InspiredResearch

Winter 2021 Issue 19

News from the Department of Computer Science, University of Oxford

Tracking Sea Turtles via Satellite

Read more – p20

Peering into the Moon's shadowed regions with AI – p22

All about EVE: AI model shows capacity to interpret human gene variants as benign or disease-causing – Page 15

AI to detect floods launches into space – p24

University of Oxford

Department of Computer Science

Inspired Research

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

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

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

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

Contributors

Robert Ashley
Jonas Beuchert
Ipshita Chatterjee
Emma Dunlop
Leslie Ann Goldberg
Amanda Matthes
Oliver Moore
Ekaterina Pesheva
Oliver Sampson
Carolyn Ten Holter
Ian Watts
Jun Zhao

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

CONTENTS

Leaving a legacy ...

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

If you have already left a gift in your will, please let us know so that we have an opportunity to thank you properly.

If you would like to know more about leaving a legacy to the Department of Computer Science, please contact Caroline Reynolds at: caroline.reynolds@devoff.ox.ac.uk

New Head for the Department of Computer Science

We are delighted to welcome Professor Leslie Ann Goldberg, our new Head of Department. Leslie is the Head of the Algorithms and Complexity Theory Research Theme and is a Senior Research Fellow at St Edmund Hall. She is the first female Head of Department.

Leslie takes over from Professor Michael Wooldridge, who has been Head of Department for six of the past seven years. Michael will remain with the Department. On his departure, Michael commented, 'I am proud to have been head of such a vibrant and active department at such an extraordinary time for our discipline. I think Oxford is the most exciting place to be a computer scientist in Europe right now, and I am pleased to be leaving the department in such rude good health. Leslie has a raft of exciting new ideas for advancing the department. The next seven years are shaping up to be every bit as exciting as the past seven.'

Leslie said, 'Oxford's Computer Science Department has gone from strength to strength under Michael's leadership. I'm delighted to work in such a wonderful research and teaching environment, and am excited about the future of the department.'

Letter from the Head of Department

Welcome to the Winter 2021 Issue of *Inspired Research*. Lots has happened in the department since the summer edition. I'd like to first highlight some of the many developments in AI. Eight members of our faculty have been chosen as Fellows of the Alan Turing Institute for the 2021/2022 academic year. Professor Michael Wooldridge has been awarded the Turing AI World-Leading Researcher Fellowship which he will use to lead a new research project on robust agent-based modelling at scale. Professor Yarin Gal was selected as one of five Samsung AI Researchers of the Year.

This issue of *Inspired Research* contains two different stories connecting departmental research, AI and space! One of these describes AI technology that was recently deployed into space to map flood events on Earth. The other describes research using deep learning to study shadowed regions on the moon. Professor Alex Rogers received a Prominent Paper award from AI Journal for his paper 'Efficient crowdsourcing of unknown experts using bounded multi-armed bandits'. The citation praised the paper for developing the first comprehensive framework for the rigorous analysis of task allocation in crowdsourcing systems.

In addition to these successes in AI, the department has produced exciting research in many other areas. Professor Ian Horrocks has received a Skolem Test of Time award for his paper 'Reasoning with Individuals for the Description Logic SHIQ'. The citation commended the paper for giving a practical and efficient satisfiability decision procedure for a very widely-applicable description logic. Professor Marta Kwiatkowska was a joint winner of this year's Van Wijngaarden Award for her research in preventing software faults. You can also read in this edition about applications of research led by Professor Alex Rogers in our Cyber Physical Systems group, developing a solution that is being used to track endangered sea turtles.

At the beginning of this academic year, we have welcomed three new faculty members: Ana Namburete whose research focuses on machine/deep learning with applications to human brain imaging; Bartek Klin, whose research focuses on rigorous methods for attributing meaning to programs and to constructs in programming languages; and Seth Flaxman whose research is in the area of Bayesian machine learning for public policy and social science. As part of our expansion in Computer Science, we are currently advertising six faculty positions including the Christopher Strachey Chair in Computer Science, and faculty positions in algorithms and complexity theory, automated verification, cyber physical systems, database systems, and systems. It is going to be an exciting year!

Professor Leslie Ann Goldberg
Head of the Department of Computer Science
December 2021

NEWS

News in brief

Congratulations to 4th year Computer Science & Philosophy Hertford College student Dapeng Gao on receiving a CadenceLIVE Master Thesis award for his work on the formal verification of a RISC-V processor with capability hardware extensions.

Human genes related to diseases can have numerous protein variants of yet unknown consequences. Usually, machine learning methods used to model predictions need expert labels, which are sparsely available & therefore not robust. Professor Yarin Gal, Oxford's OATML (Oxford Applied and Theoretical Machine Learning Group) and colleagues at Harvard and MIT achieve reliable performance without the use of labels by using generative models. Read the paper: go.nature.com/3DvBWMm

We are delighted to announce that Ana Namburete has joined Pembroke College's Governing Body as their first Rokos Fellow and Tutor in Computer Science. Ana is an internationally known expert in the use of machine learning to collate and analyse neurosonographic data. She will lead the introduction of Computer Science to the college's undergraduate course offering from 2022. Jesus College welcomed Seth Flaxman, and University College also welcomed Bartek Klin as additional Computer Science Tutors this academic year.

Professor Marta Kwiatkowska is part of a collaborative research team on the newly announced Project FAIR: Framework for responsible adoption of artificial intelligence in the financial services industry. Read more about it here bit.ly/3fytfHR

Collegiate Programming Contest's 2021 Northwest Europe regional heats

Teams from the University of Oxford showed off their skills in a programming teamwork challenge on the weekend of 20-21 November, held online, and hosted by Reykjavík University in Iceland. NWERC 2021 is an official regional contest in the International Collegiate Programming Contest. It draws students from colleges and universities throughout Belgium, Germany, Denmark, Estonia, Finland, Great Britain, Ireland, Iceland, Lithuania, the Netherlands, Norway, Sweden and Luxembourg. 134 teams took part this year. After a close finish, a team from the University of Oxford took the second place and won a gold medal.

The team is called "nwerc is bad" and consists of the following students: Radostin Chonev (University College); Viktor Kozhuharov (Magdalen College); and Harris Leung (Merton College). They will represent Oxford, and the entire Northwest Europe, in the ICPC World Finals.

Another Oxford team, called "Almost Retired", came 5th and won a silver medal: Xingjian Bai (St. John's College); Yeyuan Chen (Lady Margaret Hall); and Yichen Huang (St. Catherine's College).

In the contest, teams of up to three students had to try and solve as many programming problems as possible from a given problem set. The goal of each participating team is to solve as many problems as possible within the five-hour time limit.

The full scoreboard can be found here: https://www.nwerc.eu

Linde Hesse made Microsoft Research PhD Fellow

Doctoral student Linde Hesse has been selected as a Microsoft Research PhD Fellow. The Fellowship is a global programme that identifies and empowers the next generation of exceptional computing research talent. Microsoft recognises the value of diversity in computing and aims to increase the pipeline of talent receiving advanced degrees in computing-related fields to build a stronger and inclusive computing-related research community.

Linde is a doctoral student in the Department of Engineering but now does her research in the Department of Computer Science, where her main supervisor Ana Namburete recently moved.

Doctoral Training Centre Symposium

On Thursday 30 September, the members of three Oxford University CDTs (Centre for Doctoral Training) came together to talk about their current research at a student symposium, held at Trinity College, Oxford. The three CDTs were: Autonomous Intelligent Machines & Systems (AIMS), Cyber Security, and Health Data Science (HDS).

Students from all CDTs spoke on topics which included 'Hot takes on Bayesian modelling of COVID-19', 'The HAPPY Project: Histology Analysis Pipeline.PY', 'Self-Supervised Multi-Modal Alignment for Whole Body Medical Imaging', 'Understanding Sleep's Association with Cardiovascular Diseases Using Machine Learning for Wearables', and 'Dead Man's Switch: Forensic Autopsy of the Nintendo Switch'.

AIMS is based within the Engineering Department and Cyber Security and Health Data Science are both hosted by Computer Science.

Infectious disease: key surveys overestimate COVID-19 vaccination rates in the USA

Estimates of COVID-19 vaccine uptake in the USA based on large surveys that are used to guide policy-making decisions tend to overestimate the number of vaccinated individuals, research published in Nature suggests.

The authors of a new paper published in Nature find that in May 2021, Delphi–Facebook's COVID-19 symptom tracker (250,000 responses per week) overestimated vaccine uptake by 17 percentage points, and a survey from the US Census Bureau (75,000 responses per wave) overestimated vaccine uptake by 14 percentage points compared to benchmark estimates from the CDC (Centers for Disease Control and Prevention). These overestimation errors go orders of magnitude beyond the statistical uncertainty provided by the surveys. A survey by Axios–Ipsos also overestimated uptake, but by a smaller amount (4.2 percentage points in May 2021)—despite being the smallest survey (about 1,000 responses per week). These findings indicate that bigger is not always better when it comes to datasets, if we fail to account for data quality.

The authors note that COVID-19 vaccine uptake was not the primary focus of any of the surveys. For example, Delphi-Facebook intends to measure changes in COVID-related behaviours over time. However, the bias in estimates of vaccine uptake in the two large surveys indicates that they are not representative of the US adult population. They suggest lack of statistical representativeness may also be causing bias in other survey outcomes.

The research presented in the paper suggests that design choices in survey data collection can lead to inaccuracies that are not overcome by sample quantity; for example, the three surveys use different recruitment methods, which may introduce different biases in their estimates. The authors conclude that efforts to measure data quality and improve the accuracy of assessments of vaccine uptake are needed to better inform public policy decisions.

