Establishing a Pro-Innovation Approach to Regulating AI

A Response to the UK's AI Regulation Policy Paper

*Dr. Robert H. Wortham and Catriona Gray**
September 2022

Summary of Submission

We welcome the recent publication of the UK government's proposals on the future regulation of AI in its Policy Paper, [Establishing a pro-innovation approach to regulating AI](#). Overall, we broadly support the sector and context-specific approach which stands in contrast to the anticipated cross-sectoral [EU regulatory regime](#). We believe the **key strengths** of the Policy Paper are:

(1) Its recognition of the complex reality of AI system design, development and use;

(2) Its ambition to promote the UK as a leader in responsible innovation;

(3) Its recognition of AI assurance as an enabler for the implementation of the principles in practice.

However, we have **several concerns**, set out in more detail below. These include issues relating to:

(1) Regulatory coordination and capacity;

(2) Potentially overlapping or contradictory regulatory remits and responsibilities;

(3) Burdens facing businesses operating across jurisdictions; and

(4) Insufficient detail on oversight and monitoring.

To fully realise opportunities presented by AI, the forthcoming White Paper ought to ensure the UK maintains its position and reputation as a leader in responsible innovation.

Question Responses

Question 1

What are the most important challenges with our existing approach to regulating AI? Do you have views on the most important gaps, overlaps or contradictions?

As highlighted in the paper, the UK approach to regulating AI has been relatively piecemeal until now. At present, there is considerable legal uncertainty, and no robust and consistent principles which would allow us to characterise the general properties of AI while remaining attentive to use-specific issues and technological advancements. This situation makes current rules **difficult to navigate without costly professional advice** and leaves businesses unable

to fully assess, mitigate and insure against risks. This holds back investment, innovation and competition.

Whilst recent policy initiatives on AI assurance – particularly the roadmap to an effective AI assurance ecosystem published in 2021 – indicate the right direction of travel, the pace of this **work must accelerate to meet expected demand for reliable assurance services.**

Question 2

Do you agree with the context-driven approach delivered through the UK's established regulators set out in this paper? What do you see as the benefits of this approach? What are the disadvantages? Do you agree that we should establish a set of cross-sectoral principles to guide our overall approach?

Global Digital Foundation broadly agrees with the context-driven approach. The rationale for avoiding a rigid or decontextualised definition of AI is sound. AI systems cannot be understood simply as static products; they are better understood as dynamic systems – often developed across complex supply chains, embedded within other products and systems, and applied across varied domains. The potential risks posed are also highly contextual. AI regulation, then, should seek to govern adoption and use as well as development. A key merit of the proposed approach is its **recognition of the complex reality of AI system design, development and use.**

The context-driven approach as proposed, however, has some potential shortcomings. Whilst efforts are already underway to coordinate the work of existing regulators, the proposed approach sets regulators an immense task. The paper suggests that the government will look for ways to "support and encourage regulatory coordination" including by working closely with the Digital Regulation Cooperation Forum (DRCF) and other regulators and stakeholders. We believe that a stronger coordination framework will be required to ensure the effectiveness and coherence of the regulatory regime. In addition, the role the DRCF and individual regulators will play in the AI assurance ecosystem should be clarified. The DRCF 2022-2023 workplan includes activities on algorithmic auditing but no reference to AI assurance or to the work of the CDEI.

All stakeholders, including the public, must have trust in the competence, independence, and integrity of regulators. To meet the objectives of the regulation, the capacity and powers of some regulators will need to be expanded significantly. **The upskilling required for regulators to effectively carry out these newly delegated duties should not be underestimated.** Careful consideration must be given to achieve the right mix of broad, multi-sectoral and domain-specific knowledge and expertise.

A multi-regulator approach brings additional risks of procedural and substantive inconsistencies. **Many AI systems will span multiple regulatory remits.** A lack of harmonisation in requirements and procedures will create legal uncertainty and, inevitably, legal challenge. This is likely to give rise to a risk of 'regulator shopping' whereby AI developers and operators seek to exploit potential loopholes. Finally, the context-driven approach appears to leave **gaps in less regulated sectors**, such as recruitment, which do not fall within the remit of specific regulatory bodies. Further analysis of these gaps could prove very useful.

Overall, we agree that the application of cross-sectoral principles to be interpreted within specific regulatory contexts is a sensible approach, and that the principles suggested are sound. However, we identify some potential shortcomings.

