acrobatic partner dance club. At Flight Club, we mix several partner dance styles (modern jive, jive, rock and roll, etc) with jaw-dropping acrobatics. We take part in lots of competitions, both as individual

couples and with a group routine. Basically we’re a crazy bunch having fun while climbing each other, laughing at our fails and celebrating successes. It makes for some very good holiday pictures too!



Tokyo exchange

The Department of Computer Science has for many years had an arrangement with the National Institute of Informatics (NII) in Tokyo, whereby they will provide scholarships for our MSc and DPhil students to spend time as an intern in their lab in Tokyo.

NII is a government-sponsored research institute, located right in the centre of Tokyo, just next door to the Imperial Palace. They have research groups in a variety of areas. In particular, there is a very strong group under Ichiro Hasuo on categorical semantics and co-algebras; they also have groups in quantum computing, cyber security, cyber-physical systems, graph algorithms, verification, and social technology.

With so many research areas in common, a collaboration between our department and NII has many advantages. Jeremy Gibbons is a Visiting Professor at NII and one of his former DPhil students relates his experience as a scholarship intern at NII:

Meng Wang (now a lecturer at Bristol University) recalls:

'I visited Tokyo on a scholarship in 2008, towards the end of my second year of PhD.

In my case, the area of research I collaborated on while I was at NII was very closely related to my PhD research, which was programming languages.

Academically it was not very different to working at Oxford. We were all focused on doing good research and trying to publish good papers. Generally speaking, Japanese research labs tend to be more formally organised, with a professor on top, followed by an associate professor assisted by an assistant professor. There are typically a number of postdocs, PhDs, and MSc students. The groups tend to be very close knit, with regular social events that everyone attend.

Life in Japan is (obviously) very different from that in the UK. So there is much to explore and learn. Through a Japanese friend I got to know back in Oxford, I managed to do a house-share with three Japanese young men. It was a

fantastic experience as I was able to observe local life up close. Japanese people are generally very helpful. It is not difficult to make friends if one tries.

My NII scholarship has significantly contributed to my research, DPhil and beyond. A chapter of my DPhil thesis is based on my work done during the internship, and the academic connections fostered there have been long lasting. I have been constantly collaborating with Japanese researchers since then. As we speak, a PhD student I got to know then (who is now an associate professor in a top Japanese university) is visiting me at Bristol.

I would definitely recommend this scholarship to current students. Visiting other research groups is always a good way to generate new ideas, and Japan is a place where top-rate research is done. The social life is very colourful too.'

Read more about NII here: www.nii.ac.jp/en/research/centers/



Student Profile

Sara Dutta : A DPhil in Computational Biology who has gone on to work with the FDA and pharmaceutical and medical device companies.



What course did you study here and when?

DPhil in Computational Biology (DTC Systems Biology 2009) 2010-2014

What was your background before that?

I studied Computer Science BEng at Imperial College London and I did my high school at the Lycée François Premier in Fontainebleau (France) where I grew up.

What attracted you to studying Computer Science as a subject?

I always enjoyed mathematics and logic at school and I was fascinated by how it could be applied in Computer Science to build algorithms that could solve amazing problems.

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

I particularly enjoyed interacting and learning from my colleagues in the department, such as during the student departmental conference and having the opportunity to be a demonstrator and tutor during some undergraduate courses. Most importantly my team under Professor Blanca Rodriguez was a great source of support, fun and knowledge.

What did you enjoy about life at Oxford?

I really valued interacting with friends and peers from other subjects. Balliol MCR dinner parties were a source of exchange and fun as we were a diverse group of PhD students from a range of subjects from Physics, Computer Science, Philosophy, Literature and Economics. I also enjoyed that there were always fascinating talks, conferences, concerts and other events that were easily accessible. At times it was intimidating, but mostly I was just fascinated by all professors and students around me, and tried to learn from them as much as possible. I have particularly fond memories of being part of the organising committee for the Balliol's 750th Commemoration Ball and rowing for Balliol during Torpids and Summer Eights.

What did you do when you left Oxford?

I went to work for the U.S. Food and Drug Administration in the Washington D.C. (to be precise, Silver Spring MD) where I investigated how computational models of the heart could be incorporated in the FDA drug approval

process to improve prediction of lethal arrhythmias. I worked there for 3 years

and then moved to

Boston over a year ago to work for Dassault Systèmes (the 3DEXPERIENCE Platform company) where I help pharmaceutical and medical device companies embrace digital transformation throughout their lifecycle, one of the main ways being through modelling and simulation.

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

The work I did at Oxford has really helped lay strong foundations for all the work I have done since then. All the technical knowledge I learnt during my PhD has been directly applied during my research work at the FDA and now I use that technical expertise to help demonstrate the business value of such tools to our life sciences companies. I also have kept close ties with my team from Oxford throughout and they have always been a great source of support throughout my time at the FDA and in my current job.

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

Computer Science gives you a great set of analytical and logical skills that can be applied to virtually any field. Don't be scared to try areas a little out of your comfort zone or that seem a little too ambitious. My 18 year old self would have never thought that computational modeling of the heart existed and was something she would make a career out of.

If we went back in time and asked, what would the student you have thought about what you are currently doing – would you have been surprised, proud, amazed?

I think I would have been impressed and I would have found it hard to believe that there are so many opportunities to continue to learn and apply my knowledge outside of the traditional academic route. I really believe you can forge your own path and be surprised at where it takes you.

Detecting extremist content online

Mariam Nouh discusses new research to automatically detect radical content in social media.

