The Foundation for Science and Technology
Debate Record Note

The impact on society of machine learning - an opportunity or a threat?

Date and Location:	14th November, 2017 at The Royal Society
Chair:	The Earl of Selborne GBE FRS Chair, The Foundation for Science and Technology
Speakers:	Dr Mike Lynch OBE FRS FREng DL Founder, Invoke Capital Dr Claire Craig CBE Director of Science Policy, The Royal Society Amir Saffari Head of AI, Benevolent AI Professor Dame Wendy Hall DBE FRS FREng Regius Professor of Computer Science, University of Southampton
Panellist:	Professor Chris Bishop FRS FREng Laboratory Director, Microsoft Research Cambridge
Sponsors:	Institute of Practitioners in Advertising, Invoke Capital, Microsoft Research and The Royal Society
Audio Files:	www.foundation.org.uk
Hash tag	#fstmachinelearning

DR MIKE LYNCH said that the impact of machine learning would be profound, by replacing more advanced cognitive tasks currently performed by humans. He had worked in this field since the late 1980s, but recently there had been great leaps forward, in unravelling the complexities of visual recognition and speech recognition.

Arguably current developments meant that computer recognition could now be more accurate than human recognition in a range of tasks. This was leading to the risk of inappropriate regulation, proposed by those with inadequate understanding of machine learning. Machines could learn empathy, and he felt the world had much more to fear from evil people than from machines. The impact of machine learning on jobs would be significant, but there would also be important new skill requirements, particularly in the fields of maths and signal processing. There would be significant economic impact on some existing industries, but also scope for generating economic growth. There would also be investments in new start-ups, most of which inevitably would not be successful.

The assumption that machine learning could only operate with clean data was wrong. Machine learning needed access to data in many fields, and this raised the strategic issue of whether significant data owners, like the NHS, should share in the rewards which could be generated from their data. Although this might be contrary to existing Government policy on open data, there was a strong case for this.

Currently it was hard for machine learning to handle exceptions, such as how two buses might pass on the narrow, cliff hugging Amalfi coast road. There also needed to be thoughtful

development of the insurance market in relation to machine learning. The market could handle car accidents involving human driving, but needed to consider insurance in relation to driverless vehicles. There were also tricky areas in that machine learning might enable greater understanding of insurance risks between some groups in society and others, which could cause the insurance market to differentiate insurance premiums in ways which would be socially problematic. By contrast, other machine learning applications could soon have benefits for society, such as enabling more early stage dementia patients to remain cared for at home, through optimising use of existing infrastructure.

Culturally machine learning would hold up a mirror to people's attitudes, and this might not prove comfortable. There would be less opportunity in human society for hypocrisy. Hopefully society would come to understand probability in relation to human events better. Whilst in the coming years machine learning might not delivered all the promise anticipated, it would create a revolution, and even if some changes might take longer to deliver than expected, the full impact of its changes would be greater than currently understood. This would require a new ethical framework.

DR CLAIRE CRAIG opened by saying that the UK was well placed to lead on the opportunities for machine learning to improve human health. The Royal Society's machine learning report[1], published earlier this year, had involved extensive public engagement on machine learning for the first time, including a popular event led by Professor Brian Cox at the Royal Festival Hall. It had recommended five areas for action, and had built on a report in 2014 by the then Government's Chief Scientific Adviser, Sir Mark Walport, on managing risk in relation to innovation. This year's report had found that only 9% of the population were aware of machine learning, but more understood concepts like speech recognition. Responses had varied between those inherently comfortable with technological advances and those more concerned that machine learning challenged existing concepts of what it meant to be human. Although young people typically understood machine learning intuitively, they had had the same spectrum of responses.

The Royal Society's report had recommended that, although new frameworks could be needed to respond to modern innovation, it would be right not to develop governance for machine learning per se, but for particular areas of application such as transport. During the public dialogue some had called for an element of human involvement always to be retained in key decisions, whilst others felt, for example, that a tired and overworked GP might produce less effective diagnoses than machine learning could achieve.

The Royal Society had worked with the British Academy to produce a recent report "Data management and use: Governance in the 21st century[2]", which had included the ethical dimension. They were waiting for a response from the Government to this. A next stage was for the Nuffield Foundation to undertake longer term ethical foresight work. The way ahead would require a new wave of research, in interpretability, causality, verification and validation, human-machine interaction, privacy, security and dealing with real-world data.

AMIR SAFFARI opened by providing a historical introduction to artificial intelligence (AI) from the 1950s onwards. By the 1990s machine learning was being taken more seriously, through the development of spam filters and search engines. More recently neural networks and large data had created great interest.

Object recognition had become much more advanced, but it had only been in 2017 that error rates had become lower than the rate for humans. Speech recognition was beginning to have very diverse applications. Media coverage was still problematic, contributing to the current hype, but some was becoming more reputable.

