AI regulation: a pro-innovation approach – consultation response

About the Institution of Engineering and Technology (IET)

The IET is a trusted adviser of independent, impartial evidence-based engineering and technology expertise. We are a registered charity and one of the world's leading professional societies for the engineering and technology community with over 155,000 members worldwide in 148 countries. We work collaboratively with government, industry and academia to engineer solutions to society's greatest societal challenges, including tackling climate change and building a better digital world.

Introduction

The IET welcomes the opportunity to respond to the government's white paper on AI regulation. AI can spur innovation and productivity across all engineering and technology sectors – including construction, energy, healthcare, and manufacturing. However, it is critical that the appropriate legal and regulatory structures are in place to allow its safe development and use.

We are well-placed to offer guidance on this topic, and this response was put together with input from AI thought leaders from our Digital, Healthcare, Innovation & Skills, and Safety policy panels. In the last year, we have published several papers the development and use of AI systems – for example, an analysis of AI technologies used in healthy ageing research, and guidance on the principles of the safe operation of AI systems. Our response addresses the questions at the end of each section of the white paper, and are labelled accordingly.

Main recommendations

- **A statutory duty on regulators to implement the framework is necessary to ensure the outlined principles are taken seriously and engaged with**
- **The government should set up a regulatory oversight body to co-ordinate guidance on good practice and deliver sanctions where misuse has occurred.**
- **Funding to deliver agile short courses (micro-credentials) for AI upskilling is required, given the fast-moving pace of the technology**
- **The government should establish a centralised repository of AI resources, tools, and guidelines for businesses to access and learn from**

1. **Transparency and redress (Section 3.2.3 A principles-based approach: Questions 1-6)**

AI systems must be transparent and explainable to maintain public trust in the technology. This is especially true if an AI system is responsible for delivering public services or safety-critical applications. Transparency and explainability require that AI systems be designed and implemented to allow for oversight, including – 1) the translation of their operations into intelligible outputs, and 2) the provision of information about where, when, and how they are being used. Moreover, the rationale and benefits of using AI should be made clear where it is employed.

We support the principle of the government's suggestion that organisations should make it clear when they are using AI – this would improve transparency for both regulators and end users. However, the approach taken should be tuned to the context to

provide maximum benefit. For example, a general notification appearing in every product which uses AI may end up becoming redundant and widely ignored, like the 'accept cookies' notifications which appear on websites.

In addition, we are concerned that the framework cites nuclear fusion as successful large-scale application of AI (section 1.1, paragraph 19). This is a safety-critical application where caution should take precedent. These systems must meet accepted levels of safety according to the use case, regardless of opportunity costs.

Statutory duty (Section 3.2.4 Our preferred model for applying the principles: Questions 7-8)

A statutory duty on regulators to implement the framework is necessary to ensure the outlined principles are taken seriously and engaged with. We do not foresee any equivalent alternatives. However, whatever is in place should support exemptions subject to monitoring and evaluation. In other words, the statutory approach should be risk-based and proportionate, and balanced to avoid duplicating efforts in areas in which careful use of AI has already been evidenced.

2. **Implementation of the framework**

   a. **Central functions (3.3.1 New central functions to support the framework: Questions 9-11)**

We support the central functions outlined in Box 3.1, and agree that they would benefit the implementation of an AI regulation framework. Centralised implementation of these functions ensures consistent interpretation and the implementation of principles across the AI ecosystem, allows for better monitoring and evaluation of the framework's effectiveness, and fosters cross-sector collaboration. It also aids in making the regulatory landscape more accessible for businesses and promoting innovation.

However, it should also be made clear who is responsible for delivering these functions, and how it will be costed. Existing organisations that could deliver one or more of the proposed central functions may include The Alan Turing Institute, The Information Commissioner's Office (ICO) and Innovate UK.

In addition to the central functions outlined in Box 3.1, we also recommend:

- Providing financial incentives or support mechanisms for businesses and researchers to invest in AI innovation, particularly for small and medium-sized enterprises.

- Establishing an ethical oversight board or committee with diverse representation from stakeholders, including experts in AI ethics, to ensure that the framework adheres to ethical standards and principles.

- Developing a platform for sharing best practices, research findings, and case studies among businesses, researchers, and regulators to encourage knowledge exchange and learning.

   b. **Encouraging innovation (3.3.1 New central functions to support the framework: Questions 12-14)**

i. Businesses

We support an approach which encourages businesses to innovate and use AI confidently. Skills and knowledge exchange should be at the heart of this approach. Therefore, we recommend the provision of the following initiatives:

- **Skills –**

50% of engineering employers that view AI as 'important to growth' report that they do not have the relevant skills to take advantage.[1] **A focus on agile short-courses (micro-credentials) is required given the fast-moving pace of the technology** – these courses could be delivered by institutions such as the Open University, or other HEIs with digital platforms. The government should make funding available for these courses to increase their reach, allowing employees to upskill / reskill in AI continuously.

In addition, the government should build on its current work by expanding the funding options available for more intensive, domain-specialist retraining – for example, covering the tuition fees of master's degrees in AI.