In recent years we have seen several examples of how terrorist groups and violent extremists exploit the internet to advance their agendas and promote their actions. For instance, the Islamic State in Iraq and Syria (ISIS) creating ISIS-branded media outlets and heavily utilising the internet and popular social media platforms to broadcast propaganda hoping to spread fear and radicalise individuals. Similarly, the recent attack on a mosque in Christchurch (New Zealand) was broadcasted live on multiple online platforms, and many of the attack plans were advertised online a few days before the attack took place. Such incidents show the importance of identifying radical and extremist content online in a timely manner.

In this research we aim to automatically identify signals to detect online radical content. We mainly focus on the ISIS group and we analyse the propaganda material published in their online magazine, *Dabiq*. *Dabiq* is published by ISIS with the purpose of recruiting people and promoting their propaganda and ideology. Using this data source, we investigate what topics, textual properties, and linguistic cues exist in these magazines. Our intuition is that utilising these linguistic cues extracted from extremist propaganda should allow us to detect supporters of ISIS who are influenced by such propaganda.

An overview of our two-phased approach is presented in *Phase 1: Radical Properties Extraction*. First, we build a radical language model and a psychological profile model

from the ISIS propaganda corpus. To create the radical language model, we use different methods to extract the radical language from the corpus, such as Term-Frequency Inverse-Document-Frequency, and neural network model trained on word embeddings to capture semantic.

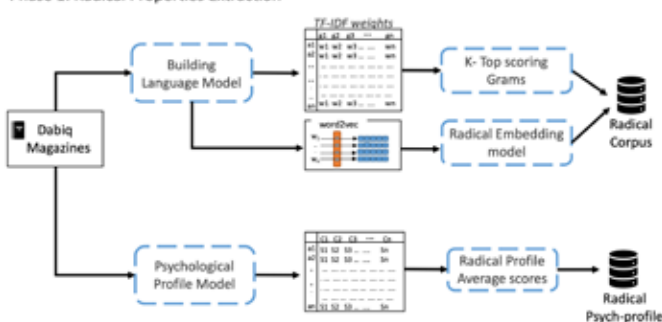
Moreover, inspired by research in fields such as linguistics, social science, and psychology that suggest that the use of language and the word choices we make in our daily communication, can act as a powerful signal to detect our emotional and psychological states, we extract psychological properties from the radical corpus in order to understand the personality, emotions, and the different psychological properties conveyed in these articles. Such psychological properties include the 'OCEAN' or 'Big Five' personality traits which are openness, conscientiousness, extraversion, agreeableness, and neuroticism. A number of studies looked at the motivating factors surrounding terrorism, radicalisation, and recruitment tactics, and found that terrorist groups tend to target vulnerable individuals who have feelings of desperation and displaced aggression. In particular, in research into the recruiting tactics of ISIS groups, it was found that they focus on harnessing the individual's need for significance. They seek out vulnerable people and provide them with constant attention. Similarly, these groups create a dichotomy and promote the mentality of dividing the world into 'us' versus 'them'. In our research, we extract psychological properties from the radical corpus in

order to understand the personality, emotions, and the different psychological properties conveyed to readers in these articles.

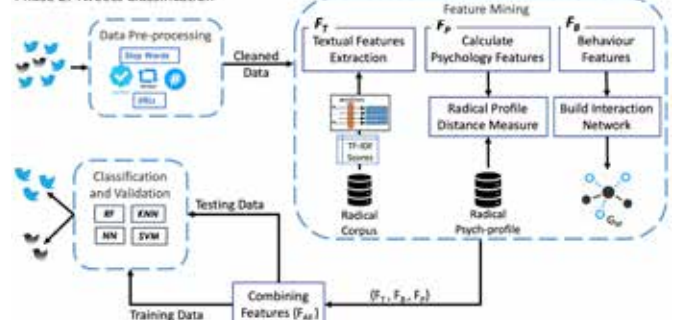
Furthermore, we make use of the radical language model and psychological model to detect pro-ISIS tweets that are influenced by their ideology. Unlike previous efforts, our approach does not only focus on lexical keyword analysis of the tweets, but also adds a contextual and psychological dimension. Using machine learning classification models, we are able to distinguish between radical and non-radical tweets. We validated our approach in different experiments and the results show that this method is robust across multiple datasets. Among the most important feature for distinguishing radical tweets is the distance between the psychological profile calculated from the propaganda magazine corpus and the Twitter user. In addition, the us/them dichotomy which looks at the total number of pronouns used (I, you, we, they) was significant.

These identified signals are indeed critical to help improve existing efforts to detect the spread of online violent extremist content and radicalisation behaviour. This system can aid law enforcement and Online Social Network companies to better address such threats and help solve a challenging real-world problem. Finally, it is important to note here that such methods should always be guided by a human analyst to oversee the results of any automated models. The complete research paper is available here: <https://bit.ly/2D1QRAY>

Phase 1: Radical Properties Extraction



Phase 2: Tweets Classification





New standards to protect children online

A requirement of the Data Protection Act 2018 (which supports and supplements the implementation of the EU General Data Protection Regulation) is the introduction of an Age Appropriate Design Code. This will set out the design standards required of online services and apps likely to be used by children in relation to how these services and apps collect and process children's data. The background to this requirement is the ubiquitous collection of information about people whenever they use online tools. This information forms 'digital footprints' and may contain revealing data about our identities, behaviours and so on.

There are growing concerns that this collection of information can threaten our privacy, data rights, and even our online safety. Children are spending more and more time online; evidence suggests¹ that in 2017, over half of 3-4 year olds were online and nearly all 12-15 year olds were.