Paper co-author Seth Flaxman (University of Oxford Department of Computer Science) comments, 'Human behaviour is key to COVID-19 pandemic response. Who is wearing masks? Have they got the message that COVID is airborne? For those who haven't been vaccinated, what reasons do they give – hesitancy, access, or lack of trustworthy information? Surveys are our primary source of information to answer these questions. But when data is collected online, especially through mobile apps, we should always ask: how do we know whether this data is representative of the population? Trump voters were famously less likely to respond to opinion polls before the 2016 election – do we see something similar with those who are vaccine hesitant? There is an urban/rural divide in terms of Internet access – and also vaccine access. Who is being left out of vaccine outreach campaigns guided by online surveys?

'Our study is the first to use a widely available benchmark – how many people have received the COVID-19 vaccine? – to evaluate the accuracy of online surveys. We find that in the US, the two largest surveys, from Facebook and the Census Bureau, significantly overestimated vaccine uptake despite enormous sample sizes. By contrast, a smaller survey from Ipsos was very accurate. Bigger isn't always better.

'Online surveys supply a plethora of data on everything from COVID-19 symptoms to in-person schooling to social distancing behaviour. We're trying to raise the alarm that data quality can be a huge issue, and even random samples – the gold standard for surveys – can still have a lot of bias because of who is willing or able to respond.'

Read the full article here: bit.ly/3yqoBUz

News in brief

We are thrilled that the *Complete University Guide* has ranked us #1 for Computer Science in their 2022 league tables.

Professor Michael Wooldridge gave evidence as the Justice and Home Affairs Committee held its second oral evidence session on new technologies for law enforcement. He was also one of the experts consulted by members of the House of Lords who are looking at the role technologies such as AI and facial recognition have in modern policing methods.

Professor Sadie Creese was featured on BBC Radio programme The Briefing Room, talking about the threat of ransomware. bbc.in/3paXXu6

Professor Marta Kwiatkowska joint winner of Van Wijngaarden Awards

The Van Wijngaarden Award 2021 has been awarded to computer scientist Professor Marta Kwiatkowska and mathematician Susan A. Murphy for the numerous and highly significant contributions they made to their respective research areas: preventing software faults and improving decision making in health. The five-yearly award is established by CWI, the national research institute for mathematics and computer science in the Netherlands, and is named after former CWI director Adriaan Van Wijngaarden. The winners will receive the prize during a festive soirée on 3 November 2022 in Amsterdam.

NEWS

Professor Ian Horrocks wins Skolem Test of Time award

Professor Ian Horrocks has won the CADE (Conference on Automated Deduction) Skolem Award for 'a CADE paper that has Passed the Test of Time, by being a Most Influential Paper in the field'. The award was in respect of the paper: Ian Horrocks, Ulrike Sattler, and Stephan Tobies. Reasoning with Individuals for the Description Logic SHIQ.

In 2014 CADE Inc. established the Thoralf Skolem Award to reward a CADE paper that has passed the test of time, by being a most influential paper in the field. Beginning with CADE-25 (2015), at every CADE the Skolem Award Committee selects a paper presented at the CADE held 10 years earlier to receive the Skolem award.

The paper was originally published at CADE-17 (2000) and the award was presented at the latest iteration of the conference, CADE-28, which took place in July 2021.

Meltwater acquires AI spinout to help build one of the largest knowledge graphs of public information

Meltwater B.V., a leading global SaaS ('Software as a service') provider of media intelligence and social analytics, has acquired AI start-up DeepReason.ai, a spinout from Oxford University's Department of Computer Science.

DeepReason.ai was founded in 2018 by Professor Georg Gottlob [right], Emanuel Sallinger and Stéphane Reissfelder and is one of the leading knowledge graph reasoning providers. Knowledge graphs are AI systems that connect entities such as companies, people, products and topics to discover new connections and insights. DeepReason's system can update and maintain complex knowledge graphs at scale and in real-time. The basic research underlying this system was led by Georg Gottlob, Professor at the Department of Computer Science and fellow of the Royal Society, and carried out at Oxford in the context of project 'VADA: Value Added Data Systems – Principles and Architecture' funded by the EPSRC research council.

Meltwater processes over 800 million documents every day and is each day growing its knowledge graph (consisting of 14 million companies, 50 million people, and 75 million topics) by approximately 2 billion connections to the conversations around those companies, people, and topics. With the DeepReason.ai acquisition, Meltwater will be in a unique position to continue expanding this knowledge graph to discover even more connections and insights.

'In the last two decades, Meltwater's proprietary search engine has been powering the insights and analytics offered to Meltwater's clients. Going forward, Meltwater's reasoning engine will unlock insights and customer value not possible to derive from traditional search engines capabilities,' says Meltwater CTO Aditya Jami.

Professor Leslie Ann Goldberg, the Head of Oxford University's Department of Computer Science, said `The department is delighted by the deep and continuing impact of the research of Professor Gottlob's team and we are very happy about the success of this spinout from our department.'

VADA project:
vada.org.uk/project-description/

Ján Pich awarded a Royal Society University Research Fellowship

Understanding the power of computation is one of the biggest challenges in science and society. Ján Pich's newly awarded Royal Society Fellowship will ask if we can automate creative and challenging tasks such as proving mathematical theorems or designing learning algorithms. Such questions, can be formalised in the language of computational complexity theory, and constitute some of the most fundamental scientific problems of our times. A famous obstacle in the centre of these pursuits, known as the P versus NP problem, asks whether it is possible to solve efficiently all problems whose solutions can be efficiently verified. A positive answer to the question would single-handedly resolve a myriad of major problems including theorem-proving and algorithm-design.

Unfortunately, the P versus NP problem has resisted solution for 50 years. In fact, it is consistent with our current knowledge that almost all problems in practice are solvable extremely efficiently with very small computational devices. The decades of research have led to two prominent but contrasting perspectives on the problem: the meta-mathematical view manifested in the theory of complexity of proofs (proof complexity) and a concrete combinatorial view dominant in the analysis of computational models such as logical circuits (circuit complexity).

Ján's fellowship will focus on bridging these approaches. The project will build on the emerging theory of hardness magnification, a counter-intuitive phenomenon which avoids previous formal barriers against the solution of the P versus NP problem, and reveal new connections between logic, learning theory and cryptography. It thus has the potential to bring us closer to a full understanding of the power of computation and logical reasoning.

Oxford Computer Science academics recognised for papers on Artificial Intelligence

Professor Edith Elkind received a distinguished paper award at the 30th International Joint Conference on Artificial Intelligence (IJCAI-21). The paper, titled 'Keep your distance: Land division with separation', was written with Erel Segal-Halevi (Ariel University) and Warut Suksompong (National University of Singapore). Just 3 distinguished paper awards were awarded, out of 587 accepted conference papers, and 4204 initial submissions.

Professor Alex Rogers received the 2021 Artificial Intelligence Journal (AIJ) Prominent Paper Award. The paper, titled 'Efficient crowdsourcing of unknown experts using bounded multi-armed bandits', was originally published in 2014 and was written with Long Tran-Thanh (Warwick), Sebastian Stein (Southampton) and Nick Jennings (Loughborough).

The AIJ Prominent Paper Award recognises outstanding papers published in the journal in the last seven years that are exceptional in their significance and impact. The paper has been identified by the awarding body as 'having developed the first comprehensive framework for the rigorous and principled mathematical analysis of task allocation algorithms in crowdsourcing systems' as well as for its work on 'the proposed bounded bandits, a new sequential decision-making model to solve task allocation problems with resource constraints'.

The committee commented that the award has been given as the work has had a 'significant impact on subsequent work carried out in both industry and academia'.

Read more: ijcai-21.org

Professor Edith Elkind made EurAI Fellow

Congratulations to Professor Edith Elkind, who has been made a Fellow by the European Association for Artificial Intelligence (EurAI). The EurAI Fellows Program honours only a very small percentage of the total membership of all EurAI member societies, and Edith will join four other members of our department (Professors Thomas Lukasiewicz, Ian Horrocks, Nigel Shadbolt and Michael Wooldridge) who are already Fellows.

Tim Rudner one of four winners of the Qualcomm Innovation Fellowship (Europe)

Qualcomm Technologies, Inc. has announced that Tim Rudner is one of four winners of the 12th edition of the Qualcomm Innovation Fellowship (QIF) Europe program. QIF is an annual program that focuses on recognising, rewarding, and mentoring the most innovative engineering PhD students across Europe, India, and the United States.

The Europe program rewards excellent young researchers in the fields of Artificial Intelligence and cyber security with individual prizes of $40,000, dedicated mentors from the Qualcomm Technologies team as well as the opportunity to present their work in person to an audience of technical leaders at the company's HQ in San Diego.

Tim has been selected for his proposal: 'A Fully Probabilistic Theory of Autonomous Decision Making'. Tim's proposal is about developing a fully probabilistic framework for reinforcement learning to provide reliable and mathematically rigorous uncertainty quantification. In contrast to previous approaches, he proposes to treat both the learning process as well as the model components, such an agent's policy, probabilistically. The approach will combine advances in probabilistic inference and modelling with probabilistic reinforcement learning.

EPSRC Early Career Fellowship awarded to Andrew Cropper

Andrew Cropper has been awarded an Early Career Fellowship for his research project, The Automatic Computer Scientist. He explains, 'Algorithms are ubiquitous: they track our sleep, find us cheap flights, and even help us see black holes. However, designing novel algorithms is extremely difficult, and we do not have efficient algorithms for many fundamental problems. This project aims to accelerate algorithm discovery by building an automatic computer scientist (AutoCS). The idea of developing machines that automatically write computer programs is a long-standing grand challenge in artificial intelligence (AI) and offers the potential to automatically build bug-free and efficient programs.

'To work towards this grand challenge, this project will build on major recent breakthroughs in inductive logic programming (ILP), a form of symbolic machine learning based on mathematical logic. ILP currently has the ability of a first-year undergraduate Computer Science student: given much guidance, it can learn simple algorithms and implement small programs.

