

UK Government White Paper on AI Regulation

PA consultation response FINAL

Deadline 21 June 2023, to be submitted online by 23:45

All respondents and answers will be published, but not linked together.

Questions & responses

Our revised AI principles

Our framework is underpinned by five principles, which we expect to guide and inform the responsible development and use of AI in all sectors of the economy:

- 1) Safety, security and robustness
- 2) Appropriate transparency and explainability
- 3) Fairness
- 4) Accountability and governance
- 5) Contestability and redress

1. Do you agree that requiring organisations to make it clear when they are using AI would adequately ensure transparency?

~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

2. What other transparency measures would be appropriate, if any? Please limit your response to 1-2 sentences. [text box]

Disclosure on the use of AI is welcome but not sufficient in itself – there must be clear and agreed methods for communicating when AI is used, which AI tools are used (including versions), and which sources the ingested input data came from. This should include full transparency of input sources and machine-readable metadata. Given the exceptionally large datasets used by generative AI systems, transparency must be built into the input stage through adherence to copyright law including appropriate acknowledgement of sources and the licensing of data which is protected by copyright.

3. Do you agree that current routes to contestability or redress for AI-related harms are adequate? ~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

4. How could routes to contestability or redress for AI-related harms be improved, if at all? Please limit your response to 2-3 sentences. [text box]

Existing routes to contestability and redress must be strengthened and the relevant bodies, regulators, and ombudsmen given additional powers and resources to enforce existing laws.

This includes redress for rightsholders where their copyright-protected material has been infringed in training AI systems, including retrospectively. A review of the legal framework should be carried out to identify gaps to redress, and ensure legal responsibility for harms throughout the AI lifecycle are understood and respected. Agreed codes of conducts should be developed for AI systems, with meaningful methods of enforcement, and confidential and responsive new mechanisms developed for AI-related harms to be flagged, investigated, decided upon, and proportionate redress provided, whether by existing organisations with additional powers or new organisations if needed.

5. Do you agree that, when implemented effectively, the revised cross-sectoral principles will cover the risks posed by AI technologies?

~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

6. What, if anything, is missing from the revised principles? Please limit your response to 1-2 sentences.

The existing principles are good in theory but, given widespread evidence of copyright infringement for ingested input data, must include 'Compliance with all UK laws, including copyright and data protection'. Much is also dependent on effective implementation, which should be explicitly detailed e.g. plans, resource commitments, accountabilities, and timeframes, as well as an analysis of whether the existing legal and regulatory framework gives regulators sufficient power to enforce them.

A statutory duty to have due regard to the principles

The AI regulation framework will be implemented on a non-statutory basis at first. However, we anticipate that introducing a statutory 'duty to have due regard' on regulators might be needed to strengthen the framework at some point. A statutory duty would create a legal obligation on regulators to have due regard to the AI principles (see section 3.2.4).

7. Do you agree that introducing a statutory duty on regulators to have due regard to the principles would clarify and strengthen regulators' mandates to implement our principles while retaining a flexible approach to implementation?

~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

8. Is there an alternative statutory intervention that would be more effective? Please limit your response to 1-2 sentences. [text box]

While it would strengthen regulators' mandates, 'a statutory duty to have due regard' is not sufficient to enable regulators to act with the speed and authority needed to match the pace and significance of AI development. Individual regulators will need additional resources,

expertise, and additional powers to meaningfully enforce the principles. Alternatively, some form of cross-sector regulator or legislation may be required to prevent inconsistencies and contradictions in how principles are applied across sectors.

New central functions to support the framework

We intend to coordinate, monitor and adapt the framework through central mechanisms that will supplement and support the work of regulators without undermining their independence or duplicating existing activities. We will bring together a wide range of interested parties including regulators, international partners, industry, civil society organisations such as trade unions and advocacy groups, academia and the general public (section 3.3.1).

9. Do you agree that the functions outlined in [section 3.3.1](#) would benefit our AI regulation framework if delivered centrally?