Despite this, online services and apps are typically designed with the adult internet user in mind and there is a clear need for specific guidance for organisations offering online tools used by children so that they can better develop websites, platforms and apps etc that protect and uphold the rights on young internet users.

The Information Commissioner's Office (ICO) has been tasked by the UK government to produce this Age Appropriate Design Code. It will ultimately recommend design standards for data processing that providers should follow when they produce online services and apps

likely to be used by children. The ICO produced a draft Age Appropriate Design Code in May 2019². This draft provides practical guidance on 16 standards of age appropriate design. The standards include:

- Primary consideration given to the best interests of the child
- Consideration of the age range of users and the needs of children of different ages
- Transparent privacy information that is easy for users to understand
- Avoidance of the use of children's data in ways that is detrimental to their well being
- High privacy settings by default
- Geolocation switched off by default
- Avoidance of data sharing unless there is a compelling reason to do so.

On publication, the draft Code was opened up for public consultation so that Individuals and groups could provide feedback on it. The work done by members of the Human Centred Computing theme at the University of Oxford takes a very strong focus on the well-being of online users and several of our projects have investigated children's use of technology. We were therefore very keen to submit a response to the draft Code.

Our response welcomed the recognition of responsible data handling as a critical issue in the contemporary digital economy and as an issue that requires regulatory attention. We also particularly welcomed the focus on children due to their need for specific protections online. Promoting age appropriate design in online services is extremely complex; there are nuances around determining and applying relevant age ranges for children, and it can be very difficult to identify and assess all potential risks and harms.

However, in our view, the 16 draft standards in the Design Code are comprehensive and well-defined. Several of them can be achieved with little practical complexity – for instance, the default for high privacy and geolocation off. Based on the findings of our own projects, our response to the draft Code argued that fully serving the best interests of children requires developers to engage with young people, parents, carers and educators. This will help them to best understand children's online experiences and form part of a community that serves to care and take responsibility for young people online.

We also argued that developers and organisations should incorporate a long-term view of what constitutes detrimental use of data. Protections should be put in place even when the harm or risk of harm cannot be discerned in the short term as the long term and cumulative consequences of profiling can be substantial.

1 <https://bit.ly/2r82XWM>

2 <https://bit.ly/2P6rBPH>

Engaging with your children:

recipes for bringing up the 'AI natives', a brief background of the KOALA project

The rapid application of Artificial Intelligence technologies is having a deep impact on our daily life, such as more accurate real-time traffic notifications or better support for our healthcare. However, AI not only makes our lives more efficient but may also influence the information we consume online. AI is becoming far more ingrained in our lives than we could have imagined.

At the heart of this new form of AI-powered digital life there are our children, the so-called 'AI natives', who grow up with smart speakers from infancy and watch more YouTube than TV. We wonder, how does this generation perceive and manage their life with these technologies? How do technologies impact their online wellbeing? To what extent can our children recognise and manage risks that they may encounter when interacting with technologies? These are the fundamental research questions that Oxford Human Centred Computing researchers set out to explore in the KOALA project (<http://koala.web.ox.ac.uk>).

The general goal of Kids Online Anonymity & Lifelong Autonomy (KOALA) is to investigate the impact of personal data collection practices used by mobile apps, upon the general well-being of young children.

The project started in 2017, and we have worked with over 40 families from 16 different local primary schools. We focus on children aged between 6 and 10 because their experience with digital technologies has not yet been extensively studied in existing literature. We are curious to learn about how this age group may perceive the role of digital technologies in their lives.

We first seek to understand children's existing use of digital technologies and what they may enjoy and struggle with. We are particularly interested to find out how children may describe any risks or barriers that they experience, so that we can communicate with children in a language that is most familiar to them [1]. We have identified that children have an amazing ability to use a wide range of language to capture their experience online (such as 'recommenders', 'youtubers' or 'autoplay').

However, we also identified their lack of ability to recognise implicit persuasive design features (such as the game/video promotions) in the technologies they interact with. This lead to them spending more time in front of screens or consuming content less appropriate to their age.

These are critical understandings that provided key inputs to the development of KOALA Hero (<https://koala.web.ox.ac.uk/toolkits>), an Android mobile app prototype that is aimed at supporting parents in engaging with their young children to discuss the use of digital technologies and to co-learn implications associated with the use of these technologies.

In early 2019, we presented our early app prototype to six local families. Our parent participants told us that they found existing the control mechanisms on the market insufficient for their needs, and they welcomed KOALA Hero's alternative approach to encourage parent-child engagement and co-learning; they agreed that children must 'learn' instead of simply being shielded from risks. Parents also appreciated the additional expert reviews about

the apps (which we collected from various information sources); but expressed that they need more support to establish trust on such information or to make use of it.

This is positive feedback to an initial prototype which tries to disrupt the existing market-leading approaches, which mainly focus on enabling parental control and monitoring. KOALA Hero's next step therefore will focus on providing better support for parent-child engagement by bringing in existing parenting theories and providing better mechanisms for parents and children to contribute and share their own experiences with technologies with their peers. We hope that in the future parents will become empowered to support each other.

In the meanwhile, we would encourage parents to:

- Engage with their children, by discussing what they enjoy during their interactions with technologies.
- Start the conversation from a young age, to develop a habit of co-learning and co-experiencing a new form of digital childhood
- Pay more attention to what children do in front of the screens, instead of how long.

More information about the KOALA project: <http://koala.web.ox.ac.uk>.

