techUK
FOR WHAT COMES NEXT

AI Adoption in the UK: Putting AI into Action

March 2023

Supercharging UK Tech and Innovation

Contents

Introduction

Artificial intelligence (AI) is changing the face of the technological ecosystem and is unlocking unprecedented opportunities for innovation in the UK, as well as on the global stage. The anticipated economic benefits over this decade are significant with estimates predicting that the UK's GDP will be up to 10% higher in 2030 from the development and adoption of AI.[1]

The good news is the UK is a global leader in AI and is ranked highly in 'AI readiness'.[2] The UK Government's National AI Strategy[3] recognises and emphasises the vital role that AI will play in the UK's efforts to build economic resilience, productivity, growth and innovation across the public and private sector. Using three key pillars, over the next 10 years the Strategy aims to increase investment for the long-term needs of the AI ecosystem; support the transition to an AI-enabled economy; and ensure effective governance of AI. With over £2.3 billion being invested into AI since 2014, the Government is keen to strengthen its position as a global superpower. The UK is also ranked third in the world as a destination for private AI investment; from 2013 to 2022, almost £10.8 billion has been invested by the private sector.[4] It is now about how the UK fully realises the benefits of AI and unleashes the vast potential of AI to create opportunities for businesses and society as a whole.

We have already started to see this happen. In a recent UK survey, AI solutions for data management and analysis were particularly prevalent, followed by natural language processing and machine learning.[5] The benefits of this adoption of AI include greater automation of business processes, enhanced IT or network performance and improved customer experiences. However, as we look to the future, the benefits of adopting and innovating through AI are surpassing mere automation and process efficiencies. Globally, organisations are using AI to solve some of society's most complex issues and two-thirds of respondents to a recent global survey said that they are aiming to or have already applied AI to sustainability-related goals.[6]

While having a National AI Strategy is key for ensuring we have a clear vision and blueprint for the development of the AI industry in the UK, it is equally important that we have a plan for the increased adoption of AI. Given the immense opportunities and benefits that come from adoption, we must accelerate implementation and encourage more companies and organisations to put AI into action. However, the extent to which these benefits can be felt by organisations in the UK, as well as the wider public, would be hindered if barriers to AI adoption are left unresolved. In order to further progress the adoption of AI, greater efforts must be made to overcome the persistent challenges that industry are currently facing across multiple areas, including access to quality data; trust in AI; organisational culture; compute infrastructure; and AI skills.

This report aims to enable businesses to put AI into action by showing how far the UK has come and what we need to do to make the country a global AI superpower. Through working alongside techUK's AI Adoption Working Group, this report is a celebration of how AI has been implemented in a variety of use cases across the UK. Using the first-hand experiences of AI leaders across industry, this report will also shed light upon the key barriers that are impacting wider adoption, as well as recommendations for businesses and policymakers on how these can be mitigated.

The current state of AI adoption

AI adoption can be defined as the use of AI capabilities, such as machine learning, in at least one business process.[7] The global growth of AI adoption is well-documented in current literature. This increase has been driven by multiple factors, including a greater accessibility to these technologies, for example through off-the-shelf solutions, as well as a broader diversity of well-established use cases[8] and improved digitisation of businesses' core processes.[9] According to IBM's 2022 Global AI Adoption Index, 35% of businesses report using AI; a 4-point increase from the previous year.[10] Another global survey by McKinsey showed similar findings, reporting a doubling of adoption since 2017.[11] However, it was also reported by the same paper that for the past few years, the proportion of organisations utilising AI has plateaued between 50% to 60%. This is indicative of AI becoming a mainstream technology, as well as the potential organisational change that is required in order to successfully adopt the technology and reap its benefits. Despite this, the gap between AI adopters and laggards is increasing, with the proportion of 'underachievers' increasing.[12] It is therefore imperative that barriers to entry are addressed to ensure that benefits are widely promoted.

The degree to which AI has been adopted varies significantly depending on the industry and company size. In the UK, the IT and telecommunications sector have the highest rate of adoption, followed by the legal sector. Sectors with the lowest rates of adoption include hospitality, health and retail.[13] Yet, company size has the most striking differences in adoption rates; a recent study showed that within the private sector, 90% of large organisations have already implemented AI, or have actively planned for its adoption, whereas this was 48% for SMEs.[14] According to a global study, larger companies are now 100% more likely to have deployed AI in comparison to smaller companies.[15] It is therefore clear that barriers to entry for SMEs are higher, in comparison to larger companies. Thus, any efforts to increase AI adoption may require a more nuanced approach for smaller organisations.

When considering international comparators, whilst the UK has a high rate of adoption compared to other advanced nations, there is a distinct gap that can be seen between other global leaders. According to a global survey by IBM last year,[16] China and India have the highest rate of adoption with almost 60% of IT professionals reporting the active use of AI in their organisations. This is a stark difference when compared with the UK (26%), US (25%), Australia (24%) and South Korea (22%).

Also, while the UK is globally recognised for having a strong reputation in incubating AI businesses, this is a competitive global market with many countries looking for ways to support and enable AI companies, particularly those looking to scale-up, to thrive and grow. This presents a potential risk of the UK losing out to larger economies and markets, and as a result missing out on the opportunity of large-scale economic benefits. Therefore, greater momentum is needed to widen adoption of AI across industries and sectors in the UK to increase market demand for AI that will also help AI businesses scale-up and grow here in the UK.

The development of AI is now more than ever at the forefront of the public conversation with the advent of OpenAI's large language model (LLM), ChatGPT. The potential opportunities for innovation using AI are endless, however, thoughtful design, guidance and ongoing monitoring is needed to ensure the responsible and ethical use of these technologies. The UK is in a strong position to capture the full potential of AI and its adoption is anticipated to continue rising.[17] The Government's Spring Budget 2023 highlighted activities to boost these efforts.