'This project aims to significantly advance ILP to the level of a Computer Science PhD student so that given little guidance it can discover novel and complex algorithms and implement large programs. As a marker of success, a key objective of this project is to use an AutoCS to discover a novel algorithm and publish it in a Computer Science journal. Such a result would be a landmark achievement for AI and would herald a new era of automatic scientific discovery.'

Eight Computer Scientists named Turing Fellows

The Turing Institute has announced its latest cohort of Turing Fellows, and they include researchers from 14 Oxford departments spanning the Mathematical and Physical Life Sciences (MPLS), Social Sciences and Medical Sciences divisions. Eight of the new fellows are Department of Computer Science academics.

Professor Sam Howison, Head of MPLS Division, commented:

'I am delighted to see the work of so many Oxford researchers recognised in this way. Their wide-ranging expertise illustrates not only the breadth of outstanding data science and AI research in Oxford, but also the critical importance of data science to so many aspects of the modern world.

The full list of fellows from our department is as follows:

– Alessandro Abate, Professor of Verification and Control

– Michael Benedikt, Professor of Computing Science

– Ani Calinescu, Associate Professor

– Georg Gottlob, Royal Society Research Professor and Professor of Informatics

– Ian Horrocks, Professor of Computer Science

– Varun Kanade, Associate Professor

– Marta Kwiatkowska, Professor of Computing Systems

– Thomas Lukasiewicz, Professor of Computer Science, AXA Chair

Professor Michael Wooldridge named as one of five Turing AI World-Leading Researcher Fellows

Professor Michael Wooldridge is one of five internationally-recognised researchers who have been appointed as the first Turing AI World-Leading Researcher Fellows to conduct ground-breaking work on Artificial Intelligence's (AI) biggest challenges.

The new fellows are:

– Professor Zoubin Ghahramani, University of Cambridge

– Professor Samuel Kaski, The University of Manchester

– Professor Mirella Lapata, University of Edinburgh

– Professor Philip Torr, University of Oxford

– Professor Michael Wooldridge, University of Oxford

The fellowships, named after AI pioneer Alan Turing, are part of the UK's commitment to further strengthen its position as a global leader in the field. Retaining and attracting some of the best international research talent in a highly competitive international environment will increase the UK's competitive advantage and capability in AI. The fellows' research will have a transformative effect on the international AI research and innovation landscape by tackling some of the fundamental challenges in the field.

Working with industrial partners Accenture Global Solutions, Facebook, JP Morgan, Oxford Asset Management, Schlumberger, and Vodafone, Michael aims to improve the agent-based AI models that are increasingly used in sectors such as financial modelling and logistics.

Michael said: 'Agent-based modelling is an old idea, but one that has really started to show promise recently. The fellowship will give me a unique opportunity to advance the science and technology of agent-based modelling along a broad range of directions.

'I will investigate languages that allow us to transparently capture agent-based models and the complex assumptions that they embody, so that these models can be more easily developed and understood.

'I'll investigate how models can be populated with realistic agent behaviours, which is essential if the models are to be of any value. Finally, I'll investigate how models can be calibrated and validated, so that we can have confidence in the predictions they make.'

The fellows are supported with a £18 million investment by UK Research and Innovation (UKRI). In addition to this, 39 different collaborators including IBM, AstraZeneca and Facebook are making contributions worth £15.7 million to the fellows' research programmes. The fellowships are being delivered by UKRI's Engineering and Physical Sciences Research Council.

Professor Michael Bronstein appointed to DeepMind Professorship of Artificial Intelligence

We are very happy to announce the appointment of Professor Michael Bronstein to the DeepMind Professorship of Artificial Intelligence. He will join our department with effect from 1 January 2022. He will also be a Fellow of Exeter College.

Michael joins from Imperial College London, where he has held the Chair in Machine Learning and Pattern Recognition for the past three years while serving as the Head of Graph Learning Research at Twitter. Previously, Michael had visiting appointments at MIT, Stanford, and Harvard universities and was a Fellow at the Harvard Radcliffe Institute for Advanced Study.

Michael is a world leader in geometrically-inspired machine learning, a field that has recently seen a broad adoption ranging from recommender systems to drug design. He is a Member of Academia Europaea, Fellow of IEEE, and a recipient of multiple ERC grants, Google Faculty Research awards, Amazon ML Research awards, Facebook Computational Social Science award, Dalle Molle prize, Royal Society Wolfson Merit award, and Royal Academy of Engineering Silver Medal.

Michael says, 'I am honoured and delighted to join Oxford, which has unparalleled expertise in AI and related fields and amazing students. I am looking forward to forging new collaborations and synergies within the Department and beyond that would allow us to develop the next generation of machine learning methods that solve real-world problems and at the same time have the trust of domain experts and the broader public.'

This chair is generously supported by DeepMind, the world-leading British AI company.

Obum Ekeke, Global Lead of University Relations and Education Partnerships, DeepMind, said 'Congratulations to Michael on his appointment as the DeepMind Professor in Artificial Intelligence at the University of Oxford. Michael is an excellent researcher with deep experience working across academia and industry simultaneously, and we're delighted to support him in this new position. I have no doubt that Michael will contribute further to Oxford's distinguished work in the field, and add enormously to the development of the next generation of AI researchers.'

Professor Leslie Ann Goldberg, Head of Department of Computer Science comments, 'We are delighted by this appointment. Michael is a world-leading researcher whose research has already had a huge impact across the increasingly important area of machine learning. His research is important both theoretically and because of its many applications. Our students will be very excited by the opportunity to learn from Michael. We are extremely grateful to DeepMind for supporting a Chair in this very significant area.

You can find out more about Michael here: imperial.ac.uk/people/m.bronstein

Professor Yarin Gal one of five Samsung AI Researchers of the Year

Professor Yarin Gal has been announced as one of five 'Samsung AI Researcher of the Year' award winners. The awards were launched last year to discover rising AI researchers globally.

The awards were announced on the first day of the Samsung AI Forum, an event which gathers leading experts to discuss the latest technology trends and research findings. The two-day event held on 1 and 2 November enabled participants to discuss applications of AI that will make a practical contribution to people's daily lives.

This year's awards went to Yarin, and also Professor Diyi Yang (Georgia Tech), Professor Jacob Andreas (MIT), Professor Judy Hoffman (Georgia Tech) and Professor Phillip Isola (MIT).

Yarin commented 'I am honoured to receive this award, and I want to thank my students and collaborators – this award is really for them'.

MPLS Impact Awards 2021: winners and commendations announced

Each year the University of Oxford's Oxford University's Mathematical, Physical and Life Sciences division (MPLS) runs a competition which recognises the impact of research undertaken by MPLS researchers at all career stages – from DPhil students to senior academics. Nominations are made across four impact categories: commercial, social, early career and public engagement with research. The Department of Computer Science has done exceptionally well this year.

Professor Kasper Rasmussen won the Commercial Impact category, for his research on 'Resolution of Multiple Critical Design Flaws in Bluetooth Standard'. Kasper lead research on weaknesses in wireless protocols which uncovered critical flaws in multiple parts of the Bluetooth standard (as implemented in billions of devices worldwide), demonstrating how both the Bluetooth session key establishment and the authentication procedures can be completely compromised. The research team coordinated with key industry bodies to disclose each vulnerability, allowing them to be remedied before they could be abused. The work led to changes to the Bluetooth Core Specification, and to mitigations applied by major manufacturers.

Professor Michael Wooldridge, then Head of the Department of Computer Science, said: 'Kasper and his colleagues discovered a jaw dropping bug in the Bluetooth protocol, used by hundreds of millions of people every day. Their discovery necessitated a change to the protocol, and firmware updates for billions of devices. It is wonderful work, and one of the clearest and most direct examples of impact I have seen in my department during my seven years as head.'

Members of the department were also awarded commendations in the Commercial Impact category as follows:

Professor Georg Gottlob: Efficient Web Data Extraction and Knowledge Processing via Datalog

Georg's research at Oxford has led to fundamental advances in efficient reasoning languages and their application to web data extraction and management. Systems developed as part of this research led to two spinout companies, Wrapidity and DeepReason.ai.

Professors Bernardo Cuenca Grau, Ian Horrocks, Boris Motik: Enabling Applications of Ontologies in Healthcare and in Industry via Reasoning Systems

The researchers have developed state-of-the-art reasoning systems that represent important advances in exploiting the potential of semantic technologies for complex data and knowledge applications. Their open source reasoning tools are enabling applications of ontologies in areas as diverse as global healthcare IT and large-scale infrastructure design. They have also developed and commercialised RDFox, a high-performance knowledge graph and semantic reasoning engine, through the spinout company Oxford Semantic Technologies (OST).

Professor Niki Trigoni: Improving Workforce Efficiency in Hospitals via Infrastructure-Free Indoor Localisation

Niki's research group developed a frictionless infrastructure-free indoor positioning solution based on smartphones, which avoids the significant cost and effort required to deploy existing indoor location-tracking solutions. The patented technology underpins the spinout Navenio, founded by Niki in 2015. Navenio has applied the technology in multiple NHS trusts to build a workforce tracking and tasking solution.

Find out more about the MPLS Impact Awards here: bit.ly/3lKNqpd

Science Together: Oxford Researchers and Communities

The Mathematical, Physical and Life Sciences Division and Medical Sciences Division at Oxford University have launched a new Public Engagement with Research (PER) pilot programme to connect researchers and local communities, so that our research is grounded in local needs, interests and priorities.

Science Together: Oxford Researchers and Communities explores what challenges can be overcome or opportunities seized by local community groups when they are given access to the world-class skills, knowledge and resources of our researchers. There are eight wonderfully diverse range of projects taking place, some examples include:

- Daybreak – exploring how technology can help people with dementia live independent lives for longer.
- Barton Community Association - running an after-school science club to help young people catch up on learning missed during lockdown
- Urban Music Foundation – exploring the interplay between science, technology and Hip Hop
- Oxford Neighbourhood Watch - conducting a data mapping exercise to understand the scale of under-reported bike crime in Oxford

Researchers from the Department of Computer Science will be working with Leys Community Development Initiative (CDI) to co-produce an app with young people. The partnership will provide hands-on opportunities for young people to build their skills and experience as well as streamlining the charity's operations by making session registration and promoting activities etc much easier.

