GROWING THE ARTIFICIAL INTELLIGENCE INDUSTRY IN THE UK

Professor Dame Wendy Hall and Jérôme Pesenti

FOREWORD

We are grateful to the Business Secretary and Culture Secretary for asking us to conduct this Review of how to grow Artificial Intelligence in the UK, in terms of those developing it and deploying it. We believe that this is the right time for the UK to accelerate on AI, and ensure that our unique history of ground breaking research bears fruit in the social and economic benefits that the technology offers.

We are at the threshold of an era when much of our productivity and prosperity will be derived from the systems and machines we create. We are accustomed now to technology developing fast, but that pace will increase and AI will drive much of that acceleration. The impacts on society and the economy will be profound, although the exact nature of those impacts is uncertain. We are convinced that because of the UK's current and historical strengths in this area we are in a strong position to lead rather than follow in both the development of the technology and its deployment in all sectors of industry, education and government.

We have a choice. The UK could stay among the world leaders in AI in the future, or allow other countries to dominate. We start from a good position in many respects but other leading countries are devoting significant resources to growing and deploying AI. The UK will need to act in key areas and to sustain action over a long period and across industry sectors, to retain its world leading status, and to grow our AI capability as well as deploying it much more widely.

If we can judge by the contributions to this Review by academic and industry experts, the UK has the expertise and the appetite to grasp the opportunity if we act decisively now.

We would like to thank everyone who helped us with this review. We consulted with a wide range of experts and organisations but because the time we had to undertake the review was relatively short, it was not possible to talk to everyone with valuable expertise and experience. We hope that many more people and organisations will take part in developing and delivering these recommendations to make the UK the best place in the world for AI companies to flourish and deploy AI across all sectors of society for the benefit of all.

EXECUTIVE SUMMARY

Increased use of Artificial Intelligence (AI) can bring major social and economic benefits to the UK. With AI, computers can analyse and learn from information at higher accuracy and speed than humans can. AI offers massive gains in efficiency and performance to most or all industry sectors, from drug discovery to logistics. AI is software that can be integrated into existing processes, improving them, scaling them, and reducing their costs, by making or suggesting more accurate decisions through better use of information.

It has been estimated that AI could add an additional USD $814 billion (£630bn) to the UK economy by 2035, increasing the annual growth rate of GVA from 2.5 to 3.9%.[1]

Our vision is for the UK to become the best place in the world for businesses developing and deploying AI to start, grow and thrive, and to realise all the benefits the technology offers.

The pioneering British computer scientist Alan Turing is widely regarded as launching and inspiring much of the development of AI. While other countries and international companies are investing heavily in AI development, the UK is still regarded as a centre of expertise, for the present at least. This report recommends that more is done to build on Turing's legacy to ensure the UK remains among the leaders in AI.

Key factors have combined to increase the capability of AI in recent years, in particular:

- New and larger volumes of data
- Supply of experts with the specific high level skills
- Availability of increasingly powerful computing capacity.

The barriers to achieving performance have fallen significantly, and continue to fall.

To continue developing and applying AI, the UK will need to increase ease of **access to data** in a wider range of sectors. This Review recommends:

- Development of data trusts, to improve trust and ease around sharing data
- Making more research data machine readable
- Supporting text and data mining as a standard and essential tool for research.

Skilled experts are needed to develop AI, and they are in short supply. To develop more AI, the UK will need a larger workforce with deep AI expertise, and more development of lower level skills to work with AI. This review recommends:

- An industry-funded Masters programme in AI
- Market research to develop conversion courses in AI that meet employers' needs
- 200 more PhD places in AI at leading UK universities, attracting candidates from diverse backgrounds and from around the world.
- Credit-bearing online AI courses and continuing professional development leading to MScs
- Greater diversity in the AI workforce
- An international AI Fellowship Programme for the UK.

[1] https://newsroom.accenture.com/news/artificial-intelligence-poised-to-double-annual-economic-growth-rate-in-12-developed-

The UK has an exceptional record in key AI research. Growing the UK's AI capability into the future will involve building on this with **more research on AI in different application areas**, and coordinating research capabilities. This Review recommends:

- The Alan Turing Institute should become the national institute for artificial intelligence and data science
- Universities should promote standardisation in transfer of IP
- Computing capacity for AI research should be coordinated and negotiated.

Increasing uptake of AI means **increasing demand as well as supply** through a better understanding of what AI can do and where it could be applied. This review recommends:

- An AI Council to promote growth and coordination in the sector
- Guidance on how to explain decisions and processes enabled by AI
- Support for export and inward investment
- Guidance on successfully applying AI to drive improvements in industry
- A programme to support public sector use of AI
- Funded challenges around data held by public organisations.

Our work has indicated that action in these areas could deliver a step-change improvement in growth of UK AI. This report makes the **18 recommendations** listed in full below, which describe how Government, industry and academia should work together to keep the UK among the world leaders in AI.

RECOMMENDATIONS

RECOMMENDATIONS TO IMPROVE ACCESS TO DATA

1. **To facilitate the sharing of data between organisations holding data and organisations looking to use data to develop AI, Government and industry should deliver a programme to develop Data Trusts – proven and trusted frameworks and agreements – to ensure exchanges are secure and mutually beneficial.**

2. **To improve the availability of data for developing AI systems, Government should ensure that public funding for research explicitly ensures publication of underlying data in machine-readable formats with clear rights information, and open wherever possible.**

3. **To support text and data mining as a standard and essential tool for research, the UK should move towards establishing by default that for published research the right to read is also the right to mine data, where that does not result in products that substitute for the original works. Government should include potential uses of data for AI when assessing how to support for text and data mining.**

RECOMMENDATIONS TO IMPROVE SUPPLY OF SKILLS

4. **Government, industry and academia must embrace the value and importance of a diverse workforce for AI, and should work together to break down stereotypes and broaden participation.**

5. **Industry should sponsor a major programme of students to pursue Masters level courses in AI, with an initial cohort of 300 students.**

6. **Universities should explore with employers and students the potential demand for one-year conversion Masters degrees in AI for graduates in subjects other than computing and data science.**

7. **Government and universities should create, at a minimum, an additional 200 PhD places dedicated to AI at leading universities. As the UK trains and attracts additional academic talent, this number should grow continually year on year.**

8. **Universities should encourage the development of advanced credit-bearing AI MOOCs and online Continuing Professional Development courses leading to MScs for people with STEM qualifications to gain more specialist knowledge.**

9. **An International fellowship programme for AI in the UK should be created in partnership with the Alan Turing Institute: the Turing AI Fellowships. This should be supported by a targeted fund for identifying and recruiting the best talent, and by ensuring that the UK is open to any and all of the eligible experts from around the world.**

RECOMMENDATIONS TO MAXIMISE UK AI RESEARCH

10. **The Alan Turing Institute should become the national institute for artificial intelligence and data science, becoming truly national and expanded beyond the current five universities, with a key stated aim that centres its mission on artificial intelligence.**

11. **Universities should use clear, accessible and where possible common policies and practices for licensing IP and forming spin-out companies.**

12. **The Alan Turing Institute, Engineering and Physical Sciences Research Council (EPSRC), Science and Technology Facilities Council (STFC) and Joint Information Systems Committee (JISC) should work together to coordinate demand for computing capacity for AI research, and negotiate for the UK research community.**

RECOMMENDATIONS TO SUPPORT UPTAKE OF AI

13. **Government should work with industry and experts to establish a UK AI Council to help coordinate and grow AI in the UK.**

14. **The Information Commissioner's Office and the Alan Turing Institute should develop a framework for explaining processes, services and decisions delivered by AI, to improve transparency and accountability.**

15. **The Department for International Trade should expand its current support programme for AI businesses.**

16. **TechUK should work with the Royal Academy of Engineering, the Digital Catapult, and key players in industry sectors, to develop practical guidance on the opportunities and challenges of successful adoption of AI across the UK economy.**

17. **Government, drawing on the expertise of the Government Digital Service, the Data Science Partnership and experts working with data in other Departments, should develop a programme of actions to prepare the public sector and spread best practice for applying AI to improve operations and services for citizens.**

18. **Government should ensure that challenges addressed by the Industrial Strategy Challenge Fund (ISCF) and Small Business Research Initiative (SBRI) are**

designed to attract and support applications of AI across the full range of challenge areas and set funded challenges which use public sector data for AI.

Contents

SECTION ONE: ARTIFICIAL INTELLIGENCE AND OBJECTIVES OF THIS REVIEW

This section covers what AI is, why it matters now in the UK, and the objectives of this Review.

What is Artificial Intelligence?

Artificial Intelligence (AI) describes a set of advanced general purpose digital technologies that enable machines to do highly complex tasks effectively.

The Engineering and Physical Science Research Council uses this description: “Artificial Intelligence technologies aim to reproduce or surpass abilities (in computational systems) that would require 'intelligence' if humans were to perform them. These include: learning and adaptation; sensory understanding and interaction; reasoning and planning; optimisation of procedures and parameters; autonomy; creativity; and extracting knowledge and predictions from large, diverse digital data.”

Examples of AI already in use include: communicating with computers in natural language, deriving new insights from transport data, operating autonomous and adaptive robotic systems, managing supply chains, and designing more life-like video games. Applied AI is already changing business practices across financial services, law, medicine, accounting, tax, audit, architecture, consulting, customer service, manufacturing and transport. More examples are shown below, but no limited set of examples can be representative. AI could improve the functioning of most digital operations, products and services. Wherever a process uses digital data, AI may enable us to use that data more effectively and in new ways.

This report uses “Artificial Intelligence” as an umbrella term to cover a set of complementary techniques that have developed from statistics, computer science and cognitive psychology. While recognising distinctions between specific technologies and terms (e.g., artificial intelligence vs. machine learning, machine learning vs. deep learning), it is useful to see these technologies as a group, when considering how to support development and use of them.

What can AI do?
Artificial intelligence holds great potential for increasing productivity, most obviously by helping firms and people use resources more efficiently, and by streamlining the way we interact with large sets of data. For example, firms like Ocado and Amazon are making use of artificial intelligence to optimise their storage and distribution networks, planning the most efficient routes for delivery and making best use of their warehousing capacity. Artificial intelligence can help firms do familiar tasks in more efficient ways. Importantly, it can also enable entirely new business models and new approaches to old problems. For example, in healthcare, data from smartphones and fitness trackers that is analysed using new machine learning techniques can improve management of chronic conditions as well as predicting and preventing acute episodes of illness.

Artificial intelligence can help both companies and individual employees to be more productive. Routine administrative and operational jobs can be learned by software agents ('bots'), which can then prioritise tasks, manage routine interactions with colleagues (or other bots), and plan schedules. Email software like Google's Smart Reply can draft messages to respondents based on previous responses to similar messages. Newsrooms are increasingly using machine learning to write sports reports and to draft articles: in the office, similar technology can produce financial reports and executive briefings.

Artificial intelligence can reduce the burden of searching large sets of data. In the legal sector, groups like ROSS, Lex Machina and CaseText are using artificial intelligence to sift court documents and legal records for case-relevant information. Other firms are using similar techniques as part of due diligence. Artificial intelligence can also offer a way of interacting with these datasets, with platforms such as IBM's Watson able to support expert systems that can answer factual natural language questions. For cybersecurity firms, artificial intelligence offers a way of recognising unusual patterns of behaviour in a network.

These examples focus on using software to do the same thing as humans but, in many cases, analysing data of volume or complexity that is beyond the analytical capability of individual humans. Indeed, artificial intelligence is not a replacement, or substitute for human intelligence. It is an entirely different way of reaching conclusions. Artificial intelligence can complement or exceed our own abilities: it can work alongside us, and even teach us, as shown by Lee Sedol's unbroken string of victories since playing AlphaGo. This offers new opportunities for creativity and innovation. Perhaps the real productivity gain from artificial intelligence will be in showing us new ways to think.

Government Office for Science 2016 report, Artificial intelligence: opportunities and implications for the future of decision-making[2]
Summary: Artificial intelligence for innovation and productivity

[2] https://www.gov.uk/government/publications/artificial-intelligence-an-overview-for-policy-makers

Why does AI matter?

AI matters because it can lead to major economic and social benefits. It holds great potential for increasing productivity in existing industries - a pressing need for the UK economy - and for creating wholly new products and services.

In one estimate, the worldwide market for AI solutions could be worth more than £30bn by 2024, boosting productivity by up to 30% in some industries, and generating savings of up to 25%.[3] In another estimate, "AI could contribute up to $15.7 trillion[1] to the global economy in 2030, more than the current output of China and India combined. Of this, $6.6 trillion is likely to come from increased productivity and $9.1 trillion is likely to come from consumption-side effects."[4]

The overall estimates can be overwhelming, so it is useful to focus on what AI could do in major business sectors. The excerpt below is PWC's overview on short, medium and long term uses of AI in Healthcare, Automotive and Financial Services (the same report also covers transport and logistics; technology, communications and entertainment; retail; energy; and manufacturing).

The applications, time-frames, benefits and obstacles will be different in different sectors, which makes it difficult to generalise or to reach confident predictions across the economy. But it is evident that AI is widely seen as having enormous potential to improve the functioning of many sectors.

[3] https://www.bofaml.com/content/dam/boamlimages/documents/PDFs/robotics_and_ai_condensed_primer.pdf

[4] PWC: Sizing the prize: what's the real value of AI for your business and how can you capitalise? June 2017. http://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf

Healthcare
Three areas with the biggest AI potential:
- Supporting diagnosis in areas such as detecting small variations from the baseline in patients' health data or comparison with similar patients.
- Early identification of potential pandemics and tracking incidence of the disease to help prevent and contain its spread.
- Imaging diagnostics (radiology, pathology).

Patient benefits: Faster and more accurate diagnoses and more personalised treatment in the short and medium term, which would pave the way for longer term breakthroughs in areas such as intelligent implants. Ultimate benefits are improved health, and lives saved.

Time saved: More effective prevention helps reduce the risk of illness and hospitalisation. In turn, faster detection and diagnosis would allow for earlier intervention.

Timing
- Ready to go: Medical insurance and smarter scheduling (e.g. appointments and operations).
- Medium-term potential: Data-driven diagnostics and virtual drug development.
- Longer-term potential: Robot doctors carrying out diagnosis and treatment.

Barriers to overcome: It would be necessary to address concerns over the privacy and protection of sensitive health data. The complexity of human biology and the need for further technological development also mean than some of the more advanced applications may take time to reach their potential and gain acceptance from patients, healthcare providers and regulators.

High potential use case: AI-powered diagnostics use the patient's unique history as a baseline against which small deviations flag a possible health condition in need of further investigation and treatment. AI is initially likely to be adopted as an aid, rather than replacement, for human physicians. It will augment physicians' diagnoses, but in the process also provide valuable insights for the AI to learn continuously and improve. This continuous interaction between human physicians and the AI-powered diagnostics will enhance the accuracy of the systems and, over time, provide enough confidence for humans to delegate the task entirely to the AI system to operate autonomously.

Automotive
Three areas with the biggest AI potential:
- Autonomous fleets for ride sharing.
- Semi-autonomous features such as driver assist.
- Engine monitoring and predictive, autonomous maintenance.

Consumer benefit: A machine to drive you around and 'on-demand' flexibility – for example a small model to get you through a city or a bigger and more powerful vehicle to go away for the weekend.
Time saved: The average American spends nearly 300 hours a year driving – think what you could with that time if you didn't have to spend it behind the wheel.

Timing
Ready to go: Automated driver assistance systems (e.g. parking assist, lane centring, adaptive cruise control etc.).
Medium-term potential: On-demand parts manufacturing and maintenance.

Longer-term potential: Engine monitoring and predictive, autonomous maintenance.

Barriers to overcome: Technology still needs development – having an autonomous vehicle perform safely under extreme weather conditions might prove more challenging. Even if the technology is in place, it would need to gain consumer trust and regulatory acceptance.

High potential use case: Autonomous fleets for ride sharing. Autonomous fleets would enable travellers to access the vehicle they need at that point, rather than having to make do with what they have or pay for insurance and maintenance on a car that sits in the drive for much of the time. Most of the necessary data is available and technology is advancing. However, businesses still need to win consumer trust.

Financial services
Three areas with the biggest AI potential:
- Personalised financial planning.
- Fraud detection and anti-money laundering.
- Process automation – not just back office functions, but customer facing operations as well.

Consumer benefit: More customised and holistic (e.g. health, wealth and retirement) solutions, which make money work harder (e.g. channelling surplus funds into investment plans) and adapt as consumer needs change (e.g. change in income or new baby).

Timing
Ready to go: Robo-advice, automated insurance underwriting and robotic process automation in areas such as finance and compliance.
Medium-term potential: Optimised product design based on consumer sentiment and preferences.
Longer-term potential: Moving from anticipating what will happen and when in areas such as an insurable loss (predictive analytics) to proactively shaping the outcome (prescriptive analytics) in areas such as reduced accident rates or improved consumer outcomes.

Time saved: The information customers need to fully understand financial position and plan for the future is at their fingertips and adapts to changing circumstances. Businesses can support this by developing customised solutions rather than expecting consumers to sift through multiple options to find the one that's appropriate.

Barriers to overcome: Consumer trust and regulatory acceptance.

High potential use case: Personalised financial planning. While human financial advice is costly and time-consuming, AI developments such as robo-advice have made it possible to develop customised investment solutions for mass market consumers in ways that would, until recently, only have been available to high net worth (HNW) clients. Finances are managed dynamically to match goals (e.g. saving for a mortgage) and optimise client's available funds, as asset managers become augmented and, in some cases, replaced by AI. The technology and data is in place, though customer acceptance would still need to increase to realise the full potential.

PWC: Sizing the prize: what's the real value of AI for your business and how can you capitalise? June 2017

Why does AI matter here and now?

Action to grow AI capability in the UK - in particular on data and skills - matters now because early leaders (businesses and countries) could gain major and lasting advantages from taking a lead in building and using AI.

AI matters here because (for now) the UK is one of a group of countries leading in AI. That advantage could be built on successfully, or it could be lost. Industry outside the UK is taking up AI. UK industry will need to, to compete.

The UK also has a unique position in terms of the history of development of AI. Public support has been highly effective in the past in enabling the UK to make advances in AI.

AI matters now because the technology has matured to the point that it can be applied widely. Even though AI technologies have been in development for decades, and have been in use in some consumer services for several years, the past five years have seen an unprecedented level of interest and investments in AI which has led to a very fast pace of new discoveries and improvements, even by the standards set by previous digital technologies.

