

UK AI COUNCIL

AI ROADMAP

The AI Council is an independent expert committee. It provides advice to the UK Government, as well as high-level leadership of the Artificial Intelligence (AI) ecosystem.

This independent report draws on the expertise of its members and those in its wider ecosystem to summarise four pillars on which to build the UK's future in AI. It invites action across government to keep the UK at the forefront of safe and responsible AI.

UK AI COUNCIL | AI ROADMAP

Executive summary	2
16 recommendations to help the government develop a UK National AI Strategy	4
Introduction	6
1. Research, Development and Innovation	10
The Alan Turing Institute: local presences and global reach	11
Challenge-led innovation	12
AI to transform research, development and innovation	14
2. Skills and Diversity	16
What everyone needs to know	16
Further, Higher and Graduate Education	17
Supporting every teacher	18
Lifelong learning and an online academy for AI	18
3. Data, Infrastructure and Public Trust	20
Trustworthy and accessible data	20
Public trust and good governance	22
Digital Twins and Living Labs	24
A globally leading role	25
4. National, Cross-sector Adoption	28
Business as smart adopters	28
Supporting high growth AI startups	29
Enabling public sector adoption	30
Health and social care	31
Climate change	32
Defence and security	33
Acknowledgements	36

Executive summary

As Artificial Intelligence (AI) becomes embedded in people's lives, the UK finds itself at a pivotal moment.

Estimates show that AI could deliver a 10% increase in UK GDP in 2030.¹ If approached correctly, this would offer huge benefits to the economy, to recovery and resilience, the environment and for people from all walks of life across all parts of the UK. This summary report by the AI Council sets out long-term ambitions and suggests near-term directions for all government departments, with the aim of cementing the UK as one of the very best places in the world to live with, work with and develop AI.

The UK starts from a place of strength in research, enterprise and regulation, and with its history of recent support for AI it stands among the best in the world. The UK government's ambition should push

for scale and reliability in areas of unique advantage. To be influential in attracting talent, shaping global markets and global governance we call on the government to build a UK National AI strategy that scales up its investment for a decade and beyond.

This report has two underlying messages. The first is that we need to "double down" on recent investment the UK has made in AI. The second message is that we must look to the horizon and be adaptable to disruption. Support for AI needs to reflect the rapid pace and evolution of the science & technology and its applications. This means staying at the forefront of the development of AI and integrating approaches to ethics, security and social impacts and planning for the next 10-50 years. The UK will only feel the full benefits of AI if all parts of society have full confidence in the science and the technologies, and in the governance and

“**The UK will only feel the full benefits of AI if all parts of society have full confidence in the science and the technologies, and in the governance and regulation that enable them**

“**A National AI Strategy is needed to prioritise and set a time frame that will position the UK for success**

regulation that enable them. That confidence will depend on the existence of systems that ensure full accountability, clear ethics and transparency. Developing the best science and the most robust applications requires commitment to an ambitious programme of investment in talent; one that promotes cutting edge skills and does so in ways that makes AI accessible in ways that are more diverse and inclusive.

Given the breadth and significance of potential AI applications, this Roadmap is concerned with overcoming the challenges of enabling change. It sets out suggested directions across four pillars: Research, Development & Innovation; Skills and Diversity, and Data, Infrastructure and Public Trust - and then, in the final section, addresses some specific measures to support adoption and the key areas of health, climate and defence.

This Roadmap and its recommendations reflects the views of the Council as well as 100+ additional experts. We recognise it will not be possible to achieve all at once, and this is why a National AI Strategy is needed to prioritise and set a time frame that will position the UK for success. In order to support the government, we are ready to convene workshops with the wider ecosystem in order to capture more detail and work together to ensure that a future National AI Strategy enables the whole of the UK to flourish.

¹ [The economic impact of artificial intelligence on the UK economy](#), PwC (June 2017)

16 recommendations to help the government develop a UK National AI Strategy

Research, Development and Innovation

- 1. Scale up and make sustainable public sector investment in AI; ensure consistent access to top talent from around the world; and find new ways to bring researchers, disciplines and sectors together.** Build on the commitments in the government's R&D Roadmap² and suggestions in the soon to be published UKRI AI review³.
- 2. Cement The Alan Turing Institute as a truly national institute, with a set of regional investments that draw on strengths from across the UK.** Provide assured long term public sector funding that will give the Turing and others the confidence to plan and invest in strategic leadership for the UK in AI research, development and innovation.
- 3. Ensure moonshots, as described in the R&D Roadmap as challenge-led, high-risk, scalable programmes, are both advancing and leveraging AI.** These could tackle fundamental challenges such as creating "explainable AI", or important goals in any area where AI can contribute strongly, such as the UK Digital Twin program or developing smart materials for energy storage in the move towards Net Zero carbon emissions.

Skills and Diversity

- 4. Scale up and commit to an ongoing 10 year programme of high level AI skill-building.** This would include research fellowships, AI-relevant PhDs across disciplines, industry-led Masters and level 7 apprenticeships.
- 5. Make diversity and inclusion a priority.** We suggest benchmarking and forensically tracking levels of diversity to make data-led decisions about where to invest and ensure that underrepresented groups are given equal opportunity and included in all programs.
- 6. Commit to achieving AI and data literacy for everyone.** The public needs to understand the risks and rewards of AI so they can be confident and informed users. An online academy for understanding AI, with trusted materials and initiatives would support teachers, school students and lifelong learning.

Data, Infrastructure and Public Trust

- 7. Consolidate and accelerate the infrastructure needed to increase access to data for AI.** Invest in the relevant organisations, link general principles to specific applications, and pursue initiatives for pump priming innovation and enabling safe data sharing for valuable uses.
- 8. Lead the development of data governance options and its uses.** The UK should lead in developing appropriate standards to frame the future governance of data.
- 9. Ensure public trust through public scrutiny.** The UK must lead in finding ways to enable public scrutiny of, and input to, automated decision-making and help ensure that the public can trust AI.
- 10. Thoughtfully position the UK with respect to other major AI nations.** Building on its strengths, the UK has a crucial opportunity to become a global lead in good governance, standards and frameworks for AI and enhance bilateral cooperation with key actors.

National, Cross-sector Adoption

- 11. Increase buyer confidence and AI capability across all sectors and all sizes of company.** Support investment for local initiatives to enable safe value-creating innovation and improve the data maturity needed for AI innovation.
- 12. Support the UK's AI startup vendor community.** Enable greater access to data, infrastructure, skills, compute, specialist knowledge and funds.
- 13. Enable robust public sector investments in AI,** building capability in the use of data, analytics and AI to ensure intelligent procurement of AI as part of projects for public benefit.
- 14. Use AI to meet the challenges of Net Zero carbon emissions.** Work on access to data, governance, to develop cleaner systems, products and services.
- 15. Use AI to help keep the country safe and secure.** Work with government departments/agencies and defence and security companies to ensure AI is available to assess and respond to modern defence and security threats and opportunities.
- 16. Build on the work of NHSX and others to lead the way in using AI to improve outcomes and create value in healthcare.** The UK's comparative advantage will depend on smart strategies for data sharing, new partnership models with SMEs and skill-building.

² [UK Research and Development Roadmap](#), UK Government (July 2020)

³ [To be published early 2021]

Introduction

Artificial Intelligence (AI) is breaking into the mainstream of societies and the global economy.