Claudine Tinsman, Thomas Serban Von Davier, Max Van Kleek and Jun Zhao kick-started the project in October 2021 with an interactive and inspiring session at Leys CDI giving young people a taster of what is involved in designing and producing an app.

A tribute to Samson Abramsky on his retirement

Samson Abramsky, head of the Quantum Computing research theme, Fellow of Wolfson College and Christopher Strachey Professor of Computing has retired having been a hugely important member of our department for over 20 years.

Throughout his career Samson has made pioneering contributions to a wide range of subjects. He has played a leading role in the development of game semantics, and its applications to the semantics of programming languages.

Other notable contributions include his work on domain theory in logical form, the lazy lambda calculus, strictness analysis, concurrency theory, interaction categories, and geometry of interaction.

More recently, he has been working on high-level methods for quantum computation and information. He introduced categorical quantum mechanics with (former member of the Department) Professor Bob Coecke. He introduced the sheaf-theoretic approach to contextuality and non-locality with Adam Brandenburger, and has contributed extensively to developing a structural theory of contextuality and its applications.

More recently, Samson has been interested in connections between Computer Science and other scientific disciplines. He says 'I believe that the distinctive methods of Computer Science, above all compositional semantics and logic, have much to offer across a broad sweep of the physical and biological sciences, and to the modelling of complex systems.'

Other achievements include:

He is a Fellow of the Royal Society (2004), a Fellow of the Royal Society of Edinburgh (2000), a Member of Academia Europaea (1993), and a Fellow of the ACM (2014).

His paper 'Domain theory in Logical Form' won the LICS Test-of-Time award (a 20-year retrospective) for 1987. The award was presented at LICS 2007.

His paper 'A Fully Abstract Game Semantics for General References' with Kohei Honda and Guy McCusker won the LICS Test-of-Time award (a 20-year retrospective) for 1998. The award was presented at LICS 2018.

He was the Clifford Lecturer at Tulane University in 2008.

He was awarded the BCS Lovelace Medal in 2013.

He received the Alonzo Church Award for Outstanding Contributions to Logic and Computation in 2017.

DeepMind Scholarships awarded:

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 scholarship programme for under-represented groups.

Four more DeepMind graduate scholarships for students wishing to pursue a master's degree in Computer Science in the Department of Computer Science have been awarded to students commencing courses in 2021–22. These were open to individuals ordinarily resident in the UK who belong to one or more of the following groups:

- identifying as female and/or
- of Black or Minority Ethnic origin and/or
- from a low-participation area of the UK identified by post code at date of taking A Levels and/or equivalent.
- One of the scholarships was open to students who are liable to pay overseas fees.

Introducing our Google DeepMind scholars:

Lili Janzer (St John's)

'I am from Hungary, I have a BA in Computer Science from the University of Cambridge, and currently, I am studying towards an MSc in Advanced Computer Science. Receiving the DeepMind scholarship allowed me to pursue a master's degree immediately after graduation without having to worry about finances. It is also a huge honour and I am looking forward to the events and mentoring that accompany this scholarship.'

Hashan Punchihewa (St Catz)

'I grew up in Essex, and I recently completed my undergraduate degree in Computing at Imperial College London. During this time, I benefitted from internships in industry and research. I also worked as an undergraduate teaching assistant. I'm grateful to have received a DeepMind scholarship, which has allowed me to pursue the MSc in Advanced Computer Science at Oxford. This is an incredible opportunity for me to develop advanced knowledge and versatile skills, that will be invaluable in my career and help me achieve my future aspirations.'

Maëlys Solal (Somerville)

'I am a post-graduate student at the University of Oxford, studying in the MSc in Advanced Computer Science. I come from France, where I studied both Mathematics and Computer Science during my undergraduate degree. This Google DeepMind scholarship opened the doors of this prestigious university to me and provides a mentorship program which will help me in the future.'

Maria Stoica (St Edmunds Hall)

'I am the daughter of Romanian emigrants to the USA. I grew up in Kansas and I graduated from Harvard with a bachelor's degree in Computer Science in 2017. Since then I have been working in finance in New York City and London at Goldman Sachs and NatWest Markets. Having the support of DeepMind has not only increased my confidence in taking up my course at Oxford but will allow me to fully focus on my studies without the burden of supporting myself financially.'

Four more DeepMind graduate scholarships for students wishing to pursue a master's degree in Computer Science in the Department of Computer Science will be made available for students commencing courses in 2022–23. Read more here: bit.ly/3IHFhf9

Alumni Profile

Ipshita Chatterjee – studied the MSc Computer Science at the Department of Computer Science in 2020 - 2021.

What was your background before joining Oxford University?
I am from New Delhi, India. Before coming to Oxford, I completed my Bachelor of Engineer in Computer Engineering from the University of Delhi, and worked for two years as a software engineer at Adobe.

What attracted you to studying Computer Science as a subject?
I have been deeply fascinated by Computer Science, ever since I took up HTML as an elective in school, at age 13. I was very captivated by how seemingly simple lines of code could create beautiful web pages. To me, studying Computer Science is empowering. It gives me the tools to create powerful solutions to everyday problems and catapults me to the forefront of creating the future through technology.

What aspects of the course you studied here did you particularly enjoy?
What attracted me most to Oxford Computer Science was the first principles approach to teaching the subject. For such a rapidly changing industry, I wanted to strengthen my skills for the future and equip myself to solve problems which may not even have presented themselves today. The focus of the course on solving problems from first principles, and not just the latest technologies, was quite commendable. Additionally, I was able take up courses presenting the bigger picture - the impact of Computer Science on society, and application areas such as biology, medicine and genomics.

I was involved with Oxford Women in Computer Science (OxWoCS) as an Outreach Officer, alongside an amazing committee. I am very proud of our work championing the voices of women in Computer Science, with Female Faces of Computer Science. And the hands-on impact we created with Challenge Club, an online computer science outreach programme to support students in growing their maths skills. I hope we were able to inspire (even a tiny bit) the next generation of women to take up STEM careers. I also represented the Department as a Student Ambassador in several outreach activities.

As part of Oxford Artificial Intelligence Society (OxAI), I worked alongside Dell Technologies to produce a whitepaper examining how AI could change the future of work, which was received very well.

Most of all, I really enjoyed the company and learning from the people I met here, with vastly different backgrounds and experiences, all incredibly kind and supportive.

What did you do after you graduated from your Oxford University course?
I have started working with Amazon in London as a Software Development Engineer with the Prime Video team, building bespoke financial systems for royalty calculation.

How has the course you studied here helped you in your current career?
While I had a Computer Science background before, this course has given me a strong understanding of technical concepts and inculcated the ability to solve ambiguous problems. The Oxford term system has taught me time management and the tutorial system has taught me the value of curiosity, asking the right questions and formulating my thoughts cogently on the fly. Networking skills I acquired at Oxford University have also helped me in my current career.

If we had interviewed you when you were a student, what might you have thought about what you are currently doing?
I think she would have been surprised, amazed AND proud of everything she has achieved during the course and beyond. She would go on to complete her MSc in a very challenging year full of uncertainties, contribute to projects and initiatives she cares about, make friends who support her, and start a job she is really excited about.

Ethical Web and Data Architecture in the Age of AI – Creating an Open and Transparent Data Common for the Digital Future

By Senior Researcher Jun Zhao

Thirty years ago, the World Wide Web was launched as an open, common, universal infrastructure for anyone with a computer and a modem. The Web was created for everyone. In recent years, however, this medium has been rapidly diverted from its initial values. Instead of being an open data platform to bring together information from the world and the voices of many, the Web is being dominated by a few platform companies, who have accumulated unprecedented concentration of power.

The consequences are twofold: information on the Web is no longer created and shared by users and citizens of different voices, but being filtered and manipulated by a few dominating platforms for profit or political gains; users are no longer simply information consumers, but being treated as commodities, as a means of value extraction, whose beliefs, interests, activities and ways of life are processed constantly and used to influence and direct their access to particular content, groups and viewpoints. Users' data are locked in company 'silos', and they lose control of what they see and how their data is being used. The opaqueness of the algorithmically-powered data analytics - the backbone of many platforms, puts extreme challenges to users' ability to make informed interpretations about the data-driven society around them, involving their education, career opportunities, or healthcare.

To address the current state of 'digital monopoly', University Oxford Computer Science researchers, led by Sir Tim Berners-Lee and Sir Nigel Shadbolt, recognised that not only a regulatory change is needed to bring the Web back to its beginning but also an innovative rethink and re-design of the current technical and legal infrastructures.

The Ethical Web and Data Architecture project (EWADA) is an ambitious three-year programme funded by the Oxford Martin School, set off in June 2021. The key ambition of EWADA is to bring back a more equitable and ethical Web for everyone, providing a digital environment that respects individuals' personal data privacy, rights and autonomy. It is led by the Oxford Computer Science Department, who have joined forces with leading scholars on data privacy and ethics (Professor Carissa Véliz), data fairness (Professor Sandra Wachter) and digital economy policies (Professor Stefan Dercon) .

The core of the EWADA mission is to explore a re-decentralised web architecture as the means for users to understand and control how their data is being used in existing and future applications. Our vision is that we will provide usable tools for end-users to manage their data and digital footprints in an informed and transparent way, and gain access to applications that are responsible and ethical. At the first step, users will take back control of data ownership: Users' data will no longer be locked in specific company platforms, and instead, they will control who can access what about them. Users will no longer need to put up with platforms that may ruthlessly collect their data without any explicit consent or make use of their data to manipulate their digital experiences. Ultimately, users will be empowered by a community of open-source developers and ethical innovators, who would develop transparent, accountable, and fair applications that are guided by best practices.