[1] Zhao et al. 'I make up a silly name': Understanding children's perception of privacy risks online. CHI 2019






New research creates computer simulations of heart attacks

Hector Martinez-Navarro, a DPhil student in Computer Science, has essentially simulated the electrical consequences of a heart attack, as part of his research in the Oxford Computational Cardiovascular Science group, supervised by Ana Mincholé, Alfonso Bueno-Orovio and Professor Blanca Rodriguez. He has developed a computational pipeline for modelling and simulation of how the electrical activity of the human heart is affected by the block of an artery during a heart attack. The study has been published in Nature's *Scientific Reports*.

Computational modelling and simulation are becoming ubiquitous in medical research. Widely exploited in many areas of engineering (eg aeronautics), computational modelling and simulation based on experimental and clinical data represents an alternative to animal and human experimenting, overcoming multiple ethical and technical limitations. Computational modelling and simulation are now also generating digital evidence for the development and evaluation of medical therapies. Both the European Medicine Agency and US Food and Drug Administration are currently developing new regulatory pathways that include modelling and simulation for both devices and drugs.

In this study, detailed mathematical models of human cardiac cells were incorporated into a whole-heart model, allowing the simulation of lethal arrhythmias and clinical signals such as the electrocardiogram. This approach was adopted by the authors to conduct highly biophysically-detailed simulations of the human heart electrophysiology under ischemic conditions, when cells in the region affected by coronary occlusion are deprived of oxygen and nutrients. These are computationally expensive, requiring the use of high performance computing, such as ARCHER (<http://archer.ac.uk/>), the UK's largest supercomputing facility.




During a heart attack, a coronary artery is occluded or obstructed, and blood irrigation of the cardiac muscle is compromised. This leads to ischemia in the affected tissue: the tissue is starved of oxygen and can eventually die. Myocardial ischemia leads to abnormalities in the electrical propagation in the heart muscle, potentially causing life-threatening arrhythmias (irregular heartbeat). Variability in symptoms and arrhythmic risk caused by ischemia remains a challenge in the diagnosis and management of patients. Some ischemic patients exhibit ECG abnormalities, whereas others show no clear changes in the ECG. The relationship with arrhythmic risk is unknown.

Changes in the simulated electrocardiogram in the virtual human heart affected by coronary occlusion. *Transmural ischemia* (when the whole thickness of the heart wall is affected) [left] causes larger electrocardiographic abnormalities than *subendocardial ischemia* (which is when the inner layer of the heart muscle is affected) [right]. In both cases, the dashed line corresponds to a control non-ischemic simulation. LAD occlusion of the Left Anterior Descending artery, the most important artery of the human heart in blood supply.

continued on next page ►

from previous page ►




Arrhythmia around the region of ischemia (figure-of-eight pattern) in transmural ischemia, affecting the whole ventricular wall from the ventricular cavity to the epicardial surface.

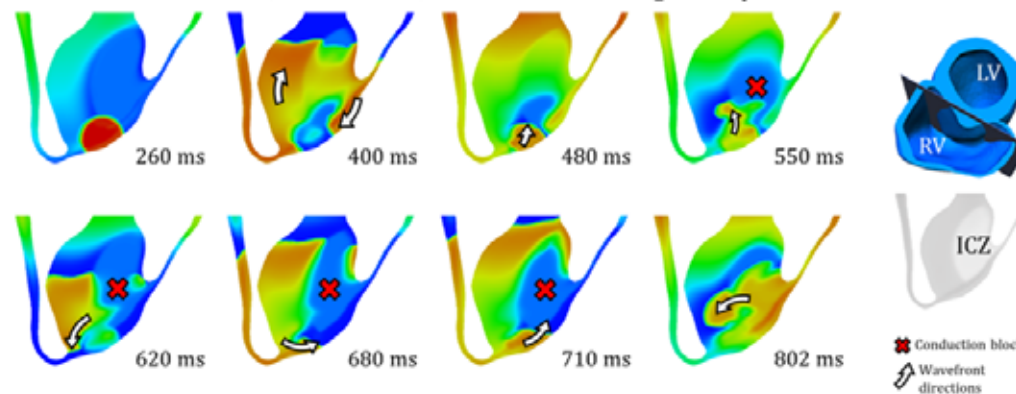
In this study, the authors create multiple virtual scenarios of ischemia reproducing the variability observed clinically, using experimental and clinical data for model construction and validation. The scenarios considered included different sizes of ischemic region, extensions across the ventricular wall, and locations through consideration of the occlusion of different coronary arteries. Fully transmural ischemia scenarios (ischemia affecting the entire myocardial wall) reported remarkable changes in the simulated electrocardiogram, and typical arrhythmic patterns around the afflicted region.

Subendocardial ischemia, on the contrary, only affects the inner part of the myocardial wall, close to the ventricular cavities (eg the endocardium). Even though this only produced mild or negligible changes in the simulated electrocardiogram, the occurrence of arrhythmias was still very high, in agreement with clinical evidence. Through detailed analysis, arrhythmias in subendocardial ischemia are shown in this study to be linked with a different mechanism, as they are sustained transmurally within the myocardial wall, rather than around the region affected by the coronary occlusion.

B Subendocardial ischemia, LAD occlusion, epicardial view



C Subendocardial ischemia, LAD occlusion, transmural view through the septum



Transmural arrhythmia in subendocardial ischemia. The top panel shows the cardiac surface, providing no information about the arrhythmic pattern. The bottom panel shows the transmural pattern with a detailed view, where the arrhythmia develops in the form of a spiral waves inside the ventricular wall.

Thereby, the simulations show that dangerous transmural patterns of arrhythmia can occur in patients with only subtle changes in the electrocardiogram. This would explain the still high arrhythmic risk reported in patients suffering from silent ischemia episodes, typically produced by incomplete occlusion of their coronary arteries.