The history of AI is characterised by cycles of boom and bust, with waves of excessive hype followed by “AI winters”. This time things are different: advances in machine learning and neural networks have led to great advances in fields such as computer vision and speech recognition, and natural language processing, with numerous applications used by hundreds of millions of people on a daily basis. Safely harnessed AI is revolutionising certain types of tasks, disrupting business models, and complementing and enhancing human capabilities in wholly new ways. These tremendous advances have come from increasing the sizes of training data sets, and the associated networks and compute power, by factors of a million or more, along with modifications to the

underlying algorithms, many of which have been with us in some form for a long time. This accounts for tens of billions of dollars in global annual value creation, and looks set to increase markedly over the coming decade.⁴

The UK has fundamental strengths in AI which reach back to the intellectual origins of the science, with figures like Alan Turing, Joan Clarke, Donald Michie and many others, to its ability to debate and enact sound governance of new technologies. Successive governments, businesses and civil groups have worked hard to make increasing amounts of data and information more available in ways that combine safety and usefulness.

In recent years, the UK has been the site of an array of initiatives aimed at developing and deploying AI across the economy for society. Building on these with focus and structure will ensure AI is

an essential part of tackling the major challenges of economic recovery following the COVID-19 pandemic, of levelling up the UK economy through the creation of jobs and prosperity beyond London and the South East, of ensuring new forms of health and social care, of reaching net zero carbon emissions, and of ensuring resilience to future economic, health, or environmental shocks. It will also create entirely new opportunities for human flourishing.

These fundamental strengths and recent initiatives have helped the UK to be amongst the global leaders in many areas of AI, from accelerating drug discovery⁵ to helping businesses make decisions that accurately factor in climate volatility.⁶

However, the next few months and years will be crucial in determining where the UK places its desired level of ambition in AI, as massive investments by governments and businesses around the world are already priming themselves to influence the global conversation. The scale of effort elsewhere must not be underestimated: Germany has committed €3.1 billion towards a national AI strategy up to 2025⁷, and France has pledged €1.5 billion up to 2022, with almost half earmarked for research.⁸ While the UK has played a leading role in the development of AI, the US has recently announced \$1 billion over the

coming five years to open 12 new institutes that will keep the country at the forefront of research in AI and quantum computing⁹, and like-minded countries such as Canada and Germany are leading on commercialisation as measured by patents.¹⁰

The Council calls for a UK National AI Strategy that creates an even stronger general base of support for AI. Scotland has already embarked on an AI Strategy and the UK strategy should be created in consultation with the devolved administrations and the wider ecosystem.

The UK cannot, and need not try to, compete equally in all economic sectors so it must shape its strategy carefully to ensure it prioritises and is able to act selectively to lead and to capture markets around the world. In doing this it will need to consider its options for collaboration and competition within the international scientific and economic landscapes.

⁴ Some estimates show that leading nations could gain an extra 20-25% in economic growth and productivity gains through AI over the next decade.

⁵ [Modelling the impact of AI on the global economy](#), McKinsey (September 2018)

⁶ [Biotech harness AI in battle against Covid-19](#), Financial Times (May 2020)

⁷ [UK tech for a sustainable future](#), Tech Nation (2020)

⁸ [Germany AI Strategy Report](#), European Commission (March 2020)

⁹ [France AI Strategy Report](#), European Commission (February 2020)

¹⁰ [The Trump Administration Is Investing \\$1 Billion in Research Institutes to Advance Industries of the Future](#), The White House (August 2020)

¹¹ [WIPO Technology Trends 2019: Artificial Intelligence](#), World Intellectual Property Organisation (January 2019)

¹² [Developing an AI Strategy for Scotland](#), Scotland's AI Strategy (2019)

The Council has considered how to diffuse the development of AI across sections of society that would help to both improve the lives of its citizens and bring economic benefits in the form of recovery and resilience. This has resulted in four pillars:

- support for research, development and innovation;
- literacy in AI across the population with enhanced AI-related skills at all levels of educational and lifelong attainment;
- sound physical, digital and virtual infrastructure, including robust and flexible regulation as a form of social infrastructure, resulting in good governance that boosts public trust;
- and pragmatic approaches to national, cross-sector adoption that meaningfully advances public and private sector applications.

This Roadmap sets out the requirements in each of these areas in more detail and builds on the 2017 *Hall-Pesenti AI Review*. Its premise is that, in most areas, what the UK needs to do is to “double down” on the best of the investments it has already made, and to accelerate what’s proved to be ‘working building’ capability across organisations and sectors. Doing this is not easy, and least of all in the face of jobs at risk that could be made more vulnerable by the twin pressures of the COVID-19 pandemic and automation.¹²

Fundamentally, it requires commitments of a size, and for sustained timescales, that are often only possible at times of disruption such as the one the world is going through now with its response and recovery from the COVID-19 pandemic. It therefore requires a degree of collective vision and of boldness, combining this leadership with the confidence to create spaces for entrepreneurship in all parts of society and the economy in order to find new ways to use AI to benefit everyone.

In a small number of important areas this Roadmap points in new directions. As AI becomes more integrated into daily life, developing trustworthy technology is vital and businesses will require a novel approach to good governance and consideration of the ethical and societal implications of AI as part of the design process, with greater multi-disciplinary and cross-disciplinary approaches throughout. The UK needs to develop a world-class enabling environment with strong but flexible legislation and well-resourced, capable regulators able to provide responsive guidance and timely oversight.

This Roadmap is not a draft strategy, and nor is it an instruction manual. This is a set of recommendations and we recognise it will not be possible to achieve all at once. This is why a National AI Strategy is needed to prioritise and set a time frame. We are ready to convene workshops across all of the areas laid out in this Roadmap, and more, to capture the full spectrum of thought across the

UK in a way that would lend itself to the development of a UK National AI Strategy. We have set out a clear direction that we believe will accelerate the UK to a positive technological future and start the necessary conversations needed to create a strategy that will enable the whole of the UK to flourish.

¹² [Who is at risk?](#) The Royal Society for Arts, Manufactures and Commerce (October 2020) finds that, broadly speaking, industries with the highest level of furlough take-up are more likely to be at risk of automation.

1. Research, Development and Innovation

AI's full potential will only be realised in the UK if the government continues to create the conditions for leading edge research, development and innovation in AI for the public and private sectors, and if it creates the conditions for AI to work in and with multiple disciplines and areas of society and the economy.

In common with some other parts of this Roadmap, what can seem like a restatement of general principles that could be taken for granted is, in fact, a challenging requirement for sustained systemic change.

The government's R&D Roadmap¹³ makes an excellent start in outlining some of the areas of systemic change that will particularly benefit AI. One of the most important commitments is to make the UK a place that is open and inviting to top talent. Demand for AI talent is at an

all-time high, and there is more to do to attract and retain the very best to the UK's shores. Skills training, both at a research and workforce level, is vital to translate research, development and innovation, and suggestions are detailed in the forthcoming section.

The government R&D Roadmap includes valuable commitments to strengthening the interactions between discovery and applied research, innovation, commercialisation and deployment. This is vital: funding either excellent basic or applied research in isolation risks the outputs and the people being easily transferred to other countries, so any competitive advantage can be easily and quickly lost, especially to other countries that have more resources than the UK. The R&D Roadmap also signifies support for entrepreneurs; creating a more place-based approach to research, development and innovation; and the

aspiration to be more imaginative in ways of ensuring that research and innovation is directly responsive to the needs and aspirations of society. This last point echoes strongly the views of the AI Council's ecosystem.

In addition to building on these broad directions, there are three areas which are more specific to AI and on which the government needs to take significant action.

The Alan Turing Institute: local presences and global reach

The first is to build on the existing success of The Alan Turing Institute as a national champion for AI research, development and innovation in the UK. In a relatively short period of time the Turing has developed to play a pivotal role in promoting AI research across sectors and in attracting talent by leveraging its growing network. Over the past year, the Turing has attracted 77 new Turing Fellows, bringing the total to 369, and 21 new research associates alongside 116 events across their research portfolio, and a student enrichment scheme intake representing 36% female and 32% from BAME backgrounds.¹⁴ It has also reached out to crowdsource best practice for data science and AI research, was pivotal in creating a new discipline in data-centric engineering, and has worked across government to provide policy advice and operational support.

