techUK
FOR WHAT COMES NEXT

Governance for an AI future: Enabling innovation and securing trust

March 2022

Governance for an AI future: Enabling innovation and securing trust

The UK is among the global leaders of AI, ranking highly in international comparisons of peer-reviewed publications and private investment.[1] AI technologies have the potential to drive economic growth, help improve many of the services we interact with daily and even contribute to solving some of the most complex social and environmental challenges facing the modern world. Yet, a recent global poll found that the population of Great Britain are among the most sceptical of AI use, with only 35% saying they trust a company using AI as much as they trust a company which does not.[2]

One way to help secure greater public trust is by adopting a clear and transparent approach to AI governance, which facilitates informed engagement. In the National AI Strategy published last year, the government established that a governance regime which "supports scientists, researchers and entrepreneurs while ensuring consumer and citizen confidence in AI technologies" is fundamental to securing the UK's ongoing position as a global AI superpower.[3]

techUK wholeheartedly welcomes the ambition to build such a regime. If successful, it will encourage greater overall AI adoption, fuelling the economy and improving the standard of living across the country – crucially, in a way that safeguards against potential adverse consequences of some AI applications. But this will only happen if we create a governance regime which encourages and enables AI innovation in the UK, and makes it easier, not harder, to navigate the responsibilities that befall those developing and deploying AI.

[1] OECD and Stanford's Institute for Human Centered Artificial Intelligence, Artificial Intelligence Index Report 2021
[2] Ipsos, Global Opinions and Expectations about Artificial Intelligence, 2022
[3] HM Government, National AI Strategy, 2021

As announced in the National AI Strategy, the government is currently preparing a white paper that is expected to set out its chosen path for the UK's future AI governance framework. To support this work, techUK has prepared this short paper setting out what we believe must be key elements included in the white paper if we are to a build proportionate, innovation-friendly and effective AI governance regime. We encourage the government to:

1. Take a risk-based approach; *tier AI governance requirements by the estimated level of risk posed by a given AI model or application, informed by clear criteria and categories.*

2. Consider the entire AI lifecycle; *clarify at what stages, from AI planning and procurement to ongoing use, risks can reasonably be expected to be addressed.*

3. Encourage and oversee the development of an effective AI assurance market; *work with industry to develop consistent and transparent requirements catering to different levels of risks and AI lifecycle stages.*

4. Acknowledge the role of existing regulation; *ensure that any potential new regulation or governance mechanisms do not replicate or contradict existing regulation.*

1. Take a risk-based approach: Scope and tiers

AI is not one thing; it is a general class of software. AI-based solutions can encompass many different technologies and approaches, including supervised, unsupervised, semi-supervised and reinforcement learning, and more traditional algorithmic approaches. This diversity of technologies requires the government's planned AI governance framework to identify clearly when AI applications will be within scope rather than apply blanket conditions to the use of any technology which may be classified as AI.

The most instrumental element in doing so will be by defining tiers of risk. There is a range of areas across industries which benefit from AI in ways that can be considered low-risk, for example within supply chain optimisation or business pricing strategies.

It would be counterproductive if AI in such contexts required extensive governance reporting as it could discourage uptake and delay the potential of AI to improve productivity. Governance requirements should therefore be based on estimated levels of risk, considering for example the risk of physical or psychological harm, infringements on fundamental rights or damage to the environment. Risk categories should factor in both the severity and likelihood of a potential harm and be subject to review at appropriate intervals. Governance requirements should focus on the highest risk uses of AI.

Recommendation:

Tier AI governance requirements by estimated levels of risk, with clearly defined categories and criteria.

2. Consider the entire lifecycle of AI systems

For a risk-based AI governance approach to be effective, it is not enough to just establish tiers of risk; it must also lay out how these risks are distributed between stages of the AI lifecycle. Responsible use of AI depends on those commissioning and planning the implementation of AI, those marketing AI solutions, those developing AI models and those overseeing its use and monitoring its impact – as well as a host of other stakeholders.

If an AI governance framework only emphasises the role of developers, it places an undue level of responsibility on AI providers. It also reduces the chances of achieving a truly safe and ethical AI environment. Developers of AI can focus on the expected model use and check key data indicators for performance and outcomes, but they cannot guarantee their products will be used as intended. They are also not in a position to oversee the iterations of each software cycle with its resulting changes, and therefore not able to predict or monitor the potential impact. These challenges are particularly salient for developers of general-purpose AI systems, who cannot predict what these will be used for or how they will be integrated or built upon.

Clarifying at what stages an identified risk can reasonably be expected to be addressed will encourage greater uptake of AI solutions, as each party can feel confident both in what has been checked at each stage of the lifecycle, and what their own responsibilities are. For this to work, there needs to be consistency in what is checked and explained for each stage. This could be enabled by an effective AI assurance ecosystem, explored further below.