The computer modelling and simulation technology and results in this study are an illustrative example of the successful application of in silico models in biomedical research. The human heart models and simulations in disease can be used as a test bed for therapies in in silico trials, to inform further experimental and clinical studies.

Read more: <https://bit.ly/2qjIBts>

Virtual Assay: user-friendly software for in silico drug testing

By Elisa Passini

In the last few years one of the research lines within our Computational Cardiovascular Science Group has been focusing on how computer models can be used to predict drug safety and efficacy for the human heart. Driven by the idea that drug testing could be performed by a computer rather than on animals, our researchers developed the Virtual Assay Software which provides a user-friendly graphical user interface to run human in silico (computational) drug trials in populations of heart models. Predictions obtained with the software will guide drug discovery in the future, identifying potential side effects for the human heart in the early stages of drug

development thus leading to better and safer medicine. Virtual Assay is also easy to use by non-experts in modelling and simulations providing a very efficient simulation engine (1 drug trial takes about 5 minutes for 100 cells, using a modern PC).

The first version of the Software was released in 2014, and after almost 5 years of iterations, guided by the feedback received from our pharma collaborators testing the software, Virtual Assay V.3.0 is now ready. A free academic licence and a commercial licence are available through Oxford University Innovation. Multiple evaluation studies have been performed using Virtual Assay, and

have demonstrated the ability of the software to predict the risk that drugs would cause abnormal heart rhythms in patients with almost 90% accuracy, whereas similar studies conducted in animals showed 75-85% accuracy.

The potential of the software for the replacement of animal experiments has been recognised by two international 3Rs prizes, awarded to Dr Oliver Britton (2014) and Dr Elisa Passini (2017). Through dissemination of our results, we have also contributed to changing the culture in pharmaceutical industries, opening their minds to computational approaches. Read more: <https://bit.ly/2Lq0yhf>

Beyond Academia: what we have learned from our experience with industry

By Elisa Passini

Since cardiovascular side effects are still one of the major problems during drug development, our research on human in silico drug trials is very relevant to pharmaceutical companies. By disseminating our results in industry-oriented conferences, we have been able to attract the attention of the people working in the safety pharmacology departments, and we have established successful research collaborations with more than seven big pharma companies in the last six years. This has given us the opportunity to appreciate many differences between academia and industry.

First, every new methodology is evaluated in terms of quality, time and cost, and strong evidence must be provided to demonstrate its advantages compared to the existing ones.

Second, a good management of time and deadlines is essential. Even before starting a job, we are required to provide an accurate estimate of how much time/work will be needed to complete it, and any publication – even a conference abstract – needs to go through an internal review process, which could take up to a month, before submission. Things are quite different in academia, where a project might end up taking

a very different direction compared to what was planned originally, and where last-minute deadlines happen all the time. Also, industry tends to take into account a larger time scale when evaluating new methodologies. They want to make sure that what is working today, will also be working in a few years, and sustainability in the long term is crucial.

Third, changes take time. The CiPA initiative (<https://cipaproject.org>), sponsored by the US Food and Drug Administration and other organisations, started in 2013, with the aim of implementing a new paradigm for drug safety assessment, also including computer models. The proposed implementation was supposed to happen by 2015 and now, 4 years after this – very optimistic – deadline, things are still work-in-progress. However, changes are slowly happening, and we can proudly say we are making our contribution.

For me – as an early career researcher – establishing good collaborations with industry has been very important. It made a big difference in terms of extending my network of contacts, opening the way for future research projects and grant applications. It gave me insights on the operational

aspects of pharma, and also the opportunity of appreciating more applied and structured style of research. It also gave me visibility and personal recognition, through the International 3Rs prize and the Technological Innovation Award (Safety Pharmacology Society). This also led to numerous invitations to present my research to national and international meetings, to different audiences with varied expertise. Managing and nurturing the industrial collaborations, as well as contributing to large industry-led projects/consortiums, gave me the opportunity of establishing my role within the group, maturing into a more senior research position. All this will have implications for future grant/fellowship applications, where I will be able to bring on board industrial collaborators, providing support and experimental data, which are essential to develop our computational models.

Finally, from a personal point of view, it is very rewarding to see that our research can achieve impact beyond academia, pioneering changes that could lead to a revolution in the way drug safety and efficacy testing are currently performed by pharma industry, with concrete implications such as the reduction of animal experiments.

Training autonomous vehicles using real-life human behaviour

By Stuart Gillespie

Driverless cars are on their way – there’s little doubt about that. But before they hit the UK’s roads, they need to be tested in realistic simulations to ensure that this transformative technology will be a safe and positive addition to our lives.

Latent Logic, an Oxford University spinout company, is helping enable this vital testing with technology that teaches autonomous systems using real-life examples of natural human behaviour on the roads.

Professor Shimon Whiteson, co-founder and chief scientist at Latent Logic, explains:

‘Autonomous vehicles must be tested in simulation before they can be deployed on real roads. To make these simulations realistic, it’s not enough to simulate the road environment; we need to simulate the other road users too: the human drivers, cyclists and pedestrians with which an autonomous vehicle may interact.

‘Latent Logic is using a machine learning technique called imitation learning to build realistic human behaviour models. These models make it possible to test autonomous vehicles quickly and robustly.’

Latent Logic grew out of an EU-funded research project that trained semi-autonomous telepresence robots to behave in a socially ‘normal’ way. Since it’s difficult to quantify what is meant by socially normal, it’s much easier to train such systems to imitate the behaviour of humans acting in a socially normal way. Shimon decided to explore the commercial potential of this technology, recruiting postdoctoral researcher Dr João Messias to be co-founder and chief technology officer. Kirsty Lloyd-Jukes then joined the company as CEO.