Concentrating on these beginnings, an enhanced and reshaped Turing will be a crucial part of the necessary research, development and innovation infrastructure but, to play that role fully, some changes are needed. The UK's research system as a whole cannot reach its full potential while the original membership model has the effect of inadvertently excluding excellent research teams across the four nations. The UK needs the Turing to have a full national reach so that it can play its part in levelling up and in supporting talent and great ideas wherever they start. To do this, it will need to be able to bring together researchers, businesses, civil groups and the public sector to make great ideas work at all geographic levels and at a sector-wide scale.

A new Turing model might, for example, extend both its virtual and physical presences to energise more regional and local activity. There are several ways to achieve such aims, building on the Turing's convening role in the UK landscape to join up and build on existing teams and infrastructure. This could be delivered through Turing acting as a headquarters and convenor of a nationally network of centres-of-excellence or hubs, which may be physical or virtual, which would not only advance AI research but also development, innovation and deployment into the regions. Across these centres, skills training will be at both the research level and, most importantly, in apprenticeships, upskilling and reskilling regional

¹³ [UK Research and Development Roadmap](#), UK Government (July 2020)

¹⁴ [Annual Report 2019-20](#), The Alan Turing Institute (2020)

workforces to exploit the opportunities of AI across sectors of the UK economy. Models for such centres might be based on, for example, UK's National Cyber Security Centre, or focus on regions to support AI research, development and innovation capability to support the local economy.

Well-funded institutes are becoming highly valued components of national AI strategies, and as a national institute of global standing, the Turing also needs the assured long term public sector funding that will give it and others the confidence to plan and invest for the medium term and strategic leadership for the UK, with the needs of academia, businesses, civil groups and the public sector reflected in its strategic direction. A well-funded UK institute would not only signal the UK's commitment to remaining a globally leading player over the coming decades, it could also promote the UK's interests through collaborations with international partners and enable it to play a significant role in building the international infrastructure discussed as part of the third pillar in this Roadmap.

Sustainable funding would enable the Turing to provide visible national leadership on matters of AI security and ethics, and use this leadership role and connectivity across the landscape to ensure the UK remains at the forefront of next-generation AI developments, such as those involving machine learning from small data sets, and machine learning as part of wholly new types of neural

interface technologies.¹⁵ This added focus on how to leap-frog current machine learning-based AI technologies would enable the UK, through the Turing and its network-trusted and neutral status, to also become the place for Explainable AI, bringing us closer to safe, responsible and trustworthy AI.

Challenge-led innovation

The significant promise of AI and its high risk, high reward, aspirational nature lends itself to goal-directed “moonshot” programmes in the style often associated with ARPA-esque models.^{16a} Properly designed and delivered moonshot programmes play to AI's strengths by requiring people to work across boundaries and existing organisational structures, and to build new relationships, networks and common languages in order to develop entirely new solutions to big challenges. In addition to this, there is a need for the UK to consider further the question of how AI interacts with people, and what it means to create effective human-AI partnerships.

One such moonshot, directly relevant to AI, might involve a major challenge aimed at developing and establishing appropriate methods for safe, ethical, explainable and reproducible AI which will accelerate its use across many sectors. Another would be the development of an innovation ecosystem enabled by the UK Digital Twin programme, as discussed later in this

Roadmap, within which AI would use these virtual worlds to explore real world optimisations that are beyond human scale.

AI should also be built into moonshots in other sectors where it could play a major part, for example to radically accelerate materials science to develop new materials that can better store energy from renewable sources, or in solving for real-time control of the increasingly complex electricity distribution network as the country moves towards net zero emissions.

Across all the potential types of challenge-based funding, programmes need to be selected and designed to draw in and accelerate AI's integration across disciplines and sectors. One suggestion would be to select broad sectors that are more mature, and thus ripe for AI adoption than others, and use the mission-based approach to focus minds. This would bring together basic and applied research and multiple disciplines across research horizons and opportunities for systems development, helping to mature AI as a discipline.

The specific vehicle for delivering moonshots is a ripe topic for discussion, but a forward-looking framework for research, development and innovation consisting of a series of labs, which conceive of and build fully engineered prototypes and platforms, could offer a flexible model of experimentation. These labs could generate vision-based

sequences of projects around moonshot challenges and could serve those with an interest in working between industry, academia and the public sector.

By providing significant signals about potential future developments, moonshots create space for informed public and business engagement with the emerging science and applications. Under these circumstances the science and technology is not simply communicated to; it is inspired by, and carried out with, individuals, communities and businesses. Such space enables all those involved to explore how AI interacts with people and what it means to create human-AI partnerships. In turn, the projects will help define responsible and ethical practices to meet society's needs and aspirations for a technology that behaves as people and businesses hope and expect: safely, securely and reliably.

¹⁵ [Human perspective: Neural interfaces](#), The Royal Society (September 2019)

^{16a} The US' Advanced Research Project Agency (ARPA) is a model designed to concentrate funding around pivotal investments in breakthrough technologies, often for national security purposes. DARPA (for Defense) is the most well-known ARPA agency in the US and is credited with developing a number of disruptive technologies such as stealth technology, the foundations of the Internet, and miniature GPS receivers.

AI to transform research, development and innovation

Capitalising on the promise of AI will continue to require investment in basic research into the fundamentals of AI itself. Researchers, increasingly working across disciplines, will need to address some of the fundamental challenges that remain, such as establishing a common language through which to develop data science projects or interoperability between datasets, which can drive smart collaboration. AI systems often rely on a narrow disciplinary base, and this needs to be broadened to cover an array of domains in an intelligent way. AI technologies need to get better at extending into areas where they do not have structured data, and to navigate effectively through that.

At the same time AI is transforming the nature of research itself, with examples in particle physics and bioinformatics.^{16b} For the UK to retain its status as a global leader in research, it needs to stay well abreast of the potentially disruptive effects of AI and other forms of data science which are beginning to be incorporated into all forms of knowledge creation, including the social sciences and the humanities. Given the availability of massive data sets and computer power, such disruptions could range from game-changing increases in efficiency in routine tasks such as searching texts or large observational data sets, to entirely new forms of enquiry brought about by AI's ability to

spot new patterns and prompt new questions and insights, often across existing disciplinary boundaries.

^{16b} [AlphaFold: a solution to a 50-year-old grand challenge in biology](#), DeepMind (November 2020)

2. Skills and Diversity

The Council's vision is for everyone to be able to live confidently with AI, and for those who go on to work with it and to build it to do so with the very best foundation.

For the most technically talented individuals we need to guarantee a best-in-class postgraduate offering to access globally attractive pathways that will contribute to the country's renowned research base.

Success will depend on increasing the diversity of people working on and with AI until those AI communities reflect the society they inhabit. It's important to focus on diversity; not as an afterthought, but because doing so will generate crucial research questions and lead to the widest and most innovative range of future technologies which will serve the needs of the people. There are wonderful examples of emergent leadership,

creativity, and energy in promoting diversity but it is not enough. We suggest benchmarking and forensically tracking levels of diversity and inclusion in AI as the NCSC¹⁷ has recently done for cyber security, to make data-led decisions about where to invest and ensure that underrepresented groups are included in funding programs. Government must do more to encourage the best of the existing initiatives where diversity is enabled, and to expand with their own schemes and campaigns, if the needle is to be sufficiently moved.

What everyone needs to know

The UK needs to set itself challenging but realistic goals to ensure that every child leaves school with a basic sense of how AI works. This is not just about understanding the basics of coding or quantitative reasoning, or to describe the mathematical concepts; nor is it simply

about ethics. It is about knowing enough to be a conscious and confident user of AI-related products; to know what questions to ask, what risks to look out for, what ethical and societal implications might arise, and what kinds of opportunities AI might provide. Without a basic literacy in AI specifically, the UK will miss out on opportunities created by AI applications, and will be vulnerable to poor consumer and public decision-making, and the dangers of over-persuasive hype or misplaced fear.