The recent upturn in performance has been driven by:
- Increasingly powerful and affordable computing capacity
- New and larger volumes of data
- Novel algorithms and applications developed by experts with very specific skills.

Objectives and scope of this Review

The objective of this Review is to identify the actions required to grow the AI industry in the UK.

In March 2017, the Government announced an industry-led Review of how industry and Government can create the conditions for the AI industry to continue to thrive and grow in the UK, as part of a broader Digital Strategy for the UK economy. This report is also a key contribution to the Government's Industrial Strategy, for which a White Paper will be published later in 2017.[5]

This report is the outcome of that Review. It offers a high level summary of current AI activity in the UK, the challenges to increasing AI activity and recommendations to address those challenges. The recommendations made are complementary and coordinated actions by industry, academia and Government which are anticipated to inform a potential AI sector deal.

Related activity is proceeding in parallel with this Review.

Industrial Digitalisation Review: This Review of AI has run parallel with a review aimed at accelerating advanced digitalisation of UK industry. As explained in more detail below, AI can only be used successfully in sectors which are digitised and making effective use of data, so the reviews are complementary.

Life Sciences: Industrial Strategy: Another parallel review, of LIfe Sciences, has published a strategy to identify ways to establish the UK as the global leader in clinical research and medical innovation.[6] One strand of the review focuses on the potential uses of health data for research, including by AI.

Robotics: AI is driving improvement in robotics and automation. It is perhaps impossible in practice to draw a clear line between AI and physical functions that AI can now manage. This Review focuses on advanced information processing, rather than physical actuation, and does not cover the development of the UK robotics industry. The Industrial Digitalisation Review also covers the benefits of deploying robotics in manufacturing.

Trust, ethics, governance and algorithmic accountability: Resolving ethical and societal questions is beyond the scope and the expertise of this industry-focused review, and could not in any case be resolved in our short time-frame.

However, building public confidence and trust will be vital to successful development of UK AI. Therefore this Review stresses the importance of industry and experts working together to secure and deserve public trust, address public perceptions, gain public confidence, and model how to deliver and demonstrate fair treatment. Fairness will be part of gaining economic

[5] In January 2017, the Industrial Strategy Green Paper invited sectors to come forward with targeted proposals to increase productivity and growth through dialogue with decision makers at the heart of Government via the sector deal process. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/611705/building-our-industrial-strategy-green-paper.pdf

[6] https://www.gov.uk/government/publications/life-sciences-industrial-strategy

benefits, and addressing ethical issues effectively to support wider use of AI could be a source of economic advantage for the UK.

Royal Society and British Academy report - Data management and use: Governance in the 21st century, June 2017: Ethical and governance questions in the broader data economy (of which AI is a part) have been addressed very thoroughly by the recent report from a major project by the Royal Society and British Academy, which sets out principles, essential functions and the role of stewardship.

The overarching principles for governance of data set out by the Royal Society and British Academy could also be - as a default - those for AI in the UK. While AI will generate some specific challenges, it would not be helpful to see AI governance as something unrelated and separate to broader data governance.

It may be appropriate for the stewardship function for data governance suggested by the Royal Society report to include AI expertise, and to oversee AI issues within its remit, rather than having the stewardship of data governance for AI placed in an additional organisation.

Among our recommendations below is the additional suggestion that this data stewardship function could also be an effective partner in developing frameworks for data-sharing for AI, described as data trusts.

Royal Society and British Academy, Data management and use – Governance in the 21st Century June 2017

Principles for Data Governance

A set of high-level principles is needed to visibly shape all forms of data governance and ensure trustworthiness and trust in the management and use of data as a whole.

The promotion of human flourishing is the overarching principle that should guide the development of systems of data governance. The four principles that follow provide practical support for this overarching principle across the varied ways data is managed and used:

- protect individual and collective rights and interests
- ensure that trade-offs affected by data management and data use are made transparently, accountably and inclusively
- seek out good practices and learn from success and failure
- enhance existing democratic governance

Essential functions and stewardship

The governance framework for data management and data use should perform three broad categories of functions. These may be carried out by a variety of public and private actors:

- **A**nticipate, monitor and evaluate
- **B**uild practices and set standards
- **C**larify, enforce and remedy

Despite the range of actors already carrying out some of these important governance functions in their specific sectors or domains, there is a clear need **for a new body** to steward the landscape as a whole, rather than being directly responsible for implementation within specific domains.

- The purpose of such a stewardship body would be to support delivery of the full breadth of critical functions in accordance with the principles set out above.
- We expect that such a body would primarily recommend actions to others, but it may also need the capacity to carry out some functions itself if they could not be performed elsewhere, being careful to not duplicate existing efforts.
- This stewardship body would be expected to conduct inclusive dialogue and expert investigation into novel questions and issues, and to enable new ways to anticipate the future consequences of today's decisions.

The characteristics of such a stewardship body are that it should be:

- Independent
- Deeply connected to diverse communities
- Expert across and beyond disciplines
- Tightly coupled to decision processes
- Durable and visible
- Nationally focused but globally relevant

AI and work: The changes which AI will make to the nature of work in some industry sectors are also out of scope for this Review. This by no means implies that this Review does not consider those important. Helping people adapt to changes in work will be necessary to protect

people and share the benefits of AI, to gain public acceptance of AI, and to ensure employers can fill the new, different roles that will be needed for working with AI. Industry, Government, business and the workforce will all need to respond to changes resulting from automation of processes, including by providing and taking up opportunities for retraining.

Understanding of the broad potential impacts of this phase of automation is at an early stage, but is developing.[7] This will be a complex, varied and evolving process, which a short-term review cannot predict or direct. Government and employers should actively monitor these developments, in particular where they see opportunities for AI to offer improvements to their functions and services.

The Review process

Over 100 experts from academia, industry and Government, were involved in this Review through a series of workshops and meetings used to inform, develop, refine and test these recommendations. This report also makes use of a wide range of industry and expert sources of evidence and insight, published and unpublished.

In April 2017, the Royal Society published a report on Machine Learning, the culmination of a major project to investigate the potential of Machine Learning over the next 5-10 years, and the barriers to realising that potential.[8] This Review builds on that work and complements it by focusing on concrete actions that UK industry, academia and Government should take now. As mentioned above, it also sits alongside the joint Royal Society and the British Academy report on Data Governance.

The review has considered reports by companies and analysts, by the World Bank, IEEE, World Economic Forum, and by other countries that are supporting AI, including US, France, China, Singapore, companies, in particular the October 2016 White House report “Preparing for the Future of Artificial Intelligence”.[9]

[7] For instance see “The future of the professions: how technology will transform the work of human experts" by Richard Susskind and Daniel Susskind, 2015.

[8] Machine Learning: the power and promise of machines that learn by example, Royal Society April 2017 https://royalsociety.org/~/media/policy/projects/machine-learning/publications/machine-learning-report.pdf

[9] https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

AI in the UK: a short history

The pioneering British computer scientist Alan Turing is widely credited with launching and inspiring much of the development of AI, and the philosophy of AI, with his 1950 paper “Computing Machinery and Intelligence”. Turing had already developed the principle of the modern computer in 1936, and played a critical role in breaking ciphers at Bletchley Park during the Second World War. In the 1950 paper, Turing explored what is meant by ‘machines’ and by ‘thinking’, and in what became known as the Turing Test, he proposed that if a machine could conduct a conversation in print that could not be distinguished from conversation with a person, it could be said that the machine was “thinking”. As his earlier work on the foundations of computing, delivered to the London Mathematical Society, had shown that all digital computers were in effect equivalent (i.e. any computer can simulate the behaviour of all other computers given enough memory and time), this thought experiment expressed a very powerful, elegant and precise concept. The paper is still widely read, discussed, cited and anthologised today.

Early pioneers in AI focused on the development of the necessary tools and techniques to explore Turing’s idea. Early approaches focused on symbolic programming (i.e. programs that can manipulate expressions in their own programming languages), as the most promising paradigm available. Many special purpose languages were written as part of that drive, most famously LISP in the US, but also including significant contributions from Britain, such as POP-2 (created by Robin Popplestone and Rod Burstall at the University of Edinburgh) and Edinburgh Prolog (by David H.D. Warren, also at Edinburgh).

In 1952 Christopher Strachey used the Ferranti Mark 1 system at the University of Manchester to write a programme to play draughts, and later programmed the generation of love letters. The performance of AI in increasingly more complex games has been an indicator of progress ever since.

Another former Bletchley codebreaker was Donald Michie, later director of the Department of Machine Intelligence and Perception at Edinburgh. His noughts-and-crosses-playing program MENACE was too complex for the computers available to him at the time, and he initially implemented it with 300 matchboxes!

By the 1960s, AI techniques were being applied to far more complex problems with more practical applications. Planning involves developing strategies for problem-solving that produce a series of actions that will approach a goal; example applications include automated reasoning, or planning proofs, as pioneered by Alan Bundy.

Understanding natural language was another important strand; for instance Karen Spärck Jones developed ways of retrieving information from documents, and Yorick Wilks’ preference semantics was a computational approach to disambiguating word senses, which not only contributed to AI but directly challenged the dominant Chomskyan paradigm in linguistics. Both of these were alumni of the Cambridge Language Research Unit, a legendary crucible of computational linguistics founded by Wittgenstein’s student Margaret Masterman.

In subsequent developments, robotics systems, such as Edinburgh’s Freddy I and Freddy II, were able to combine vision, intelligence, versatility and physical engineering to perform tasks such as assembling objects (special-purpose AI languages needed to be developed for robotics). AI systems were also influential on the discipline of cognitive psychology, as researchers including Richard

Gregory, Christopher Longuet-Higgins, Philip Johnson-Laird and David Marr realised than many human cognitive processes could be seen as a type of computation, and be modelled as computer programs.

Globally and in the UK, AI has gone through periods of development and periods of relative stagnation (often referred to as "AI winters"). One major such event followed the publication in 1973 of Sir James Lighthill's report on AI, which recommended concentrating AI funding in a smaller group of British universities. Lighthill was sceptical about the ability of AI at the time to scale up to solve the complexity of real-world problems, and indeed the dominant approach of the 1960s, to model complex reasoning as a search through a tree of possible decisions, was vulnerable to the problem of combinatorial explosion.

However, in the longer run, the advances in symbolic programming enabled greater understanding of high-level problem-solving intelligence, with especial progress in tools and techniques to simulate or support complex expert reasoning in relatively well-structured domains – ideal for applications in the workplace.

So-called Knowledge-Based Systems (KBS) combined AI techniques with other kinds of computing inference and domain-relative expertise to create systems for often quite mundane but important real-world applications. The unspectacular but practical success of KBSs helped defuse Lighthill's pessimism, and paved the way for a productive expansion of funding with the Alvey Programme. In retrospect, the AI winters that we have seen have been products of excessive hype – overclaiming by boosters leading to an erroneous impression of failure, and consequent undervaluing of the important but unsung successes of the research.

The UK's Alvey Intelligent Knowledge-Based Systems (IKBS) Programme, which ran from 1983 to 1987 was developed in response to progress in other countries, in particular Japan (whose 5th Generation Project rested upon techniques and languages, particularly Edinburgh Prolog, originating from the UK). Alvey influenced development of academic research and research capability but also encouraged industry applications, focusing on the practical problems in which progress had been made, notably natural language processing, interfaces, and KBSs.

These applications gradually coaxed the field of AI away from the idea of producing 'machines that think' (a concept that has always been philosophically problematic), towards the more measurable idea of creating machines that can produce performance that would certainly be ascribed to intelligence if produced by a human (an idea implicit in the Turing Test). Such intelligent performance might be produced by 'brute force' methods that neither mirrored, nor attempted to mirror, human problem-solving. Interestingly, the UK has produced many important philosophers who have helped uncover the concepts behind such distinctions, including, for example, Margaret Boden and Andy Clark.

Post-Alvey, AI funding dipped once more, but the promise of the field was already on an upturn, as new methods of programming, which did not rely on linear combinations of symbolic inferences, became feasible. Whereas symbolic programming is the easiest type of programming for a human to understand, there are also great gains to be had from simulations of natural techniques of inferring information from the perceived environment (e.g. information from the senses), that are often called sub-symbolic, because they do not include direct representations of declarative or propositional knowledge.

One example of inspiration from nature is the genetic algorithm, which encodes a program as a set of 'genes', and then modifies them in a way that mimics evolution, looking for a 'fit' with an ever-changing environment (pioneering work here includes Richard Forsyth's BEAGLE system for pattern

recognition). Another is neural nets or connectionist systems, in which artificial ‘neurons’ are connected in a system intended to work like the human brain, with the ‘neurons’ stimulating or inhibiting each other. As with symbolic AI, researchers have often rowed back from the aim of copying the human brain to improve performance (for instance, with techniques such as back-propagation, developed by Geoffrey Hinton), but the large-scale neural net SpiNNaker (2005-), led by Steve Furber, remains in the tradition of direct modelling of the brain. Other non-traditional computing approaches with relevance to AI include parallel processing (using multiple processors in parallel to solve problems), multi-agent systems (where many intelligent autonomous agents interact within an environment), and machine learning (algorithms which can learn to find significant structure in data, given training in identifying the interesting patterns).

Other countries and international companies are investing heavily in AI development, but the UK is still regarded as a centre of expertise in research and application of AI, for the present at least. Two of the founders of DeepMind, for instance, met as PhD students in University College, London’s Computational Neuroscience Unit, whose founding director is Hinton. The UK can continue to build on the legacy of Turing and those who have followed him, to remain one of the great centres for AI.

Dr Kieron O'Hara, associate professor and principal research fellow in Electronics and Computer Science at the University of Southampton.

SECTION TWO: OVERVIEW OF CURRENT USE OF AI, MARKETS AND SUPPORT FOR GROWTH

Fundamental factors:
- The digital context
- Hardware
- Data

AI business activity in the UK
- Global tech companies
- AI startups
- Geographical spread
- AI in established industries
- AI in the public sector
- Estimated economic impacts

UK support for development of AI

International context
- Investment
- Comparisons
- Government support

The digital context

AI is the next stage of digital evolution, adding new capabilities to digital technologies already in use, in vertical digital technology sectors and in digitised sectors horizontally across the economy. It may in the future contribute to every sector that makes (or could make) substantial use of data.

UK AI is developing on the foundation of the UK's existing technology capability, which is comparatively strong by international standards. The turnover of the digital tech sector was estimated at £170 billion in 2015, representing a growth rate of 22% over the previous five years. The UK now has 1.64 million digital tech jobs, and the growth rate of that employment market was more than double that of non-digital jobs between 2011 and 2015.

Cybersecurity is a good example of an established digital sector that will see an improvement in performance with greater use of AI. A large number of organisations face cyber threats every day. Machine learning can identify, categorize and analyse these more effectively than individual researchers. By working simultaneously on different tasks, across a large number of devices and systems, AI can help defend against large attacks. Automating some cybersecurity functions can help identify anomalous behaviour more quickly, highlight areas of concern that

can be followed up by human network engineers, and identify and patch network weaknesses before they are exploited.

Parallel AI applications are already enhancing other digital and digitised sectors.

The capability to add AI into digital functions across sectors has been supported by improvements in hardware, and by greatly increased volumes of data.

Hardware

Development and use of AI have been accelerated internationally by continued and competitive improvements in the capability and availability of hardware.

Central Processing Units (CPUs) were the standard for interpreting and executing commands in servers, tablets, computers and mobile phones. More recently the development of machine learning and deep learning has been boosted by the use of Graphics Processing Units, which have the ability to perform many calculations simultaneously, or in parallel, speeding up training processes. Google has developed the Tensor Processing Unit (TPU), a custom made chip for machine learning and has announced plans for further increases in chip capability, as has the GPU developer Nvidia, recently named the Smartest Company in the World by MIT.[10] Apple is reported to be working on a dedicated chip for AI on devices.[11]

As a result of this continual, competitive development, the cost of using high performance computing has fallen significantly, and continues to fall, making it available to a growing range of users. The plans announced by market leaders, and reports by expert analysts, suggest that this trend will continue.

Data for developing AI, AI for managing data

Since 2000 there has been an exponential increase in the quantity of data generated globally, much of it from internet and mobile personal devices. Parallel technologies including the internet of things are also contributing to the strong upward curve in the volumes of data generated. This trend is projected to continue. Cisco estimates that globally, mobile data traffic will increase sevenfold between 2016 and 2021.[12]

Developments in data economies are now powerful factors driving changes in national and global economies. According to a recent Government policy paper: “Our data economy will be

[10] https://www.technologyreview.com/lists/companies/2017/intro/ - nvidia
https://www.technologyreview.com/s/607879/google-reveals-a-powerful-new-ai-chip-and-supercomputer/
https://www.forbes.com/sites/moorinsights/2017/05/15/why-nvidia-is-building-its-own-tpu/ - 19004918347f
[11] https://www.bloomberg.com/news/articles/2017-05-26/apple-said-to-plan-dedicated-chip-to-power-ai-on-devices
[12] http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/vni-hyperconnectivity-wp.html

integral to the UK's growth and future prosperity. Analysis predicts that data could benefit the UK economy by up to £241 billion between 2015 and 2020." [13]

This vast, fast increase in data has also nurtured AI. Access to both bulk and specific data is key to successful training of machine learning algorithms. As the Royal Society's Machine Learning report explains in more detail, exposing technologies to larger datasets improves methodologies and drives continually improving outcomes over time.[14]

To use AI in a sector, it is necessary to train AI with data relevant to that sector. Without ample relevant and good quality data, AI technologies cannot develop. The increased availability of training data has made it easier to improve the accuracy of AI algorithms, and enabled it to operate in more sectors.

Increasing data flows also make AI more necessary: data flows in some sectors are now so great that only AIs have the capacity to deal with the volume and complexity.

Many organisations - public and private - hold significant quantities of data. As more functions are digitised, much more data will be generated than has been in the past.

However, there are many reasons why organisations find it difficult or impossible to share data externally, including privacy, security, commercial advantage and other considerations. Even where organisations see a case to proceed with secure data-sharing, and trust the external organisations they plan to share with, they often lack expertise and know-how to form agreements, establish trust between the parties, and manage the data sharing processes in practice.

This is explored in more detail in the first set of recommendations below.

AI business activity in the UK

The UK has AI companies that are seen as some of the world's most innovative, in an ecosystem that includes large corporate users of AI, providers large and small, business customers for AI services, and research experts. Competition for talent and investment is global, so it is useful to see UK activity in that global context.