The technology works by combining state-of-the-art computer vision

with imitation learning. Shimon says: ‘Our models extract the “latent logic” behind real-life examples of natural human behaviour. As a result, they can respond realistically even in new situations.

‘We use computer vision to collect these examples from video data provided by traffic cameras and drone footage. We can detect road users, track their motion, and infer their three-dimensional position in the real world. Then, we learn to generate realistic trajectories that imitate this real-life behaviour.’


By providing a service that enables better training and testing of autonomous vehicles, Latent Logic hopes it can hasten the safe introduction of what will be life-changing technology. Shimon adds: ‘While many of the biggest players in this market are international, there is also a lot of energy in the UK in this sector, and Oxford is a hotbed of talent and entrepreneurship in machine learning, robotics and autonomous vehicles.

‘Autonomous vehicles are improving rapidly, but we are still some way from realising the dream. It is not just about perfecting existing technology – there remain fundamental unsolved problems in building sufficiently robust autonomous systems. At Latent Logic, we believe our technology can play a critical role in addressing those unsolved problems.

‘In addition, our technology has numerous other applications. Situations in which you want socially normative behaviour are great candidates for learning from demonstration, as is robotics – from factories to warehouses to homes. You can also think about video games, where people might want to play against bots that can imitate the style of their favourite professional gamers. The sky’s the limit.’

Latent Logic grew out of a research project which trained semi-autonomous telepresence robots to behave in a socially ‘normal’ way.

Read more: www.latentlogic.com



The knob attack on Bluetooth

By Professor Kasper Rasmussen

Bluetooth is a wireless communication technology that is present on most modern devices. Every mobile phone has a Bluetooth chip and Bluetooth is used by people around the world every day, eg, to enable your phone to communicate with a headset, your car, or your smart watch. It is the protocol of choice for connecting devices together and is used in everything from kitchen appliances to implantable medical devices. If you buy a new Tesla and want to open your car with your phone (yes that is a new feature Tesla offers), your car and phone communicate via Bluetooth. Anytime two devices are meant to communicate wirelessly over a short distance Bluetooth is the obvious choice because of its ubiquity and the fact that it offers a secure connection through which to communicate.

Bluetooth offers several different modes of communication, including those that offer a confidential and authenticated link. When two Bluetooth devices want to communicate securely they first have to be 'paired'. Pairing is a procedure by which the two devices agree on a long term secret key that forms the basis for future secure communication. This key is often authenticated through an out-of-band channel, for example by a user reading digits off the screen of one device and typing them into the other device to verify that the two devices do in fact share the same key.

Using this long term key, two devices that wish to communicate, can now establish a new session. A new session is established every time two devices wish to communicate after having been disconnected. When a new session



is started a new session-key is generated in order to encrypt and authenticate the data exchanged during the session. It is in this session-key establishment procedure that the Bluetooth standard has a critical flaw. When two devices wish to establish a new session-key, they do so by first generating a preliminary session-key utilising freshly generated random numbers as well as the Bluetooth address, and the long term pairing key. These parameters are passed through a key derivation function to generate the preliminary session-key. This preliminary session-key is a perfectly good 128 bit long session-key that could be used to secure the session adequately. However as part of the session-key generation mechanism there is an additional step that requires the devices to reduce the entropy of the freshly generated session-key to a value that the devices must negotiate ahead of time. This

value can be anything between 1 and 16 bytes. The problem is that this negotiation is not protected in any way, and it therefore allows an external adversary who is present during the negotiation, to force the devices to generate a session-key with only a single byte of entropy. That means that the resulting session-key only has 256 possibilities, all of which can be quickly checked by the adversary and thus it allows the adversary to obtain the session-key.

Once the session-key has been obtained an adversary can choose to passively listen in on the session, ie, eavesdrop on the data being shared, or the adversary can choose to participate in the conversation, pretending to be either one of the endpoints. The adversary listen in, and inject arbitrary data into the connection, for the duration of the session.

continued on next page ►

from previous page ►

This attack is not against a particular implementation of the Bluetooth standard but rather against the standard itself. For this reason this vulnerability is present on every single Bluetooth enabled device in the world, which means every single mobile phone, every single laptop, every single smart watch, and most cars produced in the last 15 years. Millions and millions devices are affected. To make things worse, it is far from trivial to roll out a fix for the problem. The Bluetooth standard mandates that the protocol is handled not by the operating system of the device, but by a separate Bluetooth chip-set which hides the session establishment details from the rest of the system. This means that the only way to change the behaviour of a Bluetooth device, is to upgrade the firmware of the Bluetooth chip-set, which is not always possible, and even if a device allows this, not all users would be willing or able to perform such an update on all their devices. Even though this attack was disclosed to the Bluetooth working group, and they had eight months to come up with a workaround, many devices will continue to be vulnerable for many years to come.

In order to take advantage of this vulnerability, an adversary must be within Bluetooth range when two devices are starting a new session. This seems like it would limit the attack surface somewhat, and it does, but not as much as it might first appear. First of all, starting a new session is something that happens all the time, and if an attacker arrives late, say he arrives in a cafe and there is already a Bluetooth connection in progress,

the attacker can simply jam the ongoing transmission and wait for one of the devices to re-establish the connection. In re-establishing the connection the devices will start a new session, and the adversary can interfere with the session establishment and make the devices derive a weak key, that in turn can be immediately cracked.