We have already seen proof-of-concepts of this web decentralisation vision. Ethical Web and Data is to be built upon an open source development of the Social Linked Data (Solid) platform (solidproject.org/), developed at MIT, which enables data to stay at rest on the users' devices, namely 'Solid pods'. Solid has been piloted with the NHS to store patients data and enable faster healthcare provisions (nytimes.com/2021/01/10/technology/tim-berners-lee-privacy-internet.html). More recently, we have seen BBC R&D's announcement of their personal data stores for viewers' audience data, to stop data lock-in and create a more transparent platform for the users bbc.in/3y3YgeK

We recognise that to bring this vision to full realisation would require not only strong technical innovations but also a comprehensive re-thinking about what a decentralised data architecture may mean for users' data rights and privacy and for platforms' data governance and accountability: Who should be responsible for data privacy and security now? What is needed to enable users' trust of access to their data? What can be an effective way to empower users to control access to their data? This new vision will also raise exciting challenges to existing approaches of privacy-preserving AI algorithms that are needed to create powerful and complex user applications on top of the decentralised solid pods: Can we ensure the scalability and performance of distributed algorithmic training at a web scale? Can we still ensure sufficient algorithmic accountability when data is processed and analysed in a distributed way, instead of centrally?

EWADA is rapidly developing its initial reference architectures in application domains like children's data protection, personal media viewing data, lifelogging and IoT (Internet of Things) data protection.

Read more about The Ethical Web and Data Architecture project here: bit.ly/3GKxJ9q
Read more about the Oxford Martin School here: https://www.oxfordmartin.ox.ac.uk

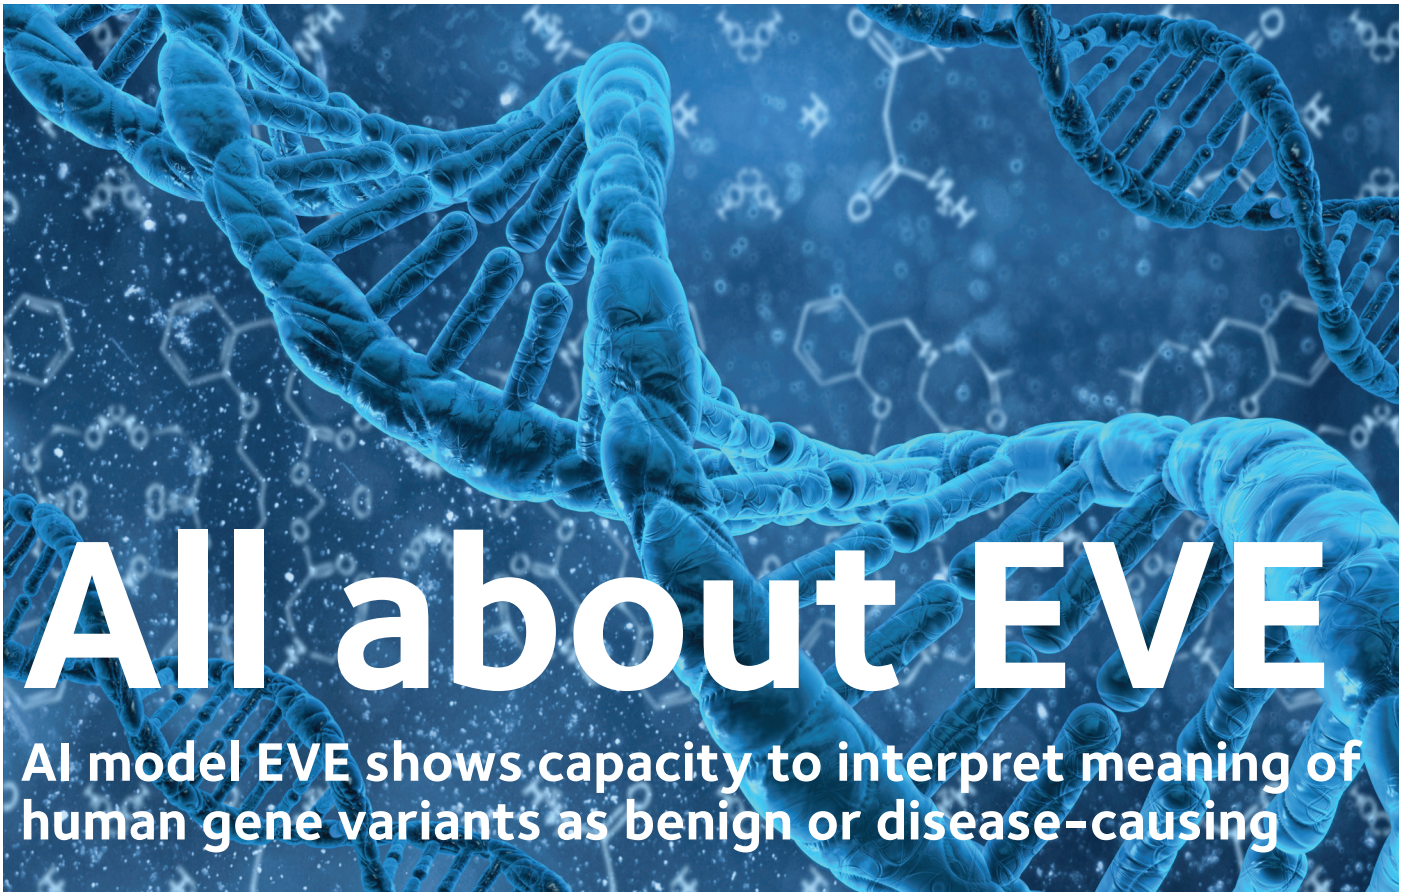

All about EVE

AI model EVE shows capacity to interpret meaning of human gene variants as benign or disease-causing

By Ekaterina Pesheva, Harvard University, our partners in this research

No two human beings are the same, a biologic singularity encoded in the unique arrangement of the molecules that make up our individual DNA.

Variation is a cardinal feature of biology, the driver of diversity, and the engine of evolution, but alterations in DNA sequences and the resulting proteins that build our cells can sometimes lead to profound disruptions in physiologic function and cause disease. But which gene alterations are normal or at least inconsequential, and which ones portend disease?

The answer is clear for a handful of well-known genetic mutations, yet despite dramatic leaps in genome sequencing technology over the past 20 years, our ability to interpret the meaning of millions of genetic variations identified through such sequencing still lags.

To make sense of it all, researchers at Harvard Medical School and Oxford University have designed an AI tool called EVE (Evolutionary model of Variant Effect), which uses machine learning to detect patterns of genetic variation across hundreds of thousands of nonhuman species and then uses them to make predictions about the meaning of variations in human genes.

In an analysis published in *Nature*, (Oct 2021) the researchers used EVE to assess 36 million protein sequences and 3,219 disease-associated genes across multiple species. The results suggest that 256,000 previously identified human gene variants currently of unknown significance should, in fact, be reclassified as either benign or disease-causing.

The tool, the researchers said, can be used to augment current clinical methods used to determine the meaning of gene variants. And, when used in combination with such tools, EVE could boost the precision and accuracy of diagnosis, prognosis, and treatment choice.

'Increasingly, people have access to sequencing their genomes, but making sense of the data is not always straightforward. There is very little information about what it even means for likelihood of disease or disease progression,' said study senior author Debora Marks, Associate Professor of Systems Biology in the Blavatnik Institute at Harvard Medical School, who co-led the research with colleague Professor Yarin Gal at Oxford University's Department of Computer Science, (and co-first authors Jonathan Frazer and Mafalda Dias at Harvard Medical School, and Pascal Notin at Oxford).

The researchers emphasise that EVE is not a diagnostic test, but its computational prowess can augment current clinical tools to make diagnoses, predict disease progression, and even choose treatment based on the presence of certain disease-causing genetic mutations.

Analysis showed that EVE outperformed other computational prediction models in predicting clinical effect and also scored as high as or better than current gold-standard high-throughput experiments that test the effect of a mutation on biologic function.

The stakes of accurately interpreting the meaning of genetic variation are enormous. Reading a benign

continued on next page ▶

from previous page ▶

variation as disease-causing could lead to erroneous diagnosis, fuelling a cascade of further testing, anxiety, and even unnecessary medical interventions. Conversely, misinterpreting a disease-fuelling change on one's DNA as inconsequential or benign could provide false reassurance when watchful observation, further testing, and preventive measures may be needed.

'What we hope this approach will do is generate powerful data that can empower the clinicians on the frontlines to make the right diagnostic, prognostic, and treatment decisions,' Yarin said.

With more data, more questions

The historic sequencing of the human genome in 2003 established a reference human genome against which newly sequenced ones are compared. Yet, this reference genome is not a standard or a baseline for a 'normal' human genome. The rapidly growing amount of data from DNA sequencing renders the reference genome less of a standard and more of a fluid baseline that shifts as researchers better grasp the meaning of genetic variation.

Relating specific changes in the human genome to disease occurrence continues to bedevil the field of clinical genetics because the number of variants in the human population dwarfs the number of scientists who can investigate. Even though only a tiny fraction of the human population has been sequenced, researchers are already seeing millions of variants whose significance and meaning are unclear. Of those variants, only 2 percent are classified as benign, neutral, or pathogenic. The remaining 98 percent of the identified gene variants are currently deemed of 'unknown significance.'

In the human genome, protein-coding regions alone account for millions of observed mutations involving the position of a single amino acid in a protein made by a gene. These so-called 'missense mutations' may have no effect on the function of a protein, or they may render the protein dysfunctional, causing disease. Every individual has many variants in their genome, compared with other people and with the reference human genome.

Adding another twist to an already complicated plot, humans inherit two versions of each gene—one from each parent. And, as people age, genes may acquire changes, known as somatic mutations.