All the major global tech companies active in the UK are developing and using AI. Some startups have been acquired by these majors, and it is likely that more will be. The majors use various routes to building expertise, for example: "Companies have adopted M&A as a way to

[13] Data - unlocking the power of data in the UK economy and improving public confidence in its use. DCMS policy paper March 2017. https://www.gov.uk/government/publications/uk-digital-strategy/7-data-unlocking-the-power-of-data-in-the-uk-economy-and-improving-public-confidence-in-its-use

[14] Chapter 3. Machine Learning: the power and promise of machines that learn by example, Royal Society April 2017 https://royalsociety.org/~/media/policy/projects/machine-learning/publications/machine-learning-report.pdf

sign up top talent, a practice known as “acqui-hiring,” for sums that typically work out to $5 million to $10 million per person.”[15]

IBM and Microsoft offer a range of AI services to business customers for key functions (predictive analytics, computer vision, language, customer service, news discovery, management of internet of things applications) and in key target sectors (financial services, health).

It is estimated that there are more than 200 startups and small and medium sized businesses developing AI products in the UK.[16] AI startups have been founded to address specific areas of industry sectors (**Your.MD** in personal health) and major technology challenges (**Darktrace** in cybersecurity).

Large established companies outside the tech sector are using AI to deliver more efficient operations and services (Ocado, GE). Public sector organisations, for example HMRC, are also using (or exploring using) AI to optimise services.

AI covers a set of rapidly evolving, complementary general purpose technologies, applied across many sectors.

Because of this range of activity and organisations, It is challenging to make a clear, current and accurate assessment - quantitative or qualitative - of AI in the UK. The rate of change is fast, uneven between businesses and sectors, and hard to measure. There are no absolute recognised distinctions between AI technologies and other big data and data science applications in use. “Artificial Intelligence isn’t a matter of any single technology or application—whether driverless cars or smartphone virtual assistants or trend detection solutions or a myriad of other examples. Artificial Intelligence is a rich and diverse field.”[17]

Where AI is in use, it is often integrated into other digital functions and cannot entirely be separated from those. Much AI activity is internal within organisations and difficult to assess from the outside. Much of it happens in international tech companies, and it is not fully clear what it is developed in the UK, and what is only used here, given that globally-sourced data and distributed innovation teams feed continual development across these companies.

No single company’s AI activity is representative. The range of business users is mixed and going to become much more mixed. The spectrum already reaches from consultant experts to global majors, and very new AI startups to very long-established businesses enhancing core functions with AI. Some figures on company formation and investment are available, and help identify those AI-first companies, but give no insights into the latter group. Larger, older

[15] Artificial Intelligence: the next digital frontier. June 2017. http://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/how-artificial-intelligence-can-deliver-real-value-to-companies

[16] https://medium.com/mmc-writes/artificial-intelligence-in-the-uk-landscape-and-learnings-from-226-startups-70b9551f3e4c#.n8898das9

[17] Accenture, Turning Artificial Intelligence into business value today, 2016. https://www.accenture.com/t20160814T215045__w__/us-en/_acnmedia/Accenture/Conversion-Assets/DotCom/Documents/Global/PDF/Technology_11/Accenture-Turning-Artificial-Intelligence-into-Business-Value.pdf

businesses are acquiring newer ones. More new ones will appear. Some small companies provide AI services to established companies. Some large established companies provide AI as a service. Over time, more established companies big and small, using AI through a variety of models of business and contractual arrangements (eg. in-house, off-the-shelf, bespoke, licensing), and of products and services.

Therefore, what follows is an overview of available analysis with illustrative examples of company activity, not a comprehensive account of all UK AI activity.

The variety of activity does make assessment difficult, but it should be seen as a strength of AI. Realising the future economic value of AI will rest on the maximum range and mix of organisations being able to make informed choices about how to use it, and being ready to use it.

AI in global technology companies: The major US-founded global digital companies are using AI now in combination with their core business, and some of this AI is being developed in the UK.

Globally, the US tech majors appear to be the majority investors in AI, though exactly how much and where geographically is not fully clear from the outside. McKinsey: "Globally, we estimate tech giants spent $20 billion to $30 billion on AI in 2016, with 90 %of this spent on R&D and deployment, and 10 per cent on AI acquisitions." In this analysis, the acquisitions, including of leading UK AI companies, which gain media and public attention and are rightly seen as indicating appetite for AI, only represent a relatively small fraction of their overall investment in AI.

IBM is a long term pioneer leader in AI development and in provision of AI as a service. Microsoft has been investing in AI for more than 25 years and has added it into several key offerings. The majors have diversified AI interests. In December 2016 its VC arm announced an investment fund dedicated to AI startups that focus on "inclusive growth and positive impact on society."[18]

Facebook, Google, Amazon, Apple, Microsoft and Baidu all use AI to develop their principal services, using the rich, continuous data streams from user interactions continually to train AIs to improve performance in face recognition, language interactions (Siri, Alexa, Cortana etc), and customer service. Cisco, Samsung and Huawei are all using AI to develop their core products.

Major software providers have added AI applications to their suites of services to industry. SAP is developing AI services to automate employee approvals, payment processing, and sales discounting. Sage has launched an AI-powered "virtual accounting assistant" chatbot to submit expenses, and track receipts and payment of invoices, integrated with messaging apps.

[18] https://news.microsoft.com/features/microsofts-ai-vision-rooted-in-research-conversations/

The addition of AI to everyday services from global companies means that consumers and business users are already using AI services without knowing it. AI-driven functions are not necessarily visible from the outside.

To date, when US majors have acquired UK AI companies, the companies and their expertise have largely stayed in the UK. This is encouraging, but cannot be guaranteed in the future. The acquiring companies are global, and can move assets to where those develop and work most effectively.

Attracting and retaining the investment and expertise of the global majors is a key part of making the UK the best environment for developing AI .

UK AI Startups: The UK has produced a number of very innovative AI companies, and companies are being formed frequently. According to a 2017 Coadec report "a new AI startup has been founded in the UK on almost a weekly basis in the past 36 months".[19] In December 2016 one study estimated that there are 226 independent, early stage AI companies in the UK.[20]

Some companies that identify themselves as AI specialists (**Swiftkey**, **DeepMind** and **Ravn**). Some of these companies have been acquired and now operate within larger global players, for instance DeepMind within Google / Alphabet. Some UK AI companies have focused on a single sector. A number are working on key challenges with the NHS.

In the UK, TechCity estimated that AI received 3% of investment in digital tech in 2016, and rising. It appears to be one of the fastest growing parts of the digital sector.[21] One in 10 applications to Startupbootcamp, a leading Financial Services focused accelerator based in London, look to exploit AI technologies.[22] However, Coadec points out that only 1 in 10 UK AI companies is in late "growth" capital stage compared to 1 in 5 in the US.[23]

[19] http://coadec.com/Coadec-Report-A-Global-Britain.pdf
[20] https://medium.com/mmc-writes/artificial-intelligence-in-the-uk-landscape-and-learnings-from-226-startups-70b9551f3e4c
[21] TechCity UK, 2017. Tech Nation 2017: At the forefront of global digital innovation. http://technation.techcityuk.com/
[22] https://www.pwc.co.uk/financial-services/start-up-bootcamp/SBC-PwC-FinTech-Trends-report-16-Final-Online-v2.pdf
[23] http://coadec.com/Coadec-Report-A-Global-Britain.pdf

AI companies in key UK Sectors

Healthcare / Life Sciences: Seen as one of the most important sectors for AI both for better services and for better efficiency. **BenevolentAI** uses AI to speed up the drug delivery process.[24] **Babylon Health** works with the NHS on trialling chatbot advice to patients via mobile. It recently raised $60m worth of funding.[25] **Your.MD**'s Personal Health Assistant is a free chatbot offering personalised and accessible healthcare advice.

Digital Marketing: The UK has AI companies across many aspects of digital marketing, sales and business development functions. **AdBrain**'s customer ID mapping platform enables marketers to target and track an individual consumer across different devices, channels and platforms, to deliver better marketing results. **Pixoneye** offers AI based image and feature analysis of mobile phone pictures to help clients better segment their customer base. **Attest Technologies** applies AI to market research, **Growth Intel** to business intelligence and business development. **Decibel Insight** specialises in web analytics.

Automotive: Connected cars and autonomous vehicles offer very large future growth opportunities for AI companies. Bristol based **FiveAi** works on safe autonomous driving. Oxford's **Oxbotica** has developed an autonomous operating system. **Selenium** uses data from lasers and cameras placed on vehicles for autonomous navigation.

Identity: The nascent RegTech (Regulatory Technology) sector has companies using AI for checking identity. **Onfido** uses machine learning to conduct global background checks for companies.

Financial Services: Behavioural analysis of financial transactions with AI can dramatically improve controls to spot fraudulent transactions.[26] Chatbots using Intelligent Voice Recording can handle telephone requests from customers.[27] **HSBC** has launched a chatbot, Olivia, for verifying customer identity in order to increase security via an individual "voiceprint".[28] In May 2017 it was reported that by 2019, exams for the **Chartered Financial Analyst Institute** would include questions on AI, robo-advisory services, and methods for analysing big data.[29]

Legaltech: AI already helps lawyers to do legal search, to identify the best standard documents, to review documents, and automatically draft documents. Pinsent Masons' **TermFrame** system identifies precedents and templates. **MarginMatrix**, a joint venture between Allen & Overy and Deloitte, auto-drafts legal documents to help banks comply with new financial regulation, reportedly cutting down drafting time from hours to minutes.[30]

Education: AI can improve the effectiveness of education, for example by assessing the success of online learning, and can support better personalisation.[31] **Gradescope** provides autograding for teachers.[32]

[24] https://www.siliconrepublic.com/start-ups/ai-startups-europe
[25] https://www.ft.com/content/1f56997a-290f-11e7-bc4b-5528796fe35c
[26] http://www.wired.co.uk/article/how-ai-is-transforming-the-future-of-fintech
[27] http://thefinanser.com/2017/01/11-fintech-trends-need-follow.html/
[28] https://www.hsbc.co.uk/1/2/voice-id
[29] https://www.bloomberg.com/news/articles/2017-05-09/cfa-exam-to-include-big-data-artificial-intelligence-as-topics
[30] https://www.ft.com/content/5d96dd72-83eb-11e6-8897-2359a58ac7a5
[31] https://edtechdigest.wordpress.com/2017/04/24/how-artificial-intelligence-will-transform-education/

As stated above, there have been a number of high-value acquisitions of UK AI companies by US-based global tech companies.[33] Evi in 2012; DeepMind in 2014; VocalIQ in 2015; SwiftKey and Magic Pony in 2016.

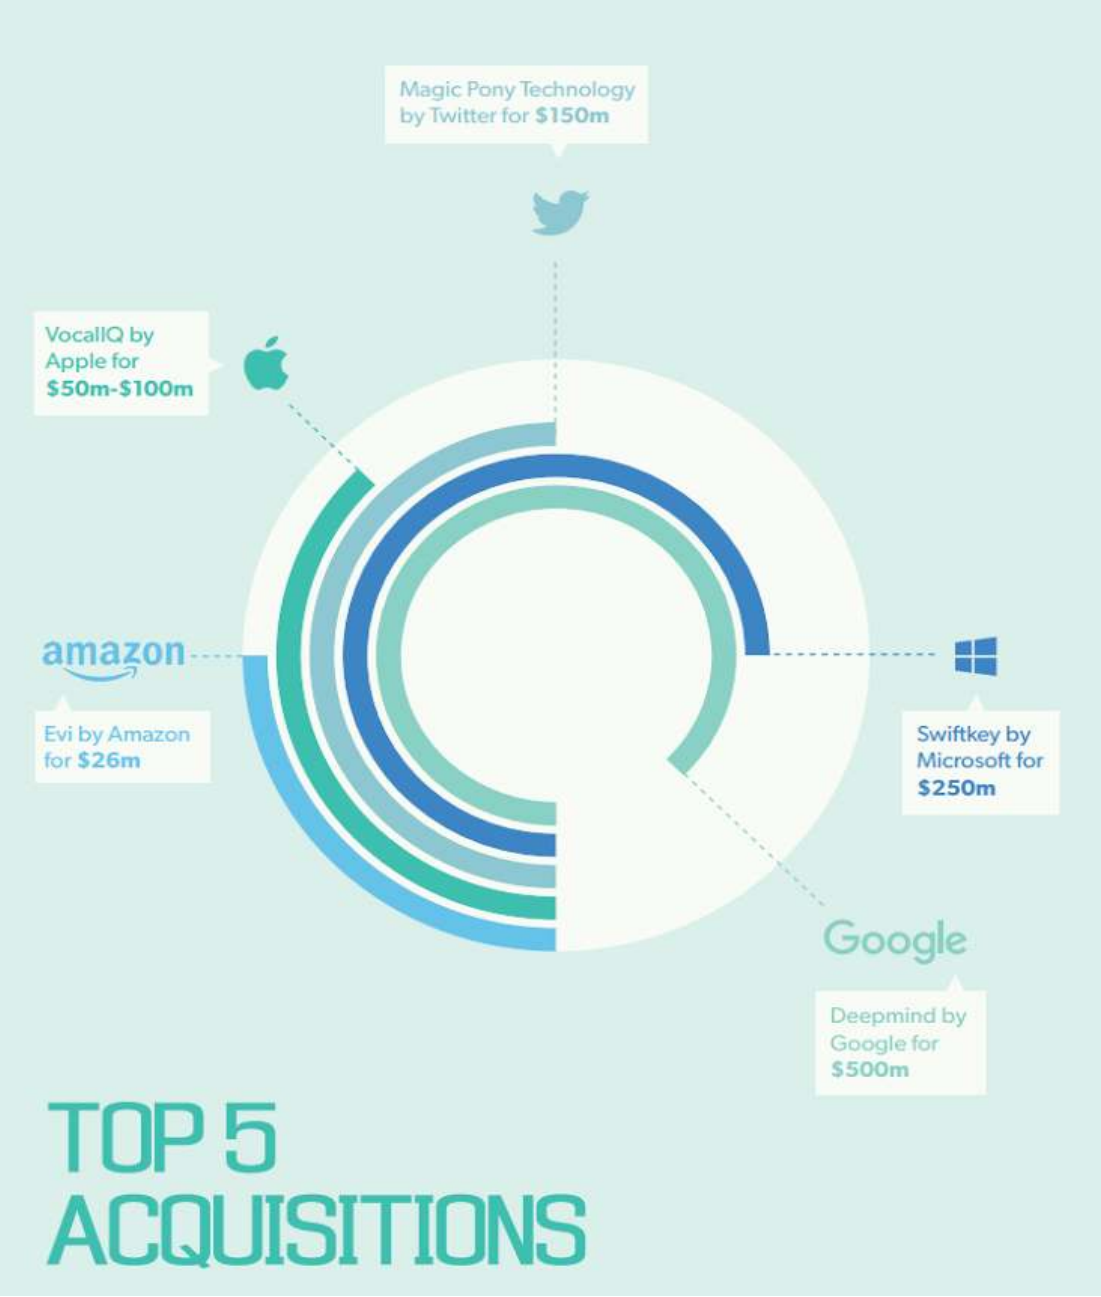

[32] https://blogs.nvidia.com/blog/2016/09/02/gradescope-brings-ai-to-grading/
[33] http://www.sonovate.com/quickview/50-hottest-uk-ai-companies

Geographical Spread of AI Companies in UK

While the majority of AI companies are based in London, there are a number of geographical clusters around the UK. The map below developed for the Open Data Institute and the Digital Catapult illustrates clustering of activity in AI based on tech events, scientific publications, and data such as local skills, business startup rates, and research and development spending.[34]

London: London is the strongest centre of AI startups and SMEs in the UK. 80% of AI companies on the UK Top 50 list[35] are based in the capital and many international companies have an AI presence as well (e.g. **DeepMind**, **Adbrain** and **BenevolentAI**). UCL, Kings and Imperial have significant AI and machine learning research groups which strengthen the cluster. **Entrepreneur First** brings together engineers and computer scientists to build companies, and has a strong focus on AI. **Cognition X** is a community market intelligence platform that provides information and research on products and resources required to build AI solutions. It provides a daily newsletter, frequent topical events, industry research and talent services. Since starting in 2016 it has run more than 40 events, has 7000 members and lists 10000 organisations supporting the AI market.

Cambridge: A range of AI startups including **Evi**, **Vocal IQ**, **Cytora**, **SwiftKey** and **Darktrace** have been created in Cambridge, often with a direct link to the University of Cambridge's Computer School and with support from local investors like Amadeus Capital. International tech companies including Amazon and Apple have an AI presence in the region too.

[34] http://imactivate.com/clusters/?options=true&datagroup=Artificial Intelligence&location=null
[35] http://www.sonovate.com/quickview/50-hottest-uk-ai-companies

Edinburgh: The University of Edinburgh has a successful track record of fostering spin-out companies in data analytics and AI, for example **Skyscanner**. Amazon has set-up a development centre which has a strong focus on machine learning. **CodeBase** is the UK's largest startup incubator, working with more than 80 of the country's most promising technology companies.

Oxford: Oxford University is a renowned centre for machine learning and deep learning and has seen the successful spin-out of companies including Dark Blue Labs and DeepMind[36].

Bristol: Bristol has large tech companies including HP, Oracle and BAE Systems and young companies working in AI. **Five AI** develops software for autonomous vehicles. **Graphcore** puts AI in low power consumer devices. The University of Bristol's Intelligent Systems Lab and internationally recognised Robotics Lab[37] develop AI talent.

AI in established industries

Uptake of AI varies significantly by sector and within sectors. Businesses and sectors that have digitised operations and services can take up AI more easily and effectively than those that have not. In particular, organisations that have good data capability (collection, retention, curation, analysis, protection) have a head-start in becoming AI-ready.

Some businesses and some sectors have taken more of a lead than others. Many of the major accountancy and law firms is have grasped the initiative, are undertaking R&D and engaging with research developments, connecting with the academy, and experimenting with AI for services and operations.

Ocado is an illustrative example of a company that has uses AI both to enhance essential functions, and to develop innovative ones, below.