Over time, AI needs to be built into the curriculum as a specialist subject. As well as being its own subject AI needs to be part of computer science, citizenship studies, and as part of new ways of doing other subjects such as geography or history. These changes may take a decade to complete but getting there is essential. For now, the government should focus energy on the most effective existing curriculum enrichment initiatives, not as alternatives to fundamental changes to the curriculum, but because they provide opportunities to accelerate change in the short term.

Further, Higher and Graduate Education

Since the AI Review in 2017, the government has invested significantly in schemes to promote graduate-level learning through Apprenticeships, Masters courses and PhD training. This is a good start, but only that. Current numbers would suggest there is unmet

demand to justify a five to tenfold increase in places compared to the numbers in the *AI Review*.

To ensure a steady supply of skilled entrants to the workforce, and to send the right signals to young people making further academic choices, the UK should commit to scaling up its programmes significantly at all levels of education and maintain this commitment for at least a decade.

Such a scale up would include AI-related Level 7 apprenticeships, Masters programmes, Centres for Doctoral Training and novel forms of PhDs. It would extend the scope of the schemes to include AI-related courses in disciplines outside those traditionally associated with data sciences, with a review of digital apprenticeship standards to identify where additional AI knowledge and skills ought to be unlocked. It must also include the postgraduate AI and data science conversion course programme announced by the government last year, whose 2,500 places for people to develop new skills or re-train. Properly implemented, this commitment would help level up access to economic development across regions. We must also ensure that no graduate student leaves university without some grounding in AI and its applications.

¹⁷ [Decrypting diversity: Diversity and inclusion in cyber security](#), NCSC (July 2020)

Meanwhile the Turing AI Fellowships programme for established researchers, an international fellowship programme for AI in the UK in a partnership between UKRI, the Office for AI, and the Turing, has already seen a successful first wave,¹⁸ alongside the first batch of Turing AI Acceleration Fellowships.¹⁹ In line with the 2017 *AI Review*, this programme must continue to be developed and expanded so that momentum can be built to attract world-leading talent to the UK.

These efforts, directed at the supply side, would allow graduates and apprentices to advance into emerging AI professions. They would need to be delivered across the country and to be available to young people of all backgrounds, ensuring that the UK's future AI workforce reflects the diversity of background and thought needed to develop world leading AI capabilities that tackle the issues which matter most. Success will depend on collaboration between universities, colleges and businesses, not only in terms of how academia can help support continuous learning and upskilling within businesses, but also through business support such as mentoring, real life use cases and data to academia.

Supporting every teacher

The key to progress in schools is to enable teachers to have confidence in the basics themselves, so they can make the best possible choices about how to engage their students. One way to make progress relatively fast would be to

extend the work of the NCCE, supporting it to provide comprehensive curriculum resources that cover AI at all Key Stages, and providing an incentive and opportunity for teachers to learn and develop their AI knowledge. A comprehensive programme aimed at all teachers and with a clear deadline for completion would enable every teacher confidently to get to grips with AI concepts in ways that are relevant to their own teaching.

As in many other subjects, learning should combine access to knowledge about AI with exercises based on tackling real-world problems in ways that encourage the skills associated with problem-solving (including creativity and resilience) and team-working. Therefore new curriculum resources need to include curated data sets that will enable teachers to create the case studies and exercises that will in turn create the familiarity that leads all young people to increased confidence and data literacy.

Lifelong learning and an online academy for AI

Looking across society, the scale of re-skilling and up-skilling that will be required is vast when you consider a population whose working lives may last more than half a century beyond the time they left school, college or university. The UK must contend with changing business models in goods and services, and the changing nature of professions and trades, and the priority of levelling up.

These challenges point to the need for significant change in the contract between the state, individual and employer and a step change in the scale of continuous professional development and the ability to move between types of work.

While these socio-economic trends cut across sectors and subjects, they represent an even more urgent challenge in AI where the required skills have recently emerged and are still changing rapidly.²⁰ The government should consider the two following major steps.

The first is extending business tax relief from investment in research, development and innovation to investment in specified forms of retraining and upskilling, including AI.

A second step would be to enable the creation of an Online Academy for AI. Such an Academy could provide a key part of the social infrastructure necessary to support systemic change. Properly designed, such an Academy would provide the opportunity efficiently and creatively to meet several overlapping needs which emerge from different parts of this Roadmap. It might start by providing entry-level material to support life-long learning. Ideally, this would be accredited in a recognisable and portable way. In doing this it could also provide an entry-point and repository for online learning aimed at teachers, complementing the work of bodies such as the NCCE, and it would encourage the

growing national community of product managers retooling for the future in AI. Either separately, or integrated into the Academy, a 'data skills observatory' to identify and address the on-going and evolving data skills deficit function that would otherwise hinder AI innovation would also be needed.

While the Academy's main objective would be to enable more people to live and work with AI, it would also be well placed to help connect and support local initiatives involving public engagement and public dialogue as part of responsible innovation in new AI deployments, including moonshots. This sustained and distributed resource, rather than, say, a single educational or promotional campaign, would be an integral part of negotiating the benefits and the risks of AI across different applications and publics and finding new forms of settlement and new narratives. It would help to ensure that trustworthy information, in relevant and accessible forms, informed public narratives on AI across sectors and applications. It would need to be able to reach into the forefront of research, development and innovation, such as through the Turing, and to connect people and information for discussions on opportunities, the distribution of benefits and of risks, shortcomings and uncertainties.

This kind of resource would in turn underpin the regulatory infrastructure that is outlined in the next section, by ensuring that public debates were consistently well-informed.

¹⁸ [Welcoming world-class Turing AI Fellows to the Institute](#), The Alan Turing Institute (October 2019)

¹⁹ [New Turing AI fellows to deliver world-class AI research](#), UKRI (November 2020)

²⁰ [Short, online courses can unblock gender and other diversity issues in digital skills education](#), Institute of Coding (September 2020)

3. Data, Infrastructure and Public Trust

The UK should aspire to be the best place in the world to access and use safe, secure and good quality data in order to develop new applications and business models.

The UK should also be the best place for the public to exercise scrutiny of, and input to, automated decision making, in order to ensure the public can trust AI. To do this, the infrastructure it will need is not only physical, such as broadband networks and high performance computing capacity, but also virtual and social, encompassing processes, people, standards, governance and practices. A systemic reimagining of data availability and governance will be required across disciplines to transform traditional practices and public acceptability.

As for other infrastructure assets of national importance, the government has a key role in funding and providing data

infrastructure. An important component to this should be in creating a cross-sectoral Information Management Framework that is AI-friendly and designed for clean, codified, real-time data so that it forms the best possible basis for private and public sector innovation.

Trustworthy and accessible data

AI thrives on data, and makes new demands on data collection, curation and use. The government should focus its plans to make more public sector data safely and securely available, being clear about which data will be increasingly available, under what conditions and on what timescale. In the private sector, while regulators have begun good work to audit AI for personal data protection compliance²¹, more work is needed to help businesses seeking to use data for AI by creating the conditions for the

deployment of suitable privacy enhancing technologies.²² This should be furthered by accelerating work on translating the intent of a data sharing agreement into an actionable legal framework, and establishing guidelines for legal frameworks around different data sharing structures such as trusts, cooperatives and contracts.

It will be important to note that while some AI systems will continue to rely on access to the largest possible data sets, volume should not be the only priority for schemes to promote access: some of the next generation of AI technologies will be designed to work with much smaller data sets. This illustrates a more general principle which is that the detailed infrastructure needed to make AI work in different markets is often very context-specific, and the UK will need a National AI Strategy that links general principles to industry-specific arrangements of governance to enable innovation and growth. One need shared by many markets, however, is in establishing data infrastructure to permit companies to safely share and access data to develop common AI-related products. The public sector should lead the way with examples of good data sharing practices.