These commitments include the establishment of an AI Foundation Models Taskforce to build sovereign capability, a yearly AI Challenge Prize, as well as an investment of £900 million into a new AI Research Resource and exascale supercomputer.[18] Currently, the UK also ranks particularly highly in AI readiness with a solid footing in research activities and start-up investment.[19] In order to fulfil the UK's ambition of being a global superpower, these activities need to be scaled up. It is therefore imperative that we act now to help businesses put AI into action.
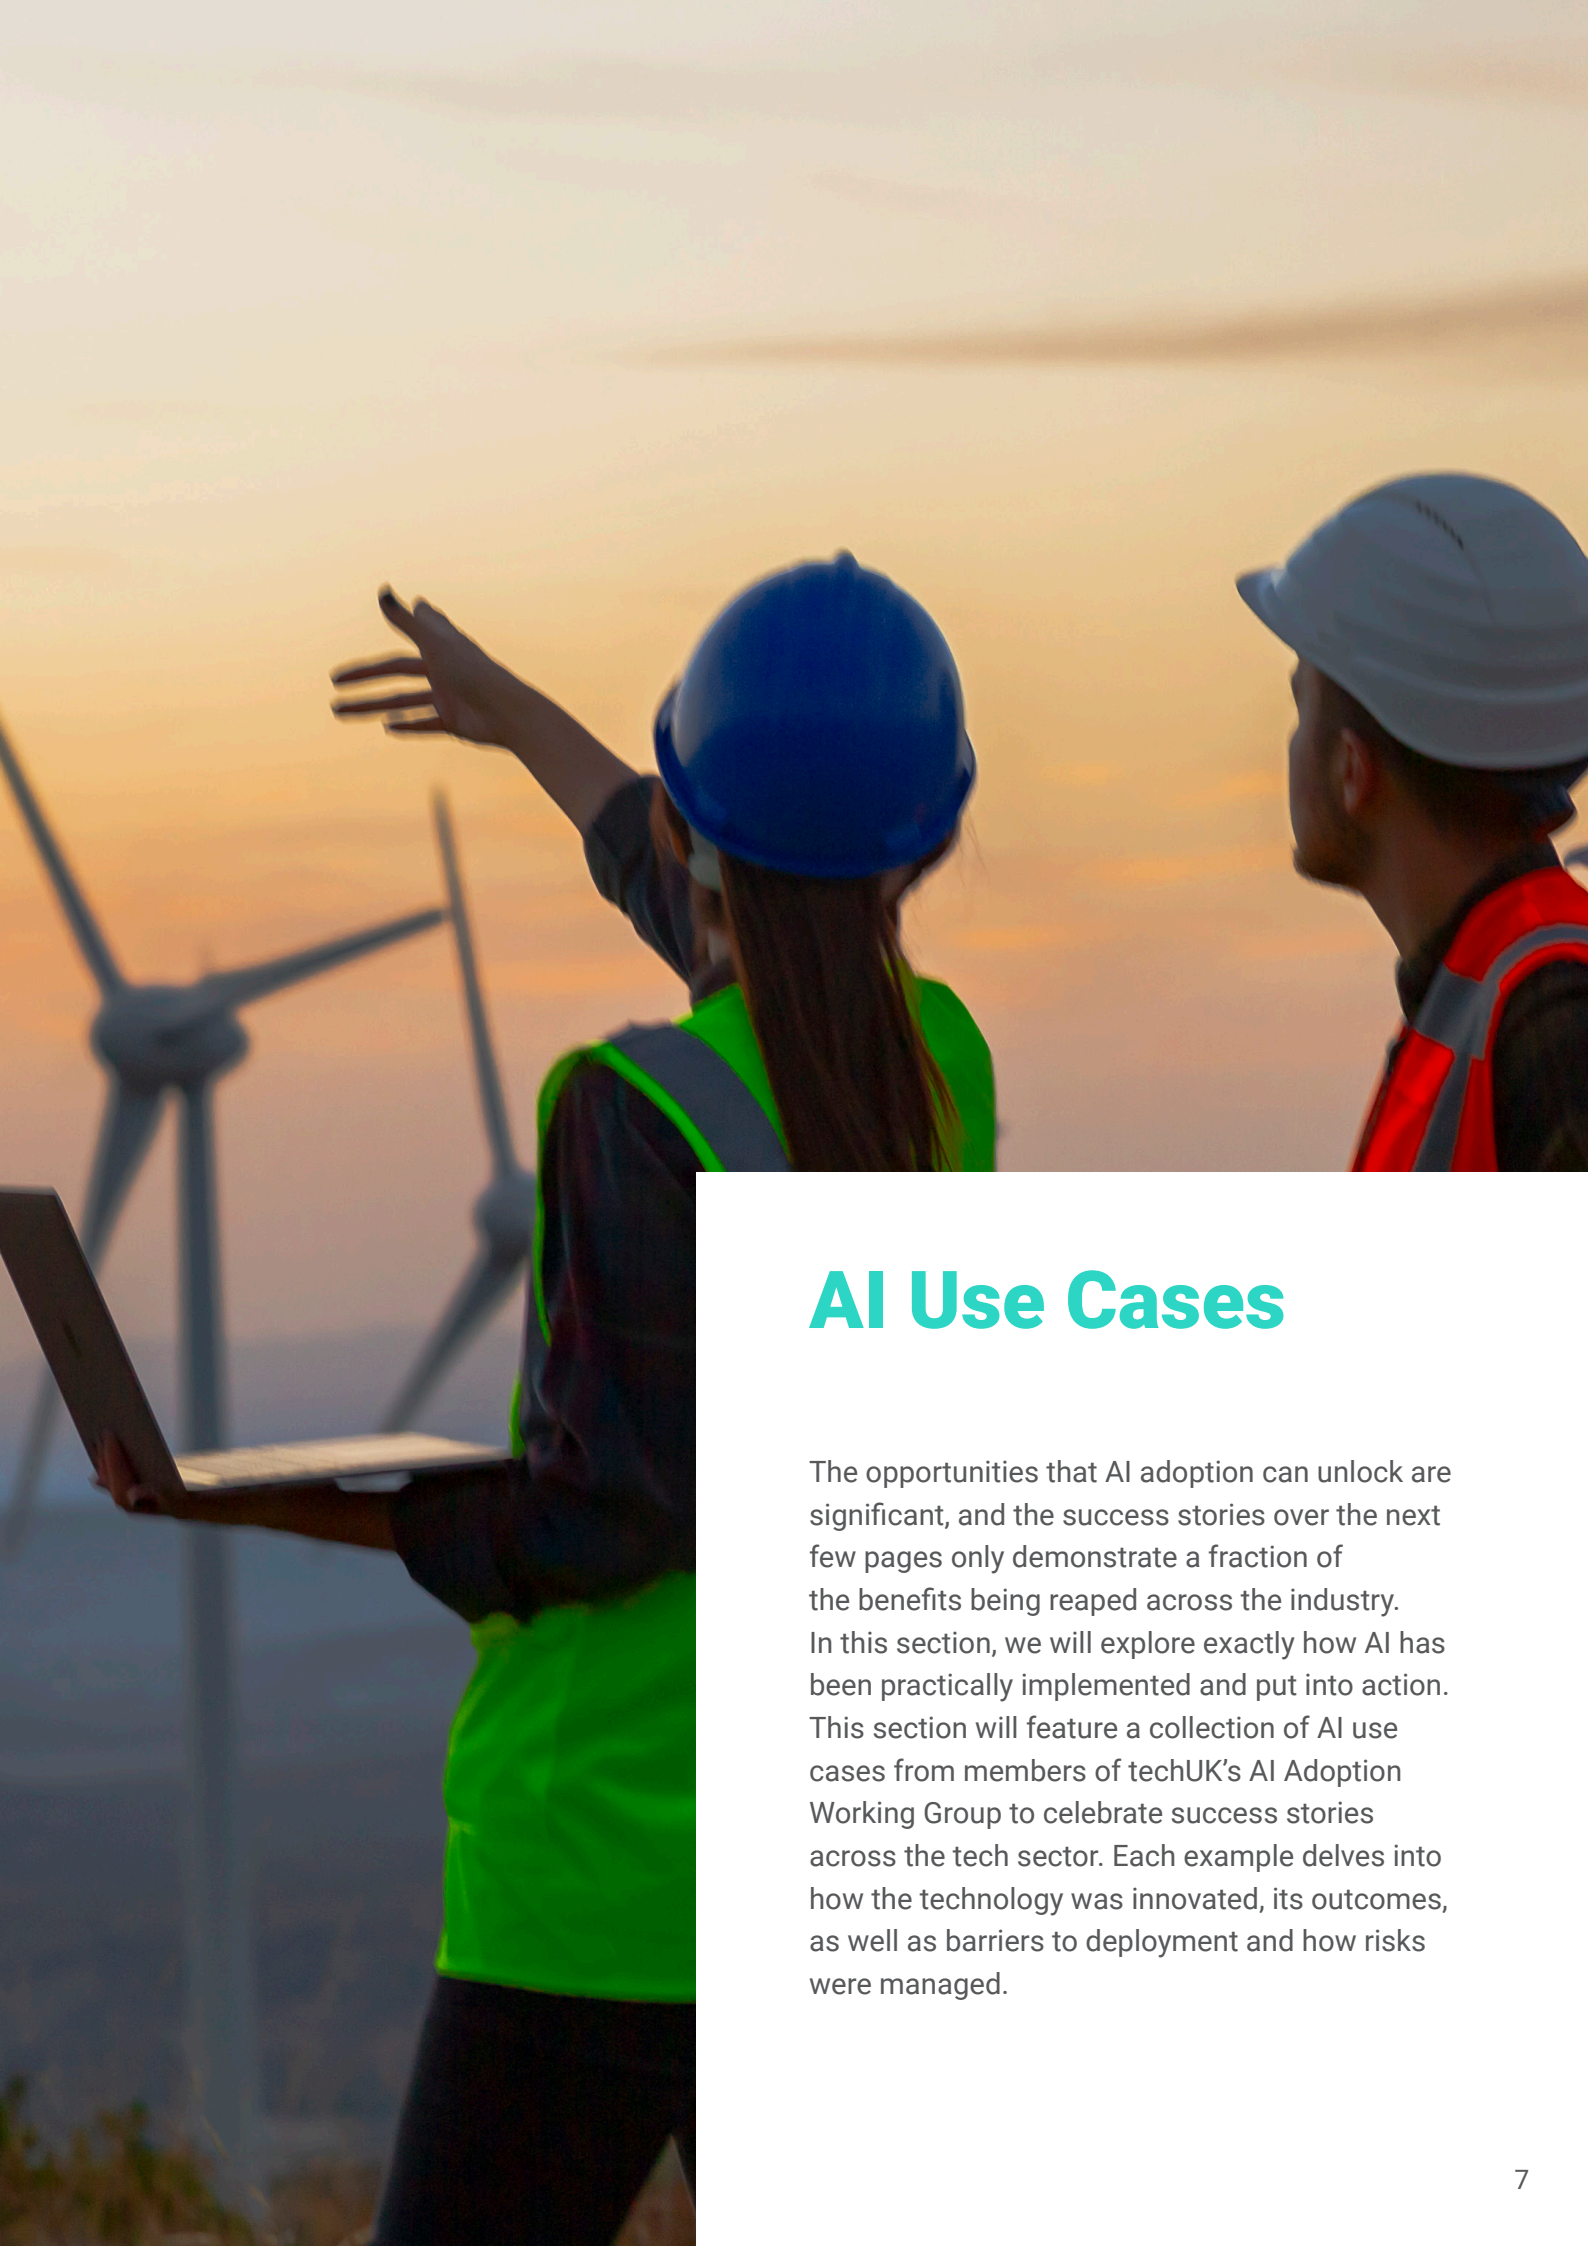
AI Use Cases

The opportunities that AI adoption can unlock are significant, and the success stories over the next few pages only demonstrate a fraction of the benefits being reaped across the industry. In this section, we will explore exactly how AI has been practically implemented and put into action. This section will feature a collection of AI use cases from members of techUK's AI Adoption Working Group to celebrate success stories across the tech sector. Each example delves into how the technology was innovated, its outcomes, as well as barriers to deployment and how risks were managed.

Wayve's 'AV2.0' pioneers a new approach to self-driving using AI and data

Challenge: Developing commercially viable and scalable AV solutions to unlock self-driving benefits faster

Wayve is a British start-up pioneering a scalable way to bring automated vehicles (AVs) to the UK and beyond. The uptake of AV technology in the UK will enable greater access to safer, smarter and more sustainable forms of transportation while contributing to the UK's long-term economic growth and international competitive advantage. Our vision is for a world where people can reclaim the time spent driving and where safe, reliable and environmentally friendly AVs are the standard for transporting people and goods.

Traditional AV technology relies on expensive sensors, HD maps and rules-based coding, which are prohibitive to scale. This makes adapting to new locations, vehicle types, use cases, or business needs time-consuming and expensive.

Wayve is using AI to build AV technology that is truly scalable with the potential to be 'adopted' by a range of sectors across different use cases to deliver real-world benefits. By designing affordable vehicle-agnostic AV technology embodied with artificial intelligence that can be applied to many different electrical vehicle platforms and adapted to any location, Wayve will give people greater access to the benefits of this transformative technology.

Solution: Building embodied AI to enable any vehicle to become self-driving

Wayve's pioneering 'AV2.0' approach uses a single deep neural network, i.e., an end-to-end machine learning system. This removes the need for HD maps and programmed rules, so AVs powered by Wayve's solution are not restricted to predefined routes and can therefore drive in 'new' areas with minimal preparation and infrastructure. This makes it highly attractive for commercial deployment as few transport services operate on fixed routes.

We expect that commercial fleet deployments will be where the early benefits of AV technology are first realised. That means AVs will be used to move people and goods, in last-mile deliveries, logistics, public transport and ride-hailing - and fleet operators can adopt Wayve's generalisable Automated Driving System (ADS) to optimise their operations and transform their service offering. AV fleets offer greater utilisation and efficiency, leading to fewer vehicles on the road and supporting efforts to decarbonise the transport network.

These are applications where the UK Government agrees there are prominent early use cases and where we already have commercial partnerships with businesses like Asda and Ocado Group to trial last-mile grocery delivery.

Barriers: Building confidence and certainty for commercial deployment

We are showcasing the promise of our approach through daily testing on UK roads with electric cars and light commercial vans, which we test with expert safety operators behind the wheel. Our testing is enabled by the UK's progressive testing framework and favourable business environment. Our on-road and virtual testing using the **Wayve Infinity Simulator** enables us to gather the evidence we need to demonstrate our systems' safety to prepare for future commercial deployment.

However, the UK risks losing this market if the Government does not bring forward legislation to enable commercial deployment of self-driving vehicles by 2025. Regulation will ensure that a safety and assurance framework is in place to build public confidence and trust in this technology while providing businesses with the certainty to invest in the long term.

Impact: Harnessing AV technology for a smarter, safer, more sustainable world

Our current trials with ASDA and Ocado Group are proving the adaptability of our ADS performance on complex grocery delivery routes in London while helping us and our partners understand the operational requirements of automating grocery delivery services. The UK has the opportunity to become a hub for AV innovation, attracting and growing new businesses in advanced manufacturing and tech, while transforming existing business models across various industries by adopting AV technology. We need regulation to keep pace with innovation and provide the certainty and confidence businesses want to see to keep investing in the UK.

Hazy's synthetic data helps enterprises unlock their data to enable safe and fast innovation

Challenge: Enabling safe and responsible use of data

The financial industry is rapidly changing. New entrants, ease of switching, digital services and macro-economic factors, mean the need to transform has never been more critical to remain competitive.

Banks such as Nationwide Building Society (NBS) are speeding up their innovation as a result and many work with third parties to innovate faster by capitalising on external tech firms' specialisms. This approach offers huge potential for banks to tap into additional capability faster and more cost-effectively but faces a significant challenge: the ability to share data quickly, safely and easily.

Solution: Synthetic data paving the way for innovation

As a result, NBS looked to synthetic data to quickly and safely employ their data. They chose Hazy for its enterprise credentials and expertise in deploying the technology in complex environments.
Hazy's software uses artificial intelligence to create new 'synthetic' data which retains the statistical properties of the real sets, unlocking data for enhanced propositions while protecting customer privacy. Synthetic data can therefore be used by internal teams and third parties freely and safely.