How Ocado uses AI

Ocado makes significant and fast growing use of AI across their e-commerce, fulfilment and logistics platform. Applications fall into four main categories:

- Predictive analytics
- Monitoring and oversight
- Managing complexity
- Real time optimisation

Specific applications include:

- Personalisation, tailoring services to fit the needs of individual customers
- Adaptive user interfaces which can respond to whether the customer is new or experienced,

[36] https://www.theguardian.com/technology/2014/oct/23/google-uk-artificial-intelligence-startups-machine-learning-dark-blue-labs-vision-factory

[37] http://www.brl.ac.uk

their different shopping styles and needs

- Instant Order that looks at a customer's historical purchasing patterns and predicts what repeat needs, enabling a (typically) 50 item order with one click
- Processing unstructured business data including voice, emails and social
- Forecasting the demand for each of the 50,000 different grocery items
- Monitoring and optimising delivery routes
- In first generation automated warehouses, managing the flows of thousands of crates to avoid congestion and respond to any hardware failures
- In new warehouses, optimising the routing of swarms of thousands of robots and the storage they are accessing
- Streaming the data exhaust from these thousands of robots to the cloud where machine learning based analytics provide oversight. For example, spotting that a robot needs servicing or that it's performance has changed in some way
- Swarm based learning, where one robot learns something that it can share with other members of the swarm (as will happen with driverless vehicles)
- Driverless vehicle trials in Greenwich in partnership with Oxbotica. Obviously this makes massive use of machine learning
- Ocado is the project coordinator for a Horizon 2020 funded project to build a humanoid maintenance robot call SecondHands which will be subjected it to a series of DARPA style challenges. It will learn by observing human engineers at work and discover how to help them with their tasks, keep them safe, do things they cannot do.

AI is increasingly being applied to improve management of infrastructure. The National Infrastructure Commission is currently conducting research on technologies including AI to improve use and maintenance of existing and future infrastructure.[38]

AI is seen as key to the next stage of improvement of quality and efficiency in manufacturing: "Using digital technologies such as artificial intelligence, sensor technologies and automation makes companies more agile and better equipped to respond to, or even act ahead of, changing consumer demands, supplier conditions and technology availability. And in today's world, agile makes competitive."[39] Though there may be scope for many more manufacturing applications. "In the Manufacturing sector, for example, there are few startups to address a substantial need. Machine learning has the potential to unlock 20% more production capacity through predictive, optimised maintenance of machines."[40]

Some common applications could be taken up across most categories of large organisation, public, private and utilities, for instance:

- More responsive and adaptive scheduling of transport
- Better predictive maintenance of infrastructure

[38] https://www.gov.uk/government/publications/national-infrastructure-commission-technology-study

[39] Professor Sam Turner, CTO at HVM Catapult https://hvm.catapult.org.uk/news-events-gallery/news/ati-guest-blog-digital-engineering-manufacturing-leadership/

[40] https://medium.com/mmc-writes/artificial-intelligence-in-the-uk-landscape-and-learnings-from-226-startups-70b9551f3e4c

- Better predictive rostering and responsive scheduling of staff resources
- Better demand management in networks
- Management of pay, payments and invoices.

Survey of 160 business on AI-readiness, for CBI and IBM, May 2017[41]

"Nearly half of firms believe the current wave of Artificial Intelligence will be transformational and widespread; fundamentally transforming the industry and markets they work in. However, only a third feel their business has the skills to adopt data-driven technologies, so the UK must act quickly to bridge the knowledge gap as international competition heats up.

- Artificial Intelligence (AI) investment is gaining momentum with 42% of companies planning to invest over the next five years (Internet of Things – 42% / Advanced Analytics – 51%). That's on top of the one in five companies (21%) who already invested in AI during the past 12 months
- The AI wave will be transformational and widespread, half (49%) of business believe it will fundamentally transform their industry/ market.
- Firms believe AI can lead to greater efficiency (78%) and profits (68%). Benefits to consumers (71%) are big leading theme but also collaboration (70%).
- Yet many businesses risk being left behind, 52% of pioneers have already invested in AI while 47% of 'followers' do not plan to invest in AI at all.

AI in the Public sector

AI should prove a very effective tool in addressing complex public sector challenges, and improving efficiency in mainstream public services.

AI is likely to have wide application in the processing of applications and submissions, including for tax, benefits, visas, passports, and other Government licences. Many of these applications from citizens can be processed more quickly, and abuse of the system spotted more accurately, using AI. A common machine learning algorithm could be used to triage the risk associated with different applications.

The very low risk applications can be processed automatically, leaving caseworkers to focus on the more difficult applications and on applicants who need more assistance. The time it takes applications to be processed would fall, the associated costs would also fall, and the system would become more secure. There will be technical and policy challenges to overcome in relation to specific uses, as with any operational innovation, but the outcome would be a smarter and more responsive Government.

Successful uptake of AI in the public sector will depend on many of the same things as successful uptake in established private sector businesses: a good basis of digitisation already,

[41] May 2017. http://www.cbi.org.uk/insight-and-analysis/adopting-the-future/

leaders who can understand opportunities and authorise use of AI (including authorising use of privileged data), and the digital and data skills and capability to deliver.

UK Government departments and agencies are already developing and deploying AI applications. The Government Digital Service (GDS) uses machine learning to help automate and process user comments from surveys on gov.uk, and for predicting peak traffic demands to the most popular content searched for by the public. GDS works with the Pensions Regulator to improve efficiency using predictive algorithms for future pension scheme behaviour and HMRC uses AI to help identify call centre priorities. There are plans for the Digital Government Partnership to experiment with machine learning applications in Government.

The Data Science Campus at the Office for National Statistics acts as a hub, bringing together data and digital expertise and leadership. The Campus aims to gain practical advantage from the increased investment in data science capability, and help cement the UK's reputation as an international leader in this field. The Data Science Accelerator and Government Data Science Partnership train data scientists across government in advanced data analytics including machine learning techniques to gain new insights into live departmental services and processes. Knowledge and best practice relating to use of AI to tackle policy and operational challenges are widely shared across the government data science community through conferences such as the recent Government Data Science Conference. ONS has an MSc in Data Analytics for Government and Data Science apprenticeships.[42] The Data Science Advisory Board will work to align cross-Government efforts to leverage the potential of data science, with a particular focus on its value as an input to broader policy making processes.

The Cabinet Office Government Commercial Function is running a procurement process to bring in a strategic partner to help promote Robotic Process Automation (RPA) and accelerate uptake. This will be a vehicle through which Departments can identify, develop and purchase RPA solutions. The Cabinet Office is also exploring how central commercial arrangements and government standards could support use of robots and AI.

There are also existing channels for collaboration on innovation in addressing public needs, notably the Small Business Research Initiative (SBRI).

Some other countries are also using AI in the public sector, including in: chatbots in Singapore and Australia; back office services in Finland and Japan; pattern detection in local law enforcement in the United States; and the processing of grant applications in Denmark.

Estimated future impacts of AI on the UK economy

The majority view is that impacts will be positive, large, and widely spread across sectors, with uneven rates of uptake. The fact that AI could be used in most industry sectors over time makes

[42] https://datasciencecampus.ons.gov.uk/2017/03/28/how-are-we-building-data-analytical-skills-locally-and-internationally/
https://www.ons.gov.uk/news/news/onsrecruitingtheuksfirstdataanalyticsapprentices

overall numbers hard to reach, and leads to differences between different assessments. There are also various views about the pace of change.
As above, Accenture has estimated that AI could add an additional USD $814 billion to the UK economy by 2035, increasing the annual growth rate of GVA from 2.5 to 3.9 per cent.[43] PWC's recent analysis of UK impacts is summarised below.

These are broad projections, covering the whole economy and over a long period. The persuasive force of these analyses is not so much in the precise numbers, but in the consensus that a very wide range of industry activity is expected to be changed by application of AI.

PWC June 2017[44]

UK GDP will be up to 10.3% higher in 2030 as a result of AI – the equivalent of an additional £232bn – making it one of the biggest commercial opportunities in today's fast-changing economy.

The impact over the period will come from productivity gains (1.9%) and consumption-side product enhancements and new firm entry stimulating demand (8.4%).

There will be significant gains across all UK regions, with England, Scotland, Wales and Northern Ireland all seeing an impact from AI in 2030 at least as large as 5% of GDP, and extra spending power per household of up to £1,800-£2,300 a year by 2030.

It is important to note that the impact of artificial intelligence will not be concentrated in any one sector of the economy and will not be limited to the firms that develop and produce the AI technologies. The uptake of AI will have direct impacts in the sectors in which this uptake occurs, both through the automation and augmentation of process and the enhancement of product offerings for consumers. Furthermore, the total economic impact includes the potential indirect and induced impacts that are likely to be felt by firms and consumers throughout the economy.

It is possible to see AI becoming a utility, improving the functioning of all digital applications and digitised functions. In this context, where proven applications become ever easier to access, the businesses that benefit most will be those that have made the best preparation to be "AI-ready".

Kevin Kelly, 2014: "The AI on the horizon looks more like Amazon Web Services—cheap, reliable, industrial-grade digital smartness running behind everything, and almost invisible except when it blinks off. This common utility will serve you as much IQ as you want but no more than you need. Like all utilities, AI will be supremely boring, even as it transforms the Internet, the global economy, and civilization. It will enliven inert objects, much as electricity did more than a century ago. Everything that we formerly electrified we will now cognitize. This new utilitarian AI will also augment us individually as people (deepening our memory, speeding our recognition) and collectively as a species. There is almost nothing we can think of that cannot be made new, different, or interesting by infusing it with

[43] https://newsroom.accenture.com/news/artificial-intelligence-poised-to-double-annual-economic-growth-rate-in-12-developed-economies-and-boost-labor-productivity-by-up-to-40-percent-by-2035-according-to-new-research-by-accenture.htm
[44] The economic impact of artificial intelligence on the UK economy http://www.pwc.co.uk/services/economics-policy/insights/the-impact-of-artificial-intelligence-on-the-uk-economy.html

some extra IQ. In fact, the business plans of the next 10,000 startups are easy to forecast: Take X and add AI. This is a big deal, and now it's here."[45]

Google CEO Sundar Pichai, Google Founders' Letter 2016: "Google started in the cloud and has been investing in infrastructure, data management, analytics, and AI from the very beginning. We now have a broad and growing set of enterprise offerings: Google Cloud Platform (GCP), Google Apps, Chromebooks, Android, image recognition, speech translation, maps, machine learning for customers' proprietary data sets, and more. Our customers like Whirlpool, Land O'Lakes and Spotify are transforming their businesses by using our enterprise productivity suite of Google Apps and Google Cloud Platform services.

As we look to our long-term investments in our productivity tools supported by our machine learning and artificial intelligence efforts, we see huge opportunities to dramatically improve how people work. Your phone should proactively bring up the right documents, schedule and map your meetings, let people know if you are late, suggest responses to messages, handle your payments and expenses."[46]

Microsoft statement, Democratizing AI: "Every walk of life has changed because of our ability to create knowledge and distribute knowledge. But one thing has remained constant and scarce: time.

In the midst of this abundance of information, we're still constrained by our human capacity to absorb it. The question is, how can we use all we have in terms of computational power to solve this fundamental constraint? To make better sense of the world? That's the essence of what AI is. It's not about having AI that beats humans in games, it's about helping everyone achieve more — humans and machines working together to make the world a better place.

We're taking a four-pronged approach to how we think about Microsoft AI and how we pursue this bold ambition to democratize AI for all:

- We're going to harness artificial intelligence to fundamentally change how we interact with the ambient computing, the agents, in our lives.
- We're going to infuse every application that we interact with, on any device, at any point in time, with intelligence.
- We'll make these same intelligent capabilities that are infused in our own apps — the cognitive capabilities — available to every application developer in the world.
- We're building the world's most powerful AI supercomputer and making it available to anyone, via the cloud, to enable all to harness its power and tackle AI challenges, large and small."[47]

[45] https://www.wired.com/2014/10/future-of-artificial-intelligence/
[46] https://www.blog.google/topics/inside-google/this-years-founders-letter/
[47] https://news.microsoft.com/features/democratizing-ai/ - EfQesLtOujFRMv12.97

UK support for development of AI

There are already support mechanisms for UK AI, for academic research and for businesses.

Funding for Research - Engineering and Physical Sciences Research Council: AI is a priority for EPSRC, which supports 143 relevant research grants under one of its key objectives: to deliver intelligent technologies and systems. EPSRC believes that multi-disciplinary research involving social scientists will enable AI tools and technologies to be acceptable, usable and ethical.

EPSRC aims to deliver:

- A portfolio of AI research and training in data science that complements work undertaken at the Alan Turing Institute, with links to underpinning statistical and theoretical sciences.
- A supply of people highly skilled in AI technologies able to work across a wide range of domains (e.g. the future of healthcare delivery)
- Researchers combining development of new methodologies and applications (e.g. by working alongside research enablers like research engineers, translational researchers and collaborators with application expertise)
- A portfolio that contains AI-enabled technologies co-created with other disciplines (e.g. robotics, human-computer interaction, computer vision and the social sciences). This should assess how intelligent systems interact with humans, and consider their dependability, safety and security.

The UK is currently ranked fourth in terms of the volume of academic papers published relating to AI, behind China, the United States and Japan, with 10,000 papers published between 2011 and 2015 (China published 41,000 over the same period).[48]

Research in Data Science - the Alan Turing Institute: The Alan Turing Institute is the national institute for data science, headquartered at the British Library. Five universities – Cambridge, Edinburgh, Oxford, University College London and Warwick – and the EPSRC created the Institute in 2015, investing £42 million in total.

The Institute brings together researchers in mathematics, statistics, computer science, social science and data ethics, software engineering, machine learning and AI to generate world class research in data science. It applies its research to real-world problems, working with partners in industry, Government and third sector. Core areas of research include defence and security (with GCHQ and other government agencies), health and well-being (working with a range of partners), data-centric engineering (with Lloyd's Register Foundation), computational technology (with Intel) finance (with HSBC) and smart cities. Further core areas of the Institute's remit are training the next generation of data scientists, and shaping the public conversation around data and its powerful impact on science, society, the economy and our way of life.

[48] https://www.timeshighereducation.com/data-bites/which-countries-and-universities-are-leading-ai-research

Turing Data Study Group: The Alan Turing Institute regularly holds Data Study Groups, in which top public and private organisations bring major data science problems to be worked on by researchers at the Institute. Over a week-long period, researchers apply their cutting-edge data science expertise and techniques to the problem, culminating in a presentation and report at the end of the week containing their recommended approaches to tackling the challenge. Participating academics gain the opportunity to work on real-world industry problems and datasets, and industry participants benefit from intensive practical research into their problem, with viable business solutions at the end of the process. Participating companies to date include Siemens, Shell, National Grid, Defence and Security Technology Laboratory, Tata Steel and Thomson Reuters.

Open Data Institute (ODI): The ODI's mission is "to build a strong, fair and sustainable data economy by helping businesses and governments get data to people who need it." ODI is an independent, non-profit, non-partisan company based in London, with a national and international reach, bringing together commercial and non-commercial organisations and governments around specific sectors to address global challenges with data.

As above, the **ONS Data Science Campus,** the government **Data Science Accelerator** and **Government Data Science Partnership** all support government data expertise and data applications.

The **National Innovation Centre for Data** in Newcastle will bring together industry, the public sector and world-leading academics to develop the skills, ideas and resources needed to exploit the opportunities offered by the explosion in digital data.

Royal Statistical Society (RSS) Data Science Section: Recently established, this includes representatives from business, industry, government and academia. The section will organise meetings for a broad range of attendees and generate outputs that are aimed at:

- Supporting the Data Science community throughout the UK
- Promoting good practice by addressing what good Data Science looks like (with exemplars) and what it does not look like
- Promoting the statistical aspects of Data Science and reinforcing the statistical framework
- Being a trusted voice on Data Science for employers, including inputting to consultation exercises
- Supporting the pipeline and career development of data scientists and statisticians by elevating skill sets to work in the modern world
- Supporting important emerging topics such as ethics, privacy, algorithmic responsibility and personalization - lifting the quality of the conversation
- Fostering multi-disciplinary connections and the exchanging of ideas.

Support for commercialisation of digital R&D - Digital Catapult: The Digital Catapult is a technology and innovation centre for advanced digital technologies, created to accelerate access to new digital markets and carry out applied research and development to identify new applications of emerging technologies. AI is one of four key technology layers. It occupies a neutral position between multinational corporates, investors, startups, government departments,

academia and other Catapults (each focused on a high-value technology area), with active partnerships across those organisations. The Catapult has worked with DSTL on autonomous agents, with SwissRe on smart conversational interfaces and with Seagate on metrology for complex production lines, in each case bringing together academics, experts and SMEs. The Catapult is developing a machine learning computation lab to help early stage companies reduce the costs of training machine learning models.

Industry representation - TechUK: The tech industry representative body techUK sees AI as significant driver of change across the UK economy and society, and works to encourage better understanding of benefits and challenges. In April they brought together leading industry and academic including Luxoft, Artificial Solutions and UCL to discuss risks, challenges and barriers that could stand in the way of realising the potential of AI in Fintech, and how these might be overcome. In May they hosted discussions on AI applications in healthcare with IBM, University of Cambridge and DeepMind. This focused on raising awareness of AI's applications in healthcare and how to build public trust in the technologies. In May 2017, Tech UK ran a campaign week featuring the opportunities and benefits AI can bring to the UK more broadly, including improving productivity and economic growth.[49]

NMI, the body for the UK Electronic Systems & Technology Industry, is now working to support "UK Deep Tech", as TechWorks.

Additional non-competitive collaboration in AI happens through the **Society for the Study of AI and Simulation of Behaviour**.

The **Leverhulme Centre for the Future of Intelligence** aims to build "an interdisciplinary community of researchers, with strong links to technologists and the policy world, and a clear practical goal: to work together to ensure that we humans make the best of the opportunities of artificial intelligence as it develops over coming decades.[50]

An **All-Party Parliamentary Group on Artificial Intelligence** has been established to explore the impact and implications of AI including machine learning, and to improve understanding and engagement among Parliamentarians and other policy-makers.

Support for digital scale-ups: Tech City UK. Tech City UK support growth in the UK digital tech ecosystem by building a pipeline of scale-ups to generate more late-stage pioneers and create jobs nationwide. Tech City UK is working on a number of sector-specific programmes building on Upscale, a six-month mentoring programme for fast growth scaling companies.

International activity

Global investment

[49] https://www.techuk.org/insights/news/item/10673-techuk-s-ai-week
[50] http://lcfi.ac.uk/

Investment is rising too quickly to track definitively, but there are indications of the potential future value perceived by industry globally.