Monitoring and evaluating the framework as a whole	Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know
Assessing and monitoring cross-economy risks arising from the use of AI	Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know
Scanning for future trends and analysing knowledge gaps to inform our response to emerging AI	Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know
Supporting AI innovators to get new technologies to market (see section 3.3.4 for more detail)	Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know
Promoting international alignment on AI regulation	Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know

10. What, if anything, is missing from the central functions? Please limit your response to 2-3 sentences. [text box]

The central functions are welcome in principle as a means of reducing regulatory inconsistency across sectors and increasing coordination on cross-sector issues. However, detail is needed on how they will be deployed and what powers they will have to operate effectively. For example, effective monitoring and evaluation will rely on the central function having the power to compel AI companies to transparently report on key data such as input sources, use cases, and user volumes. An additional function is required to analyse the legal framework in relation to AI, identify gaps, ensure legal responsibility is respected throughout the AI lifecycle, and ensure compliance, particularly with IP law.

11. Do you know of any existing organisations who should deliver one or more of our proposed central functions? Is there, for example, an academic research group that conducts AI horizon scanning or a think tank that gathers evidence on regulatory impact.

Yes (please describe) [text box]

An additional central function should be established to ensure compliance with UK law. For IP law, this can be delivered by the Intellectual Property Office (IPO) playing a more proactive role, for instance, in ensuring AI developers do not infringe copyright and seek appropriate licences in respect of the data they use to train LLMs.

~~No.~~

12. Are there additional activities that would help businesses confidently innovate and use AI technologies? Please limit your response to 2-3 sentences.

Yes (please describe) [text box]

- Clear guidance on, and enforcement of, UK copyright law in relation to AI. Compliance with copyright law by AI developers will give them, and the businesses employing AI systems, confidence to invest in AI on the basis that the systems are legal, the data sources known, reliable and licensed, and the role of human creators has been respected. Education, training, and clear guidance about copyright law for AI developers will ensure they can pursue innovation without infringing IP rights.
- Detailed analysis of the existing legal framework in relation to AI, ensuring that legal responsibility throughout the AI lifecycle is understood and respected.
- Clearer definitions of AI technologies, linked to their use cases, including clarity on what 'open' and 'closed' AI systems are, how they operate and interact, how they use and reproduce data, and the extent to which datasets can be ring-fenced within them.
- Development of expertise in regulators and links to research organisations to stay at the forefront of technologies.
- Pilot projects and benchmarks standards to test AI systems' safety and compliance with the framework principles.

~~Unsure~~

[12.1. If so, should these activities be delivered by government, regulators or a different organisation? For regulators only] [text box]

N/A

13. Are there additional activities that would help individuals and consumers confidently use AI technologies? Please limit your response to 2-3 sentences. [text box]

Yes (please describe) [text box]

Individuals' confidence and trust in AI systems will be built by those systems being safe, legally compliant, transparent in their processes, built on robust and reliable data, respectful of human creators, and with clear lines of redress. From a rightsholder's perspective, this can be assisted by AI developers respecting UK copyright law, being open and transparent about the data used in their systems (particularly to train LLMs) and licensing the curated, robust and reliable data streams that only rightsholders can offer. More broadly, consumers must have clear lines of redress and the application of AI systems in higher-risk scenarios should be subject to more stringent regulation.

~~No~~

~~Uncure~~

[13.1. If so, should these activities be delivered by government, regulators or a different organisation? For regulators only] [text box]

N/A

14. How can we avoid overlapping, duplicative or contradictory guidance on AI issued by different regulators?

The central functions must be given sufficient power and resource to fill gaps and prevent overlap, including through close communication and co-ordination between regulators and businesses. Certain standards should be established that apply across all regulators, including a defined set of criteria for regulators to assess risk levels and minimum standards for each principle. Consistent and co-ordinated guidance can also be achieved by ensuring all regulators and other relevant bodies demand compliance with existing law, including IP law. The need for a single AI regulator should be kept under review. It is critically important that the IPO takes the lead in making clear across the board that the training of AI systems on copyright protected content requires a licence from the rightsholder.

Monitoring and evaluation of the framework

We will need to monitor the implementation of the framework closely to make sure that it is working as designed. We will monitor the regime to ensure it aligns with 6 key characteristics, these being: pro-innovation, proportionate, adaptable, trustworthy, clear and collaborative (see box 3.2).