NBS's dataset contained a rich variety of characteristics and patterns that needed to be learned in order to produce a synthetic dataset that was fit for purpose. To verify that the synthetic data successfully preserved these characteristics, the teams used Hazy software's built-in metrics as well as some bespoke metrics for evaluation. This enabled the Nationwide team to finetune the data model, compare real vs synthetic datasets, and prove the value of the technology internally.

The project centred around specific, measurable business outcomes, along with clear implementation and adoption planning.

Barriers: Developing trust with a nascent technology

Barriers to deployment of AI for generating synthetic data include:

- the nascent nature of the technology,
- a lack of understanding regarding legal requirements around data access,
- a lack of understanding about methods of synthetic data that reduce the risk of leakage.

Deploying AI technologies in enterprise firms with complex IT infrastructures and layers of regulation is a complex process. Hazy's team is composed of technical experts who have spent six years deploying this technology side by side with customers.

It is important to note that mere use of synthetic data alone does not protect the original records from linkage. But, unlike some providers, Hazy mitigates the risks by embedding differential privacy technology in the process of synthetic data generation. Differential privacy solves the risk of re-identification by adding a carefully tuned level of 'noise' to the original dataset. Additionally, the software is deployed on-premises, meaning that the source data never leaves the source environment, further reducing the risk of leakage.

That said, the potential benefits are great. Hazy's enterprise software enabled Nationwide to generate synthetic data in a way that both preserved the statistical properties of the original data and simultaneously prevented exposing the business to regulatory risk.

Impact: Accelerating transformation

Using synthetic data, Nationwide has reduced the time needed to provision data for validating vendors from six months to three days.

As a result of faster vendor analysis, Nationwide is able to execute on its transformation programme much faster whilst reducing risk of a data breach from using customer data outside of production environments. In essence, synthetic data has de-risked innovation at Nationwide; mitigating against data breaches, whilst also advancing the speed and volume of projects that can be completed each year.

Throughout the implementation of Hazy, Nationwide has educated a wide audience of internal stakeholders, teams, and partners on the benefits of synthetic data. This has resulted in a comprehensive list of use cases for adoption across the building society, which are currently in progress with the Hazy and Nationwide teams.

Intuit's AI automates repetitive tasks, allowing people to be more productive and focus on the work that they love

Challenge: Beating the boredom of repetitive tasks

Repetitive tasks can cause productivity to decline due to boredom and a lack of stimulation, which can negatively impact on an individual's ability to work efficiently and effectively. Many repetitive tasks can be automated, freeing up time for people to focus on more complex and higher-value tasks.

Solution: A suite of tools helping people to do their jobs better

Intuit have built a range of tools using AI that works hand-in-hand with people to help them do their jobs better. Much of the AI functions as decision support systems, helping both customers and financial experts make insightful financial decisions.

Expert matchmaking: When it comes to filing taxes or needing accounting advice, some people prefer to handle it by themselves. Other people prefer to let the experts take care of everything. For the people somewhere in the middle, Intuit has services called **TurboTax Live** and **QuickBooks Live** – which utilizes AI to match customers with tax specialists or bookkeepers who will help guide them through the process and address their unique financial situation to review and discuss documents together.

Cashflow prediction: AI helps identify patterns beyond what our human minds could identify. Intuit's cash flow prediction tool uses a small business customer's bill payments, projected invoicing, and other transaction data to predict their cash flow 30 days and 90 days out. This can empower a small business owner to determine if they can weather a drop in business, if it's the right month to invest in a new piece of equipment or help them spot a payroll or invoicing issue early, and provide small business owners with the confidence to make the right financial decisions at the right time for their business.

Error checker: Every small business owner is shadowed by a pile of paper documents waiting for their data to be entered into the system. It's tedious, time-consuming, and error-prone. Intuit's document understanding platform uses a combination of computer vision, and machine learning technology to capture, classify, and extract images and data from structured and unstructured user documents of all kinds - smartphone photos, PDFs, forms, and so on. This helps eliminate data entry requirements and reduce errors, saving time and improving accuracy.

Auto-classification: Small businesses and self-employed individuals have to track their expenses for tax purposes. We applied advanced machine learning to a dataset of over a billion transactions of our QuickBooks users to automatically categorize our customers' transactions. This means every time you buy a coffee from your favourite coffee shop, Intuit's algorithm can categorise it automatically, saving you time and money and increasing accuracy.

Monitoring impact through a multi-disciplinary approach to governance

We're operationalising responsible AI through governance practices. We take a multi-disciplinary approach to governance. Our governance includes internal review, executive level oversight through our AI Governance Committee, employee training, and channels for stakeholder feedback.

We've designed a thoughtful approach to reviewing our use of AI that helps implement responsible AI practices. It's a risk-based approach that takes into account factors, such as the intended use/what's at stake with respect to customers, applicable laws, and the nature of the data involved. It's designed to identify and address potential issues around the proposed use of AI, including legal, compliance, ethics, privacy, etc. For example, models that could have a potentially significant impact on a person, such as AI used in the context of lending would be reviewed. As we learn from our reviews, we will continue to build on those learnings and iterate to evolve our approach.

We consider laws including but not limited to employment, credit and fair lending, anti-discrimination, privacy, and other laws. We recognise laws may not always keep pace with technology. That's why we consider our values, the benefits to our customers, and other ethical considerations when reviewing our AI models.

The Benevolent Platform™ augments scientific ingenuity with AI to navigate an ever-expanding data landscape to map complex diseases

Purpose: AI-enabled drug discovery to understand complex and underserved diseases

The primary use of AI at Benevolent is to understand complex disease biology, particularly in the context of a clinically-defined cohort of patients, to find new medicines for untreated and underserved diseases. To do so, we have built a robust data architecture and Knowledge Graph that captures a multidimensional representation of human biology across all diseases, which scientists can explore using our suite of AI predictive tools in the Benevolent Platform™. The tools reason across all of this information to suggest novel biological targets that could evolve into a drug discovery programme if experimentally validated, deemed safe and able to provide an opportunity to differentiate should the programme be successful.

At BenevolentAI, we invest much time in ingesting and harmonising a wide variety of biomedical data from more than 85 different data sources. We ingest pseudonymised or anonymised patient-level data, including electronic health records, genetic data, multi-omics data and scientific journals and literature. Our data is then represented in the form of a knowledge graph comprising nodes representing biological entities such as diseases, disease mechanisms, genes, cell types, and tissue types that are linked if they are related in some way.

BenevolentAI also uses large language models (LLM) to identify biological targets. Our models are trained to encode the meaning of literature sentences that refer to potential biological targets and are deployed to receive questions that ask for target suggestions for a particular disease, generate a ranked list of targets and provide evidence that backs up the target suggestions.

Barriers: Building trust in methodology and preventing algorithmic biases

One of the key barriers to AI deployment is trust in the methodology. Modern AI methods are typically black-box and provide limited rationale behind a prediction. As the typical user is unfamiliar with how the underpinning algorithms work, it is often difficult for them to act on a model suggestion, particularly when the supporting evidence they can understand (e.g. the literature) is sparse. The way to overcome these barriers is through investment in uncertainty quantification, explainable AI approaches and methods that look to identify causal relationships between objects of interest.

Another issue that we must continually be aware of is biases in data. Data bias can come in many forms. For example, biomedical literature is heavily biased towards a small number of diseases that are typically of interest to the developed world.

We have to consider these biases when making predictions and have used data normalisation strategies to down-weight this type of information. Another source of data bias is at the patient level. Biobanks and other data sources provided by either Government or commercial organisations are heavily biased towards patients of European descent, with other ethnicities often underrepresented. We have developed tools to quantify diversity in data and have set up BenevolentAI's Data Diversity Initiative to raise awareness of these issues.

Impact: Empowering scientists to discover drug targets for patients in need

Today, our Platform is scientifically and commercially validated. We have built an in-house pipeline of 15 named and 10 exploratory programmes, all Platform-generated, spanning multiple therapeutic areas in diseases affecting millions worldwide. Our most advanced asset is an atopic dermatitis topical compound in Phase 2a clinical development, and we also have a novel asset for ulcerative colitis planned to enter Phase 1 studies this year.

Secondly, our Platform supports a number of commercial and non-commercial partnerships, including a successful multi-year collaboration with AstraZeneca, who have already validated and selected a total of five novel targets identified by our Platform for their drug development portfolio in two disease areas: Idiopathic Pulmonary Fibrosis (IPF) & Chronic Kidney Disease (CKD).

Trilateral Research provides Ethical AI solutions to tackle complex social problems and benefit humanity

Challenge: A threat that can erode the dignity and liberty of millions of people

Trafficking is not only a crime that erodes the dignity and liberty of over 24.9 million people worldwide; it also presents a significant threat to international security, fuels criminality and corruption, and exacerbates conflict and violence. Combatting human trafficking remains an acute challenge, not only because of the scale of the problem and the difficulty in spotting the signs, but also because it's a highly complex process. It isn't merely a single isolated event, but a series of interrelated interactions, transactions, and experiences across time and space.

Solution: Countering human trafficking by driving insights through Ethical AI

Trilateral Research's Ethical AI application, **STRIAD:Honeycomb**, offers powerful ways to help public, private, and third sector organisations improve their ability to understand and tackle human trafficking. STRIAD:Honeycomb's use of Ethical AI offers opportunities to unpack, explore, and analyse large volumes of data in ways that would take a human being many hours, days, or even weeks to achieve – while ensuring its AI-driven insights are clear, understandable, explainable, and digestible so they can easily be used to support good decision-making.

In particular, STRIAD:Honeycomb's Natural Language Processing (NLP) tools – co-designed with the very people tasked with understanding and tackling human trafficking, and trained and developed from the outset using Trilateral's in-house subject matter experts – help ensure that in-depth human expertise, professional judgement, and the reality of human trafficking survivors' experiences can be integrated directly into analytical workflows, planning processes, and resourcing decisions.

The Ethical AI application, CESIUM®, is a software application co-designed by Lincolnshire Police and Trilateral Research that augments Multi-Agency Child Exploitation (MACE) decision-making with Ethical Artificial Intelligence (AI) to transform safeguarding arrangements for children.

As part of its ethical approach to developing AI, Trilateral Research integrated its teams of ethicists and data protection advisors into the co-design process. The development team created several front-end XAI features to show the relative importance of the different features of the algorithm's output. Through co-design, Trilateral worked closely with end users to create different visualisations. The development team used a "multi-modal" explainability approach which groups several different kinds of explanations together to facilitate understanding and informational uptake by different individuals who may have distinct, or even conflicting, explanatory needs.

These visualisations include infographics, descriptive sentences and the placement of information on the screen to determine the optimal ways for explaining the relevant information in a meaningful way. The result is a system which MACE partners can scrutinise for assured decision-making, and that operates within the limits of AI.

Impact: Accelerating the prevention of harm and benefitting humanity

In terms of operational impact, an October 2022 exercise validated CESIUM's prioritisation scores against historic decisions to assess their alignment with professional judgement. This exercise identified 16 children up to six months before they were referred by a safeguarding partner. Where this first workshop was a hindsight review, a subsequent foresight workshop identified two children for consideration at MACE pre-screening. While these workshops reflect a single agency deployment, this second workshop also identified a potential capacity uplift of 400% in a multi-agency deployment.