- **Knowledge exchange –**

We recommend that the government establishes a centralised repository of AI resources, tools, and guidelines for businesses to access and learn from. Knowledge exchange should also be encouraged at the local / regional level – for example, the North of England Robotics Innovation Centre, at Salford University, a hub for small to medium companies in the North West wanting to design test and validate innovation in robotics and automation.

ii. Consumers

In addition, there are several initiatives that could help consumers to use AI with confidence:

- A government-led public awareness campaign to educate individuals on the potential risks and benefits associated with AI technologies.

- Easy-to-understand resources and guidelines on AI technologies, their applications, and their implications for consumers – developed by regulators.

- Establishing a platform or helpline where consumers can report concerns, ask questions, or seek advice about AI technologies and their usage.

3. Monitoring and evaluation (3.3.1 New central functions to support the framework: Questions 15-17)

The white paper suggests that the government will 'develop and monitor' metrics to evaluate the framework's performance. These metrics should be co-developed alongside academic experts, professional bodies such as the IET, coalitions of end-user industries, and representatives of the technology industry.

---

[1] The IET, *Skills for a Digital Future survey*, 2023

Compliance alone is not sufficient to measure the impact of the framework – rather, it just shows that processes are being followed. Evidence of innovation and risk-management would suggest that the framework is performing as intended.

4. Role of regulators (3.3.1 New central functions to support the framework: Question 19 + L1)

The government's current approach of leaving implementation to existing regulators creates several challenges. Firstly, this approach increases the likelihood that regulators of different industries will produce contradictory guidance. Secondly, it may be challenging to identify appropriate regulators for all industries.

Therefore, we recommend that a regulatory oversight body be set-up to co-ordinate guidance on good practice and deliver sanctions where misuse has occurred. One option to deliver this is to set-up and fund a regulatory oversight body within the Health and Safety Executive (HSE), which has a track record of excellence, impartiality, credibility and accountability. This is necessary to ensure AI is used safely and help prevent incidents from occurring – this is fundamental to maintaining public trust, which underpins the economic and social benefits AI can bring.

The development of regulatory guidance should be a holistic, cross-sector process that accounts for perspectives in industry, academia, professional organisations, and the general public. **The government should initiate and fund the development of guidance for the safe use of AI via a trusted institution** – such as the BSI or the HSE. This guidance could take the form of British Standards Institute (BSI) standards or HSE Approved Codes of Practice (ACOPs). Whilst not legally binding, both approaches would provide robust benchmarks for users to adhere to. A starting point for such guidance could be the IET's AI-Safety policy position, which sets out ten key pillars that support the safe development and operation of AI systems in safety-critical applications.[2]

5. Role of government / legislation (Section 3.3.2 Government's role in addressing accountability across the life cycle: Questions L2-L3)

New legislation should clarify responsibility and accountability for the safety and security of AI systems, and outline powers to sanction misuse. In human-operated systems, someone can often be held accountable for negligence in these situations. However, in AI systems this is not always the case – the overseeing operator may not have been negligible. Therefore, victims of damage have fewer obvious paths to legal action. There should be greater legal clarity about responsibility for the safety and security of AI systems, and on powers to sanction their misuse.

Greater clarity is also required on the legal use of data in the research and development of AI systems. Many organisations do not understand their legal obligations in terms of permission to use data, or taking decisions based on that data. This lack of clarity hinders research and innovation. While there is much openly accessible data that could be used to 'train' and test algorithms in their early stages, it is unclear the extent to which such data is legally allowed to be used, especially when the data is not accompanied by explicit terms and conditions. Greater legal clarity on these issues would give researchers and

---

[2] The IET, *AI and functional safety*, 2022. (The IET will publish a more detailed version of this guidance in late 2023.)

companies more confidence to develop AI systems, supporting the UK's pro-innovation approach.

Furthermore, we encourage the government to undertake an immediate review of how AI-generated content is governed in the metaverse. The hyperrealism of these environments coupled with the freedom of movement is enabling new forms of abuse online. The quality of 'deep fake' imagery is now at the point where a realistic VR 'experience' of an abusive scenario featuring a child known to an offender in real life can be ordered from a supplier. The production of these models and scenarios is already a lucrative commercial venture on a scale of concern.

6. **Foundation models (Section 3.3.3 Foundation models and the regulatory framework: Questions F1-F3)**

Regarding question F2., measuring compute capacity would not be a good tool for the governance of foundation models. Newer models are increasingly efficient, and therefore the threshold for compute capacity in such a tool would be evermoving.

More generally, greater transparency is needed around the operation and use of Large Language Models (LLMs). Most publicly available LLM interfaces offer only a brief, incomplete description of the limitations. Our recommendation is that LLMs offer comprehensive, visible guidance to ensure that users can leverage these technologies effectively.

Greater transparency is also needed around the data used to train LLMs. The use of LLMs for enterprise poses legal risks if the intellectual property rights of the data cannot be traced. In addition, there is a lack of clarity about the use of open-source data for training these systems. The UK government should address this lack of clarity as legislation to regulate AI is introduced.