Recommendation:

Provide clarity on the split of responsibilities between stakeholders in the AI lifecycle.

3. Encourage and oversee the development of an effective AI assurance market

AI assurance refers to a process of providing trustworthy information about how a product is performing on issues such as fairness, safety or reliability, and it can also be used to ensure compliance with any relevant standards.[4] The government recently published a roadmap to an effective AI assurance ecosystem,[5] and while it is difficult to assess the roadmap before the government has announced its chosen approach to AI governance, it seems clear that the provision of high-quality AI assurance tools and services will be crucial to the framework's success. We therefore strongly encourage the government to focus on the necessary steps to make sure these tools and services are available as soon as possible.

As outlined above, techUK proposes that the approach to AI governance should be based on tiers of risk. For the AI assurance system to support a risk-based framework, it needs tools and services catering to different levels of assurance needs. This could involve third-party auditors, and a critical component will be clear baseline requirements which create consistency across different providers of assurance services. Industry should inform discussions about what such baseline requirements should be, contributing insights gained from their experience in quality and safety assurance so far.

[4] Centre for Data Ethics and Innovation Blog, The need for effective AI assurance, 2021

[5] Centre for Data Ethics and Innovation, The roadmap to an effective AI assurance ecosystem, 2021

There may be lessons to draw from sectors with long-standing risk assessment processes. For example, in cyber security, the Network and Information Security Directive requires businesses identified as providers of essential services to complete a risk assessment framework, tailored to the estimated level of risk. Services classified as critical national infrastructure need to go through a framework overseen by a relevant competent authority. The cyber security sector also has well-established issue-reporting procedures, aiding consistency across stakeholders.

Just like businesses or service providers that are either required to or interested in assessing cyber security risks, it would be beneficial for those involved at any stage of the AI lifecycle to have tools or services to assess potential AI-associated risks, and clear ways of reporting any issues. This may increase both their own confidence and that of their stakeholders.

However, the assessment would have to be in line with our recommendations above to be proportionate – that is, only be required for AI systems in the high-risk tier and at the stages of the lifecycle where such risks can reasonably be expected to be addressed.

Finally, the role of international industry-driven standards should be considered as the AI assurance ecosystem is developing. As AI is a complex, global and evolving topic, the ongoing voluntary standardisation work developed by international standardisation organisations will have increasing relevance to AI governance across the world.

Recommendation:

Work with industry to foster a high-quality AI assurance market, catering to the needs of varying risk categories and stages of the AI lifecycle, based on consistent and transparent requirements.

4. Acknowledge the role of existing regulation

Lastly, the framework must make clear how new AI governance measures will interact with existing regulation across different sectors and levels of the supply chain. Many AI products and services are already regulated at a sectoral level, for example through the Medicines and Healthcare products Regulatory Agency in healthcare and the Financial Conduct Authority and Prudential Regulation Authority in financial services. Such existing regulation is often based on deep domain-specific expertise which addresses the risks of AI in the context it is used.

In addition, organisations are also subject to the Equalities Act 2010 and the Data Protection Act 2018 (including any forthcoming reform to the legislation following the Data: A new direction consultation). These will continue to apply to the development, deployment and use of AI regardless of AI-specific governance measures.

A situation where potential new AI governance measures could result in duplication and possible confusion about how AI-specific governance relates with existing regulation must be avoided. Any new measures must therefore be created to work alongside existing regulation, and regulatory bodies need to collaborate to ensure a consistent approach across sectors.

The role of the Digital Regulation Cooperation Forum will continue to be crucial in creating such consistency, which is particularly important to SMEs where the burden of confusing documentation requirements would have the most severe impact and may discourage uptake of AI altogether.

Recommendation:

Ensure that any potential new regulation or governance mechanisms do not replicate, confuse or contradict existing regulation.

Conclusion

Finding the right approach to AI governance is a question of balance. The government's ambition as set out in the National AI Strategy is the right one; to position the UK as the best place to live and work with AI. To enable the AI innovations of the future that will improve the economy and people's lives, businesses need clarity and certainty but they also need established ways of checking, as well as proving, that the way AI is created and deployed is responsible.

If the forthcoming white paper on AI governance adheres to the recommendations laid out in this paper, techUK believes it will bring us one step closer to a flourishing and trustworthy AI environment. Innovation will be encouraged through the recognition that many applications of AI pose low or no risk of harm, responsibility will be distributed fairly through the AI lifecycle, trust will be built through clear and consistent assurance mechanisms and bureaucracy minimised through clarity on the impact of existing and potential new regulatory measures.

techUK looks forward to working with the government throughout the process of finalising and implementing its approach to AI governance. If we get this right, it will enable an economy and society powered by responsible AI, working in the interest of the entire country.

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