NVDIA partnered with King's College London to deliver the Synthetic Brain project to help diagnose neurological diseases through AI-powered images

Challenge: Diagnosing neurological diseases more effectively

The Synthetic Brain Project, a collaboration between NVDIA and King's College London, is focused on building deep learning models that can synthesise artificial 3D MRI images of human brains. This is helping researchers to accelerate their work by dramatically increasing the amount of data available. In turn, this aims at developing AI models to help diagnose neurological diseases based on brain MRI scans. They also could be used for predicting diseases a brain may develop over time, enabling preventative treatment. The use of synthetic data has the additional benefit of ensuring patient privacy and gives King's College the ability to open the research to the broader UK healthcare community.

Solution: synthetic data to make brain images more accessible

King's College London used a computer network inspired by VQ-VAE to compress brain images to less than 1% of their original size while maintaining their quality. This network outperformed the previous state-of-the-art method. They used a transformer model to learn the latent space of the data and generate new images based on clinical variables such as age or disease. Transformer models use attention to weigh the significance of different parts of the data and understand its sequence length.While it is artificial, synthetic data reflects real-world data, mathematically or statistically. Research demonstrates it can be as good or even better

for training an AI model than data based on actual objects, events or people. The use of synthetic data in healthcare research helps address some key risks in applying AI for healthcare, such as the availability of large amounts of high-quality labelled data, preserving data privacy, ensuring data represents a diverse cross-section of individuals, and enabling the creation of plentiful data on rare diseases. As part of the synthetic brain generation project from King's College London, the code and models are open-source. NVIDIA has made open-source contributions to improve the performance of the fast-transformers project, upon which The Synthetic Brain Project depends.

Barriers: Lack of access to diverse sets of medical images

Synthetic data is annotated information that computer simulations or algorithms generate as an alternative to real-world data. This technique already widely used in computer vision for consumer and business apps. These fields tend to have access to open datasets with millions of real-world images. By contrast, medical images are relatively scarce, typically only available to researchers connected to large hospitals, given the need to protect patient privacy. Even then, medical images tend to reflect the demographics the hospital serves, not necessarily the broader population. This has been a major barrier preventing researchers from applying AI to their work.

Impact: AI-created brain images to support research

The UK Biobank maintains multiple brain images taken from more than 50,000 participants, curated at an estimated cost of $150 million. The researchers responsible for the Synthetic Brain Project have already tripled this resource with their donation of 100,000 synthetic brain images to Health Data Research UK, a national healthcare data repository which will host the images. The AI models will also be available so researchers can create whatever images they need. A fortunate feature of this AI approach is it can make images to order, helping researchers ensure their research is based on diverse datasets. Though the brain images are simulated, they are highly useful because they preserve key biological characteristics.

Creating generative brain images that are eerily similar to real life neurological radiology studies helps understand how the brain forms, how trauma and disease affect it, and how to help it recover. Instead of real patient data, the use of synthetic data mitigates problems with data access and patient privacy.

AVEVA worked with ENEL Green Power to produce AI-powered predictive analytics for performance optimisation

Challenge: Optimisation of the ENEL's energy plants and processes

ENEL Green Power worked with AVEVA to build and deploy a complete diagnostic process, including teams from Engineering and Construction, as well as Innovation and Global Technology. The objective of the project was to develop analytics that support the diagnostic process and the performance optimisation of the fleet, exploiting features of thermal monitoring and diagnostic rooms, and combining the expertise of global and national teams to drive accelerated performance through the use of advanced AI capabilities and prognostic approach to operational efficiency of their power network.

Solution: AI to power green analytics and identify risks

As part of their effort to meet and exceed greenhouse gas emissions and their corporate sustainability targets, ENEL partnered with AVEVA to implement this AI solution. The AI solution is focused on identifying risks within ENEL's power generation fleet and providing mitigation instructions to minimize adverse impacts to major assets and production operations. The AI itself is proven to be accurate by "learning" from real-world production data that ENEL stores, and it continues to learn in real-time as operations progress.

This increased efficiency leads to improved sustainability by allowing them to burn less fuel to achieve required levels of power production and, hence, emit lower levels of greenhouse gases.

Barriers: Security and human trust

Two barriers to AI deployment in power generation are 1) security and 2) human trust. AVEVA has a long history of developing mission-critical, secure software and has leveraged this experience and various industry-standard processes in developing this AI solution to be extremely secure, and able to pass security audits. To improve human trust, ENEL's process data was used to train the AI. The software itself is very mature, incorporating a high degree of usability features and details of how conclusions and guidance were determined. All of this provides users with a highly usable solution that is not merely a "black box."

Impact: Increased operational efficiency, lower costs and energy savings

Dozens and dozens of issues have been detected, allowing ENEL to improve their operational efficiency, lower maintenance costs, and ensure maximum power generation.

For example, the AI solution was deployed after discovering a circumferential crack in the HP rotor of a steam turbine. The steam turbine transient model used model training data to improve efficiencies and identify the fault.

Smith Institute uses AI to support the transition to Net Zero

Challenge: Identifying communities' needs to support the green transition

To support Scottish and Southern Electricity Networks (SSEN) Distribution with their Vulnerability Future Energy Scenarios (VFES) project which helps deliver a just transition to net zero. In particular AI was being applied to identify drivers of vulnerability of their customers which would then enable SSEN to forecast how those communities and needs may change in the transition to net zero, and to support them now and through the transition.

Solution: Using Machine Learning to improve decision-making

The purpose of applying ML in this instance was to identify trends and patterns that were not known to humans, due to the scale of the data and the complexity in the correlations between datasets. The results are feeding back into business operations to support human understanding and impact future decision making, with the customers' well-being at the forefront.

Barriers: geographically limited data

With SSENs data, it was well managed and labelled, meaning that it was very amenable to AI. The data was limited to the areas served by SSEN which has two consequences. Firstly, data at the boundary where the community was served by multiple providers was not exclusive to SSEN and thus needed to be excluded from the analysis to avoid inaccurate results. Secondly, this is not applicable to the entire UK, due to being biased to the areas the model has been trained on. To improve the model further for any future application, the full UK dataset would need to be accounted for.

Otherwise, most risks are mitigated by the fact that this is not an automated and deployed model. The output is provided to humans and scrutinised, then used in conjunction with other data and insights to inform decision making. The model itself has been built by mathematical experts in collaboration with the domain experts of SSEN, ensuring that the data and models have been designed and implemented responsibly, for the use case and requirements of today.

All data used was compliant with GDPR. Datasets were not assessed at an individual level, but at a community level.

Impact: Mapping out social vulnerabilities for better support

The model indicated that certain social indicators are strong drivers in vulnerability, either increasing or decreasing it, but other features were found to be far more prominent in one cluster than any of the others. Interesting results arise from inspection of the clustering results. For example, the number of households in private rented dwellings strongly influences vulnerability in one cluster but is not observed as a main driver in any of the other clusters. Many other demographic features arise in multiple clusters, for example old age, provision of care and household sizes, however the average vulnerability in each cluster differs.

The results of our work identified distinct groups of locations that SSEN service, with each group having similar demographic features driving vulnerability, and each requiring distinct strategies to ensure that those who would be particularly adversely affected by the risk of frequent power cuts will be served by secure, stable, and resilient future networks. These groups cover different locations across Central Southern England and the north of Scotland: a robust strategy for improving infrastructure and services in Scotland and another for the south of England is therefore required to address the varying causes of vulnerability that customers experience.

KPMG use AI to understand and improve A&E waiting times and support operational decision making

Challenge: Understanding and improving A&E waiting times

The issue being tackled is understanding and improving patient waiting times in hospital A&E. KPMG developed a machine learning algorithm to predict the likelihood of A&E patients waiting more than the 4-hour target and identifying the primary clinical and operational factors influencing the breach.

Solution: AI models and data to support operational decision making

KPMG developed a machine learning algorithm to predict the likelihood of A&E 4-hour target breaches and derive the actionable insights to reduce A&E waiting time. Anonymised patient level records and dynamic features regarding hospital operational conditions were used for model training. Data was stored in Azure blob storage. Extensive feature engineering was performed along with class imbalance techniques for better predictions. The final model was selected after multiple iterations and implemented in Azure ML workspace for predictions.

The AI model aims to provide insights to support operational decision making. They should never be used as an automation or decision-making tool as they are.

Barriers: AI skills and training for healthcare staff

There are couple of barriers to AI adoption in healthcare domain. The successful AI solution requires innovation in integrating software systems to hospitals and involved new processes for doctors, pharmacists, and others to interact which is quite challenging. However, this can be reduced by proper training and user-friendly & easy to understand tools/processes.

Typically, advanced models are neural network based which are considered as black box which are difficult to understand and interpret. This lack of transparency can reduce trust in AI and reduce adoption by health care providers, especially considering that doctors and hospitals will likely be held accountable for decisions that involve AI. Interpretable and Responsible AI might increase trust by eliminating the black box problem, allowing health care workers to understand how AI reaches a certain recommendation.

There are various regulations in healthcare industry like, privacy regulation on patient's sensitive information and the regulatory approval process for a new medical technology takes time. KPMG has been compliant with these regulations and built these models without compromising patient rights or quality of care.

Impact: Successfully identifying patterns in A&E waiting times

The prediction model for A&E patient waiting time successfully identified 79% of the patients who will breach the 4-hour waiting time target at the point of A&E arrival. Model performance was monitored through area under Receiver Operating Characteristic (ROC) curve, precision and recall.

Model explainability analysis highlighted the important factors impacting these breaches. From patient side, patient age, previous number of admissions, month of arrival, and hour of arrival are the top characteristics contributing to the likelihood of breach. From hospital side, number of breaches in previous hours and the grade and specialism of the attending clinician are the top factors.

For example, the risk of breach is increased by the patient having one or more previous admissions and age. Specifically, infants under 5 years are less likely to breach in contrast to older patients (85+ years) are significantly more likely to breach compared to younger patients. Suggesting a difference in the clinical treatment pathways in A&E for these polemical age ranges.

The model helped the hospital to mitigate the risk of these breaches, improve service delivery and thus patient experience, coupled with reduced financial penalties due to A&E breach.

Barriers to widening AI adoption

In order for the UK to retain its position as a leader in AI and achieve its ambition to become an international superpower, it is fundamental to understand the barriers to greater AI adoption that exist in industry and how they can be mitigated at a time of unprecedented change and growth. techUK's AI Adoption Working Group identified five key thematic issues that could limit the greater implementation of AI within industry.

In no particular order, these include:

1. Inconsistent data quality & accessibility
2. Lack of trust in AI
3. A limiting organisational culture & understanding of AI
4. Insufficient compute infrastructure
5. Gap in AI skills

These issues will be discussed in more detail below, along with industry perspectives on these key issues. Recommendations to mitigate these barriers to put AI into action have also been suggested for industry leaders and policymakers.