"Startups specializing in AI applications received US$2.4 billion in venture capital funding globally in 2015 and more than US$1.5 billion in the first half of 2016. Government programmes and existing technology companies add further billions. Leading players are not just hiring from universities, they are hiring the universities: Amazon, Google and Microsoft have moved to funding professorships and directly acquiring university researchers in the search for competitive advantage."[51]

McKinsey also sees a sharp recent upturn in investment: "Companies at the digital frontier - online firms and digital natives such as Google and Baidu - are betting vast amounts of money on AI. We estimate between $20 billion and $30 billion in 2016, including significant M&A activity. Private investors are jumping in, too. We estimate that venture capitalists invested $4 billion to $5 billion in AI in 2016, and private equity firms invested $1 billion to $3 billion. That is more than three times as much as in 2013. An additional $1 billion of investment came from grants and seed funding."

IDC forecasts worldwide revenues for cognitive and AI systems of $12.5 billion in 2017, an increase of 59.3% over 2016, and projects rising investment leading to revenues of more than $46 billion in 2020.[52]

International comparisons

The UK and other countries are generally seen as behind the US and China in terms of scale of AI investment and activity. The UK still lags far behind the United States in terms of global deal share, with 62% of investment deals in 2016 estimated as going to startups in the US and only 6.5% going to UK-based startups.[53] Only 5% of the value of global VC fundraisings for AI companies went to UK businesses in 2010-2016 (see table below). More UK investments appear to be at earlier stages, with three quarters of the total number of UK AI companies seeking seed or angel investment, compared with only half of US companies, and just one in 10 UK AI companies looking for growth capital, compared with one in five in the US.[54]

AI –value of VC fundraisings among international competitors, 2010-2016[55]

Country	2010	2011	2012	2013	2014	2015	2016	Total
United States	£112m	£171m	£228m	£399m	£843m	£1,503m	£1,578m	**£4,833m**
China	£6m	-	£1m	£15m	£55m	£124m	£199m	**£401m**
United Kingdom	£6m	£9m	£24m	£18m	£19m	£67m	£152m	**£294m**
Canada	£3m	£17m	£11m	£4m	£2m	£23m	£11m	**£71m**

[51] World Economic Forum Global Risks Report 2017. http://www3.weforum.org/docs/GRR17_Report_web.pdf
[52] http://www.idc.com/getdoc.jsp?containerId=prUS42439617
[53] https://www.cbinsights.com/blog/artificial-intelligence-startup-funding/
[54] https://medium.com/mmc-writes/artificial-intelligence-in-the-uk-landscape-and-learnings-from-226-startups-70b9551f3e4c
[55] Analysis based on Pitchbook data.

Germany	£3m	£8m	£8m	£0m	£0m	£7m	£9m	**£36m**
France	£3m	£1m	-	£1m	£1m	£9m	£15m	**£31m**
Total	**£132m**	**£206m**	**£272m**	**£438m**	**£920m**	**£1,733m**	**£1,964m**	**£5,666m**

International Governments' support for AI

Other leading digital economies are also acting to grow their national AI capabilities and subsequent market shares. The UK will need to raise its level of investment if we are to match the support that our global competitors are giving to their AI sectors. These are only some examples of headline actions by governments in other leading countries.

France: Launched an AI Strategy in March 2017. Key recommendations include: establishing a strategic committee to implement the strategy's recommendations; a programme to identify, attract and retain AI talent; funding a mutualised research infrastructure; a public-private consortium to identify or create an AI centre; ensuring that AI is a priority for all innovation in public bodies; investing €25m (£20m) in ten startups within five years.

Singapore: The National Research Foundation (NRF) is investing up to S$150 million (£85m) into a new national programme aimed at boosting Singapore's AI capabilities over the next five years.

United States: The government invested US$1.1 billion (£850m) on unclassified R&D for AI systems in 2015 and an estimated US$1.2 billion (£950m) in 2016. The Information and Intelligence Systems department of the National Science Foundation and the programs related to AI from the DARPA are reported to have been around US$300m-$400m (£250m- £300m) a year for the last 15 years. The 2016 White House reports included a National Artificial Intelligence Research and Development Strategic Plan.[56]

South Korea: The government announced that it will invest 1 trillion won (£700m) in AI research over the next five years, a 55% increase in annual funding for AI.

Germany: The Research Center for AI (DFKI) was founded in 1988 and has an annual budget of €41m. It is one of the world's largest AI labs, with nearly 500 researchers.

Canada: is funding a Pan-Canadian AI Strategy for research and talent. The funding is worth C$175m (£100m) and is aimed at attracting and retaining top academic talent in Canada.

China: has a stated ambition to create a US$15 billion AI market by 2018, and is reported to be preparing a comprehensive AI strategy.

56

https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/national_ai_rd_strategic_plan.pdf

RECOMMENDATIONS

As described above, several key factors have combined to increase the capability of AI in recent years, globally and specifically in the UK, in particular increased access to data, supply of advanced skills, research to develop AI, and increased computing capacity.

To realise future AI opportunities and meet growing demand, the UK will need to expand supply in the same areas.

Looking ahead, continued growth in computing power and availability of it is anticipated at least in the medium term. Data volumes are expected to continue to grow, but this does not mean all that data is easily accessible for AI. Demand for skills is expected to continue to rise. The need for further research to extend and improve what AI can do will continue to grow as the technology improves and spreads.

Through engagement with industry and academic experts and by developing their ideas as well as examining analysis of UK and global development of AI, **this Review has identified the major challenges below, and set out recommendations to address those.**

As summarised in the market overview above, use of AI is not confined only to certain sectors or sizes of company. Successful UK uptake will mean it is used much more widely than it is already.

Therefore many of the recommendations address factors that can grow UK AI capability broadly, rather than within only certain sectors or certain types or sizes of organisation.

SECTION THREE: RECOMMENDATIONS TO IMPROVE ACCESS TO DATA

Access to Data for AI

Growing the AI industry in terms of those developing it and deploying it requires improved access to new and existing datasets to train, develop and deploy code.

The move to more open access publishing has improved access to research data, and the Digital Economy Act 2017 included new provisions on using data for the public good, but there remain many areas where action is needed to maximise use of data in practice. Broadly, more data can be made available by:

- Making more data open
- Improving machine readability
- Managing trust and access to sensitive data which cannot and should not be open.

This section addresses challenges in access to data and makes recommendations for action in key areas.

- Open data
- Access to sensitive data
- Benefits of trusted data sharing arrangements
- Data Trusts
- Access to research data for AI
- Copyright and text and data mining
- Explainability of AI-enabled uses of data.

Open data

Very simply, more open data in more sectors is more data to use with AI to address challenges in those sectors, increasing the scope for innovation.

The 2016 US White House report “Preparing for Artificial Intelligence” recognised the importance of encouraging the public sector to make data available for AI to improve services, and the same is true for the UK. “Many uses of AI for public good rely on the availability of data that can be used to train machine learning models and test the performance of AI systems. Agencies and organizations with data that can be released without implicating personal privacy or trade secrets can help to enable the development of AI by making those data available to researchers. Standardizing data schemas and formats can reduce the cost and difficulty of making new datasets useful.”[57]

[57] p.14, October 2016. https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

Wherever there are no risks (including in terms of privacy, security, commercial confidentiality), data held by public organisations should be made open. The UK has been a pioneer in open data, and was recently ranked first in the world on Government performance on Open Data.[58]

The Open Data Institute (ODI) was founded in 2012 "to connect, equip and inspire people around the world to innovate with data", and now has a global network. As well as opening data in the public sector, the ODI and its network have already persuaded hundreds of businesses in the banking, agriculture and sports sectors to publish open data.

There is considerable continuing work to open more data, of which data for AI is only one element, but the additional value that can be derived from using open data as AI training data should give additional impetus to that work.

To be readily available for use, open data also needs to be machine-readable in standard formats, with clear rights information where applicable.

Access to sensitive data

Much data cannot be made open. There are many reasons why it is difficult or impossible for organisations to allow others to access their data, including privacy, security, commercial advantage and other considerations.

As described above, even where organisations see a case to proceed with secure data-sharing, and trust the external organisation they plan to share with, they often lack expertise and know-how to form agreements, establish trust between the parties, and manage the data sharing processes in practice. As a result, much data that could be used to develop specific new AI applications in high value sectors cannot be used for that in practice, and AI cannot be applied to challenges in these sectors.

This is particularly apparent in healthcare, which is also an area where AI can add great value. "The obstacle to AI implementation in healthcare is not technological but access to data. Research is hampered by difficulties in accessing large medical datasets, for legal or other reasons. It's particularly tough for startups in the field; larger players already have access to such data."[59]

Practical challenges around access to data - Matteo Berlucchi, Your.MD

Your.MD is building an AI system designed to give everyone in the world the best possible health information they need when they need it free of charge. This ambitious (and noble) challenge requires

[58] http://www.publictechnology.net/articles/news/uk-retains-pole-position-open-data

[59] Yann LeCun, Director AI Research, Facebook, quoted in Artificial Intelligence in the Real World, Economist Intelligence Unit Briefing Paper, 2016.
https://www.eiuperspectives.economist.com/sites/default/files/Artificial_intelligence_in_the_real_world_1.pdf

a profound understanding of each individual person's medical profile. Therefore, access to reliable and consistent data sets of anonymised personal health records would give us a tremendous boost towards achieving this goal.

We have tried to approach the NHS to see if there was a way to access some of this data but we have struggled to even find the right person to talk to. Navigating a complex organisation like the NHS is an unfathomable task for small startups like Your.MD.

I strongly believe that the levelling and opening of the access to such vital data is fundamental for the creation and development of ground-breaking AI services in the healthcare sector. Government is uniquely positioned to unlock this potential by creating the appropriate data sharing environment.

But while healthcare is a particularly acute example, there are challenges in any area where data cannot be made open for sound reasons.

Managing access to sensitive data has two key aspects:
- providing trust and confidence
- reducing transaction costs to sustainable levels.

Trust and confidence in use of sensitive data: In many sectors (in particular those serving individual people directly), data needs to be protected, for reasons of privacy, security, confidentiality, and commercial sensitivity. Data-holders need trust and assurance, to be confident in sharing data with AI developers. Some of the areas where data is most sensitive, commercially or personally, may also be ones where the greatest benefits are.

This is not about removing necessary protections, for individuals, businesses, or market competition. As the Information Commissioner's Office describes it, "It's not big data ***or*** data protection, it's big data ***and*** data protection. The benefits of big data, AI and machine learning will be sustained by upholding key data protection principles and safeguards."[60]

In these sectors there is a need for secure, managed access to data for AI, offering agreed sharing of benefits, and retaining protections on repeatable standard terms. These security challenges can be overcome by agreements, but at the moment agreements are made on an ad hoc basis, incurring large transaction costs and making it difficult for smaller companies to compete. Existing programmes have been complex to arrange, and some have attracted criticism, such as DeepMind's collaboration with the Royal Free Hospital.[61]

By systematically improving data access through trusted and repeatable mechanisms, the UK can help create a competitive and innovative UK market that serves the interests of holders of large datasets, and of companies who could develop AI services by leveraging that data.

Developing trusted solutions would help companies work on sensitive data in public organisations to find new solutions to major public challenges, for instance in helping long-term unemployed or socially excluded people. The potential for proven and trusted public-private data sharing is not all one way either. Public health organisations could work with data from

[60] https://ico.org.uk/for-organisations/guide-to-data-protection/big-data/
[61] http://www.bbc.co.uk/news/technology-39301901

supermarkets, transport providers and town planning, to identify how to better encourage healthy lifestyles. Public sector bodies could significantly improve their decision making if they were able to use commercially-held information, notably for the purpose of public health policy, natural and technological risk management, spatial and urban planning, managing energy supply grids or protecting the environment.[62]

Reducing transaction costs of accessing data: Solutions that provide trust and confidence also need to be affordable and sustainable. If legal and procedural costs remain high for each transaction, a particular solution for data-sharing will not be taken up widely. This is in part a matter of enabling small companies to access data as easily as very large ones, whose experience and resources can manage transactions and transaction costs more easily.

Certain large companies have developed major advantages in AI technologies because of the abundance of data they generate and have access to. Google and Facebook use data generated through their other products and services continually to improve those offerings, for instance by predicting customer preferences based on vast volumes of search data.

There is a risk that AI may further advance the position of incumbents because of this virtuous circle: more data leads to better AI technologies which drives more usage and data. The large resources of these companies also give them an advantage in managing transactions for access to other datasets, for instance in healthcare.

Even with data they do not hold internally, the larger companies have the advantages of experience, expertise and mechanisms for striking agreements with data-holders to access their data.

Comparably, for small companies, negotiating agreements and establishing practices can present major obstacles and costs. These conditions could make it difficult for new companies to enter some markets, potentially to the detriment of outcomes for the public. These barriers for small companies in particular could restrict the focus of AI innovation to areas that are core to the major incumbents' services, and draw resources away from other areas of major public benefit, including innovation in public services or medical research.

While the large tech companies contribute very significantly to innovation, a successful future for UK AI should include a broad range of small and large companies working across a wide spectrum of sectors.

Benefits of trusted data sharing arrangements

Standardised, repeatable terms for access to data would unlock value in many sectors, making possible many applications which are not economical today. Some mutually beneficial data-sharing agreements are being used in practice already, so there is a strong case for surfacing and spreading successful approaches.

[62] The European Digital Single Market Strategy has addressed the value of private sector data to public bodies. https://ec.europa.eu/digital-single-market/en/policies/shaping-digital-single-market

Action in this area would need to be a step by step transition, not a single step to a standard that can be used in all circumstances. At the moment there is relatively little widely shared understanding of even the questions that organisations should consider when approaching data-sharing for AI. As one attendee of a workshop in this Review said: “Even a list of the right questions would be good.”

Therefore action in this area should proceed by exploring what agreements are already in use, what common elements they include, and work to codify where possible and develop questions and analyses to help users proceed where codification is not possible or appropriate.

Data Trusts

Recommendation 1: To facilitate the sharing of data between organisations holding data and organisations looking to use data to develop AI, Government and industry should deliver a programme to develop Data Trusts – proven and trusted frameworks and agreements – to ensure exchanges are secure and mutually beneficial.

To use data for AI in a specific area, data holders and users currently come together, on a case by case basis, to agree terms that meet their mutual needs and interests. To enable this to be done more easily and frequently, it is proposed to develop terms and mechanisms for these parties to form, between them, individual “data trusts” to enable AI to be developed to meet the needs of the parties involved and allow data transactions to proceed with confidence and trust.

These trusts are not a legal entity or institution, but rather a set of relationships underpinned by a repeatable framework, compliant with parties’ obligations, to share data in a fair, safe and equitable way.

This programme could progress through stages, delivering benefits at each.

1. Evidence-gathering on challenges to data-sharing and solutions already in use

2. Development of frameworks through which the holders and users of data would come together to form Data Trusts to share data.

3. Trialling in one or two sectors where there are known challenges, but risks are manageable, to test the concept.

Transport may be a good candidate, as a sector where many organisations could benefit from sharing data, but where the benefits and practical processes for sharing data for AI have not

been fully trialled. ODI, Deloitte and the Transport Systems Catapult have reported on the case for supporting data-sharing in Intelligent Mobility.[63]

This proposal would not be the first attempt to develop a repeatable framework for data-sharing. Development of it should learn from existing, functioning models, and from where past models have succeeded, up to what points they have succeeded, and at what points they have failed to meet users' needs. It should address practical obstacles experienced in transactions for data for use by AI, and issues around using shared data for AI (for example novelty of objectives, complexity, uncertainty over assigning outputs and allocating benefits).

A support organisation, the Data Trusts Support Organisation (DTSO) could be developed, which would lead on the development of tools, templates and guidance for those who want to share and use data, so data owners and consumers can come together to form data trusts as and when they wish to do so. This role could be provided by a neutral and expert organisation such as the Royal Society, the Royal Academy of Engineering, the Digital Catapult or the Open Data Institute, all of which have relevant expertise. Given the importance of trust, this role should be taken by an organisation which is already trusted both for security and operational effectiveness in relation to data.

The DTSO would by default act as a trustee, a third party that helps manage a data trust. **Key functions** of the DTSO include:

- Provide the framework to define the data, or stream of data, that the parties agree to share
- Broker the purpose of sharing the data and its intended uses, including what analytics will be used
- Agree upon the mechanism of data transfer and storage
- Determine the conditions upon which commercial value generated will be distributed.

A data stewardship body as proposed by the Royal Society and the British Academy could provide advice and oversight. Developing data trusts could also give that body valuable operational experience in resolving important current data challenges and realising public benefits.

The data trust and its support functions could become instrumental in enabling AI competition based on some commonly shared data (see ISCF for AI), and these competitions could be the driver behind the creation of the initial data trusts.

Government role: Government agencies already have experience in secure data-sharing for researchers which may be relevant to the development of data trusts, for instance the Office for National Statistics Virtual Microdata Lab and 'Five Safes' framework.

In order to facilitate the development and adoption of data trusts, a set of the DTSO's supporting functions should be supported and authorised by Government through the AI Council proposed below.

[63] https://s3-eu-west-1.amazonaws.com/media.ts.catapult/wp-content/uploads/2017/04/12092544/15460-TSC-Q1-Report-Document-Suite-single-pages.pdf

These functions would initially include acting as default trustee, and provide free templates of legal contracts that could be used to establish data trusts.

If the Data Trust is successful in facilitating more data-sharing, functions would develop to include advice and guidelines for legislators to convert the concept of data trust into an official legal construct and allow other third party organizations to be used as trustees, if industry sees a need for that development.

Proposed governance structure: In the first phase, the DTSO will act as the single body with responsibility for overseeing the creation of the data trusts. This could have strong strategic relations to a data stewardship or ethics body. In addition, the Digital Catapult, ICO, Open Data Institute, the industry-led AI Council (proposed below) will all be important stakeholders who can help support the creation and promotion of data trusts with industry. Subsequently, other organisations would be accredited by the DTSO to deliver these functions.

Role of the DTSO: In order to encourage the creation of data trusts, the DTSO could initially work for free for the sharing of any kind of data coming from government entities or non-profits, or for data that is deemed benefiting the public good. For other kind of data sharing, the organization will not operate for profit, but will charge for its services should commercial value be generated from exploitation of the data. The support organisation will not be a body for governance or for setting case law on ethics of the data use.

The DTSO could be further enhanced by:

- Developing the technical, business and financial framework upon which multiple parties feel confident to engage.
- Offering guidance on anonymising, disaggregating and securing the data for sharing and transactions
- Hosting technical experts that can independently advise on data valuations
- Acting as a trusted advisor on GDPR and other data regulation where applicable, advising on fair and ethical use of data.