Debora comments, 'You may have a different variant on one copy of a gene and, as we age, there are all sorts of somatic variations that occur—not only related to cancer development but to neurodegeneration, both of which are age-related processes driven by mutation.'

There are a number of disease-associated genes for which researchers have identified mutations that carry high risk of clinical disease, but even those genes have shown other unstudied mutations, the significance of which remains unclear. All of this creates an urgent need to clarify the significance of genetic variations in humans—a process in which computation is going to play an increasingly important role in providing answers.

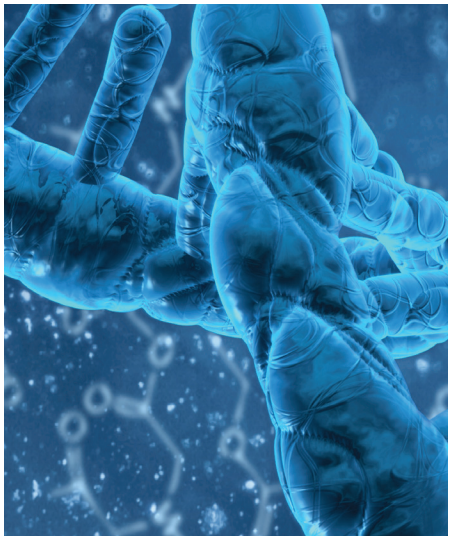

Enter AI

A defining feature of neural networks is their capacity to continually reassess and update the probability of a hypothesis as new data become available. This means that neural networks can re-evaluate evidence using new knowledge and can detect patterns and meanings missed by traditional methods.

In the current study, the researchers used a type of analysis known as unsupervised machine learning, a form of artificial intelligence that is not based on predefined parameters and rules but instead involves adaptive learning. What this means is that when presented with new data, a machine learning algorithm will become better at recognising patterns over time.

One advantage of unsupervised models is that there is no chance of biasing their learning by feeding them pre-labelled data. Also, they can adapt as the data change to perform more complex analyses. Most current computational methods used to assess the significance of gene variants employ supervised training based on clinical labels, which may bias these tools and cause inflated accuracy of prediction in the real world.

Clues from our evolutionary relatives

In this work, the researchers hoped that by studying genetic variation across multiple species, they might glean clues about the significance of variation in humans.

Evolution tends to preserve features that are critical, or at least important, to function and survival across species. Thus, amino acid arrangements that recur across species are markers of biologic importance, indicating they are important to an organism's function and its evolutionary fitness. Alterations to such highly conserved sequences are linked to pathogenicity.

EVE looked for evolutionarily conserved patterns to draw conclusions. It analysed data from 140,000 species, including endangered and extinct organisms. Scientists have used comparative genetics for many years to detect regions of similarity across DNA or protein sequences to draw meaning. The Harvard-Oxford team used a neural network to do so on a much greater scale.

Training EVE

After training on 250 million protein sequences, EVE estimated the likelihood of each single amino acid variant to be either benign or pathogenic. To determine whether EVE was making accurate

continued on next page ▶

from previous page ▶

predictions, the researchers compared its scores with established human mutations whose significance is known. The tool's results were remarkably consistent with the clinical data, the team found.

Next, the researchers applied EVE to a set of 3,219 human genes associated with disease. EVE made the right call on whether the mutation was pathogenic or benign across all genes, including 60 'clinically actionable' genes, the researchers said. When researchers compared EVE's performance with that of other tools, it showed notably greater accuracy of prediction.

But how would EVE's predictions fare compared with findings made from actual clinical experiments, the gold standard of assessing how a genetic mutation affects physiologic function?

The team compared EVE's scores against results from clinical experiments involving well-studied mutations in five genes, among them genes related to various forms of cancer, several cancer syndromes, and heart rhythm disorders. EVE's predictions overlapped with current labels from experimental data.

Debora said, 'It seems that by simply training a model to fit the distribution of sequences across evolution we extract information which enables us to make unexpectedly precise predictions about disease risk arising from a given genetic variant.'

A matter of trust

A notable advantage that EVE has over current methods is that it assigns a continuous score rather than a binary score. This is because even when gene variants are labelled as benign or pathogenic, how a mutation might manifest physiologically is more nuanced.

Debora said, 'The continuous score is very important for predicting what the level of pathogenicity is. Does the mutation mean I am going to get pain in my little toe, or am I going to die tomorrow?'

Another important aspect of the tool is that it assigns a confidence-of-prediction score on a gene-by-gene basis. For each genetic variant, EVE tells the expert how much they can trust its call. This is a matter of trustworthiness, of confidence in the model, the researchers said.

'We're not providing clinicians merely with a number but also giving them the degree of uncertainty that comes with it,' Yarin said. 'This is something that the expert can take and use in the decision-making process. The tool can say, 'I think that variant belongs to that pile, but I've never seen any variants like that before so take that with a grain of salt.' Or the tool can also say, 'I think that that other variant belongs to this pile, and I've seen very similar variants to that in the past, and I saw them belonging to this pile and therefore I'm going to assign it to this pile with high confidence.' Building trust between the tool and the expert is an important aspect of this work.

Looking to the future

This type of modelling is still in its infancy, and it's clear that evolution and genetic variation still have much to teach us about disease. The researchers plan to extend the work to other parts of the genome beyond protein-coding regions.

In the immediate future, however, the urgent task is to make clinical use of the genetic variation for which we do have some understanding. To do so, the researchers have already teamed up with a genome-sequencing company and are collaborating with various groups.

The team is also participating in the Atlas of Variant Effects Alliance, a global research effort whose mission is to map the effects of variation across the genome and create a comprehensive atlas of all possible human gene variants and their effects on protein function and physiology. The ultimate goal of the effort is to improve the diagnosis, prognosis, and treatment of human disease.

Study co-authors included Aidan Gomez of Oxford University and Joseph Min and Kelly Brock of Harvard Medical School.

This abridged version is reprinted with gratitude from an original by Harvard University.

Full article here: hms.harvard.edu/news/all-about-eve

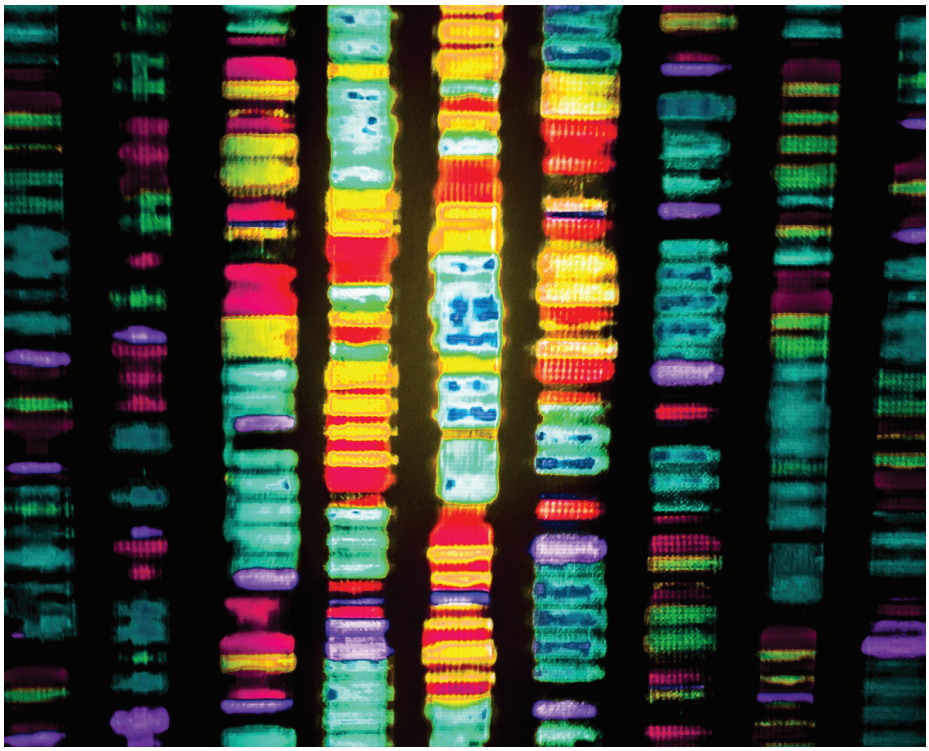

Carbon footprints and communications technologies – adjusting to the 'new normal'

By Carolyn Ten Holter, Doctoral Student and Research Assistant

The climate crisis is now a fixture on the agendas of national and local governments, industry, and research institutions. Its critical importance is reflected in the priorities of funders and researchers, and in numerous initiatives across the University and at a college level.

At University level there is now an Environmental Sustainability Subcommittee (ESS) of the Planning and Resource Allocation Committee (PRAC), chaired by David Prout, Pro-Vice-Chancellor Planning and Resources. The ESS has representatives from across the University and student body and set targets in March 2021 for the University to achieve net-zero carbon, and biodiversity net gain, by 2035. These new objectives show the importance the climate question has come to assume in the priorities of institutions across the globe.

As well as ensuring we examine our own impact on climate, these issues have also contributed to a new research agenda across many of the University's divisions. In 2020 the Engineering and Physical Sciences Research Council (EPSRC) began to highlight the importance of research into digital sustainability. A community of ICT researchers proposed an additional theme for an EPSRC call for proposals for Digital Economy Sustainability. EPSRC developed these priority areas, one of which - 'Sustainable Digital Society', was completely new for EPSRC.

Within the Department we were delighted to receive funding for a project proposal within this Sustainable Digital Society priority, co-led by Professor Marina Jirotka and the University of Lancaster. The PARIS-DE project, which will be located at the new Responsible Technology Institute (rti.ox.ac.uk) focuses on the carbon emissions of ICT technology with the goal of ensuring alignment with Paris targets to limit temperature increase to 1.5 degrees Celsius. The carbon emissions of ICT, as the most optimistic estimates suggest, will at best flatten out, so a profound change of course is necessary in order to reduce ICT emissions in a meaningful way. This project is looking at ways to ensure that future innovations in the digital economy are Paris-compliant by design.