Inconsistent data quality and accessibility

Data foundations have been defined in the National Data Strategy[20] as data that is: fit for purpose; recorded in a standardised format on modern and future-proofed systems; and, findable, accessible, interoperable and reusable. It is clear that strong data foundations are critical for developing AI systems that are robust; a study by EY found that all organisations that claimed to have high AI adoption also had high data foundations adoption.[21] High quality datasets are essential for training and testing AI systems to ensure that results are accurate and unbiased.

The availability of sector-specific data is essential for the commercialisation of AI research and development (R&D). Examples of such data include the **UK Biobank**, the NHS genomics dataset [22], as well as the Home Office's open datasets on policy recorded crime and outcomes.[23] However, there is still more work to be done on widening access to data, and increasing the availability of high quality, open-source datasets, particularly for SMEs. Unsurprisingly, data privacy and security contribute to this paucity, particularly in industries such as healthcare, financial services and Government where sensitive and personal data is involved. To ensure that the value of data is equally distributed across the economy and society, Government must take the lead to provide the right framework, tools, and resources to support organisations in sharing data ethically and responsibly.[24]

Additionally, in terms of access to data, there is often a disconnect between data owners/gatekeepers and businesses that need to deliver analytics projects or derive insights from the data.

This can make it difficult to obtain the necessary data for projects and could lead algorithms to be trained on datasets that are incomplete, leading to potential biases or inaccurate outcomes.

There are also concerns about competitors using the same datasets, which can raise issues around intellectual property, liability, and regulatory responsibility. Harmonising data assets can be challenging, as access agreements with datasets may sometimes prevent the use of certain data in combination with other data, thereby creating barriers to data availability.

Innovative approaches to responsibly promote data sharing and access to high quality data are therefore imperative to help address concerns, although no single method in isolation will act as a silver bullet. Data intermediaries, such as data trusts and trusted third parties can be one way to facilitate responsible access to or sharing of data, while protecting individual rights and privacy to minimise risk.[25] NHS Digital's Secure Data Environment (SDE) service is an example of an intermediary platform that provides approved researchers and analysts with de-identified health data.[26] New models of data consent that consider the role of civil society in decision-making could also assist data transfers. For example, a combination of data stewardship approaches, including collective and individual decision-making could facilitate greater transparency from an ethical perspective, therefore providing a means to access data.[27]

Technological approaches to data access and sharing through privacy enhancing technologies (PETs) are becoming increasingly popular. PETs are a suite of tools that can help to navigate commercial sensitivity and privacy concerns. For example, federated learning can be used to keep raw sensitive data safe in its original location whilst training machine learning models, and has been shown to be particularly promising in its practical use, although examples of end-to-end implementation remain scarce.[28] Synthetic data is another PET that can optimise privacy by representing the statistical properties of data, without using the original data values However, the use of PETs is still at an early stage and there are barriers to its widespread use. Knowledge and awareness of PETs still remains low, particularly in the public sector, and a lack of clear use cases, as well as standardisation have been cited as challenges.[29]

In summary, ensuring that businesses have the necessary digital foundations and access to the data they need is crucial for the wider adoption of AI. Data is the foundation of AI, and without the right data, businesses will struggle to build accurate and effective models. Therefore, Government-backed initiatives that support SMEs and start-ups in accessing data would be a welcome move. The UK could take example from **Data Pitch**, an EU-funded programme which brings together larger corporate and public sector organisations that have data with start-ups and SMEs that innovate using data.

Recommendations

Recommendation for Government:

1. **Facilitate collaboration between industry and Government to share public sector data and enable AI innovation that can create benefits for society.**
Trustworthy and responsible mechanisms to share public sector data for specific purposes should be considered to encourage innovation that can benefit the whole of society. These data-sharing arrangements must adopt an ethics-by-design approach, relying on existing solutions such as data intermediaries and PETs, to protect privacy.

Recommendation for Government and industry:

2. **Engage in Government and multi-sectoral collaboration to understand the challenges related to data quality and develop a set of industry-driven standards to address these barriers.**
There is a key role for standards, particularly for data sharing, in order to address poor data quality or a lack of consistency in how data is collected and stored. Industry and Government should work together to define a set of common standards, being inclusive of all stakeholders including large industry and small industry players, as well as civil society.

Recommendation for industry:

3. **Prioritise the robustness of data foundations and digitisation in the organisation.**
The value that can be derived from AI is dependent upon the quality of the data that has been used to build the model or application. A strong digital backbone and digitisation of existing data will reap dividends in the long-term in being able to extract insights from data and ultimately running successful AI projects that can be commercialised at scale.

Industry Perspectives
Why do AI projects fail?

Written by Mark Powell, Partner, Data & Analytics, EY

In the past 20 years, significant sums have been invested in AI and data programmes across UK industry. However, Gartner predicts that 85% of these programmes will ultimately fail to deliver successful outcomes.[30] Why is this? There are many reasons but, in my experience, these two may resonate more than most:

1. Organisations lack the necessary data foundations.

2. Organisations fail to connect their AI initiatives with business or market imperatives.

Build better data foundations

We obsess with the idea that data by themselves have intrinsic value, and we tend to hoard them, accordingly. Unfortunately, idle data doesn't deliver a return on investment. Instead, value for the business primarily depends on the beneficial outcomes that can be enabled by insights derived from data. One of the foundational principles of successful AI activities is that data only become valuable once their analysis has a strategic rationale.

In addition, it's all too easy to overlook problems with data quality. In an EY study of data foundations and AI adoption in the UK, carried out for the Department for Digital, Culture, Media and Sport in 2021, 41% of organisations said that quality was the most challenging characteristic of their data.[31] Unlike some conventional technology programmes, successful AI implementations disproportionately depend not only on 'gold standard' data sets, but also on data and algorithms that are fair and unbiased – qualities that are difficult to curate but provide payback for the up-front effort.

Finally, we underestimate the latent value that can be released by linking together different data sets – often data that has never been connected before. This 'integrational' value of data should help organisations to explore new directions of analysis and, potentially, drive enhanced value.

Focus on business and market imperatives

AI programmes often achieve little more than proving technological concepts that, while interesting, are never going to achieve transformational outcomes or even pay back the original investment. There's a lack of focus on operationalising the concepts or delivering cashable outcomes.

This problem may arise because those responsible for the AI (or the data) are usually not sufficiently close to the business to understand what is needed, while those responsible for the business have insufficient familiarity with AI to understand how it can help. Consequently, the services developed at best go round in circles or, like much of the data, simply sit idle.

Like many technologies that have come before, AI requires integration and industrialisation to succeed.

AI works best across departmental siloes and business functions by linking data, analyses, and insights, but often this is hard for businesses to achieve without appropriate investment in business process change alongside enhanced stakeholder engagement.

If you find yourself looking at the myriad ways AI appears to be changing the world but wondering why your organisation is missing so much of the value:

- Ask about the strength of your data foundations – How well is our data matched to our strategy? What data quality issues need to be overcome? Where are the opportunities to link disparate data together?

- Ask whether your technology initiatives are genuinely aligned to business need – What are we seeking to prove? How will we industrialise the concepts?

- Ask whether your people understand – What more can be done to help our technology teams understand the business and our business teams understand the technology?

- Ask whether your teams, structures, and processes are the right ones – How are we changing and reconfiguring our operating model to scale and maximise the opportunities?

The bottom line is that success in AI projects requires a pivot in our approach. We need to stop thinking of AI as a tool and instead think of it as the next wave in business transformation. Until organisations embrace this reality, they will continue to fail to deliver the benefits.

Lack of trust in AI

Trust in AI is multifaceted and can be affected by a range of factors, including the perceived competition between humans and AI, as well as how AI systems hold up against ethical principles. It is widely acknowledged that alongside diverse benefits, the deployment of some AI use cases could cause unnecessary harm, therefore responsible innovation and standards to cultivate trust is key. Despite the importance of trust in AI, a study by IBM[32] in 2022 found that 74% of organisations had not taken steps to reduce bias in their AI systems, and 61% were not able to explain AI-powered decisions. Moreover, another survey showed that only half of organisations had developed an AI risk management framework and the majority (70%) had not implemented ongoing monitoring and controls.[33]

To build trust in AI, organisations must take steps to ensure that their systems have the relevant safeguards in place and are transparent, accountable, and have a reduced risk of bias. For example, auditing protocols for AI would need to be built into risk management taxonomies in order to assess systems, map out risks and recommend measures that can be taken to mitigate them. Building in these safeguards is particularly prudent for high-risk applications, though it should be noted that not all deployments of AI are inherently risky. Building trust in AI is therefore essential for the widespread adoption of the technology in industry and acceptance by the wider public.

One of the key considerations in building trust in AI is the need to ensure that the technology is both built and used safely and ethically. This is particularly important in high-stakes sectors such as healthcare, financial services, and transportation, where poor product execution could have serious consequences. To mitigate these risks, the purpose of building or deploying a system should be properly examined and should be trained and tested against ethical standards before being deployed. Civil society could also be consulted where appropriate to assess the potential risks against the benefits. Furthermore, the integrity of AI requires ongoing oversight and cannot only be judged according to a single point in time.

Efforts to address unintended harms and increase public confidence have resulted in a wealth of responsible AI standards, frameworks and guidelines that are built on strong ethical principles, such as accountability, explainability, transparency and fairness. The tech industry has made progress in governance approaches and are continuing to develop more concrete and shareable measures, such as algorithm audits. This progress should be celebrated and encouraged.

However, ensuring that non-experts can understand how AI systems make decisions is also crucial in building trust and confidence in the technology. This can be partly achieved through the use of AI standards, which provide some guidance on how an AI system has been designed and implemented. As part of the National AI Strategy, the **AI Standards Hub** was launched in 2022 to provide an online platform to create practical tools for organisations to develop and benefit from global standards. This is a welcome step in improving the overall governance of AI in the UK.

Building a robust assurance regime would also be integral to increasing public trust and confidence in AI systems by evaluating and communicating reliable evidence about their trustworthiness according to relevant criteria. The Government has published a roadmap to an effective AI assurance ecosystem,[34] and techUK is working with the Centre for Data Ethics and Innovation to identify assurance practices already taking place within industry, to understand what is already working well and where further support to develop a well-functioning market may be required, such as the provision of high-quality AI assurance tools and services.

The issue of legal and regulatory uncertainty can also impact trust in AI. For example, in Europe, non-adopters of AI have been reported to feel concerned about the liability risk for damage caused by AI and the reputational risk associated with using the technology.[35] Efforts to increase certainty about the use of AI through policy, best practice guidance, and regulatory initiatives are increasingly being seen across the world. In the UK, techUK welcomes the actions that

UK Government departments and institutes are taking to support the growth of responsible AI, for example, through the creation of the **AI Standards Hub** and the Algorithmic Transparency Standard.[36]

There is broad agreement amongst techUK members that greater consistency and simplicity in regulatory approaches would be helpful, both to ensure high standards across industry whilst promoting innovation, and to give people whose lives are affected by AI a chance to understand the safeguards in place to protect them. Reducing uncertainty and providing clarity could help build trust in AI amongst the public which is key if the full potential of AI is to be realised. However, it is also seen as important that regulatory oversight is focused upon the use cases of AI and how it is deployed and used, rather than the AI technology itself. This will ensure that the UK maintains an innovation-friendly environment for businesses adopting AI, whilst protecting individuals impacted by the system from potential harms.