Access to research data

Machine readability is also a key issue in access to data for AI. As AI increases the capability to work with more data from more sectors and research domains, making it machine readable by default will greatly support using AI for research, and with prior research.

Published scientific research is a major source of data for subsequent research. AI offers great potential to realise new value from that data, but only if AI can access the data as training data. Researchers and research organisations indicate that while much progress has been made in opening research, some data that could be made accessible is currently not.

Recommendation 2: To improve the availability of data for developing AI systems, Government should ensure that public funding for research explicitly ensures publication of underlying data in machine-readable formats with clear rights information, and open wherever possible.

Enabling text and data mining of research

At the moment, some data cannot be extracted from published research because access to that data can be restricted by contractor or copyright, making it unavailable as training data for AI. This restricts the use of AI in areas of high potential public value, and lessens the value that can be gained from published research, much of which is funded by the public.

The volume of research published continues to grow, machine reading is increasingly necessary in order to use these growing volumes of research. Therefore, machine reading technologies are increasingly standard and essential tools for research, and machine readability should be the default standard for research publication. To enable extraction of data for AI, publication of research should allow extraction of data by text and data mining (TDM) for AI, so long as the results are not a substitution for the original works.

Easier access to a broader range of data in published research, with reduced transaction costs, would enable innovative businesses to develop more and better AI applications in more sectors.

Many research organisations support change to liberalise text and data mining, and some see that as necessary to international competitiveness in research: “The development of new content mining services by commercial companies is in the interest of the research community and society in general.”[64]

To date, assessments of the value of text and data mining of research and for new research do not appear to have included the potential value that can come from using data for AI. Therefore future assessments should take account of uses of data by AI, adding to the overall estimation of benefits from liberalising use of text and data mining. Government should commission research into how much value could be added to the UK economy by making data available for AI through TDM, including by commercial businesses. Government should recognise this value when forming copyright exceptions.

Recommendation 3: To support text and data mining as a standard and essential tool for research, the UK should move towards establishing by default that for published research the right to read is also the right to mine data, where that does not result in products that substitute for the original works. Government should include potential uses of data for AI when assessing how to support for text and data mining.

[64] Science Europe Briefing Paper 2015. https://www.scienceeurope.org/wp-content/uploads/2015/04/SE_Briefing_Paper_textand_Data_web.pdf

SECTION FOUR: RECOMMENDATIONS TO IMPROVE SUPPLY OF SKILLS

Global demand for skills: There is intense global competition for advanced skills to develop AI. To remain competitive, the UK needs to make a step-change in the training of these skills in the UK. The UK also needs to remain an attractive destination for the best talent from around the world, at levels from undergraduate upwards, and including through visa categories and numbers that meet industry needs. This is an opportunity to show that the UK is open to international talent, and committed to staying among the global leaders in AI.

This section covers the supply of skills across the entire pyramid and makes recommendations to increase it.

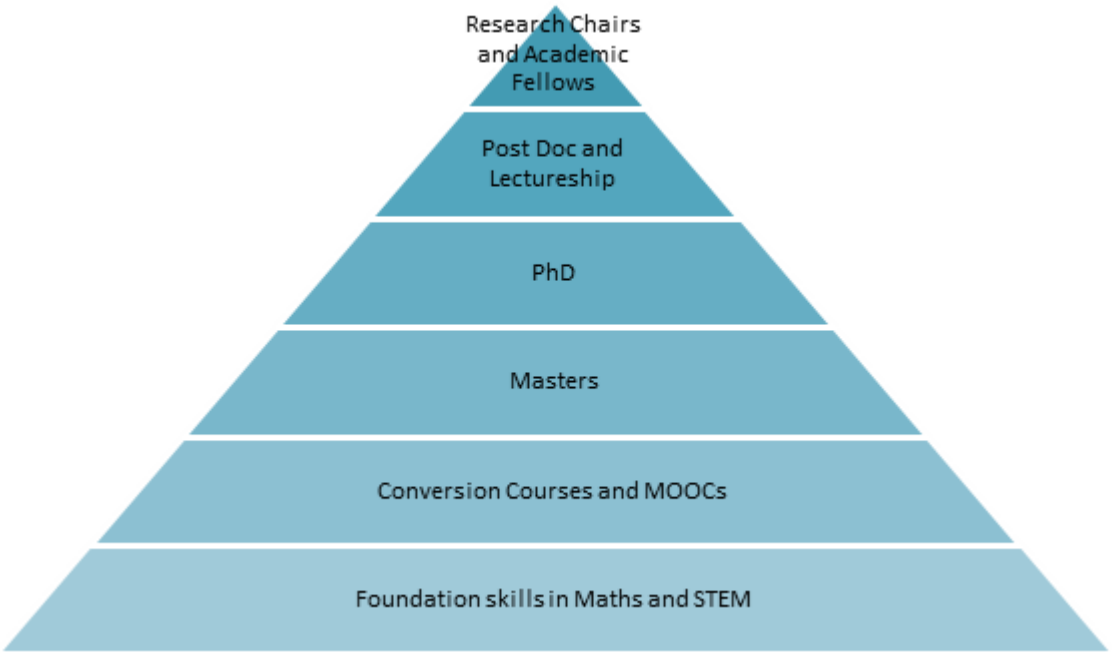

High level skills: The development of AI has been made possible by people who have been trained at a high level in computing, data science and (particularly more recently) machine learning. Development requires deep expertise, and has generally happened in places (including the UK) where that long term development of skills has been supported.

The UK now needs AI experts in greater numbers to develop more applications in more sectors. The intense demand from employers and the constrained numbers are at the high end: skilled

people who have graduated from Masters and PhDs, and ideally who also have added practical experience to that training.

Shortage of talent is a challenge for most high technology fields, but it is especially acute in AI because it is a specialized subfield at the junction of two already supply-constrained fields, software engineering / computer science and mathematics / statistics / data science. This view was expressed frequently in the workshops held for this Review, in discussions with employers and would-be employers, and in such data as are available, including data on salaries.

This review recommends a major step-change in UK development of high-level skills for AI, but this is not something that can be done in one move. Developing more expertise at PhD and above depends on the supply of graduates and candidates with the right ability, and also on having enough academic experts and institutions able to teach and support at all of these levels. The UK cannot simply increase the number of PhD places, without increasing supply of candidates and supervisors.

Therefore the recommendations below are interdependent, and will enable the capacity of the system to increase in the short term (1-3 years), but continue to grow progressively beyond that.

Foundational Skills

Those high-level skills depend on a spectrum of skills at different levels. The more people who have the right foundational STEM skills, the more can train in the higher skills, but also more will be able to work in adjacent roles: working in and around AI rather than developing it at the most complex levels.

As DeepMind's evidence to the recent House of Commons Science and Technology Select Committee Inquiry into Robotics and Artificial Intelligence stated, "one of the most important steps we must take is [ensuring] that current and future workforces are sufficiently skilled and well-versed in digital skills and technologies, particularly STEM subjects".[65]

AI can be applied in a wide variety of industry sectors and application areas, and that range is only going to grow. This means that there is growing demand for professionals who are not core specialists in AI, but will be needed to successfully add AI into functions in those sectors. There will be a need for support skills for AI across including in:

- data preparation, curation, protection,
- explaining AI functions to staff and customers,
- managing reporting, accountability, liability.

There will also be a need for professionals who can use AI tools successfully in specific domain areas, including:

- research scientists
- maintenance technicians

[65] http://data.parliament.uk/writtenevidence/committeeevidence.svc/evidencedocument/science-and-technology-committee/robotics-and-artificial-intelligence/written/33005.html

- surgical technicians and healthcare professionals working with assistive technology
- mechanical engineers in manufacturing and transport
- and "applying AI" roles in service sectors – insurance, advertising, design, creative, retail, entertainment, financial.

There have been several reports in recent years that developed and evidenced the case for improving education and training in maths, computing, data science, and the full range of digital skills. Including but not limited to:

- Digital skills for the UK economy: paper by Ecorys UK for the Departments for Culture Media and Sport, and Business Innovation and Skills. [66]
- Shadbolt Review of Computer Science Degree Accreditation 2016[67]
- UK Digital Strategy, 2017[68]
- Analytic Britain, 2015 paper developed by Nesta and Universities UK (see box).[69]

Further research and policy activity is going on now.

Mathematics: Numerous reports and commentators have stated that more, longer and better mathematics education would better prepare UK students for a data driven world and jobs in it.

Professor Sir Adrian Smith has undertaken a study of the feasibility of compulsory mathematics study for all pupils to 18. To add to the other potential benefits of more maths education that his work has identified, extending compulsory maths education would very significantly improve the foundations for skills to develop, understand, and work with AI.[70]

Embedding data science and AI widely in education: In time, AI could positively affect every area of STEM education, by providing an engaging and interesting example of application of STEM skills. Therefore there is a case for embedding understanding of AI across STEM education, up to Masters level.

Teacher Training: The addition of computer science to the National Curriculum is an excellent step, but will only deliver fully if there are enough teachers who can teach it well. The British Computer Society has expressed concern that not enough students are taking up computer science, and suggested that as many as 70% of secondary school computer science teachers could be lacking a relevant computer science background to teach at GCSE level. Therefore, more and better teacher training in computer science would improve outcomes.

Careers advice: Similarly, AI should be more fully represented in careers advice, as it will increasingly have impacts on career opportunities, both those that grow and those that shrink. There is a statutory duty on all schools to provide students with independent careers advice. Given the new opportunities likely to be created by AI, better understanding among careers advisors could substantially improve their service to students.

[66] https://www.gov.uk/government/publications/digital-skills-for-the-uk-economy

[67] https://www.gov.uk/government/publications/computer-science-degree-accreditation-and-graduate-employability-shadbolt-review

[68] https://www.gov.uk/government/publications/uk-digital-strategy

[69] http://www.nesta.org.uk/sites/default/files/analytic_britain.pdf

[70] https://www.gov.uk/government/publications/smith-review-of-post-16-maths-report-and-government-response

The application of AI also has the potential to improve careers advice, by helping to provide students and advisors with much more sophisticated, timely and local information from data about business growth and movement, changes in demand for skills, and salaries. However, successful use of AI would depend on the quality of the data available.

AI for education: In education as in other areas, AI could offer personalisation at scale, improving support for learners at different levels and with different needs.

ANALYTIC BRITAIN: SECURING THE RIGHT SKILLS FOR THE DATA–DRIVEN ECONOMY

2015 Policy Briefing by Nesta and Universities UK
http://www.nesta.org.uk/sites/default/files/analytic_britain.pdf

Schools and Colleges
1. We need stronger teaching of mathematics and statistics in schools and colleges
2. More and better information about analytical career prospects and role models in schools and colleges
3. Embed data analysis in other subjects
4. Support the development of extracurricular data activities

Universities and Vocational Education
1. Increase the visibility of strong data analytics courses
2. Embed quantitative analysis across disciplines
3. Boost the business and soft skills of graduates from data analytics courses
4. Increase the supply of high–end analytical talent
5. Foster interdisciplinary research and skills development programmes

Labour Market and Industry
1. Create a cross–cutting taskforce around data analytics
2. Actively convene industry analytics networks
3. Support innovative interventions enabling local authorities to boost local analytical skills
4. Raise awareness of the value of data for business
5. Deliver innovative solutions for data analytics training

Current supply of higher skills for AI

26 UK universities offer undergraduate courses in AI[71] and there are more than 30 graduate programmes running across 20 universities[72].

The number of students enrolled in higher education modules related to AI has been rising modestly, led by higher degree research students (see graph). The AI related modules were

[71] http://search.ucas.com/search/providers?Vac=1&AvailableIn=2017&Query=artificial intelligence
[72] https://www.findamasters.com/search/courses.aspx?CID=GB&JD=8&JS=810

identified using the Joint Academic Coding System (JACS) employed by the Higher Education Statistics Agency (HESA). The data is aggregated into three levels of study: i) First Degree + other undergraduate, ii) Higher degree (taught) + other postgraduate, iii) Higher degree (research).

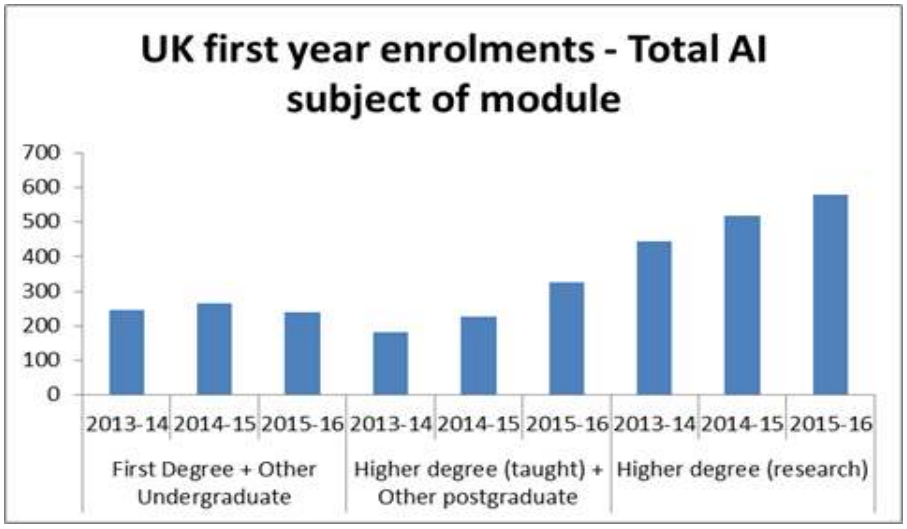

This demonstrates a steady increase in the number of enrolments, at Masters and Doctorate level at least, however, estimates both of potential and predicted growth in the use of AI in the UK would require significant increases in numbers at both levels to be realised in practice.

Applying forecast worldwide percentage growth rates in AI to UK enrolment numbers between now and 2020 suggests that a significant increase would be needed. The low, medium and high growth rate scenarios used are 15%,[73] 36%,[74] and 62%,[75] respectively. This methodology is crude, but goes some way to illustrating the scale of demand, see table below.

Demand for talent already outstrips supply, and average remuneration for data scientists and machine learning experts has increased substantially.

[73] No low scenario forecast report was found, thus 15% is an estimate based on the medium and high scenarios.
[74] http://www.transparencymarketresearch.com/pressrelease/artificial-intelligence-market.htm
[75] http://www.prnewswire.co.uk/news-releases/artificial-intelligence-market-report-2016---global-forecast-to-2022-artificial-intelligence-ai-market-is-expected-to-be-worth-usd-1606-billion-by-2022-at-a-cagr-of-629---research-and-markets-606786556.html

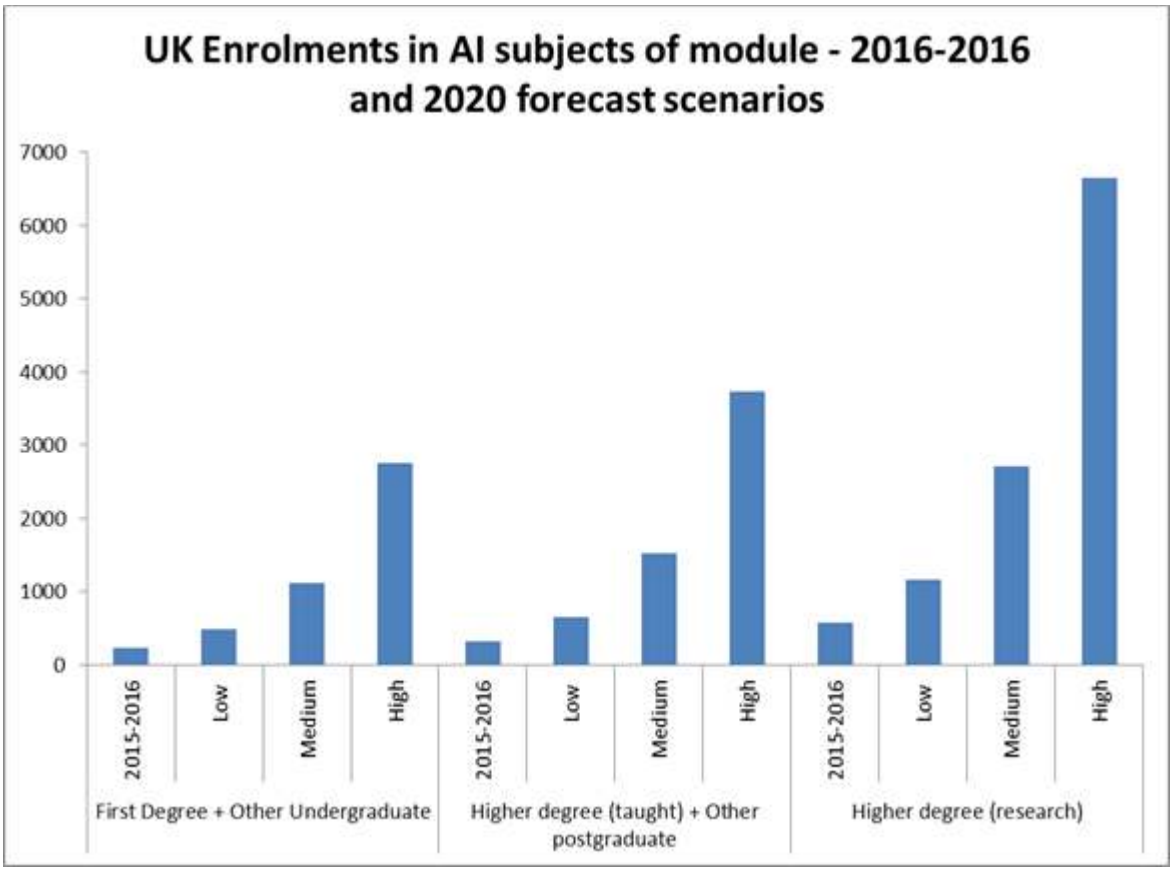

Increasing supply of skilled workers for AI

Supply can be increased by:

- Creating additional courses and places for generating new talent in the UK
- Incentivizing and improving the capacity to teach AI in the UK
- Improving the responsiveness of the skills training system to changing demand
- Attracting the best talent from other countries to the UK
- Reducing the gap between industry and academia
- Reducing the diversity gap by creating opportunities for women and other underrepresented groups.

As above, AI in the UK is already used in a broad range of organisations and sectors. However, to fully realise the potential of AI in the UK, additional sectors and different categories of organisation, with a mixed ecosystem of AI provider companies, small, medium-sized and large. All of these different organisations will need access to similar sets of skills, whether by hiring directly or contracting for services.