The proposal to task individual regulators with defining and operationalising notoriously elusive concepts such as 'fairness' may present problems. Some of these concepts and terms are already the subject of technical standards related to AI – for example, ISO/IEC TR24027:2021 on Bias in AI systems and AI aided decision making, as well as standards currently under development by ISO/IEC JTC1 SC42. Steps should be taken to **avoid unnecessary proliferation of definitions.**

The principle of allocating a defined legal person responsibility for AI governance, as it is currently stated, could be further clarified. Given the complexity of actors, inputs and interacting systems involved in the operation of many AI technologies, this requirement may place additional financial burdens and risks on some actors. The availability of **mature and reliable AI assurance services** will be essential for the effectiveness of this provision.

Question 3

Do the proposed cross-sectoral principles cover the common issues and risks posed by AI technologies? What, if anything, is missing?

The current principles could be further developed and refined in some places. Safety is rightly a guiding principle. However, AI systems can not only directly cause harm and injury but often form key components in systems to manage and monitor other safety risks. More holistic consideration needs to be given to the risks of using AI in safety critical environments across different sectors.

Whilst the principle that AI should be transparent and explainable is welcome, the discussion of this principle contained no mention of any duties to inform users or purchasers that a system incorporates AI technology. Whilst in some cases, there will be negligible risk of adverse outcomes, in others there may be good reason to expect candour from operators and vendors.

Many of the risks posed by the development and use of AI overlap with data governance issues. To support businesses in complying with any regulatory changes and guidance, the **interaction of the AI regulatory regime with the UK's data protection regime should be as seamless as possible.**

The proposal lacks detail on how "real, identifiable, unacceptable levels of risk" will be determined. Although the paper states that its approach is "risk-based" there is in fact little evidence of this. Further **elaboration is needed on the kinds of risks anticipated** and how these will fall within the remits of different regulators. The language of "hypothetical risk" is confusing and implies an unacceptable ex-post approach to risk management.

Finally, the Policy Paper contains **no reference to prohibitions on particularly harmful or dangerous AI applications.** This stands in contrast to the EU AI Act which expressly prohibits systems which pose an unacceptable level of risk, including those which enable social scoring by governments and those which manipulate users' behaviour. At present, it is unclear whether, and to what extent, regulators will have powers to ban unacceptably high- risk uses. The Policy Paper states: "In some high risk circumstances, regulators may deem that decisions which cannot be explained should be prohibited entirely - for instance in a tribunal where you have a right to challenge the logic of an accusation." This gives a great deal of responsibility to individual regulators and could result in contradictory practices and difficulties in monitoring the implementation of policy.

Question 4

Do you have any early views on how we best implement our approach? In your view, what are some of the key practical considerations? What will the regulatory system need to deliver on our approach? How can we best streamline and coordinate guidance on AI from regulators?

As stated in our response to question 2, **capacity building for regulators will be imperative for the effectiveness of the proposed approach.** The legal and institutional conditions, including clear and timely procedures for regulated entities to challenge decisions, must be put in place.

The role of end users ought to be concretely addressed in the forthcoming White Paper. Those affected by adverse outcomes related to AI systems, including end users, should have affordable and timely access to redress. We believe consumers have an important role to play in giving effect to and improving policy and regulation. The possibility of **establishing a specific consumer body to help scrutinise and monitor the regulatory regime**, and to make recommendations for reform, should be explored.

The White Paper should include a clearer and more comprehensive mapping of stakeholders across the AI lifecycle. Greater attention ought to be given, in particular, to regulatory compliance in procurement – one area in which existing regulators may currently lack specific expertise. The specific responsibilities and challenges faced by public authorities undertaking regulated activities using AI systems, and how this will be dealt with by regulators, should be detailed.

Question 5

Do you anticipate any challenges for businesses operating across multiple jurisdictions? Do you have any early views on how our approach could help support cross-border trade and international cooperation in the most effective way?

Businesses are **increasingly concerned about divergent approaches emerging across jurisdictions.** UK AI developers and businesses which operate within the European Union will eventually need to comply with the EU AI Act, and other relevant legislation, as a condition of accessing the EU market. Canada is also set to reform its legal and regulatory regime and has so far indicated an approach closer to that of the EU than the UK. Appropriate mechanisms and agreements must be put in place to ensure that UK businesses wishing to operate or expand into these markets are **not held back by burdensome and duplicate processes.** There is also an important role for AI assurance services in helping to clarify and streamline different requirements for businesses.

Question 6

Are you aware of any robust data sources to support monitoring the effectiveness of our approach, both at an individual regulator and system level?

We are not aware of any relevant sources.