Recommendations

Recommendations for industry:

4. **Integrate controls for the development and deployment of AI in the organisation's existing risk management taxonomy.**
   Many organisations already have well-established risk management taxonomies in place. AI should be incorporated into these systems, including auditing, in a way that allows that ongoing oversight of the system.

5. **Engage with AI assurance tools and services to verify the trustworthiness of AI systems.**
   AI assurance tools and services can provide the necessary foundation for the monitoring and verification needed to evaluate and communicate reliable evidence regarding the trustworthiness of AI systems. This would be performed against standards, regulations and other principles or guidelines that exist, enabling other stakeholders in the ecosystem to build trust in the development and use of these systems.

Recommendation for Government:

6. **Work with industry to foster a high-quality AI assurance market.**
   Encourage and oversee the development of an effective AI assurance market, by working with industry to develop consistent and transparent requirements catering to different levels of risks and AI lifecycle stages.

Industry Perspectives
Responsible AI is maturing from theory to real-life

Written by Maria Axente, Responsible AI and AI for Good Lead, PwC

As the risks and potential harms associated with AI become more widely recognised, responsible development of AI has become a pressing concern for organisations across all sectors. Around the world, responsible AI is transitioning from being a "best practice" to a set of high-level principles and guidance that are essential for driving systemic change. Regulators have also taken note and are calling for regulatory frameworks that incorporate measures such as enhanced data protection, governance, risk management, and accountability. While organisations new to AI may seek technical solutions to concerns like bias, explainability, and privacy, those with more experience prioritise a comprehensive, values-driven approach to AI development and use, which involves addressing issues of ethics, risk, and governance in a holistic and sustainable manner.

Ethical AI adoption is growing

As industries increasingly turn to AI to improve human productivity and decision-making, there is a growing concern about whether the benefits to the bottom line outweigh the potential impact on society. In order to apply AI responsibly to real-world problems and address identified moral implications, developers, users, and organisations require clear guidance and principles. While many organisations have risen to this challenge and published codes of ethics and principles in the last five years, the next step is to operationalise them. However, those initiatives are often launched in isolation, from AI codes of conduct to ethical boards or frameworks, which can limit their effectiveness. It is essential to integrate these initiatives more comprehensively to achieve their intended impact.

For the development of 'human-centric' AI that promotes the 'common good,' ethical frameworks for AI must be in line with international human rights law, ensuring both legal and moral accountability. Furthermore, instead of a fragmented approach to AI ethics that focuses on individual initiatives such as ethical boards, codes of conduct, ethical training, and impact assessments, a more comprehensive approach is to adopt a systematic overview that incorporates a range of ethical AI interventions, as depicted in the figure on the next page.

	Principles and codes of conduct	The ethical principles defined by organizational values should be translated in organizational policies, codes of conduct and frameworks to allow for operationalizing those principles.
	External ethics boards	Ethical boards are part of the ethical decision-making through which ethical issues can be escalated, tensions can be managed and precedents can be set.
	Culture of ethics	Cultures are at the heart of this change and where ethical skills, knowledge and behavior should be recognized, rewarded and appreciated. Proper incentives and rewards schemes should be in place to stimulate ethical behavior.
	Education and training	Formal ethical training programs and curriculas should be embraced, along with other activities that will educate individuals about ethical thinking, analysis and reasoning. These include community practice, events, book clubs, team debates and hackathons.
	Reporting/ advice channels	Having appropriate means and ways for employees to receive advice regarding ethical dilemmas or to report breaches around AI and data can help identify potential ethical issues and solve them before they escalate.
	Product development and design	Ethical decision-making and actions should be operationalized at product level, with development process ethically aligned, ethical pit stops, and check and balances embedded at every step of the process to allow for the translation of principles into norms and the norms into design and governance requirements.
	Periodic assessments	Periodic audits are necessary to assess the performance of AI in terms of fairness, safety and reliability, and where relevant areas comply with internal and external standards.

Businesses Prioritise the Risks of AI

The rise in awareness of AI risks can be attributed to the misuse of AI and the ongoing societal discussion around its ethical implications. Our research reveals a growing trend among organisations to prioritise the identification and accountability of AI risks, often through an enterprise-wide approach to risk mitigation. A **PwC study** has found that a significant number of companies, 37%, now have a communicated strategy and policies in place to address AI risk, representing a substantial increase from previous years when only 18% did so. Additionally, 25% of companies have adopted an enterprise-wide approach to AI risk that not only encompasses standardisation, but also clear communication. Many organisations have established a comprehensive taxonomy of AI risks, which include application-level risks such as performance, control, and security, as well as broader ecosystem-level risks like enterprise, societal, and economic risks.
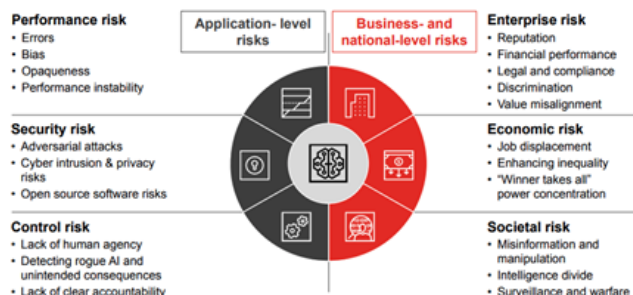
It is important to deploy an enterprise-wide risk management procedures to identify, evaluate, mitigate and monitor the AI risks over time. Educating and creating awareness of various AI risks, how they can occur and mitigating them should be raised within the whole organisation. This will create a culture of risk awareness through the teams task with building, using and maintaining AI systems.

Adoption of AI Governance Reaches a New Stage of Maturity

A sound governance framework and accountability are imperative for organisations to navigate changing regulatory environments, manage risks, and ensure ethical alignment with data and AI principles. Although technical solutions can help mitigate bias, improve explainability, and monitor the system's robustness, it's important to note that these tools have varying levels of maturity and may not entirely align with ethical standards. Therefore, it's essential to adopt a comprehensive approach to governance that emphasises processes, policies, and standards over technology.
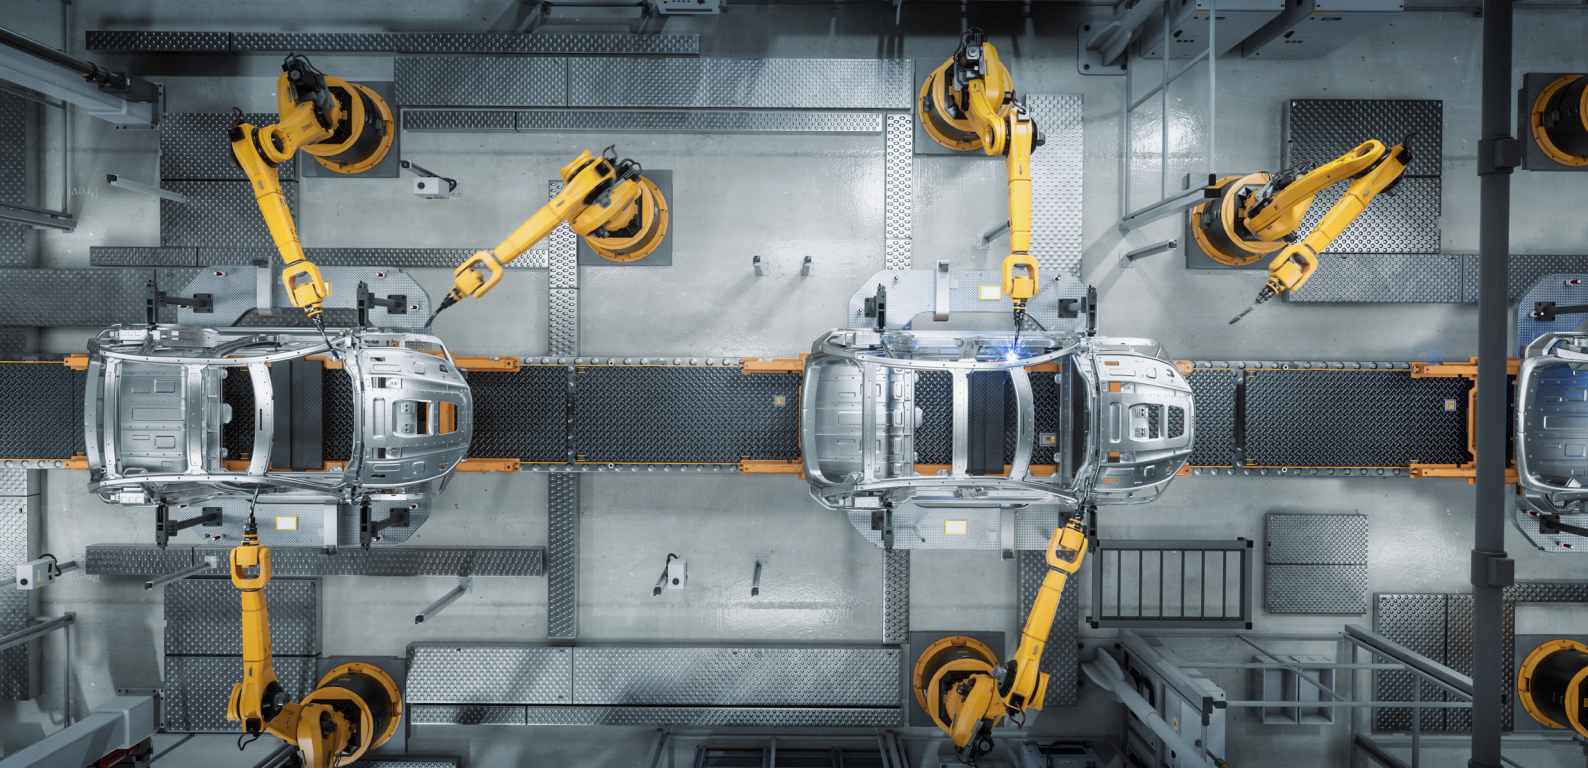
To implement a holistic governance approach, an organisation should first define its strategy, including the intended uses and expectations for data, analytics, and AI. During the planning phase, the organisation should identify its priorities

and create programs for model development and data usage. The ecosystem phase involves sourcing the necessary technology and personnel to achieve the targets established in the strategy phase. The level of governance should be proportional to the application context to avoid imposing overly burdensome requirements on development teams and stifling innovation. Governance requirements should consider factors such as the system's risk level, data privacy, novelty, and the need for new governance mechanisms.

Organisations should therefore define their guidelines and standards for governance that business units can leverage, including provision for internal and external AI assurance and validation. They should also convene diverse stakeholders across all parts of the organisation and beyond to collaborate on how to apply responsible AI practices consistently and effectively.