15. Do you agree with our overall approach to monitoring and evaluation?

~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

16. What is the best way to measure the impact of our framework? Please limit your response to 1-2 sentences. [text box]

The impact of the framework can only be assessed effectively if the central monitoring and evaluation function has sufficient power to collect the information it needs e.g. compel AI developers to share data relating to breaches of the principles. A review should be undertaken as to whether this can be achieved without new legislation or a single AI regulator. There also need to be clear metrics for what success of the principles look like for each regulator and regular dialogue with businesses.

17. Do you agree that our approach strikes the right balance between supporting AI innovation; addressing known, prioritised risks; and future-proofing the AI regulation framework?

~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

18. Do you agree that regulators are best placed to apply the principles and government is best placed to provide oversight and deliver central functions?

Yes

To a degree. Regulators will be well-placed to deal with issues in their sectors. However, government should be mindful of the significant additional resource, expertise and powers that will be needed by regulators to effectively enforce the principles. Given the growing intersectoral nature of businesses, especially in the digital space, central functions are also required to prevent contradictions and uncertainties across regulatory boundaries.

~~No (please expand) [text box]~~

~~Uncure~~

Regulator capabilities

While our approach does not involve extending any regulator's remit, regulating AI uses effectively will require many of our regulators to acquire new skills and expertise.

[19. As a regulator, what support would you need in order to apply the principles in a proportionate and pro-innovation way? For regulators only] [text box]

N/A

20. Do you agree that a pooled team of AI experts would be the most effective way to address capability gaps and help regulators apply the principles?

~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

Tools for trustworthy AI

Assurance techniques and technical standards will play a critical role in enabling the responsible adoption of AI and supporting the proposed regulatory framework. These techniques include impact assessment, audit, and performance testing along with formal verification methods (see part 4).

21. Which non-regulatory tools for trustworthy AI would most help organisations to embed the AI regulation principles into existing business processes? Please limit your response to 2-3 sentences. [text box]

Other tools could include; ethical guidelines and best practice examples such as full transparency of data sources, clear guidance that training of AI systems on copyright-protected content requires a licence from rightsholders, third-party audits and certification to show principles are being met, technical tools and checklists to assess AI products against regulations and principles, and a clear classification of AI risk and suggested risk assessments.

Final thoughts on the framework

22. Do you have any other thoughts on our overall approach? Please include any missed opportunities, flaws, and gaps in our framework. [text box]

The primary flaw is the total exclusion of intellectual property (IP). Given the exponential rise of generative AI, its reliance on a sustainable stream of the latest high quality data, and its significant implications for rightsholders and creators, IP should be at the heart of the global AI debate and addressed expressly in this consultation and the proposed regulatory framework. The UK's £109 billion creative industries are world-leading and it is imperative that, in taking steps to accommodate the expansion of AI, they are not 'traded away' or put at risk. It will be challenging for tT voluntary code of practice for AI and IP that the IPO is currently drafting to promptly and fully address the full range of issues we are highlighting in this response, especially given growing evidence that industrial-scale IP infringements have already taken place by AI firms. We encourage the government to take a clear and firm stance on UK laws, including copyright laws.

IP is essential to the development of many AI systems, particularly generative AI models which are dependent on the large datasets on which they are trained. Any regulatory framework must ensure compliance with UK IP law, which requires rightsholders' permission is sought for use of their content and in some cases adequate payment made in return for a license. This will ensure creators and rightsholders continue to produce and

invest in the datasets needed by AI. Anything else will diminish the supply of high-quality data that AI requires.

The responsible and lawful use of IP in AI will also serve to fulfil the cross-sector principles, including appropriate transparency and explainability, fairness, and accountability and governance. A robust regulatory framework, in which AI developers meet the legal requirement to license the IP used by AI models, will deliver transparency around training data, confidence in the robustness of the data used, and greater consumer trust in AI.

The government's proposed regulatory system can be amended to include IP with the addition of a cross-sector principle requiring compliance with all UK law, including IP law, and a central function tasked with ensuring compliance, analysing gaps, and ensuring legal responsibilities are properly understood and respected.