Human augmentation machine learning technologies were now emerging substantively. Four areas for this were highlighted - accelerating scientific research in drug discovery work; creation of authentic images such as in car chassis design; creating polyphonic music; and developing personalised and interactive education.

PROFESSOR DAME WENDY HALL said she had been asked in March 2017 to co-chair a group producing a report as part of the preparation of the UK Government's industrial strategy. This had been called "Growing the Artificial Intelligence industry in the UK[3]", and had had to be delivered in four months. It had been published by the Government on 15 October, and had been focussed on job creation, economic growth, and whether a coherent AI industrial sector might emerge in the UK. The report had been prepared alongside the joint Royal Society/British Academy re-

1 https://royalsociety.org/topics-policy/projects/machine-learning/?gclid=EAIaIQobChMI5cfDkIHQ1wIVCEAbCh0YIQNaEAAYASAAEgJSKfD_BwE

2 www.britac.ac.uk/data-governance

3 www.gov.uk/government/publications/growing-the-artificial-intelligence-industry-in-the-uk

port mentioned by Dr Craig, and it had not had time to cover the important issues of ethics and accountability.

The report had recommended the creation of data trusts, to improve trust and ease around sharing data, especially to benefit the growth of smaller companies. It had also made a series of recommendations to increase the UK's skills base in AI, including an industry-funded Masters programme, 200 more PhD places and the development of on-line courses for Continuous Professional Development. The report had identified the particular challenge of increasing the depth of capabilities in UK universities to deliver all this extra training, when the individuals capable of doing this were highly marketable themselves. The Alan Turing Institute should become the national institute for AI and data science. An AI Council should be set up to coordinate the necessary initiatives, with Government providing AI challenge funds so that the UK Government itself became a key adopter.

Her personal reflections were that she agreed with the call by Dr Mike Lynch for a new ethical framework. A mix of Government regulation and self-regulation would be appropriate.

Gender diversity in this emerging sector needed to be tackled urgently. The best quick approach would require involving inter-disciplinary teams with better diversity than the computing sector had recently managed, from the beginning. Otherwise it would be dangerous if all the key algorithms were established by one section of society. It had to be recognised that some people would lose their jobs and not have the capabilities to be re-trained to participate in this sector. The welfare state would need to evolve in the world of AI.

PROFESSOR CHRIS BISHOP opened the discussion by saying that there was little he had heard with which he disagreed. There was substance in the advances of machine learning, but they still remained far short of replicating general human intelligence. What was transformative was the capability of software to allow machines to be adaptive. Rather than be programmed to solve the problem, they are programmed to learn how to solve the problem using available data. Several factors had combined to generate the current potential of machine learning. There were tricky issues about who was able to use large data sets, and how the outputs from their use might be regulated. If this was not addressed there was a risk that the public could turn against AI. He commended The Royal Society for their recent work on public engagement. The challenges which had been mentioned about skills were critical. There was currently a serious skills shortage. The UK had a unique opportunity to make progress because of its great universities. He was optimistic about gender diversity, because the biggest challenges lay in areas where multi-disciplinary working was essential.

DISCUSSION

The subsequent discussion started with a debate about whether any of the estimates produced for the extent of jobs which would be lost or affected by machine learning were robust. Broad estimates had been considered by the Council for Science & Technology, but none could be regarded as definitive.

Several agreed with the proposal that, if NHS data was used in new ways, the NHS itself should receive benefit, at least by obtaining access to the relevant products at a discount. One idea was that individuals who obtained their own health data should be able to contribute that data into curated records, which could be used for wider benefit. It was suggested that the manner in which the Met Office had been forced to offer its data free to competitors had weakened it in an unhelpful way.

There was healthy overlap between AI and traditional neuroscience, which went back to the origins of AI. These synergies should continue to be fruitful.

There was debate about whether developments in AI might widen global inequalities, and whether this should be considered further by the World Health Organisation. One view was that the deployment of healthcare advances to patients might be widened if AI could contribute to widespread use of AI driven apps on phones, which could make good quality medical advice more readily available to communities lacking access to doctors.

Understanding of mathematics and probability would be key to the pace of further progress, but this needed to be coupled with understanding of how real data could incorporate bias, and the development of new techniques to avoid bias.

Although machine learning could eliminate the need for human involvement in many tasks, human understanding of the outcomes from AI would still be fundamental. Those responsible for the expansion of apprenticeships should be pressed to consider how routes to jobs in AI could be established. There were differing views over the contribution which universities might make to degree apprenticeships in AI. Further thinking was needed about how to enthuse youngsters about the potential of AI. Recent initiatives to increase the teaching of computer science in schools

were most welcome. The impact on employment might be most significant for those aged 40-60, so the potential to retrain as many as possible of those in this age group should also not be neglected.

Machine learning was already being used in copywriting and social media advertising. Ethical issues and regulation needed to be considered further in this area.