Therefore the recommendations here focus on increasing UK development and delivery of skills for the full range of potential employers, rather than narrowing the focus to, for instance, PhD skills for startups only.

Diversity

As above, ensuring that all, and not just some, people have the opportunity to work in AI it is necessary to create the largest and most talented potential workforce for AI, and to ensure that everyone has access to opportunities.

Diversity is particularly important for AI as the output quality of the algorithm depends on the assurance that the inherent bias of programmers does not transfer to code. A diverse group of programmers reduces the risk of bias embedding into the algorithm and enables a fairer and higher quality output.

Currently, the workforce is not representative of the wider population. In the past, gender and ethnic exclusion have been shown to affect the equitability of results from technology processes. If UK AI cannot improve the diversity of its workforce, the capability and credibility of the sector will be undermined.

Harvard economist Professor Iris Bohnet has researched unconscious bias and the structural opportunities to avoid poor and biased decisions in organisations[76]. Among other insights, her work indicates the benefits of anonymising and reviewing code.

To develop and apply AI and gain the widest set of social and economic benefits, it will be increasingly important to ensure that algorithmic biases are avoided in the selection of training data, design of algorithms and networks, and delivery of products and services. A diverse workforce is key to this.

While it is important to address the risks, AI also offers opportunities to support diversity and help ensure equitable treatment. AI can embed biases in systems, or it can reveal and disarm them. AIs can be developed that can detect biases, both in new AI-supported functions, but also in existing, historical systems that still influence decision-making in different sectors. AI can address the challenges faced by individuals because of unconscious bias, by bringing these to the surface more effectively than has been done in the past.

Ways in which the AI sector in the UK could improve diversity include :-

- **Demonstrating the advantages** that diversity brings to further development of AI, highlighting how to avoid algorithmic bias, and what the benefits are.
- **Breaking down stereotypes**: ensuring promotional material, course content and career opportunities are appealing to underrepresented groups.
- **Embedding unconscious bias training in selection** processes across universities offering AI related courses to ensure the application process is consistent and fair.
- **Embedding unconscious bias training in industry** for management and programming staff.

[76] https://scholar.harvard.edu/iris_bohnet/home

- **Corporate diversity initiatives,** for example offering mentoring programmes to underrepresented groups.
- **Offering additional support to women**, particularly at a later stage in their career to reduce dropouts. E.g. Flexible working arrangements and providing onsite childcare.
- **Leveraging the branding of the AI council** to ensure underrepresented groups have role models in academia and industry. The council should ensure it contributes to the function of breaking down stereotypes through all its engagements.

Recommendation: Government, industry and academia must embrace the value and importance of a diverse workforce for AI, and should work together to break down stereotypes and broaden participation.

When candidates are selected for the Masters, PhD and Fellowship programmes we advocate below, it is critical to use diversity-supporting criteria. The example of the Royal Academy of Engineering's Research Fellowship scheme may be one that could be replicated. The scheme limits applications from host institutions to two or three. However, host institutions are entitled to submit an additional candidate, if one of their candidates is considered to belong to a group that is persistently underrepresented within the engineering profession in the UK. Groups that have been evidenced as persistently underrepresented within the engineering profession are women, black, and minority ethnic.

Masters Level Skills

The challenge here is to train enough Master's level experts to meet industry demand.

Recommendation: Industry should sponsor a major programme of students to pursue Masters level courses in AI, with an initial cohort of 300 students.

The model proposed is a 15 month programme, designed to become recognised as a best-in-class training platform for machine learning graduates, funded by industry and providing skills directly to funding businesses. The first 12 months would be university training, assessed and accredited. The last 3 months of the programme would be an internship with one of the sponsoring businesses. A matching algorithm, like the deferred acceptance algorithm, will be used to ensure that each business gets a number of interns proportionate to their funding without having to identify the interns initially. The programme will be geared to attract an even ratio of genders, ensuring that the marketing and public appeal campaign is targeted for women.

Universities will apply to become members of this scheme and a panel of experts organised by the AI Council, industry and EPSRC will judge which universities are selected. Initial exploration suggests that 300 places could be funded in the first stage, expanding thereafter. We see the potential for such a programme to operate with 3000 or more places after successful implementation of the first stage, where a demonstration of the value to students and businesses is clear.

These places should be additional to the programs in existence but draw from the same pool of candidates, with the best candidates being selected for the studentships. Studentships should increase the pool of candidates, guaranteeing that the overall quality of the students to be maintained or improved despite their increase in numbers.

Each participating university must also maintain and improve teaching standards, as managed by the AI council. The participating universities will be reviewed periodically by the council to increase competition and allow upcoming institutions the opportunity to participate.

As the program matures, more industry partners and more universities can be invited, and a greater number of places funded, possibly with different tiers of funding, as long as the programme keeps a reputation for excellence. The programme could be extended to a wider range of companies including the “usage” sector, and if demand is sufficient, to sponsoring undergraduate places on relevant courses. These places should be created at qualifying universities across the UK.

Sponsoring businesses will benefit from the programme in the short and long terms, including from increased supply of high quality interns, effective leverage of the apprenticeship levy, and further on, an increased supply of higher quality candidates as the programme itself grows the market and attracts more students into AI.

In order for the businesses to leverage the apprenticeship levy, the AI council could apply on behalf of the sponsoring businesses to the Institute for Apprenticeships for approval of a new standard in machine learning/AI to support apprenticeships including up to Master’s level. The programme could draw on experience of the Data Analytics apprenticeships run by the Office for National Statistics and Welsh Government.

Masters’ Conversion Courses for Graduates

As AI is taken up more widely, there will be more and more diverse roles for people with some AI expertise, combined with other subject expertise. Graduates in subjects other than those directly involving computing have many of the fundamental skills for working on and with AI, and in particular for working on application of AI to their core specialties.

Skills in, for instance, “biomedical sciences plus AI” are likely to improve the employability of graduates enough to make this an attractive option for both employers and graduates. While they would not have as deep knowledge as those who have studied computing or data science throughout their training, there could be many opportunities for these combined skillsets.

Master’s courses directed at graduates in a wider range of subjects could also help increase diversity. There is a much higher proportion of women among graduates in biomedical sciences than in computer science. Courses aimed at this graduate cohort could help both to improve diversity in AI, and to accelerate its application in high-potential areas of science, technology and engineering.

The development of roles "working with AI" is at an early stage, and it is difficult to determine even the current state of demand. Therefore, this is best approached as market research on whether employers see value in this approach, and in what areas in particular. The objective is to determine whether such courses would contribute to meeting industry's demand for expertise, and whether graduates would respond to the opportunity.

Recommendation: Universities should explore with employers and students the potential demand for one-year conversion Masters degrees in AI for graduates in subjects other than computing and data science.

STEM graduates will have more of the fundamental skills, but a wider range of graduates will be needed in the AI workforce as it increasingly overlaps with ethics and social sciences.

Doctorate Level Training

The key challenge here is training enough doctorate level researchers to expand research capability. The UK has some of the best universities and some of the most respected academic experts in the world. Growing this capability will require additional investment in Doctorate level education. IT Jobs Watch, Tracking the IT Job Market noted that while demand has more than doubled since 2015, training provision has not.[77]

At a national scale, it is critical to ensure that sufficient investment is made into university research. Government support to academic institutions, through the research councils, is required to balance the demand and supply constraints of academia. Most of the UK's leading Universities note that the quality of the applicants for doctoral places is generally exceedingly high. Some have already co-authored publications. Demand from businesses and academia is such that Universities could support at least three times more. However, the Universities are limited both by practical constraints and funding available, so can only accept the small top fraction.

The UK faces strong competitive demand from overseas Universities for the brightest and best doctoral candidates. Unsuccessful applicants who cannot do doctoral training in the UK, because there are insufficient places, leave for opportunities overseas, often to institutions in the USA, including Stanford, NYU, MIT, Berkeley.

> Cambridge University's "Machine Learning" group see steadily rising numbers of applicants year on year. Apart from their own PhD program, they also run a joint PhD program with the Max Planck Institute for Intelligent Systems in Tübingen, Germany. This year (starting dates of Oct 1st 2017) they received about 107 PhD applications for the group, and 94 applicants for the Cambridge/Tübingen program. They are likely to give places to 15 applicants only

In 2015, the number of EPSRC doctoral students coded against AI research is 436. Of these, 80 are categorised as 100% AI. 212 are categorised as doing AI at least 50% (with the other 50%

[77] https://www.itjobswatch.co.uk/

in other areas, in particular robotics), with the other 224 doing AI less than the 50% of the time and often in more applied areas.

An approximate analysis by EPSRC suggests that demand for PhD places in machine learning and technical fields relating to AI is running at over 500 nationwide, with approximately 40% of the applications deemed to meet the minimum standard for acceptance.

As well as training and retaining UK nationals, UK universities should also actively seek to attract overseas students to study at this level. The UK as a whole will benefit from the best of the next generation of international AI experts being based here during and after their postgraduate research.

The number of available candidates can be increased by:

- Increasing supply of graduates and Masters students with degrees that will enable them to work in the AI industry.
- incentivising graduates to undertake AI PhDs in preference to other career options
- attracting the best candidates for AI PhD places from anywhere in the world.

Therefore this review recommends an immediate increase in supported PhD places to the full capacity possible now, and scaling up continually over time as the other measures below build the capacity of the system in terms of future candidates and additional academic experts able to teach them.

Attracting this talent and retaining it in the UK would make a transformational difference to the UK's capability to develop AI in new areas, and to our capacity to train the next generation of experts for the next wave of research.

Recommendation: Government and universities should create, at a minimum, an additional 200 PhD places dedicated to AI at leading universities. As the UK trains and attracts additional academic talent, this number should grow continually year on year.

Government and universities should continue to increase PhD places, supported by increases in numbers of qualified supervisors, institutions and candidates, aiming to exceed 1000 new UK PhDs in AI-related subjects by 2025.

In order to achieve such an ambitious increase in the number of PhD places in the UK, changes will need to be made in immigration, student funding and incentive schemes to ensure we can match supply with demand. This includes:

- An appropriate immigration measure to not only facilitate student entry into the UK but also incentivize graduates to remain in the UK to either continue postgraduate education and research or enter employment with a UK company.
- Changes to the rules for funding PhD places that would allow international candidates to fairly compete for grants and scholarships in the UK, this is coupled with the encouragement that charitable institutions such as the IET and The Royal Academy of Engineering expand current schemes to attract international talent to the UK

- Incentive programmes for UK students to proceed to postgraduate education. This may be in the form of awareness campaigns on the opportunities postgraduate education in AI offers, scholarships for students with financial challenges, debt relief and potential employer relations (through industry funded places).

Massive Open Online Courses and Continual Professional Development in AI

As evidenced by the *Introduction to Cybersecurity*, Massive Online Open Courses (MOOCs) are an excellent way to increase awareness of and teach basic skills in digital technology areas where there is high demand for skills. Over 80,000 people[78] have now completed the *Introduction to Cyber Security* which was developed with the Open University and was the first of its kind anywhere in the world to gain government support.

Accessible training that helps people with the right basis of knowledge to make this transition would help expand supply of AI professionals, and could help to develop understanding of how AI can deliver value among a much wider group.

Additional accredited Continuing Professional Development courses can also help people improve expertise in flexible ways that suit their personal circumstances and ambitions.

Recommendation 7: Universities should encourage the development of advanced credit-bearing AI MOOCs and online Continuing Professional Development courses leading to MScs for people with STEM qualifications to gain more specialist knowledge.

The key to success of this recommendation will be for a pool of key universities to agree on developing and assign credits to the same set of courses.

Attracting and Retaining Academic Talent: International Turing AI Fellowship Programme for the UK

Rising industry demand for highly specialized AI skills - and rising salaries - can lead academics to move into businesses. This places a strain on the resilience and capacity of the academic network to continue blue sky research and to train talent.

A solution that allows academics to freely choose between contributing to industry and academic progression, by levelling the playing field, is required. The UK must support its

[78] https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/516331/UK_Cyber_Security_Strategy_Annual_Report_2016.pdf

academics to thrive and be recognised in a global market, enabling them to contribute further to research and academic collaboration.

A national academic fellowship programme in AI would incentivise academics both locally and internationally to engage in collaborative research and higher education. The fellowship would have ties to research focused industry partners, allowing fellows to freely transfer between organizations to engage in solving challenging problems faced by industry and supporting academic research. The fellowship would include a financial stipend to bring academic salaries to compete with those of industry and will closely emulate the Canadian Institute for Advanced Research (CIFAR) Global Scholars Programme.[79]

Recommendation 9: An International fellowship programme for AI in the UK should be created in partnership with the Alan Turing Institute: the Turing AI Fellowships. This should be supported by a targeted fund for identifying and recruiting the best talent, and by ensuring that the UK is open to any and all of the eligible experts from around the world.

The Turing should scope the number of Fellowships to begin with and develop the programme over time. The aim is to create a fellowship programme that is globally respected and attractive for researchers around the world to congregate in the UK.

The programme would be linked to AI businesses to help guide and set the strategic aims of the fellowship. The fellowship programme would encourage and enable, through standardised negotiations, mixed positions for academics whose term are already defined. These mixed positions would allow experts to hold position in both industry and in academia, with the freedom to move between the two. The fellowship programme would be jointly funded by industry sponsors and Government. Incoming fellows would receive a Tier One (Exceptional Talent) Visa.

While developing the scheme, Government and partners should explore whether the best value could be derived from funding salaries not only for the senior researchers, but also for key members of their teams and other research costs.

To ensure the programme is internationally recognised in AI, the Alan Turing Institute would act to align the programme with their strengths and their ambitions for AI.

[79] https://www.cifar.ca/global-academy/

SECTION FIVE: RECOMMENDATIONS TO MAXIMISE AI RESEARCH AND COMMERCIALISATION

This section recommends targeted actions to grow AI research and commercialisation:

- The Alan Turing Institute: National Institute for Artificial Intelligence
- Access to lower-cost computing capacity for research
- Improving transfer of Intellectual Property.

Alan Turing Institute: National Institute for Artificial Intelligence and Data Science

As stated in the two previous recommendations, the Alan Turing Institute can play a key role in making the UK a powerhouse of AI research. The name “Alan Turing” has a unique significance for AI, as Turing is widely considered the father of Artificial Intelligence.

We believe, in agreement with the leadership of the institute, that the institute would be able to deliver even more effectively if it responded to this development by giving AI a central role in its focus, brand and position in the sector. At the same time, the institute should be expanded beyond the five founding universities, to bring in additional AI expertise and to grow into a truly national institution.

Recommendation 10: The Alan Turing Institute should become the national institute for artificial intelligence and data science, becoming truly national and expanded beyond the current five universities, with a key stated aim that centres its mission on artificial intelligence.

The AI institute should coordinate with leading AI institutes in other countries (eg the Vector Institute in Canada) to explore the potential for future collaboration, for instance on the model of CERN. The Turing Institute could also develop engagement with industry through the AI Council recommended below.

Improving transfer of Intellectual Property

A key component that drives the creation (and success) of new businesses in AI is the ability and capacity for ideas and technologies to spin out of the university network, or be licensed from it, and be commercialised. Universities have a valuable role to play by nurturing and

supporting those with entrepreneurial talent and contributing to the ambition to make the UK the best place in the world to start and grow a business.

As part of their mission to deliver impact for society, including the economy, universities should ensure that the primary objective of their approach to commercialisation of research in AI is the exploitation of intellectual property (IP), not just its protection.

Current spin out practices and processes can be complicated and long in duration. Different universities employ different approaches and provide varying levels and quality of support. Allocation of equity is a particularly complex and contentious process. There is a clear need to improve processes in this area to reduce barriers to commercialisation and facilitate innovation in AI.

Recommendation 11: Universities should use clear, accessible and where possible common policies and practices for licensing IP and forming spin-out companies.

Clear processes with guidelines that promote transparency of the conditions upon which IP can be commercialised would greatly aid AI innovation. This includes universities being clear about their equity requirements, provision of support and continued relationship with the business after the spin out.

There may be approaches to improve this situation using incentives. Corporate investors in AI research could develop terms and shared criteria to be used in all relevant agreements, in order to be eligible for investment. For example, the university will take no more than a 5% golden share equity stake which cannot be diluted up to an investment value of a certain value. There could also be a time limit set for negotiations.

There could also be a place for university participation in publicly-funded accelerator and incubator schemes to be dependent on a set of criteria on transfer of IP. AI Enterprise Fellowships could be established, with eligibility for the Fellowship requiring certain conditions to be met.

Universities often do not have sufficient resource to work at the same pace as businesses, preventing them from fully realising the potential of research outputs. Universities also often lack the technical sector expertise. Access to a shared group of technical experts in AI commercialisation could help universities through the process.

Access to lower-cost computing capacity for research

The most significant cost for research organisations investigating the potential applications of AI is high performance computing power. Reducing the cost and friction of accessing this computing power would help to maximise the return on investment made by the UK government through the research councils in AI research, ensuring that grants are spent as efficiently as possible. In an increasingly competitive international research environment, the availability of relatively low-cost computing power would also ensure that the UK remains a particularly attractive location for the world leading researchers to base themselves, and encourage collaborations between UK institutions and centres of expertise overseas.

Additional investment in public-sector high performance computing power would be welcome, but a reduction in the price of computing power does not necessarily have to come at an additional cost to the public funding. There may also be opportunities to aggregate demand across institutions in a way that maximises their bargaining power with commercial providers. There is an opportunity for the UK, through a lead institution to manage the high performance computing requirements institutions, gaining capacity at lower cost.

Collective action could also improve planning for future capacity for broader UK research sector.

Recommendation 12: The Alan Turing Institute, Engineering and Physical Sciences Research Council (EPSRC), Science and Technology Facilities Council (STFC) and Joint Information Systems Committee (JISC) should work together to coordinate demand for computing capacity for AI research, and negotiate for the UK research community.

SECTION SIX: RECOMMENDATIONS TO SUPPORT UPTAKE OF AI

In order for the economy to realise the benefits of AI, more needs to be done to develop and apply solutions across a wide range of applications. This section recommends actions to broaden and increase uptake.

Supporting Coordination and Collaboration among AI leaders in the UK

New industry sectors lack coordination and representation. This is natural. Technology that develops fast leaves businesses with little or no spare resource for engaging with organisations outside their immediate sector or customer base, or with policy-makers.