The 3 lines of defense ▶▶▶▶▶

Enterprise-level

External
Regulators
External auditors

External audit
Regulatory requirements and legal action
External standards development

3rd Line of Defense
Internal Audit
Ethicists

Internal audit review
Ethics review

2nd Line of Defense
Compliance
Legal & Privacy
Leadership
Independent Review Group
SME Model Approver

Standards and Specs
Independent Review and Challenge
Final Approval
Legal, Data & Privacy Governance requirements

Application-Level

1st Line of Defense
Model devs
Model owners
Ops support

Risk Assessment
Model Sourcing and/or Development
Data Collection and Use
Process Documentation

Infrastructure
Deployment
Ongoing Monitoring

A limiting organisational culture and understanding

As we have seen in the development and adoption of other emerging technologies, the culture of an organisation can be a catalyst for innovation, or a barrier. A culture that is not conducive to technological innovation may mean that it is difficult to implement an AI strategy that can effectively operationalise and apply AI to business goals.

To successfully deploy AI, organisations need a clear approach to data and AI governance. However, a key prerequisite to having an organisational strategy is ensuring that leaders are aware of the benefits and opportunities that AI can bring. A lack of leadership capability in organisations, particularly regarding the potential of AI, can lead to the benefits of AI not being fully recognised. It has been found that a lack of executive buy-in and proving business value are two significant challenges when starting AI projects.[37] It is important for business and technology leaders to establish key AI enablers, such as top-management sponsorship, a portfolio view of AI opportunities, and a data strategy, to avoid missing out on the AI opportunity.

Scaling AI's impact across an organisation is also essential. While many companies have already deployed AI to some extent, few have integrated it into standard operating processes in multiple business units or functions. A recent survey found that around one third of organisations were piloting the use of AI, and highlighted the risk of organisations becoming stuck in "pilot purgatory".[38] Putting AI into action and achieving results at scale requires the diffusion of AI capabilities across the organisation, as well as a commitment from leaders to drive large-scale change and a focus on change management, not just technology.

Recommendations

Recommendation for industry:

7. **Understand how AI can fit into the organisation's existing business strategy.**
    The way that AI is developed and used will differ across organisations. Therefore, it is important for businesses to consider how AI will help them to achieve their existing goals. This work could then grow into an individual AI strategy; however, it is important to ensure that this is in alignment with other business areas.

Industry Perspectives
Why you need an
AI Strategy

Written by Andrew Burgess, Founder of Greenhouse Intelligence

Many organisations know that they should be exploring the opportunities that AI could give them, but most do not know what AI can do for them, nor where to start. Some jump straight in with a proof-of-concept but find that they've chosen the wrong use case because it's too difficult (e.g. from poor or insufficient data, or lack of available skills) or doesn't give them any meaningful outcomes ('so what?'). Even when the proof-of-concept is successful, they then have to work out what to do next: how to get it into production, how to make sure it is ethical, plus softer aspects such as ensuring the model actually gets used by the people it is designed for. Many initial forays into AI grind to a halt for these reasons, thus wasting time, resources and, most importantly, missing out on the real opportunities that AI could bring to their organisation.

Creating an AI Strategy is one of the most important stages in the journey to exploiting the huge potential of artificial intelligence. An AI Strategy describes the 'art of the possible' for AI in the organisation, providing a roadmap to create a long-term, sustainable AI capability that, crucially, aligns to the business's objectives.

But how to know what the AI opportunities are? This is done (usually with external experts) by exploring where the different AI capabilities (e.g. image recognition, prediction, optimisation, clustering) can bring value, either by improving existing capabilities or finding completely new ways of doing things. The end result is a 'heat map', which maps each area of the business

to the identified AI opportunities based on the benefits that will be delivered, the ease of implementation and any other relevant criteria. At the core of this is a prioritised list of all of the identified AI opportunities. Business cases can be built for the highest priority opportunities, and ethical reviews carried out.

The Roadmap evolves from the prioritised opportunities, driven by aspects such as where there are the most benefits to be gained, how common certain types of AI technologies are, and which opportunities are the simplest to deliver. But there are also the softer aspects, such as where there is most enthusiasm for AI in the organisation and what ambition the organisation has for AI.

Measuring the level of AI ambition is a good way to identify the overall scope of the required change, but also helps articulate the vision to senior management and staff. The evaluation of maturity levels is a subjective one and can be assessed through interviews and the review of evidence. Many people use a third-party consultancy to carry out the evaluation to ensure consistency but also provide independence from any internal politics there may be.

The AI Strategy therefore brings together the prioritised list of AI opportunities, the heat map, roadmap, business cases and AI organisational ambition to describe the 'art of the possible' for the organisation and, most importantly, the best way to achieve the benefits for the long term.

Insufficient compute infrastructure

Compute infrastructure is an important consideration when it comes to implementing many AI solutions. Although there are many AI solutions that do not require a lot of processing power and can be built and used on lightweight hardware, access to powerful compute resources, such as supercomputers and data centres, can help to advance AI adoption. However, compute infrastructure is often overlooked and can be a significant barrier to the successful development and implementation of AI if it is not in place. Insufficient compute infrastructure has been found to be a common issue in terms of AI adoption, with non-AI adopters reporting a higher degree of this issue compared to adopters (74% vs 68%).[39]

The UK is a global leader in AI, but its large-scale computing capacity (LSC), particularly high-performance computing (HPC), lags behind other countries. The UK's HPC capacity has shrunk by three-fifths over the past five years, falling to 2% in 2019.[40] The relationship between HPC and AI is critical for supporting AI model training or scaling workloads, therefore, a lack of investment into HPC could hamper the UK's growth in AI. Without the appropriate compute infrastructure needed to further AI research, there is the added barrier to attracting and retaining talented data scientists, as people are often drawn to world-class compute infrastructure for the opportunity to work with cutting-edge technology.

The UK's largest HPC facility; DiRAC2, was established by the Science and Technology Facilities Council (STFC) and has been operating since 2012 to provide computing resources for industry and academia. However, the DiRAC facility is at risk of becoming uncompetitive as it has remained static in terms of overall capability since 2012.[41] A significant uplift in DiRAC's compute capability is needed in order to enable the facility to remain competitive. Bolstering the UK's HPC capacity will need to be a collaborative effort between industry and Government; the European HPC Joint Undertaking (EuroHPC JU) is an example of how this has worked particularly well in developing supercomputing across the EU.[42]

The recent Independent Review of The Future of Compute[43] emphasises the need to build a wider compute ecosystem in the UK that includes the role of emerging technologies, such as AI. The Government's new plans to set up an Exascale supercomputer facility, as well as an additional dedicated compute capacity for AI research[44] is therefore a highly welcomed move in responding to the Review.

Recommendations

Recommendation for Government

8. **Collaborate across industry, Government and academia to develop a National Large-Scale Computing Strategy to support AI development and implementation.**
    This strategy should set out a roadmap to support the AI industry and users of AI technologies to secure access to LSC infrastructure over the next decade, examining factors such as skills, investment, routes of access, data policy, and supply chain resilience. This strategy should be reviewed regularly to keep pace with international peers.

Industry Perspectives
Sustainable, sovereign, hybrid: why getting compute infrastructure right is critical to the democratisation of UK AI

Written by Bea Longworth, Government Affairs EMEA, NVIDIA

Compute infrastructure is the engine of AI. Just as the output of "traditional" industries, such as manufacturing, needs factories that are fit for purpose, AI researchers, scientists, and data scientists need the hardware and software tools to do their work.

The prospect of selecting and investing in AI infrastructure can be daunting, especially for smaller organisations and those who have not yet begun their AI adoption journey. Headline-grabbing large AI models like **ChatGPT**, boasting billions of parameters and demanding high-end compute horsepower, imply that impactful AI requires access to a supercomputer. Geopolitical tensions have given heightened focus to digital sovereignty and the value of storing and processing sensitive data locally. And the climate crisis requires that AI implementation must manage the environmental impact of compute resources to achieve a "twin transition". So far, so complicated. Fortunately, as adoption drives economies of scale, it's becoming much easier for organisations to match their AI aspirations to an infrastructure model that fits their budget, workloads and emissions targets.

The OECD defines AI compute capacity as "specialised stack of software and hardware (inclusive of processors, memory and networking) engineered to support AI-specific workloads or applications."[45] For those seeking to solve the most complex challenges, access to dedicated local AI supercomputing facilities remains critical. For example, Linköping University in Sweden recently unveiled an extension to its Berzelius supercomputer, tasked with accelerating national innovation in machine learning, robotics and healthcare.[46] The system has already enabled Sweden to develop its own sovereign large language model (LLM) for Swedish and other Nordic languages.

However, AI infrastructure doesn't mean supercomputing or nothing. It can take the form of an on-premises datacentre, cloud-based resources, or even a data science workstation or laptop. Each option has its advantages and disadvantages. An on-premises datacentre, whether accessible to industry, academia, the public sector or offered as a mixed multi-tenancy environment, can offer the ultimate in secure, dedicated resources which allow sensitive data to be stored and processed locally.

Alternatively, AI compute resources provided by a local or global CSP (cloud service provider) provide a high level of flexibility which has proved attractive.[47] And it doesn't have to be an either/or scenario. Many organisations are now blending on-premises and cloud resources to provide a hybrid AI compute environment that offers the best of both worlds.

Advances are also being made to ensure that reaping the benefits of AI doesn't come at the expense of climate targets. Greening AI is the responsibility of the whole ecosystem, with chip design and manufacturing to implementation through techniques such as "using pre-trained models, adopting greener forms of compute to maximise compute density given space constraints and designing more sustainable data centres with renewable resources."[48] The current focus on managing energy costs, and the possibility of incoming international regulation and certification schemes[49] to promote energy efficient datacentres, will give such efforts a further boost.

All these developments mean barriers to adopting the right AI compute environment are falling. That's a critical step in democratising this transformative technology and making it accessible to organisations of every size, from large enterprise to SMEs and startups, and from universities to individual developers.

AI skills gap

In the UK, demand for AI skills has more than tripled over the last decade, making it one of the leading nations in the world in terms of employer demand.[50] A lack of specialist skills to implement AI is a significant hindrance to wider adoption and has consistently appeared to be a challenge for organisations in recent global surveys.[51,52,53] If this gap between supply and demand is left unaddressed, the adoption of AI could be compromised in many organisations.

It is important to remember that it is not only data scientists that are needed to drive AI projects; in order to commercialise AI, organisations need teams that are multi-skilled to combine technical expertise with other capabilities, such as legal knowledge and commercial business experience.[54] Other roles within an organisation, such as business analysts, can also be upskilled as they already have the domain understanding to exploit data responsibly.

By empowering these individuals through upskilling and providing the right tools and platforms to collaborate with data scientists, this could accelerate the uptake of data and analytics projects more broadly. Equally, significant benefits would also be achieved by recruiting from cross-disciplinary backgrounds, including business and economics, as well as the social sciences and humanities. Focusing solely on technical skills risks losing sight of the bigger picture surrounding AI, such as ethics, the legal landscape and its commercial deployment.