Good AI research, development and innovation relies on a bedrock of good data practices and there is much to do to improve the current position. This includes embedding data principles such as FAIR (Findable, Accessible Interoperable, Reusable), adopting open

licences where possible, working better with multi-modal data (big data, small data, noisy data, unstructured data) and enforcing an agreed supporting framework on what it means ethically and for society to share data for research purposes, and who owns it.

Ultimately, developers of AI applications ought to consider the opening and sharing of research, analytics and data at the fundamental design stage. This should include encouraging digital research, development and innovation infrastructure, software and tools to be open and standardised; providing access to large-scale compute and data resources tuned to AI applications; and supporting the development and deployment of Trusted Research Environments, which enables collaboration and data access across several organisations.

While previous work around best practices in data management and use has been helpful in setting the direction of travel, there remains a gap between the aspirations set out in these principles and their implementation, creating an additional barrier to the development of effective data access frameworks. Further work is needed to build a common language amongst data practitioners, building understanding of the nature of the datasets being used and how best to manage them. In many cases, data can outlive the systems that collected it, and conscious investment in trustworthy systems that can outlive the lifecycle of

²¹ [Guidance on AI and Data Protection](#), ICO (2020)

²² [Privacy Enhancing Technologies](#), The Royal Society (March 2019)

our standard modes of data management will reap many benefits.

The emergence of a small number of big players in several digital markets has sparked a timely debate on the proper role of government and regulators in preserving and enhancing competition across the digital economy. The role of data as a driver of concentration and potential barrier to entry for smaller firms has been a central theme of recent reviews.²³ The government has also recognised the growing importance of data access and openness to promoting competition and innovation, having established a dedicated Digital Markets Unit to advise on the functions, processes and powers which may be needed to address the challenges posed by fast-moving digital markets.²⁴

Whatever the government's final decision, clear and flexible regulation will be critical to supporting good data practices, and can help to level the playing field for new entrants into digital and data-driven markets. There remains much to be done in interpreting existing data rules to support research, development and innovation and well-resourced, independent regulators can assist in providing guidance.

In addition to these systemic changes, the government might consider enhancing the nations' approach to data governance: making sure data is collected, used and shared in responsible and trustworthy ways. Not appropriately

addressing data governance frameworks would maintain a fundamental barrier to the adoption and diffusion of AI technologies and inhibit the quality of wider sources of data. We recommend a coalition of institutions and experts work together to develop standard data sharing agreements compliant with data protection and intellectual property law but providing practical exemplars, especially for business-to-business data sharing that could be adapted and adopted across the economy. Without a strong system of data governance, existing tenuous public trust in data use and sharing will persist, impairing the UK's long-term leadership ambitions.

Public trust and good governance

Threaded throughout discussions about AI is the need for adaptive and informed regulation and standards as part of an enabling environment of good governance, which includes public engagement, transparency measures and responsible research and innovation on the part of businesses through the lifecycle of AI operations. Achieving such a robust system of good governance is crucial to building the public trust and confidence in AI that will be essential to its uptake and long-term public acceptability.

Recent events have demonstrated that public and private sector AI and algorithmic applications continue to be plagued by public scepticism and lack

widespread legitimacy. Developing and deploying trustworthy AI will be dependent on the UK strengthening its governing environment in a manner that both provides guidance and confidence to businesses to innovate and adopt AI, and reassures the public that the use of AI is safe, secure, fair, ethical and duly overseen by independent entities. This is beginning to be addressed in the National Data Strategy,²⁵ but further exploration and scrutiny of appropriate governance mechanisms, particularly concerning public sector adoption of AI, is overdue. We highlight one example in this section - a public interest data bill - that may serve as a method of increasing accountability through public engagement.

It is clear that there remains a fundamental mismatch between the logic of the market and the logic of the law. Technology markets extract value from collective data while laws respond to individual opportunities and threats. This has a critical impact on the public acceptability of data as infrastructure - data supported by people, processes and technology - and will only worsen if it is not sufficiently addressed.

One method that could begin to combine new forms of innovation and accountability is to involve the public in considering new ways to complement individual rights-based approaches such as consent. These would include new ways of ensuring public scrutiny of automated decision-making and the types of transparency that lead to

accountability: revealing the purposes and training data behind algorithms, as well as looking at their impacts. They would include public engagement, including in algorithmic impact assessments. And they would look to ensure that existing regulations and regulatory bodies had not only the capacity, but also the capability to fully consider the implications of AI in areas such as labour, environmental, and criminal law. These three tenets: (1) clear transparency about automated decision making, (2) the right to give meaningful public input and (3) the ability to enforce sanctions could be encapsulated in a Public Interest Data Bill.²⁶

Fundamental to the development of AI systems is access to data. The UK's response to the COVID-19 pandemic has highlighted gaps in the UK's capabilities in this space, in particular around the governance of 'happenstance' data, generated by individuals in the course of their daily activities.²⁷ Trustworthy access to such data can help generate insights for public policy and research, enabling the government to better understand the impact of its policy interventions in rapidly-evolving circumstances. Their expertise in the collection and management of data means the Office for National Statistics is well-placed to lead the development of governance frameworks for these new data sources. Government should enhance their role in this respect through the National Data Strategy and associated policy frameworks.

²³ ['Unlocking digital competition'](#), Report of the Digital Competition Expert Panel, UK Government (2019)

²⁴ [New competition regime for tech giants to give consumers more choice and control over their data, and ensure businesses are fairly treated](#), UK Government (November 2020)

²⁵ [National Data Strategy](#), UK Government (September 2020). Mission 2 regards a pro-growth and trusted data regime, and Mission 3 concerns transforming the government's use of data.

²⁶ [The Data Delusion: Protecting Individual Data Isn't Enough When The Harm is Collective](#), FSI Stanford (July 2020)

²⁷ [Data readiness: lessons from an emergency](#), DELVE Report No 7 (November 2020)

Making the UK the best place in the world to research and use AI requires it to be world-leading in the provision of responsive regulation and governance. While fundamental principles of good governance may be fixed and apply across sectors, the details will not, and this means that the systems of governance themselves need to be ready to respond and adapt more frequently than has typically been true of systems of governance in the past. We therefore suggest commissioning an independent entity to provide recommendations on the next steps in the evolution of governance mechanisms, including impact and risk assessments, best practice principles, ethical processes and institutional mechanisms that will increase and sustain public trust.

Public trust in the applications of AI and algorithms depends on the trustworthiness of these applications and those that develop and implement these. Good governance also needs to encourage good culture amongst the individuals and organisations that are responsible for the development and application of advanced information technology. At the same time, the UK needs to continue investing in data-specific regulatory capacity and algorithmic accountability, ensuring that institutions are well funded to cover the advisory, enforcement and regulatory landscapes as sector digitalisation continues at pace. UK governance should combine sector-specific expertise with cross-sector insights, bringing regulators

in different sectors together to connect with each other and begin to experiment with new ways of encouraging safe AI innovation across and within their remit.

Digital Twins and Living Labs

Applying AI systems responsibly inevitably requires a great deal of development before getting anywhere near the physical world. Digital twins (a realistic representation of assets, processes, systems or institutions in the built, societal or natural environments) are increasingly recognised as important in many fields as they can provide insight into how complex physical assets and citizens behave in real time and can help organisations improve decision-making and optimising processes. An expanded programme to advance digital twin capability, with interoperable data standards, could help developers explore the potential for applications in the physical world but cheaper, faster and with less risk, as well as a better understanding of the implications and impacts to society and the economy. AI thrives when stitching digital twins together, using machine learning algorithms to analyse complex patterns and make decisions between different processes and systems alongside generating new insights with newly hybridised datasets.