The discussion had largely focussed on the potential benefits of machine learning. Further understanding of the negative consequences was needed too.

The implications of AI for security and for tracking individuals had not been covered in the discussion. The consequence of AI being used for nefarious purposes was an important area.

John Neilson

USEFUL URLS - CLICK ON URL TO GO TO SITE

CONTEXT

Alphago
https://deepmind.com/research/alphago/

Alphago - Silver D et al, Nature 550, 354–359 (19 October 2017) doi:10.1038/nature24270
www.nature.com/articles/nature24270

Deepmind
https://deepmind.com/

Department for Business, Energy and Industrial Strategy Press Release Reference to funding for AI
www.gov.uk/government/news/record-boost-to-rd-and-new-transport-fund-to-help-build-economy-fit-for-the-future

Hall and Pesenti Report to the Government on AI
www.gov.uk/government/publications/growing-the-artificial-intelligence-industry-in-the-uk

Machine Learning videos and background information from The Royal Society site
https://royalsociety.org/topics-policy/projects/machine-learning/videos-and-background-information/

Machine Learning Education
https://www.coursera.org/learn/machine-learning

Nesta on the fourth industrial revolution
www.nesta.org.uk/blog/how-can-fourth-industrial-revolution-be-made-good?

The Royal Society project on machine learning
https://royalsociety.org/topics-policy/projects/machine-learning/

The Royal Society and British Academy Report on data governance
https://royalsociety.org/~/media/policy/projects/data-governance/data-management-governance.pdf

White House Report on AI
https://obamawhitehouse.archives.gov/blog/2016/10/12/administrations-report-future-artificial-intelligence

SOFTWARE FOR MACHINE LEARNING

Mathematica
www.wolfram.com/mathematica/new-in-10/highly-automated-machine-learning/

MATLAB
https://uk.mathworks.com/campaigns/products/trials/targeted/mal.html?s_eid=ppc_51241087401&q=matlab%20machine%20learning

Nvidia
https://developer.nvidia.com/dli/onlinelabs

R MIT: Using R for machine learning
https://ocw.mit.edu/courses/sloan-school-of-management/15-097-prediction-machine-learning-and-statistics-spring-2012/lecture-notes/MIT15_097S12_lec02.pdf

RESEARCH COUNCILS:

Arts and Humanities Research Council
www.ahrc.ac.uk

Biotechnology and Biological Sciences Research Council
www.bbsrc.ac.uk

Engineering and Physical Sciences Research Council
www.epsrc.ac.uk

Economic and Social Research Council
www.esrc.ac.uk

Medical Research Council
www.mrc.ac.uk

Natural Environment Research Council
www.nerc.ac.uk

Science and Technology Facilities Council
www.stfc.ac.uk

UKRI
www.ukri.org

COMPANIES, DEPARTMENTS, RESEARCH ORGANISATIONS AND ACADEMIES:

Association of Innovation, Research and Technology Organisations (AIRTO)
www.airto.co.uk

Academy of Medical Royal Colleges
www.aomrc.org.uk

Academy of Medical Sciences
www.acmedsci.ac.uk

BenevolentAI
www.benevolent.ai

British Academy
www.britac.ac.uk

Catapult Programme
www.catapult.org.uk

Council for Science and Technology
www.gov.uk/government/organisations/council-for-science-and-technology

Department for Business, Energy and Industrial Strategy
www.gov.uk/government/organisations/department-for-business-energy-and-industrial-strategy

Department for Communities and Local Government
www.gov.uk/government/organisations/department-for-communities-and-local-government

Department for Education
www.gov.uk/government/organisations/department-for-education

Department of International Trade
www.gov.uk/government/organisations/department-for-international-trade

Francis Crick Institute
www.crick.ac.uk

Google
www.google.co.uk

Government Office for Science
www.gov.uk/government/organisations/government-office-for-science

IBM
www.ibm.com/uk-en

Innovate UK
www.gov.uk/government/organisations/innovate-uk

The Institute of Practitioners in Advertising
www.ipa.co.uk

Invoke Capital
www.invokecapital.com

Knowledge Transfer Network
www.ktn-uk.co.uk

Learned Society of Wales
www.learnedsociety.wales

Microsoft Research Ltd
www.microsoft.com

National Physical Laboratory (NPL)
www.npl.co.uk

Phrasee
www.phrasee.co

Royal Academy of Engineering
www.raeng.org.uk
The Royal Society
www.royalsociety.org
The Royal Society of Edinburgh
www.rse.org.uk

Russell Group
www.russellgroup.ac.uk

University Alliance
www.unialliance.ac.uk

Wellcome Trust
www.wellcome.ac.uk

WPP Oglivy and Mather
http://www.wpp.com/wpp/companies/ogilvy-mather-worldwide/office/ogilvy-mather-london/

UNIVERSITIES:
For a full list of UK universities go to:
www.universitiesuk.ac.uk