An attacker can basically force victim devices to start a new session whenever he wants, however he still has to be within Bluetooth range to make the attack work. This range is often listed as around ten metres, and while that would expose a connection to attack from anyone in the same cafe, office or apartment complex, it does limit the attack surface. Or it would, if the ten meter range was really true. The Bluetooth range is ten metres only when using a normal antenna and normal equipment, however with a very simple directional antenna, that anyone can construct from a tin can and a couple of wires, the range of Bluetooth can be increased up to a kilometre or more without much effort. This of course assumes that the adversary is able to point their directional antenna accurately but it certainly increases the potential range of the attack.

In response to our findings the Bluetooth working group has released an errata to the standard and will no doubt make sure that future revisions of the standard either removes the entropy negotiation all together or mandates that devices choose the maximum value of 16 bytes (128 bits). The International Consortium for Advancement of Cybersecurity on the Internet (ICASI) — an industry led coordination body founded by Intel, Microsoft, Cisco, Juniper and IBM — has spent the last eight months trying to develop a workaround for the various devices and operating systems that these companies produce. That means that while the attack itself is not prevented, the operating system can try to detect if the Bluetooth connection was subject to such an attack, and if so, tear down the connection and try to reconnect. We are currently working on an analysis of the effectiveness of this workaround.

For the foreseeable future, it would be wise to treat any Bluetooth connection as insecure, thus if you need a secure connection you must provide your own security on the application layer.



Bayesian deep learning for all humankind



By James Parr, Director - Frontier Development Lab (FDL) Europe
Hosted at Kellogg, Oxford

AI research is exploding. So too are the kind of problems it is inherently good at tackling; surfacing insight from large or noisy datasets; spotting anomalies; fusing data and reducing the time and complexity of making informed decisions.

Deep learning/machine learning has emerged on the world stage at just the moment the world needed it the most. However, the journey to deployed and trusted AI workflows is only just beginning.

The barriers are often systemic (such as the perennial challenge of AI ready data, compute costs and access to talent) but some barriers are also about the fundamental nature of neural nets. Can we trust them? Can we explain how they've reached their conclusions? Can we quantify their uncertainty?

In relations to the questions we are seeking to explore, the Frontier Development Lab (an applied AI research lab working in partnership with European Space Agency and Oxford in Europe; and NASA in the US, alongside leaders in commercial AI, such as Google Cloud and Intel) was very fortunate to meet Yarin Gal: an Associate Professor of Machine Learning at Oxford's Department of Computer Science. Yarin's particular specialty is at the intersection of neural nets and probability theory - Bayesian Deep Learning - key to making AI a useful applied tool.

Eight years after the 'Deep Learning Big Bang' in 2012, when neural nets running on GPUs were shown to match human capabilities in image recognition for the first time, the field is starting to branch off into specialty areas such as reinforcement learning and adversarial techniques. In response to the relentless pace of innovation, AI labs have started to specialise. Yarin and the Oxford team, helped steer FDL's specific niche into creating deep learning/machine learning solutions that could be explained - a crucial competitive edge when creating solutions for the exacting expectations of a wary space industry: Is it safe? Can we interpret what's happening? Can we be confident in the outcome?

FDL was established to apply cutting edge machine learning for the good of all Humankind. However, it is the 300 year old Bayes Theorem which is the key to making this dream viable.

FDL's success has been due to the hard work of Oxford Researchers such as Atılım Güneş Baydin, Steve Reece,

Bradley Gram-Hansen, Lewis Smith, Tim Rudner, Tim Seabrook, Richard Strange and many other talented researchers from around the world who have pushed the envelope of machine learning application.

For all Humankind

During FDL, the teams work intensely over a period of eight weeks, supported by world-class mentors and industry partners as well as the space agencies - with remarkable results. A particular focus has been climate and disaster response, where earth observation data is analysed by machine learning to give improved situational awareness - in other words, more timely information to humanitarian agencies, governments and communities in case of natural disasters.

Rapid damage assessment remains a problem in disaster response. FDL teams have used machine learning techniques on Earth observation data to make substantive improvements on inundation prediction (time to flood) and prototype the use of machine learning on board a spacecraft in orbit to map a flood extent, enabling maps to be computed in space and sent to first responders on the ground. They have used AI's ability to fuse data to assess building damage, but before all that, you need to know where the people are. It is hard to believe that in 2019 we still don't have access to detailed maps of vulnerable communities.

Using machine learning to create a census of Informal Settlements

One-third of the world's urban population lives in informal settlements; these are the most socially and economically vulnerable communities on Earth. Sometimes they are located within city limits. Sometimes in very rural areas. People living in these areas often have no security of tenure, they often lack basic services - such as medicine or immunisation, or infrastructure, such as waste disposal or clean water. Housing rarely complies with building safety standards. Moreover, due to zoning issues and

continued on next page ►

from previous page ►

other oversights, often these settlements are simply overlooked - becoming hotspots for poverty and disease.

FDL used machine learning to combine spectral information from satellites and ground truth information on building materials to automate mapping of informal settlements on an unprecedented scale and frequency (referred to as temporal resolution).

The result of this approach is a tool that can provide census information on a country, or even continental scale at virtually no cost, with a very high degree of temporal frequency (every 5 days). This work will hopefully contribute to projects such as WorldPop (<https://www.worldpop.org>). These capacities are not only crucial for aid agencies such as UNICEF and the WHO, but also the Red Cross and the Disaster Emergency Committee (DEC), who often have poorly defined maps of disaster prone areas.