The National Digital Twin Programme, currently being driven by the Centre for Digital Built Britain, aims to create an Information Management Framework to

enable the building of digital twins across the built environment. Building on this work, a more ambitious cross-sectoral UK Digital Twin programme, governed by the Gemini Principles²⁸, would be a moonshot challenge and would encourage more widespread use of open synthetic data to develop AI applications in many specific potential applications where the risks to individuals or society might otherwise be large. AI could also use these digital twins to explore real world optimisations that are beyond human scale.

In developing those AI applications and AI models and testing with synthetic data, there comes a point when you must begin to model and rehearse in the real world with messy, unpredictable and diverse environments. To achieve this in a safe way without causing inadvertent harm, the UK needs to create a nationwide network of living laboratories, physical environments that start as representations of the real world but more constrained.

AI will provide the embedded intelligence within all manner of new smart machines, many with autonomous capabilities, that can work collaboratively with humans to transform their productivity, safety and quality of life. By enabling the fast and safe development, testing and demonstration of new AI-based goods and services, including distributed software systems and intelligence embodied in smart machines and products, businesses, researchers and

citizens can gather the data and experience necessary to inform future policy, regulation, products and services.

A globally leading role

Finally, the UK should urgently and sustainably build upon its strong international position. It is one of the top three countries in terms of research publications in AI.²⁹ This strength combined with the UK's diplomatic weight can be a catalyst to shape international discussions, but our leading international position is not guaranteed. We should therefore strengthen existing partnerships with other like-minded countries, such as France and Canada as well as consider carefully how to increase bilateral cooperation with countries like the US and China, such as through the recent declaration signed with the US to increase collaboration on AI R&D,³⁰ to enable effective strategic action to shape a global AI landscape. One immediate way forward is to leverage our founding membership in the Global Partnership for AI (GPAI), a scalable club of like-minded governments ready to be led by experts on tackling the most pressing priorities for collaboration, where the UK is already leading on data governance.³¹

The UK has a pivotal opportunity to play an influential role in the coming year as it adopts the Presidency of the G7 and the 26th UN Climate Change conference. Over this period the UK should use this opportunity to leverage global collective

²⁸ The Gemini Principles were developed by the Centre for Digital Built Britain to guide the responsible development of digital twins and a National Digital Twin for the public good. [The Gemini Principles](#), Centre for Digital Built Britain (2018)

²⁹ [Global AI Talent Report 2020](#), (October 2020); [AI Research Rankings 2019](#), Medium (December 2019)

³⁰ [Declaration of the USA and the UK on Cooperation in AI Research and Development](#), UK Government (September 2020)

³¹ [An open call for input from GPAI's Data Governance Working Group](#), OECD (November 2020); we understand the recommendations from this working group will be incorporated into the National Data Strategy.

intelligence to advise on the biggest challenges AI presents and position the UK as a centre to operationalise recommendations, including those relating to the UN's Sustainable Development Goals.

Alongside global leadership on shared challenges, the UK ought to consider its own strategic advantage as we endeavour to become less dependent on critical technologies from non-allied sources of supply. Advancements in AI products and services underpin the capacity to compete economically but is also the domain of contests between overseas power bases and other nation states, which are sometimes not well aligned to the UK interest. One way to respond to this shift in geopolitical context is to invest strategically in research, development and innovation, choosing areas of national strengths such as AI and supporting them heavily with the aim of becoming globally competitive.

If we align this strategic investment with the construction of broad coalitions of interest on the global stage, the UK can increase our sovereign capabilities and attract researchers and technologists with the skills to develop leading products in a number of competitive sectors. Government should also consider how to redefine procurement and intellectual property to encourage a creative technology ecosystem to achieve these aims.

4. National, Cross-sector Adoption

The preceding sections outline the pillars that would provide the basis for sustained uptake of AI across the country and the economy as a whole.

This section proposes transitional and other more specific measures to boost initial take up and to evolve existing markets in the crucial areas of business adoption, startup & scale-up support, public sector adoption, health and social care, climate change, and defence.

Business as smart adopters

Across society and the economy, and since the first weeks of 2020, the COVID-19 pandemic has accelerated the adoption of digital systems such as e-commerce platforms and marketplaces, and tools to enable remote working. The UK therefore has an opportunity to build on this momentum to showcase and reward the uptake of AI by encouraging

further investment and enabling business, especially SMEs, to be smart consumers for AI systems.

Creating confidence so that business customers increasingly seek out AI solutions, knowing that they will be able to get sound advice and robust products, is crucial. There is further need to build on grants to access impartial advice on how to use digital and AI solutions to help the growth of SMEs that have been exacerbated by the COVID crisis. Scotland's Data Lab is a good example of an initiative that fosters trusted collaboration between industry and AI expertise, having invested in 120 Collaborative Innovation Projects at approximately £12 million value over the last 5 years, and recently launched a Torch Business Advisory Service to support SMEs on their data journey.

The largest companies that are able to realise the potential of AI through business transformation find themselves increasing productivity, reducing costs and innovating faster to remain competitive. By ensuring the enabling foundations described in the first 10 recommendations are in place, the UK can also be a fertile ground for its existing FTSE 100 and the future 100.

Supporting high growth AI startups

As these businesses initiatives stimulate demand for AI products, we need to ensure that the AI startup and scaleup vendor ecosystem thrives. Supporting AI-first businesses by providing the necessary physical infrastructure to offer services to the largest and smallest local and multinational companies would help these businesses scale, develop and become the next multi billion dollar companies.

The government should encourage more initiatives like Tech Nation's Applied AI Program³² to directly support smaller startups to grow across the UK. The community welcomed an additional £200 million from the British Business Bank which increased VC funding to high growth companies,³³ and there is a compelling argument to increase this funding, or alternatively to consider public investment in promising companies that seek to advance government priorities in a similar style to the National Security Strategic Investment Fund. For scaleups,

the funding gap that exists at series B+ means that AI companies struggle to stay under UK ownership. Government may consider supporting at this stage to ensure companies can take foreign investments in these rounds while also staying in the UK as they grow.

The UK government needs to assist startups and scaleups to compete most effectively for the best talent. The Study Visa program may need to be extended now we have left the European Union. To this end, there is a prominent role for the government to position the UK as the go-to AI startup location in the region. It's important to spread the message that the UK has more capital and talent than anywhere else³⁴ and a unique proximity of buyer to vendor as well as being far easier to incorporate and run a business here than in Europe.

To increase the number of SMEs and build the right entrepreneurial infrastructure for AI-first companies, it is also worth exploring the possibility of building on Innovate UK's Knowledge Transfer Partnerships (KTPs), in which the government could introduce a series of AI-specific KTPs to stimulate market ready AI research in areas of direct relevance to businesses to commercialise research.

The UK currently lacks a prominent cohort of professionals who understand the complex sales and leadership skills needed to build, articulate and sell AI tools and develop leaders of AI research labs, and there is concurrently a specific

³² [Applied AI](#), Tech Nation (2020)

³³ [Extra £200 million backing for British businesses](#), British Business Bank (April 2019)

³⁴ [Tech Nation Impact Report](#), Tech Nation (2020); [UK Tech For A Changing World](#), Tech Nation (2020)

demand to increase the number of skilled product managers in AI. We recommend a marketing campaign and programs to encourage international product managers to come to the UK, and for UK citizens to focus their education or re-skilling in this direction. This would help turbo charge the UK's ability to effectively commercialise the world class research coming out of our universities.

Enabling public sector adoption

The UK public sector will be a priority area when it comes to AI innovation. Improving policy-making processes (for example, ensuring it is more resilient to shocks to the system), and the provision of more efficient and effective public services (for example, using data-driven 'real-time' economic forecasts), built on strong ethical foundations, will become essential to improve the security and well-being of citizens.