It aims to achieve this through a systemic approach to developing a digital sustainability framework featuring consideration of evidence and reflections around responsible innovation. The framework will be instantiated in the cloud as a virtual design 'lab', featuring reusable digital tools for assessing carbon impacts in balance with social impacts for digital economy projects. The project will build on the rich heritage of socio-technical design and extend it to embed sustainability and consideration of planetary boundaries.

Additionally, as part of its own sustainability effort, the PARIS-DE project will be the first to our knowledge that will measure its own carbon footprint. Within the project there are funds allocated to work with a climate-impact measurement company that will assess the project in terms of its carbon mitigation efforts (such as reducing flights, cutting down on purchase of hardware, travel and conference attendance). Although these items can be difficult to measure, and there is not a consensus around, for example, how and where to measure 'embedded' carbon, the project will document its procedures and processes in order to provide

continued on next page ▶

from previous page ▶

a benchmark and potentially a framework for other projects to use. These sustainability concerns and records will form one of the project's outputs as well as the outcomes of the research effort.

The PARIS-DE project is the latest development in the work of Marina's research group, which is focussed on the application of responsible innovation principles in numerous disciplines and research fields. Marina was part of the Advisory Board responsible for the inclusion of sustainability within the Digital Economy theme at EPSRC and has been invited to give numerous conference keynotes over the last two years that focus on sustainability questions. Some of this reflects the work carried out for the 100+ Brilliant Women conference in 2019 (see page 9, *Inspired Research*, winter 2019), which remains Oxford's only certified carbon-neutral conference. This event not only highlighted the importance of not creating carbon in the first place (as opposed to seeking to offset it), but also the disproportionate impact of flights on the carbon footprint. The organisers gathered details during the registration process of every attendee's starting location, mode of travel, class of flight, and accommodation. The 26 flights that participants took to attend the conference generated 96% of the conference's total carbon, highlighting how crucial this issue is for academic conferences.

Going forward we are also looking at how our ways of working can change. During the last year and a half, as conferences have been forced to move to online delivery, the academic community has had the opportunity to experiment with new modes of operation, giving keynotes and panels from home, participating in workshops virtually, and delivering papers over online platforms. Any attendee of one of these virtual conferences will know that they have advantages and disadvantages. Learning a new platform or spending an entire day on these online platforms can be draining and unrewarding, and despite the best efforts of platform developers, it is nearly impossible to replicate the social and 'serendipitous encounter' elements of in-person conferences. On the other hand, it is possible to listen to a keynote speaker giving a talk many miles away while, for instance, fitting this around day-to-day activities, such as caring responsibilities or other meetings crucial to the operation of projects.

These factors, combined with carbon and budget savings, may make the availability of 'virtual' conference attendance likely for the foreseeable future. With improved data becoming available on digital tools and services such as streaming, from projects like PARIS-DE, we will also be better able to judge the relative impacts of virtual versus in-person conference carbon creation, and to factor this in as part of a wider understanding of the research community's responsibility towards sustainability goals. Ultimately, objectives such as the University's targets for carbon reduction and biodiversity gain are dependent on the availability of good data, agreed standards, and collective endeavour.

Through these initiatives, and many others, we are looking at how the way we work, the research that we do, and the resources we use can help to mitigate the climate crisis.

Can we locate endangered sea turtles using twelve milliseconds of noisy satellite signals?

By Amanda Matthes and Jonas Beuchert (EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems)

As part of our research with Professor Alex Rogers, we have developed SnapperGPS, a low-cost, low-power wildlife tracking system based on satellite navigation. In summer 2021, we deployed it for the first time on wild animals: endangered loggerhead sea turtles in Cape Verde.

Location tracking devices are an important tool for biologists to study animal behaviour. Usually, they use global navigation satellite systems like the GPS for this. However, existing devices are often expensive and come with heavy batteries for long-term deployments. One tag can easily cost more than £750, which prohibits studies with many animals. That is why we developed SnapperGPS.

SnapperGPS aims at being a cheap, small, and low-power tracking solution. Its core idea is to make the hardware as simple and as energy efficient as possible by doing as little signal acquisition and processing on the device as possible. To do this, we created a web service that processes the signals in the cloud. This allows us to build a bare-bone receiver for less than £22, which runs for more than ten years on a coin cell.

The concept that we employ is known as snapshot GNSS. Its advantage is that a few milliseconds of signal are enough to locate the receiver. With SnapperGPS we face the particular challenge that the hardware records signals at a much lower resolution than existing receivers. To address this problem, we developed and implemented three alternative algorithmic approaches to location estimation from short low-quality satellite signal snapshots, which are all based on probabilistic models.

In summer 2021 we were able to deploy SnapperGPS on nesting loggerhead sea turtles (Caretta caretta) on the island of Maio in Cape Verde. Loggerhead sea turtles spend most of their life in the ocean, but every two to three years mature females come to a beach to nest. They lay several clutches of eggs separated by roughly two weeks, which makes it possible to recover the hardware and any data it captured.

Navigation satellite signals cannot travel through water, but sea turtles regularly come to the surface to breathe. These short windows of opportunity may not always be enough for traditional GPS methods to resolve the position of the receiver. But a snapshot method only requires milliseconds of the signal which makes them ideal for such marine applications.

continued on next page ▶

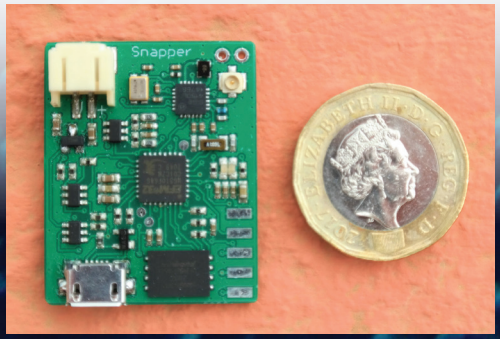

Figure 1: A SnapperGPS board next to a £1 coin. It measures 3.5 cm x 2.8 cm.

Figure 2: Once configured, SnapperGPS regularly collects snapshots of satellite data. After retrieval of the device, the collected data is uploaded to the cloud where the location track is computed.

The SnapperGPS tags are glued to the turtle's carapace. The epoxy comes off naturally if the tag is not recovered.

from previous page ▶

For this turtle deployment, the SnapperGPS tags were placed into custom-made enclosures that were tested to be waterproof to at least 100 metres.

Due to the COVID-19 pandemic, we had to deploy the tags late in the nesting season which negatively affected our recovery rate as many turtles were already laying their last nest when they were tagged. We deployed twenty tags and recovered nine. Some devices experienced unexpected technical failures but we were able to capture several location tracks that showed unexpectedly diverse behaviour among turtles. This data provides novel insights into the loggerhead sea turtle population on Maio. We also learned about the specific challenges of deploying SnapperGPS on a sea turtle and will work on an improved version for next year's nesting season.

Wildlife location tracking data can inform conservation policy decisions that help protect habitats and prevent human-wildlife conflicts. In the case of loggerhead sea turtles, understanding their movements can inform where to direct anti-poaching measures and it can identify important marine habitats that may need special protection.

SnapperGPS is supported by an EPSRC IAA Technology Fund. Additionally, Amanda and Jonas receive support from the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems (AIMS CDT). The fieldwork was made possible through a cooperation with the Maio Biodiversity Foundation and the Arribada Initiative.

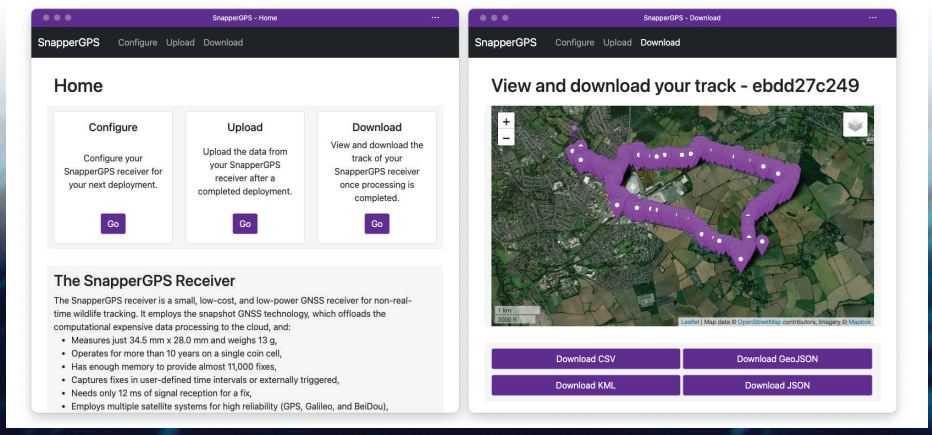

Figure 3: A web application serves as the front end for configuring SnapperGPS devices, uploading data logs and downloading computed location tracks.

Figure 4: A location track of a loggerhead sea turtle captured by a SnapperGPS tag

Peering into the Moon's permanently shadowed regions with AI

By Robert Ashley, Oxford University Public Affairs Directorate

The Moon's polar regions are home to craters and other depressions that never receive sunlight. Permanently shadowed lunar craters contain water ice but are difficult to image. An AI algorithm now provides sharper images, allowing us to see into them with high resolution for the first time.

A group of researchers led by the Max Planck Institute for Solar System Research (MPS) in Germany, supported by the University of Oxford (including members of the Department of Computer Science) and the NASA Ames Research Center, have taken a closer look at some of these regions and presented the highest-resolution images to date covering 17 such craters in the journal *Nature Communications*.