Using machine learning to improve disaster response

During a hurricane, earthquake or flooding event, the experience on the ground can be compared to the ‘fog’ described in battle zones - where a combination of stress, breakdown in trusted information services, uncertainty and panic, conspire to exacerbate the immediate challenge of the disaster itself. Moreover, despite our century’s rapid acquisition of peta-scale geospatial data - even in 2019, the translation of this data into insight on the ground during a disaster is still woefully slow. Often it is only after the Disaster has past do first responders really get an understanding of the damage to property, infrastructure and sadly, human life. Lack of information about affected communities and the level of damage also hinders efficient response coordination (such as funding) by authorities and insurance companies. So the question is can machine learning help?

The FDL team was able to show the potential of data fusion to enrich high-resolution optical satellite images



Fig 1: A future Disaster Response Dispatcher based on AI / Earth observation approaches(Source: FDL)

with medium resolution, multi-spectral optical and radar satellite imagery (SAR) to automate the creation of earth observation-based disaster impact maps (Fig 1) for first responders, affected communities and aid/assistance coordinator.

This machine learning (called Multi3Net) fuses optical and radar imagery to penetrate clouds and identify changes on the ground. The method consists of multiple streams of encoder-decoder architectures that extract the damage from medium-resolution optical images, radar data and spatial information from high-resolution images. The resulting representations are combined on a single map that accurately captures damage assessment on the ground.

Multi3Net exceeds state-of-the-art methods for assessing flood, earthquake or hurricane damaged structures, moreover this can be performed using freely available medium-resolution data instead of significantly more detailed (and expensive) very high-resolution data. This approach - in theory - enables damage assessment to happen during the disaster itself, helping governments to better coordinate medium and long-term financial assistance programs for affected areas. The bottleneck remains the frequency of satellites providing global coverage, but plans are in pace for mega constellations of satellites that will allow the ‘just-in-time’ insight discussed here.

The future: A Mission Control for Planet Earth

The concept of a ‘Mission Control’ for Earth emerges as an aspirational endgame for these technologies, with ML enabling data fusion, classification, prediction and change detection. This vision of a joined-up ‘planetary stewardship’ still requires a number of innovations in workflow to optimise AI techniques, before the human race sees the benefits. FDL, with the help of Oxford Computer Science, is already well on its journey to making this happen.



Fig 2

About Time: Changing security behaviour through smartwatch games

By Meredydd Williams

Smartwatches have grown in popularity over the past decade. Whether Apple, Android or Fitbit, millions of users are strapping functionality to their wrists. Indeed, smartwatch sales have increased by more than 60% in 2019. These devices provide connectivity and convenience in a mobile form factor. However, as with most digital technologies, these benefits come with a risk.

To provide such functionality, smartwatches access a range of data. Some of this, such as texts and contacts, may come from a synched phone. Other apps might be installed directly on the device. As watches are strapped to the wrist, they can also collect GPS and heart-rate data. Some models do include security settings. However, despite the risk, users tend to avoid this protection. After all, what threat could my watch pose? The challenge is persuading individuals to guard their devices.

Awareness campaigns tend to be the default response. Cyber Streetwise adorned tube stations and TV screens across the country. However, despite a million-pound investment, the campaigns had little effect. The reason is clear: awareness is not enough. Individuals also require education and incentivisation. Without the former, they lack the protective skills. Without the latter, they will be apathetic to act. Serious games take an active approach, embedding incentives through positive reinforcement. Meta-analyses suggest that apps have greater efficacy than awareness campaigns. Therefore, we developed the first game for smartwatch security.

Our app was designed through best practice from learning science and behavioural psychology. However, due to the scarcity of smartwatch games, we were in uncharted territory. For this reason, we began by developing an online prototype. Our game challenged users to perform tasks on a simulated smartwatch interface. Players navigated their way through several levels on a map. En route, they encountered thief characters who set them security challenges. If users updated their settings correctly, they won extra points. This design highlighted issues, enabled practice and incentivised protection.

To assess and refine the prototype, we recruited 504 smartwatch owners. At study commencement, all users indicated their degree of settings protection. One half, our treatment group, played the game one week later. The control half did not, as to mitigate confounding factors. Another week after gameplay, all participants

indicated their final behaviour. By comparing the pre-test and post-test results, we explored our game's potential. In the treatment group, individuals increased their usage of passwords and permissions. Control-group behaviour failed to change, and remained unprotective. This suggested an opportunity for smartwatch security games.

While these results were encouraging, self-reporting is prone to bias. Therefore, we chose to explore behaviour change empirically. To achieve this, we implemented our app on an Android smartwatch. This game also included thieves and security challenges. However, these tasks were customised based on each user's apps. Gameplay also updated in response to real-life behaviour, providing a positive feedback loop.

For a comprehensive evaluation, we conducted an eight-week longitudinal study. At the start, we gave 10 students an Android smartwatch. In the 18-day pretest period, we recorded their baseline use of settings. This was logged through an ethically-approved tracking app. In the gameplay phase, five treatment participants received the security game. To mitigate bias, the control group were given a similar app with generic challenges. Gameplay was followed by an 18-day post-test period, enabling analysis of behaviour change.

Prior to gameplay, all participants failed to protect their data. Indeed, no users disabled GPS tracking or restricted their permissions. However, the groups greatly differed after gameplay. Treatment participants increased password usage and decreased GPS usage. This suggests that games can indeed encourage smartwatch protection. In contrast, control group behaviour failed to change. Throughout the eight-week period, 80% of control participants did not restrict a single setting. Therefore, it appears that users will take little action unless informed and incentivised.

Awareness is not enough; individuals must receive education and encouragement. We suggest the use of serious games as a complement to security awareness campaigns. Read more: <https://bit.ly/2sgeXpx>