Greater efforts also need to be made to promote the representation and inclusion of individuals from disadvantaged backgrounds, particularly within technical roles. In the UK, only 22% of data and AI professionals are women and this reduces to only 8% for researchers who contribute to high-profile conferences on machine learning.[52] Whilst the proportion of people from ethnic minorities and with disabilities in AI roles is understudied, it is known that they are underrepresented in the tech sector more broadly, particularly within senior management. Almost three-quarters of boards and 70.5% of senior executive teams in the top tech companies have no members from ethnic minority backgrounds and less than 10% of UK tech workers are disabled, compared to almost a quarter 23% of the wider population.[56]

As AI continues to mature, long-term strategies to increase the supply and diversity of AI skilled professionals should be considered. However, short-term solutions to help bridge the skills gap are also needed to help organisations realise and deliver the benefits of AI today. For example, the use of citizen data scientists is a potential short-term approach for this that could be explored. A citizen data scientist is an individual whose primary job function is outside the field of statistics and analytics but is able to create or generate models that leverage prescriptive or predictive analytics.[57] Empowering these individuals could help to deliver AI projects at a greater pace.

Clearly a combination of innovative short- and long-term strategies to address the UK's AI skills gap are therefore essential to supporting the UK's AI industry and helping organisations to increase their adoption and use of AI. techUK welcomes the steps that the Government has already made in collaborating with industry to accelerate AI upskilling and reduce the skills gap. For example, The **AI and Data Science Conversion Course scholarship programme** is already starting to address many of these issues. By funding the development of courses from non-STEM backgrounds, the programme has created a pathway for a wider pool of talent to join the AI workforce, particularly for those from underrepresented backgrounds.

Other pioneering programmes include **Industry funded AI Masters courses** and investing **£117 million to train doctoral researchers in AI** and **£46 million in Turing AI Fellowships**. However, given the gap between supply and demand that exists more action in this area is needed. These issues were recently discussed in depth at a techUK roundtable on AI skills.

techUK Ministerial Roundtable on AI Skills

Recently techUK held a discussion between Government and industry, through the AI Skills Ministerial Roundtable. Hosted alongside the Department for Science, Innovation & Technology (DSIT) (formerly the Department for Digital, Culture, Media and Sport (DCMS)) and chaired by the Minister for Tech and the Digital Economy Paul Scully MP, and Sue Daley, Director of Tech & Innovation at techUK, the roundtable welcomed leaders from across the tech industry for a discussion about the most pressing challenges they face when looking to hire AI talent and how collaboration with Government could help businesses find solutions. The roundtable was a prime example of how Government and industry can convene to open a dialogue about this issue and discuss ways in which we can work together to alleviate the skills gap.

Roundtable attendees reported a variety of barriers affecting the hiring of AI talent in the UK. For example:

- **The UK's education pipeline produces highly talented graduates, but more could be done to make it AI-ready.** Data-driven technologies rarely feature in the classroom, and a negative culture around maths could be turning people away from a career in tech.

- **Pathways into the AI workforce can be a limiting factor.** The number of entry-level roles is too low, and for the roles that do exist, there is often a lack of awareness about career pathways.

- **A smaller talent pool is impacting salary expectations.** SMEs find it particularly challenging to compete with big tech companies on the salaries they can offer, making it difficult to attract and retain talent.

The roundtable also discussed how Government and industry could work together to unlock the benefits of AI and possible solutions to address these challenges. The following ideas were proposed during the roundtable discussion as possible areas for activity:

- **Initiatives to actively encourage diversity.** Having role models and creating initiatives such as mentorship programmes and hackathons can inspire individuals from all walks of life.

- **Design an education pipeline that is flexible and agile.** The current education pipeline could be reviewed to see how it can be improved to meet the needs of an 'AI-enabled economy'.

- **Challenge recruitment practices and build career pathways.** Tapping into non-traditional pools of talent, on-the-job training and competency-based hiring practices were discussed as ways for businesses to develop the talent they need.

- **Upskilling at scale**
  Options could be developed to promote the upskilling of talented individuals at scale in a practical and tangible way, for example through a National Reskilling Framework.

- **Learning from best practice**
  Businesses can learn from organisations and countries leading on digital transformation and exchange knowledge on best practice.

Recommendations

Recommendations for industry:

9. **Create internal career pathways to develop AI skills and bridge talent gaps found within the organisation.**
One approach that could bridge skills gap within businesses in the short-term is the upskilling of existing employees with domain knowledge of data analytics. It is also important to ensure that talented individuals are supported in their career growth to cultivate strong AI leadership in organisations.

10. **Showcase innovative recruitment practices that encourage the uptake of individuals from diverse backgrounds by sharing best practices between industries.**
Efforts to increase the representation and inclusion of disadvantaged groups should be considered. Businesses should consider how to adopt more inclusive hiring practices including open and competency-based recruitment. Organisations should also share how they are tackling talent and diversity gaps within their organisation; this could help to inspire others who are earlier in their journey.

Conclusion

We are only at the beginning of seeing the full impact of AI adoption on businesses and society. It is therefore now more important than ever to encourage and support more businesses across all sectors and industries to put AI into action for the UK to establish itself as a global leader in AI and reap the full economic and societal benefits that this technology has to offer.

In this report we have outlined a variety of use cases which have shown how business processes can be transformed through adoption of AI and how AI technologies can provide real benefits for society as a whole. However, this is only a snippet of the AI opportunity and its future potential.

We have also set out the following recommendations for industry leaders and policymakers which we believe will help organisations to put AI into action and widen adoption. However, if we are to now move forward, industry and Government must work together to address the challenges to adoption we have outlined in the report to help businesses put AI into action.

Summary of recommendations

Barriers	Recommendations	
Inconsistent data quality and accessibility	**1**	Facilitate collaboration between industry and Government to share public sector data and enable AI innovation that can create benefits for society.
	2	Engage in Government and multi-sectoral collaboration to understand the challenges related to data quality and develop a set of industry-driven standards to address these barriers.
	3	Prioritise the robustness of data foundations and digitisation in your organisation.
Lack of trust in AI	**4**	Integrate controls for the development and deployment of AI in the organisation's existing risk management taxonomy.
	5	Engage with AI assurance tools and services to verify the trustworthiness of AI systems.
	6	Work with industry to foster a high-quality AI assurance market.
Limiting organisational culture	**7**	Understand how AI can fit into the organisation's existing business strategy.
Insufficient compute infrastructure	**8**	Collaborate across industry, Government and academia to develop a National Large-Scale Computing Strategy.
AI skills gap	**9**	Create internal career pathways to develop AI skills and bridge talent gaps found within the organisation.
	10	Showcase innovative recruitment practices that encourage the uptake of individuals from diverse backgrounds by sharing best practices between industries.

References

1. PwC, The economic impact of artificial intelligence on the UK economy, 2017
2. Oxford Insights, Government AI Readiness Index, 2022
3. HM Government, National AI Strategy, 2021
4. Zhang et al., The AI Index 2022 Annual Report, Stanford University, 2022
5. Capital Economics, AI Activity in UK Businesses, 2022
6. IBM, Global AI Adoption Index, 2022
7. McKinsey, Global survey: The state of AI in 2021, 2022
8. IBM, Global AI Adoption Index, 2022
9. McKinsey, Notes from the AI frontier: AI adoption advances, but foundational barriers remain, 2018
10. IBM, Global AI Adoption Index, 2022
11. McKinsey, Global survey: The state of AI in 2022, 2022
12. Deloitte, State of AI in the Enterprise, 2022
13. Capital Economics, AI Activity in UK Businesses, 2022
14. EY, Data foundations and AI adoption in the UK private and third sectors, 2021
15. IBM, Global AI Adoption Index, 2022
16. ibid
17. Capital Economics, AI Activity in UK Businesses, 2022
18. HM Treasury, Spring Budget 2023, 2023
19. McKinsey, Artificial intelligence in the United Kingdom: Prospects and challenges, 2019
20. HM Government, National Data Strategy, 2020
21. EY, Data foundations and AI adoption in the UK private and third sectors, 2021
22. Genomics England, National Genomic Research Library grows to over 110,000 clinically linked genomes, 2020
23. Home Office, Police recorded crime and outcomes open data tables, 2023
24. techUK, Data Sharing: Getting the UK back on the right track, 2022
25. CDEI, Unlocking the value of data: Exploring the role of data intermediaries, 2021
26. NHS Digital, Secure Data Environment Service
27. Ipsos, NHS AI Lab Public Dialogue on Data Stewardship, 2022
28. Open Data Institute, Federated Learning: an introduction, 2023
29. The Royal Society, From privacy to partnership, 2023
30. Gartner, Gartner Says Nearly Half of CIOs Are Planning to Deploy Artificial Intelligence", 2018
31. EY, Data foundations and AI adoption in the UK private and third sectors, 2021
32. IBM, Global AI Adoption Index, 2022
33. Accenture, From AI compliance to competitive advantage, 2022
34. CDEI, The roadmap to an effective AI assurance ecosystem, 2021
35. Bruegel, What is holding back artificial intelligence adoption in Europe?, 2021

37. Deloitte, State of AI in the Enterprise, 2022
38. McKinsey, Notes from the AI frontier: AI adoption advances, but foundational barriers remain, 2018
39. Bruegel, What is holding back artificial intelligence adoption in Europe?, 2021
40. GO-Science, Large-scale computing: the case for greater UK coordination, 2021
41. UKRI, DIRAC-3 Operations 2019-22 - UCL
42. EuroHPC JU, Discover EuroHPC JU
43. DSIT, Independent Review of The Future of Compute: Final report and recommendations, 2023
44. DSIT, Plan to forge a better Britain through science and technology unveiled, 2023
45. OECD, Measuring compute capacity: a critical step to capturing AI's full economic potential, 2022
46. https://www.di.se/digital/hans-aktie-rusar-pa-ai-hajpen-natt-en-anvandbar-punkt/ - non-Swedish speakers can use their favourite AI-powered automatic translation service to read the article
47. Gartner, Gartner Says Worldwide IaaS Public Cloud Services Market Grew 40.7% in 2020, 2021
48. OECD, Can I help save the planet, 2022
49. For example: https://digital-strategy.ec.europa.eu/en/policies/green-cloud
50. Lightcast, Artificial Intelligence in the UK, 2022
51. Bruegel, What is holding back artificial intelligence adoption in Europe?, 2021
52. IBM, Global AI Adoption Index, 2022
53. McKinsey, Global survey: The state of AI in 2022, 2022
54. Oxford Insights, Understanding UK AI R&D commercialisation and the role of standards, 2022
55. The Alan Turing Institute, Where are the women? Mapping the gender job gap in AI, 2021
56. Inclusive Tech Alliance, Inclusive Tech Alliance Report, 2018
57. Gartner, How to Use Citizen Data Scientists to Maximise Your D&A Strategy, 2021

![techUK logo]
techUK
FOR WHAT COMES NEXT

About techUK

techUK is a membership organisation that brings together people, companies and organisations to realise the positive outcomes of what digital technology can achieve. We collaborate across business, Government and stakeholders to fulfil the potential of technology to deliver a stronger society and more sustainable future. By providing expertise and insight, we support our members, partners and stakeholders as they prepare the UK for what comes next in a constantly changing world.

linkedin.com/company/techuk

@techUK

youtube.com/user/techUKViews

info@techuk.org