Craters of this type could contain frozen water, making them attractive targets for future lunar missions, and the researchers focused further on relatively small and accessible craters surrounded by gentle slopes. Three of the craters have turned out to lie within the just-announced mission area of NASA's Volatiles Investigating Polar Exploration Rover (VIPER), which is scheduled to touch down on the Moon in 2023.

Imaging the interior of permanently shadowed craters is difficult, and efforts so far have relied on long exposure times resulting in smearing and lower resolution shots. By taking advantage of reflected sunlight from nearby hills and a novel image processing method, the researchers have now produced images at 1-2 meters per pixel, which is at or very close to the best capability of the cameras.

The Moon is a cold, dry desert. Unlike the Earth, it is not surrounded by a protective atmosphere, and water which existed during the Moon's formation has long since evaporated under the influence of solar radiation and escaped into space. Nevertheless, craters and depressions in the polar regions give some reason to hope for limited water resources.

'Near the lunar north and south poles, the incident sunlight enters the craters and depressions at a very shallow angle and never reaches some of their floors', MPS-scientist Valentin Bickel explains. In this 'eternal night', temperatures in some places are so cold that frozen water is expected to have lasted for millions of years. Impacts from comets or asteroids could have delivered it, or it could have been outgassed by volcanic eruptions, or formed by the interaction of the surface with the solar wind.

Measurements of neutron flux and infrared radiation obtained by space probes in recent years indicate the presence of water in these regions. Eventually, NASA's Lunar Crater Observation and Sensing Satellite (LCROSS) provided direct proof: twelve years ago, the probe fired a projectile into the shadowed south pole crater Cabeus. As later analysis showed, the dust cloud emitted into space contained a considerable amount of water.

However, permanently shadowed regions are not only of scientific interest. If humans are ever to spend extended periods of time on the Moon, naturally occurring water will be a valuable resource – and shadowed craters and depressions will be an important destination. NASA's uncrewed VIPER rover, for example, will explore the South Pole region in 2023 and enter such craters. To get a precise picture of their topography and geology in advance – for mission planning purposes, for example – images from space probes are indispensable. NASA's Lunar Reconnaissance Orbiter (LRO) has been providing such images since 2009.

However, capturing images within the deep darkness of permanently shadowed regions is exceptionally difficult; after all, the only sources of light are scattered light, such as that reflecting off the Earth and the surrounding topography, and faint starlight. 'Because the spacecraft

continued on next page ▶

from previous page ▶

is in motion, the LRO images are completely blurred at long exposure times,' explains Department of Computer Science Doctoral Student Ben Moseley.

At short exposure times, the spatial resolution is much better. However, due to the small amounts of light available, these images are dominated by noise, making it hard to distinguish real geological features. To address this problem, the researchers have developed a machine learning algorithm called HORUS (Hyper-effective nOise Removal U-net Software) that 'cleans up' such noisy images.

It uses more than 70,000 LRO calibration images taken on the dark side of the Moon as well as information about camera temperature and the spacecraft's trajectory to distinguish which structures in the image are artifacts and which are real. This way, the researchers can achieve a resolution of about 1-2 meters per pixel, which is five to ten times higher than the resolution of all previously available images.

Using this method, the researchers have now re-evaluated images of 17 shadowed regions from the lunar south pole region which measure between 0.18 and 54 square kilometres in size.

In the resulting images, small geological structures only a few meters across can be discerned much more clearly than before. These structures include boulders or very small craters, which can be found everywhere on the lunar surface. Since the Moon has no atmosphere, very small meteorites repeatedly fall onto its surface and create such mini craters.

'With the help of the new HORUS images, it is now possible to understand the geology of lunar shadowed regions much better than before,' explains Ben. For example, the number and shape of the small craters provide information about the age and composition of the surface. It also makes it easier to identify potential obstacles and hazards for rovers or astronauts. In one of the studied craters, located on the Leibnitz Plateau, the researchers discovered a strikingly bright mini-crater.

'Its comparatively bright colour may indicate that this

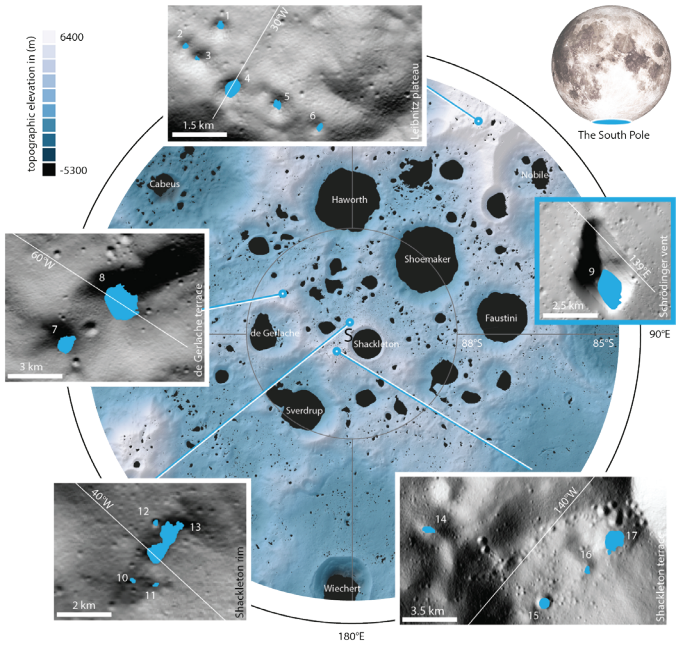

crater is relatively young,' says Valentin. Because such a fresh scar provides fairly unhindered insight into deeper layers, this site could be an interesting target for future missions, the researchers suggest.

The new images do not provide evidence of frozen water on the surface, such as bright patches. 'Some of the regions we've targeted might be slightly too warm,' Valentin Bickel speculates. It is likely that lunar water does not exist as a clearly visible deposit on the surface at all - instead, it could be intermixed with the regolith and dust, or may be hidden underground.

To address this and other questions, the researchers' next step is to use HORUS to study as many shadowed regions as possible. 'In the current publication, we wanted to show what our algorithm can do. Now we want to apply it as comprehensively as possible,' says Valentin.

This work has been enabled by the Frontier Development Lab (FDL.ai). FDL is a co-operative agreement between NASA, the SETI Institute (seti.org) and Trillium Technologies Inc, in partnership with the Luxembourg Space Agency and Google Cloud.

Read more: bit.ly/3BfSoje

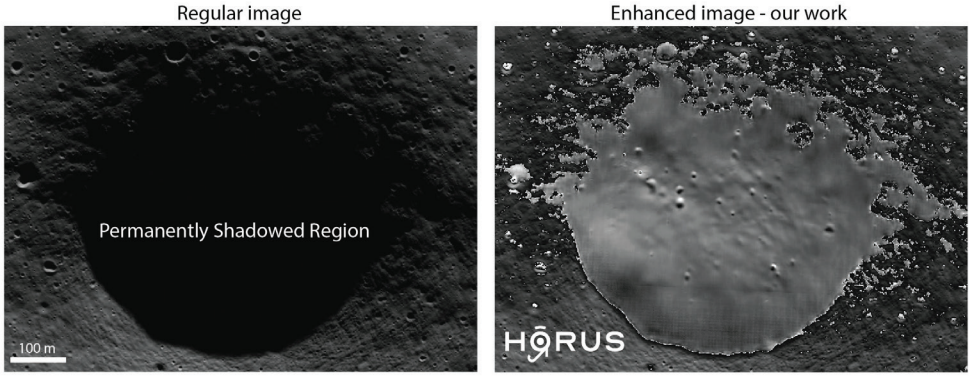

An as-yet unnamed crater (region 1 in figure above) *near the Moon's south pole and close to the proposed landing site of NASA's Volatiles Investigating Polar Exploration Rover. Left shows an image taken by the Lunar Reconnaissance Orbiter. Right shows the same image after our image processing. credits: left: NASA/LROC/GSFC/ASU; right: MPS/University of Oxford/NASA Ames Research Center/FDL/SETI Institute*

AI pioneered at Oxford to detect floods launches into space

A new technology, which was developed by Oxford researchers in partnership with the European Space Agency's (ESA) Φ-lab, will pilot the detection of flood events from space. It was deployed on hardware on D-Orbit's 'Wild Ride' mission launched by SpaceX's Falcon 9 from Cape Canaveral on 30 June 2021. The work is a first step towards relaying real-time information from space to disaster response teams.

The Oxford team has developed a machine learning model called 'Worldfloods'designed specifically for deployment in specialised hardware in space on low-cost satellites in Low Earth Orbit.

The model is a flood segmentation model whose purpose is to detect flood events so as to significantly improve disaster response operations. It has major implications in bringing down the cost of such technologies and making them accessible to low-income countries.

Atılım Güneş Baydin, based at the Departments of Engineering Science and Computer Science, University of Oxford, said: 'This will be the first time a machine learning model for this type of task will be actually deployed in space. It's a very significant step in bringing machine learning and artificial intelligence operations to space.'

The international team of eight people conducted the research as part of the 'Frontier Development Lab (FDL) Europe', a partnership between the University of Oxford, European Space Agency's Φ-lab (ESA ESRIN), Trillium Technologies and leaders in commercial AI, such as Google Cloud and Intel.

The team includes Professor Yarin Gal, Gonzalo Mateo-Garcia, Joshua Veitch-Michaelis, Lewis Smith, Silviu Oprea, Guy Schumann, Atılım Güneş Baydin, and Dietmar Backes. Three members are Oxford researchers: Atılım Güneş Baydin (Engineering Science and Computer Science), Professor Yarin Gal (Computer Science), and Lewis Smith (DPhil student in the Department of Computer Science).

The two-month FDL programme was hosted at the ESA ESRIN Φ-lab (Italy) and at Kellogg College (Oxford) in the summer of 2019. The team has continued to work on the project and kept developing the system until the present day.

The Project has had input from UNICEF due to the humanitarian aspects of the project.

An article recently published in Scientific Reports covers the details of the work: go.nature.com/31GUs7h