AI in the UK appears to lack effective levels of non-competitive collaboration both between businesses and between businesses and academia. The sector lacks clearly identifiable leadership. Unlike many traditional sectors that sell to a narrow and clearly defined market, AI opportunities exist right across the economy, making it more difficult for the businesses to identify and present themselves as AI businesses.

More and better collaboration by UK AI businesses could offer many benefits, including partnering with academia to develop the skills pipeline, and common engagement with government and institutions to develop and share successful means to manage data. Without it, the sector will be poorly represented to the rest of industry, to the public, and to the public sector. Coordination can help an emerging sector progress from success in niches to widespread uptake and acceptance.

As noted above, AI can also create new situations with new implications for fairness, transparency and accountability. AI could also change the nature of many areas of work.

AI in the UK will need to build trust and confidence in AI-enabled complex systems. There is already collective activity to work towards guidelines in ethics for automation, but we can expect this field to grow and change.[80] A publicly visible expert group drawn from industry and academia, which engages with these issues would help to build that trust and confidence.

Experts with insights into emerging developments would also be able to provide timely advice on trends in AI development and application, enabling planning of upskilling and reskilling.

In addition, for the recommendations in this report to have an enduring relevance, implementation will need to be supported by industry, academia and government.

Recommendation 13: Government should work with industry and experts to establish a UK AI Council to help coordinate and grow AI in the UK.

[80] BS 8611:2016 Robots and robotic devices. Guide to the ethical design and application of robots and robotic systems http://shop.bsigroup.com/ProductDetail?pid=000000000030320089

The intention is that this Council will operate as a strategic oversight group, establishing an open and non-competitive forum for coordination and collaboration between industry, the public sector and academia, ensuring that skills deficits are identified early and addressed, academic challenges framed in terms which have commercial implications, and that commercial funding flows into academia to bolster the large-scale current and future government investment in this discipline as and when required. The Council will be closely engaged with the Alan Turing Institute to draw on its expertise and to maximise contribution to common objectives. The Council should seek a Cabinet-level champion in Government, perhaps as a co-chair.

The council will promote cooperation within industry, improve understanding of AI in the UK, and identifying any cross-cutting barriers to growth and innovation encountered by AI businesses. As and when any barriers are identified, it will also play an important role in bringing industry, academia and government together to develop appropriate solutions on an ongoing basis.

The Council would act as an expert leadership group, rather than as a trade association for the AI sector, but could act as the catalyst and partner for industry representation.

Government has partnered with the cybersecurity sector in a number of connected initiatives, and successfully helped the sector to grow and to attain an international profile for UK businesses and experts. The Council should work with Government to explore how that model of partnership with companies and experts in AI could develop a similar programme to develop capability and visibility and increase uptake.

The Council would take a leading role in dialogue and activity to ensure that the UK becomes a leader for application of responsible innovation in AI. It should also seek to ensure that global collective action, for instance by the Partnership for AI, is supported and driven by UK organisations.[81]

The council will also be responsible for a number of specific deliverables including:

1) Driving coordinated action to deliver the recommendations of this Review, where stakeholders agree to take them forward.

2) Driving action to improve supply of skills, including coordinating the funded Masters studentship, the National AI fellowship programme and development of MOOCs. To succeed in this, the Council will need to develop a thorough overview of skills flows: through the education and skills development system; in industry; and geographically.

3) Advising policymakers on emerging trends and opportunities.

4) Oversight of the data trust programme to ensure that it remains purposeful, remaining impartial, relevant and technically strong.

[81] https://www.partnershiponai.org/

5) Providing industry expertise into discussions on issues of fairness, transparency, accountability and diversity in relation to AI, including foresight of likely developments.

The Council could provide input for UK AI in discussions about ethics, as the Turing Institute should for AI researchers. Neither should be ultimately responsible for delivering the governance and oversight functions proposed by Royal Society and the British Academy, as that would lead to conflicts of interest.

The Council should work with existing bodies that already deliver support to digital businesses in key areas, including Tech City UK and the Digital Catapult.

Explainability of AI-enabled uses of data

There is a growing consensus that decisions which affect people and are made on the basis that decisions based on data analysis data should be fair, and should be demonstrably fair. As mentioned above, the Royal Society and British Academy sets out a principle for data governance to "ensure that trade-offs affected by data management and data use are made transparently, accountably and inclusively".

Many AI innovators are among the advocates for ethical principles in data-driven activities. The Partnership on Artificial Intelligence to Benefit People and Society was founded by Google, DeepMind, Amazon, IBM and Microsoft to share and promote ethical application, and to demonstrate the commitment of the major companies in that area.[82]

This is already an area of interest to Governments, regulators and sector experts. The ICO is involved in an EU group exploring this area. The Alan Turing Institute has an interest group on fairness, transparency and privacy. Some in the commercial world are already advocating proactive self-regulation.

Demonstrating the transparency of AI-driven decisions can be difficult. When a machine learning application has trained itself iteratively on datasets, it can be hard to provide an accessible explanation of the precise factors for one single decision. "Neural networks, especially with the rise of deep learning, pose perhaps the biggest challenge—what hope is there of explaining the weights learned in a multilayer neural net with a complex architecture?[83]

The General Data Protection Regulation (GDPR) will take effect on 25 May 2018, and establishes substantial new conditions for access to and use of data. The regulation gives certain rights to data subjects (people whose personal data is used) including in relation to explanation of automated data processes (Articles 12, 13, 22). However, some commentators

[82] https://www.partnershiponai.org/ - s-founding-partners
[83] "European Union regulations on algorithmic decision-making and a "right to explanation"'. Bryce Goodman, Seth Flaxman. August 2016. https://arxiv.org/pdf/1606.08813.pdf

suggest that implementing and complying with the regulation will encounter challenges in practice.[84]

To help address these challenges, AI businesses would benefit from common, practical guidance on how to provide assurance to the public, customer businesses and regulators. There is a risk of confusion and a chilling effect on broad application of AI, if it is not clear how to comply with provisions in relation to AI-supported processes.

It would be advantageous for the UK to show a lead in achieving and demonstrating compliance with these provisions, supporting development of overarching principles and common, repeatable practices for explainability.

Recommendation 14: The Information Commissioner's Office and the Alan Turing Institute should develop a framework for explaining processes, services and decisions delivered by AI, to improve transparency and accountability.

When exploring solutions for explainability, ICO and the Turing Institute should engage with related activity within Government and standards and industry bodies (including BSI and IEEE) and with any data governance body that is developed in the future (see the Royal Society / British Academy work noted above).

Further on, it is possible that new applications of AI may hold solutions on transparency and explainability, using dedicated AIs to track and explain AI-driven decisions.

Access to Global Markets

In order to gain the maximum benefit from the current comparative advantage in the development and application of AI systems, and support the longer-term growth of the AI sector, the UK should seek to be a major exporter of products and services, and to compete with other countries to attract other world-leading companies to the UK.

The Department for International Trade already supports inward investment and trade and promotes the UK as a centre of AI excellence, but more can, and should, be done.

Recommendation 15: The Department for International Trade should expand its current support programme for AI businesses.

Including:

- Additional emphasis on AI in the Global Entrepreneur Programme, adding a new 'dealmaker' dedicated to supporting AI companies, and encouraging the creation of a

[84] Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation https://www.academia.edu/31045353/Why_a_right_to_explanation_of_automated_decision-making_does_not_exist_in_the_General_Data_Protection_Regulation?auto=download

dedicated AI seed competition, to be added to the list of 'approved endorsements' for entrepreneur visa applications.

- A series of AI focused missions for specific sectors and markets, including the US, China, India and Japan, using these as opportunities to both sell products and services, and make initial contact with companies who might be amenable to moving to the UK.

- Running a major international event as a showcase for UK firms, demonstrating the products on services on offer to major international investors and potential customers for products and services developed by UK companies.

- Building on the success of the Global Entrepreneurs Programme, Sirius and Women in Technology to develop and launch a way of attracting female business leaders, academics and technical experts to the UK.

Signposting the Opportunity of AI for UK Industry

For most businesses in most sectors, AI is new, perceived as complex, and not well understood. It is difficult for most businesses to understand how AI can improve their productivity, products and services.

McKinsey's survey of 3,000 businesses around the world "found that many business leaders are uncertain about what exactly AI can do for them, where to obtain AI-powered applications, how to integrate them into their companies, and how to assess the return on an investment in the technology."[85]

> "Business leaders are asking: What impact will AI have on my organisation, and is our business model threatened by AI disruption? And as these leaders look to capitalise on AI opportunities, they're asking: Where should we target investment, and what kind of capabilities would enable us to perform better? Cutting across all these considerations is how to build AI in the responsible and transparent way needed to maintain the confidence of customers and wider stakeholders."
>
> **PWC Report June 2017 - Sizing the prize: What's the real value of AI for your business and how can you capitalise?**[86]

There is a role for better information, generic and sector-specific, to improve understanding of opportunities.

[85] Artificial Intelligence: the next digital frontier. June 2017. http://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/how-artificial-intelligence-can-deliver-real-value-to-companies
[86] http://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf

Tech UK has the capability and experience to bring together resources and information from AI experts and providers. Action to develop resources should also bring in expertise of organisations that work with target industry sectors, including the Royal Academy of Engineering and the CBI.

Recommendation 16: To help grow demand, TechUK should work with the Royal Academy of Engineering, the Digital Catapult, providers and representative organisations in industry sectors, to develop practical guidance on the opportunities, challenges and practical actions around successful adoption of AI across the UK economy.

The best approach may be to begin with one or two sectors where AI appears to have high potential, but is currently poorly understood.

Coordinated activity could include:

- Guidance on end-to-end assessment of where and how AI could add value, identifying opportunities and prioritising investment
- Scoping and publishing skills and roles for implementing AI and working with AI, including skills needed by senior decision-makers
- Guidance on anticipating, monitoring and managing changes to working practices
- Guidance on building trust, among decision-makers, staff, customers and partners.

Coordination in this area should also connect with bottom-up activity already going on. As above, some accountancy and law firms have engaged very positively with AI, and may have lessons to share on what has successfully enabled that engagement. The Knowledge Transfer Network has extensive experience of improving understanding of emerging sectors.

Government as a User and Customer

As above, some Government departments and agencies are already using AI, but AI has the potential to make much greater contributions to public services, by:

- improving the flow of information and analysis in services, including early and better insights from economic and other data
- improved decision-making in complex areas
- managing new and increasing data resources, for example from internet of things and smart city applications
- improving the efficiency, effectiveness and usability of services for users.

These improvements can help deliver better outcomes for service users, reduce costs in regulatory compliance, and generate significant efficiencies through reduced overheads and back office costs.

Use of AI services by the public sector can also provide an important stimulus for providers, developing confidence amongst investors and other potential customers around the classes of products being procured, and challenging industry to reach beyond current capabilities in pursuit of truly innovative solutions to some significant social challenges.

Smart government procurement, using tools designed specifically to encourage and support small and medium sized enterprises, to reward commercial investment in research and development, and encourage businesses to start in communities right across the country, can help grow UK AI. Some existing initiatives (including the Small Business Research Initiative, the Innovation Partnership Procurement scheme, and the G-cloud procurement portal) can be adapted to support developing AI applications.

The challenge may be to spread "AI-readiness" beyond existing pockets of innovation and expertise, and right across the public sector. To help the public sector to realise the benefits offered by AI, Government will need to ensure that senior decision-makers and operational staff across public bodies have the skills and practical tools needed to recognise and realise how AI can help them deliver their responsibilities and objectives. Much of the core work in government still happens in-house, aimed at incrementally improving systems and is limited by the capability available to departments. Senior decision-makers will need to make informed decisions on costs, benefits and risks in using AI in their operations and services, and will need effective tools to support decisions. Opportunities for AI in the public sector will continue to evolve, and Government will need to identify and spread practical guidance and examples of best practice iteratively.

Recommendation 17: Government, drawing on the expertise of the Government Digital Services, the Data Science Partnership and experts working with data in other Departments, should develop a programme of actions to prepare the public sector and spread best practice for applying AI to improve operations and services for citizens.

This programme could include:

- Guidelines on making data ready for use, including managing privacy and security, preventing biases in training data
- Technical codes of practice for using AI
- Guidelines on ensuring equitable treatment, transparency and reporting of AI-driven functions especially when used to support decision making
- Frameworks for testing for diversity in AI applications and datasets
- Demonstrator projects, including to demonstrate how AI can assist protection of individuals' interests
- Sandbox environment for secure testing of approaches, which may include AI and machine learning
- Information for senior decision-makers on where AI can add value, and on assessing costs, benefits and risks
- Scoping roles and skills for delivering, managing, oversight of AI-driven functions, including for working with AI
- Monitoring and preparing for changes to working practices; monitoring developing skills needs, and support for skills development
- Supporting demonstrator projects to identify and codify best practice
- Outreach events and guidance that help ensure government becomes an intelligent customer of AI services supplied through the Digital Marketplace

- Supporting transformational departmental policy programmes that have been identified as potentially benefiting from using AI machine learning algorithms as part of the solution e.g. the Data Enabled Change Accelerator Programme (DECA)
- Manage flexible access to computing resources for AI for the public sector
- Publicly championing examples where AI has improved outcomes for the public
- Mechanisms to connect AI experts with datasets held by the public sector
- Supporting Masters and PhD students, and senior academics in AI to undertake placements and project work with public organisations
- Extending the Digital Marketplace to include services from industry partners providing specialist AI and data science services.

The programme should offer support to the broad UK public sector, including Devolved Administrations and the local public sector.

There could be a role for a dedicated public sector AI Innovation Fund, into which Departments could bid to fund for proof of concept AI demonstrator projects, available to UK SMEs only, to help develop the national pipeline of expertise. An additional mechanism like this could enable policy-makers and public service leaders to test applications in a supported system without having to compete for internal budgets.

Stimulating AI Innovation for Public Challenges and with Data held by Public Organisations

AI presents a transformational opportunity to solve the grand public challenges of our time, from the rising cost of healthcare, stalling productivity, managing growing cities, and cybersecurity risks. AI also offers potential for new dimensions of performance in research, and in personalisation of products and services.

Stimulating innovation by inspiring industry and academia to focus on solving a public challenge is a well respected and proven method. Public funds can have a key role here, and so can data held by public bodies.

These advances will not be made by pure research (though that is a UK strength, vital to continued progress and to preserving the UK's advantages in AI), as much as by research on applied AI. Therefore funded challenges on applying AI to major challenges hold the most promise for supporting UK innovations in high-value areas. The UK has mechanisms for funding applied R&D.

The Industrial Strategy Challenge Fund (ISCF) was created within the government's Industrial Strategy to ensure that the UK's strengths in research and innovation deliver even more tangible results with economic and public benefits. ISCF will help innovative businesses develop products and services that become UK strengths in large and growing global markets.

The fund will support challenges on a distributed and local basis, helping incentivise innovations and solutions to issues that may not be funded in the private sector without government encouragement, but are of high public value, for example:

- Improving the efficiency of social care provision by actively mapping capacity, logistics, demand and forecasting
- Improving the resilience of UK industry from cyber attracts using active search, natural language processing and automated code and security integrity verification methods.

The fund can be used to incentivize and stimulate the use of data held by public organisations to solve challenges across a wide range of areas. The challenges will be identified through a consultation process by Innovate UK and EPSRC, bringing in a range of public service bodies with ownership of data sources with academia and industry to solve the challenge.

Areas where funded challenges applying AI could address public needs and build on UK strengths include:

- cybersecurity
- personalised and integrated healthcare
- personalised education and training
- integrated transport for smart cities
- improving efficiency of infrastructure
- personalised public services
- digital manufacturing in key sectors such as pharmaceuticals and aerospace.

Targeted challenges could involve combined use of data in public organisations with data held in commercial organisations but of great relevance to resolving public challenges, eg data from transport providers and utilities, or relating to planning, or to dietary patterns. There could be a role for incentives for private sector partners to contribute data to resolve challenges. There may be lessons to be learned from Nesta's Open Data Challenge Series, when designing challenges.

The Small Business Research Initiative (SBRI) offers an additional, proven approach to support development of innovative industry solutions to public sector challenges. SBRI challenges can improve public services and generate new business opportunities for companies, provides small and medium-sized enterprises (SMEs) a route to market for their ideas. By doing this it can also bridge the seed funding gap experienced by many early stage companies.

As AI applications develop, there may also be a place for developing competitions which support experimental uses of data to identify new objectives that could be served.

Recommendation 18: Government should ensure that challenges addressed by the Industrial Strategy Challenge Fund (ISCF) and Small Business Research Initiative (SBRI) are designed to attract and support applications of AI across the full range of challenge areas and set funded challenges which use public sector data for AI.

CONCLUSION

This report outlines the areas that AI experts in business and in research have identified as key to growing UK AI capability and realising its opportunities: managing access to data more efficiently, growing the talent pool, maximising the output of UK AI research, and supporting the uptake of AI through funded challenges and better coordination and understanding of AI among potential users.

There are great opportunities for AI to do much more for the citizens and the economy of the UK, the actions recommended in this report are just a first step towards realising that potential. Sustained collaboration between Government, academia and industry will be necessary to implement these recommendations now, and to continue delivering the amazing potential of AI in the future.

Acknowledgements

This Review would not have been made possible if it weren't for the strong support we received from a wide variety of people at many levels in the Department Business, Energy and Industrial Strategy and the Department for Digital, Culture, Media and Sport – including Secretaries of State from both departments. Officials tend to remain nameless and faceless to the public but particular thanks is due to Dev Amratia (BEIS) and Ben Hawes (DCMS). They planned and delivered our ambitious engagement process, catered for our challenging schedules and were a real pleasure to work with.

We are also grateful for the contributions to this Review from the following organisations and their employees. They provided the substance of this report through workshops, meetings and ad hoc exchanges:

Accenture
Airbus
Alan Turing Institute
Amadeus Capital
ASI Data Science
Atomico
Aviva
Barclays
BP
British Academy
BUPA
CBI
City of London
CognitionX
DeepMind
Digital Catapult
Entrepreneur First
EPSRC
Facebook
GCHQ
GSK
IBM
Imperial College
InnovateUK
Malvern Instruments
McKinsey & Co.
Microsoft
Nvidia
Ocado
Onfido
Open Data Institute
Pfizer
Prof. Richard Susskind
PwC
Rolls Royce
Royal Academy of Engineering
Royal Society
Sage
SageSainsburys
Shell
Siemens
Switch Concepts
TechCity UK
techUK
UCL
University of Cambridge
University of Edinburgh
University of Oxford
University of Sheffield
University of Southampton
Your.MD