There is an urgent need to have a digitally and data enabled local and national government. To leverage the opportunities of AI in the public sector, there will need to be new ways of ensuring that data can be accessed and evaluated across public sector organisations. The government should ensure there is a central place for civil servants to access information about all of the AI-based tools and projects being undertaken across departments. AI should be used in projects that are outcome-based and only where

absolutely appropriate, and departments should learn rapidly from each other about what works well and what needs to evolve. In time, the focus should be on linking public sector data with private sector data, rapidly enhancing the insights that can be gained through diverse datasets and a more sophisticated adoption of AI.³⁵

Government should fund interdisciplinary pathfinder projects – requiring collaboration between industries and across government departments, with different academic disciplines – as a means of demonstrating the potential of AI technologies to support pathfinder projects and building capability in the civil service to deliver such projects. Importantly, the application of AI in the public sector should not fall explicitly on one department to lead. It may be that a Chief Information Officer embedded in the heart of government would work well to provide the strategic direction and leadership required for such a multi-stakeholder task.

The huge benefits of data sharing will be worth little unless government departments are enabled to intelligently procure products that suit different contexts and services. More work is needed to demystify algorithmic decision making in government, with training programs for officials so they can suitably assess AI systems for themselves and see how it would be deployed as part of an increasingly efficient civil service. The ONS is working towards a target of 500

data analysts across government trained in data science techniques this year,³⁶ and to enable a more thorough upskilling of AI and machine learning analysis, work should intertwine with the Turing, the Royal Society, the Royal Statistical Society and others in determining how data science and AI practitioners may be professionally standardised.

Public sector bodies wishing to adopt AI often face challenges in building procurement capability. Governments have tried to enable public procurement to better support innovative goods and services, while ensuring value for money and relatively low risk. This government should continue to build on initiatives that have worked, such as GovTech and the guidelines for AI procurement,³⁷ to allow the public sector to effectively purchase emerging tech such as AI. Greater AI adoption in the public sector through public sector procurement creates markets for data science and AI technologies to be developed, tested and deployed by government departments, to enable a pull-through of innovation from research, often via SMEs.

It is important to note that evidence shows the human rights of the poorest and most vulnerable are especially at risk when public services are automated.³⁸ As such we suggest this work is done carefully and with independent ethics

committees tasked to consider ethics, including social and economic impacts.

Health and social care

The demands on healthcare are changing rapidly as populations drift, age, and their needs evolve, and with a huge reduction in the cost of measurement. With this ubiquity of such data, informed and improved interventions through prioritisation and personalisation can maximise benefit.

Given the social and economic importance of health and the UK's potential to build on market opportunities created by the NHS, the sector has been, and should continue to be, at the forefront of AI adoption. The importance of the link between health, mental health and social care has been especially clear during the COVID-19 pandemic, and we have also seen areas in which AI could have been better prepared to deal with that link if we had considered the promise of AI in this sector earlier. AI can help manage the complex information flow that is needed for the integration of these systems. NHSX, in delivering the Health Secretary's Tech Vision³⁹, is taking the lead in unblocking issues including (but not limited to) governance, data access and protection, and inclusion. The Council, who are supporting NHSX with their AI Lab⁴⁰ and AI Awards,⁴¹ looks to NHSX to work with others across the research,

³⁵ There is a fundamental need to consider transparency and data bias as public sector adoption increases, taking into account analysis and recommendations from, among others: [Review on Artificial Intelligence and Public Standards](#), The Committee on Standards in Public Life (February 2020); [CDEI review into bias in algorithmic decision-making](#), Centre for Data Ethics and Innovation (November 2020)

³⁶ [Looking to the future at the Data Science Campus](#), Office for National Statistics (March 2019)

³⁷ [Guidelines for AI Procurement](#), UK Government (June 2020). Developed in collaboration with the World Economic Forum Centre for the Fourth Industrial Revolution.

³⁸ [Statement on the Visit to the United Kingdom, by Professor Philip Alston, United Nations Special Rapporteur on extreme poverty and human rights](#), United Nations Human Rights Office of the High Commissioner (November 2018)

³⁹ [The future of healthcare: our vision for digital, data and technology in health and care](#), UK Government (October 2018)

⁴⁰ [The NHS AI Lab](#), NHSX (2020)

⁴¹ [Artificial Intelligence in Health and Care Award](#), NHSX (November 2020)

development and innovation landscape to continue building towards these opportunities and developing best practice guidance for AI adoption in the healthcare sector.

As the adoption of AI increases outside of the hospital setting, so too will the ‘datafication’ of networked digital technologies which create new health insights for individuals. This carries abundant opportunities for the preventative health of citizens, for example in stimulating new and more efficient ways to diagnose and prevent disease beyond clinical boundaries, but there are also significant concerns if not governed correctly. Individuals’ health quickly becomes legible to actors outside medical and clinical settings, for example inferred sensitive personal information held by online platforms, and appropriate oversight ought to be considered in how these technologies develop.

We also encourage further development in areas with others in the sector where the UK can build a competitive advantage. Three particular areas stand out: the first is in creating a data strategy for health and care, which achieves clean, codified and real-time data with a strong sectoral data governance framework. The second is the opportunity to develop new models of partnership and incentives for partnerships between SMEs and the NHS that flourish and scale beyond ideation. The third opportunity would be to create a better understanding of the skills and knowledge required for citizens and

healthcare professionals living and working with AI technologies, building especially on the Topol review⁴² and NHS Digital Academy.

Finally, building on the UK’s world-leading life sciences industry and increasingly digitally mature NHS, there is an opportunity to establish the UK as the go-to-place for biomedical research. With the appropriate governance in place, and through closer collaboration of the life sciences industry and the NHS, AI has the potential to empower large national trials more efficiently, including through the seamless adaptation of subjects to treatments.

Climate change

AI technologies are an essential part of the toolbox for innovating to reduce greenhouse gases in the atmosphere, and to reduce the environmental impacts of goods, services and human activities. We are already seeing AI contributing to a greater grasp of complex environmental and sustainability systems, from forecasting supply and demand at real-time⁴³ to combating illegal deforestation⁴⁴ and understanding Arctic sea ice loss.⁴⁵ One way to accelerate the use of AI for these purposes would be to build it into relevant moonshots, such as that on new materials for energy storage and renewables and create incentives for AI companies to address these and other Net Zero challenges.

The challenges to increased adoption of AI in related sectors, particularly in energy, has been explored elsewhere⁴⁶ but echo similar constraints in other sectors, notably in the accessibility, interoperability and labelling of datasets that can contribute to cross-sector cooperation on energy and climate initiatives. The reality is also that the continued development of digital technologies, including AI and complementary technologies such as 5G and the Internet of Things, create an environmental impact that also needs to be monitored.

As a part of the UK’s globally leading role in AI, there is an opportunity to carve out a leading path on the global stage where others see a similar need to increase the role of AI in energy and climate for public good. The UK should consider drawing up a plan that ensures responsible AI is suitable for the planet and helps build broader awareness, inform regulation, and incentivise researchers and developers to make more conscious and accountable decisions in these areas. Ahead of COP26, the UK should consider how to develop and showcase the UK’s ability in cleantech, which spans the entire country and represents a successful, innovative path to Net Zero,⁴⁷ and use its spotlight to encourage an open, international digital ecosystem of data, algorithms and insights that can generate

reliable information in real-time about the state of the environment and interactions between the economy, society and the environment.

Defence and security

AI is central to a broad range of defence and security applications, which are becoming increasingly data rich and networked. They require multiple sources of information to be brought together to understand a rapidly evolving situation, to determine a timely course of action and to oversee the enactment of the response. The information involved is often multi-modal (text, audio and video) and is invariably uncertain, incomplete and contradictory. Sometimes it is also intentionally misleading. AI methods are necessary to make sense of this volume, variety and velocity of information. However, further developments are required to scale the underpinning fusion techniques, to have them work in closer partnership with the humans involved, and to ensure the ensuing decisions and plans can be understood and justified.

The increased connectivity of all aspects of our society means that digital assets are increasingly common, and valuable and physical assets are increasingly accessible remotely. These bring tremendous benefits, but also open up

⁴² The Topol Review advised the government on implementing technologies such as AI in the NHS and to enable NHS staff to make the most of them to improve services.

⁴³ [The Topol Review](#), NHS (February 2019)

⁴⁴ [Solar nowcasting with machine vision](#), The Alan Turing Institute (current research)

⁴⁵ [The fight against illegal deforestation with TensorFlow](#), Google (March 2018)

⁴⁶ [Understanding Arctic sea ice loss](#), The Alan Turing Institute (current research)

⁴⁷ [AI Barometer Report](#), Centre for Data Ethics and Innovation (June 2020)

⁴⁸ [UK Tech For A Changing World](#), Tech Nation (2020)

opportunities for interference and disruption, leading to a requirement for a secure, connected intelligence. AI techniques are necessary to monitor the complexity of these networked systems and to determine when they are operating normally and when they are not. As cyber-attacks become increasingly sophisticated, AI is needed to detect and respond to them.

We welcome the commitment to create a new MOD agency for AI, and look forward to further detail in the Integrated Review in 2021.

We are now living in a challenging economic and geopolitical environment in which global tech companies hold sway over national strategic advantages. By strengthening our ties with like-minded nations, the UK can increase our capabilities in co-creating technological developments, drawing innovations from oblique application areas and developing leading products which are then adopted by our allies. The broader AI community will be required to engage in this space as the security sector becomes more pervasive, and the government will need to find ways of working across demanding security frameworks to enable wider collaboration.

Part of opening up wider collaboration on this front will require new levels of internationally competitive investment in advanced hardware technologies that serve the UK ecosystem, as the driving force behind fast and secure innovation.

Co-creation – at scale – of technological developments in an application agnostic environment (benefiting from innovations from other application areas) could be a potential model. This would also provide additional benefits to improved technological export capabilities for Britain, and further benefits to attracting world-class technologists and innovators in this space, not seen in Britain since the 1980s.

Acknowledgements

AI Council

Sir Adrian Smith FRS

*Institute Director and Chief Executive,
The Alan Turing Institute
President Elect of The Royal Society*

Alice Bentinck MBE

Co-founder, Entrepreneur First

Ann Cairns

Executive Vice-Chairman, MasterCard

Christopher Bishop FRS FRSE FREng

*Laboratory Director,
Microsoft Research Cambridge
Professor of Computer Science,
University of Edinburgh*

Claire Craig CBE

Provost, The Queen's College Oxford

David Lane CBE, PhD FREng FRSE

*Professor and Founding Director,
Edinburgh Centre for Robotics*

Kriti Sharma

Founder, AI for Good

Máire O'Neill FREng

*Professor, Centre for Secure Information
Technologies, Queen's University Belfast*

Marc Warner

Chief Executive Officer, Faculty

Sir Mark Walport FRS FRCP FRCPATH FMedSci FRSE

*Chief Executive, UK Research and
Innovation, 2017 to 2020*

Martin Tisné

Managing Director, Luminata

Neil Lawrence

*DeepMind Professor of Machine Learning,
University of Cambridge*

Nick Jennings CB, FREng

*Vice-Provost for Research and Enterprise,
Imperial College London
Professor, Artificial Intelligence,
Imperial College London*

Dame Patricia Hodgson DBE

*Member of the Board,
Centre for Data Ethics and Innovation*

Paul Clarke CBE

*Chief Technology Officer, Ocado Group
2012 to 2020*

Pete Burnap

*Professor of Data Science & Cybersecurity,
Cardiff University*

Priya Lakhani OBE

Founder and CEO, CENTURY Tech

Rachel Dunscombe

*CEO, NHS Digital Academy
Visiting Professor, Imperial College London*

Tabitha Goldstaub

*Co-founder, CognitionX
UK AI Business Champion
Chair, AI Council*

Dame Wendy Hall DBE, FRS FREng

*Regius Professor of Computer Science,
University of Southampton
UK AI Skills Champion*

Wider AI ecosystem

In October 2019 the Council put out a call for any interested individuals to be part of a wider ecosystem, to contribute to the Council's work. Over 450 people responded, and the Council is grateful to them for their contributions over the last few months, including for their views on this Roadmap. This includes:

Adrian Joseph OBE, *Managing Director,
Group AI and Data Solutions, BT*

Adrian Weller, *Program Director for AI,
The Alan Turing Institute*

Alexandra Mousavizadeh, *Director,
Tortoise Intelligence team*

Andrew Hopper CBE FRS FIET FREng,
Vice-President, The Royal Society

Brent Hoberman CBE, *Executive Chairman and
Co-founder, Founders Factory*

Carly Kind, *Director, Ada Lovelace Institute*

Demis Hassabis CBE FRS FREng FRSA,
CEO and co-founder, DeepMind

Frank Kelly CBE, *Professor of the Mathematics of
Systems, University of Cambridge*

Gerard Grech, *CEO, Tech Nation*

Gillian Docherty OBE, *CEO, The Data Lab Scotland*

Hayaatun Sillem CBE FIET, *CEO,
Royal Academy of Engineering*

Sir Ian Diamond FBA FRSE FACSS DL,
UK National Statistician

Indra Joshi, *Director of Artificial Intelligence, NHSX*

Jeni Tennison OBE, *Vice-President,
Open Data Institute*

John Spindler, *CEO, Capital Enterprise*

Julia Adamson, *Director of Education,
British Computer Society*

Lila Ibrahim, *COO, DeepMind*

Professor Dame Lynn Gladden DBE FRS FRSC
FInstP FREng, *Executive Chair, EPSRC*

Matthew Taylor CBE FACSS, *CEO, The Royal
Society for the Encouragement of Arts,
Manufactures and Commerce*

Professor Dame Ottoline Leyser DBE FRS,
Chief Executive, UK Research and Innovation

Sir Patrick Vallance FRS FMedSci FRCP,
UK Government's Chief Scientific Adviser

Philip Colligan, *CEO, Raspberry Pi*

Roger Taylor, *Chair, Centre for Data Ethics
and Innovation*

Saul Klein, *Co-founder, LocalGlobe*

Simon McDougall, *Director, Technology &
Innovation, Information Commissioner's Office*

Simon Peyton Jones FRS MAE, *Chair,
National Centre for Computing Education*

Stan Boland, *CEO, FiveAI*

Stian Westlake, *CEO, Royal Statistical Society*

Stuart Russell, *Professor of Computer Science,
University of California, Berkeley*

Lord Tim Clement Jones CBE FRSA, *Co-chair,
AI All-Party Parliamentary Group*

Vanessa Lawrence CB HonFREng FRICS,
*Non-Exec Director, Satellite Applications
Catapult and Non-Exec Director and Trustee,
The Alan Turing Institute*

Zoubin Ghahramani FRS,
*Professor of Information Engineering,
University of Cambridge*

About the AI Council

The AI Council, an independent expert committee, provides advice to the UK Government and high-level leadership of the Artificial Intelligence ecosystem.

The Office for Artificial Intelligence (a joint BEIS-DCMS unit) is secretariat to the AI Council.

www.gov.uk/government/groups/ai-council

BEIS 1 Victoria Street, London, SW1H 0ET
DCMS 100 Parliament Street, London, SW1A 2BQ



© Crown copyright 2021

You may re-use this information (excluding logos) free of charge in any format or medium, under the terms of the Open Government Licence v3.0. To view this licence, visit www.nationalarchives.gov.uk/ogl or email psi@nationalarchives.gsi.gov.uk. Where we have identified any third party copyright information you will need to obtain permission from the copyright holders concerned.

Published January 2021