More broadly, we are concerned about whether regulators have sufficient expertise, resources and powers to fully implement the principles outlined, whether implementation will be effective without a strong central function, and whether the framework has taken sufficient consideration of potential harms.

Legal responsibility for AI

We recognise the need to consider which actors should be responsible and liable for complying with the AI principles. The ideal distribution of legal responsibility for AI may not be the same as the burden under current legal frameworks.

L1. What challenges might arise when regulators apply the principles across different AI applications and systems? How could we address these challenges through our proposed AI regulatory framework? Please limit your response to 3 sentences. [text box]

There will be challenges around consistency, including gaps and overlaps between regulators, and regulators applying principles in an inconsistent manner across sectors. There is also a significant need for global co-ordination and efforts to ensure that harmful AI systems are not permitted to develop abroad but affect UK rightsholders and consumers. AI platforms are already training AI systems in specific foreign territories in a blatant attempt to circumvent the UK legal and regulatory framework, and seek to influence its future direction at the expense of UK content industries.

L2.i. Do you agree that the implementation of our principles through existing legal frameworks will fairly and effectively allocate legal responsibility for AI across the life cycle?

~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

L2.ii. How could it be improved, if at all? Please limit your response to 1-2 sentences. [text box]

Clear and firm communication of the requirements of existing UK law, including the need for licensing, and routes to contestability and redress when the law is broken by AI developers are needed. For example, lack of understanding of, or respect for, the current legal framework and routes to redress mean that the UK's IP laws are not being enforced where datasets subject to copyright are being used to train AI systems without appropriate permission or compensation. A detailed analysis of the current framework is needed to identify gaps in enforcement and redress, and to ensure that legal responsibility is properly understood and respected throughout the AI lifecycle.

L3. If you are a business that develops, uses, or sells AI, how do you currently manage AI risk including through the wider supply chain? How could government support effective AI-related risk management? Please limit your response to 3 sentences. [text box]

Publishers have contributed to ethical frameworks specific to their sectors and monitor the use of AI within their products to ensure compliance with best practice.

Foundation models and the regulatory framework

Foundation models are an emerging type of general purpose AI that are trained on vast quantities of data and can be adapted to a wide range of tasks. The fast-paced development of foundation models brings novel challenges for governments seeking to regulate AI (see section 3.3.3).

F1. What specific challenges will foundation models such as large language models (LLMs) or open-source models pose for regulators trying to determine legal responsibility for AI outcomes? Please limit your response to 2-3 sentences. [text box]

Foundation models pose a profound challenge for regulators trying to enforce legal responsibility for AI outcomes. The inextricable link between the data, methods of processing, and outputs in LLMs means identifying the causal link between the outcome/harm and the AI actor will be complex. This is exacerbated by the developers of LLMs currently using datasets that infringe existing copyright law and failing to be transparent about what datasets they have used. A robust IP framework, in which the requirements of existing UK IP law (that content subject to copyright is licensed for use and properly attributed) are fully enforced will ensure greater transparency and means of redress at the input stage of LLMs.

F2. Do you agree that measuring compute provides a potential tool that could be considered as part of the governance of foundation models?

~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

F3. Are there other approaches to governing foundation models that would be more effective? Please limit your response to 1-2 sentences. [text box]

AI developers must be compelled to comply with IP law, including the need to license content subject to copyright, to ensure the models are legally compliant, and introduce greater transparency and redress at the input stage. More broadly, foundation models established abroad must be required to comply with UK law to provide their services in the UK (in line with similar requirements in the EU AI Act). The government could consider a regulatory regime similar to that of the Digital Markets, Competition and Consumers Bill for the most dominant LLM developers, in recognition of the likely imbalance in power and resources between big tech developers and many rightsholders and creators.

AI sandboxes and testbeds

Government is committed to supporting innovators by addressing regulatory challenges that prevent new, cutting-edge products from getting to market. To deliver an effective sandbox, we would like to understand more deeply what service focus would be most useful to industry.

S1. Which of the sandbox models described in [section 3.3.4](#) would be most likely to support innovation?

Single sector, single regulator (support innovators to bring AI products to the market in collaboration with a single regulator, focusing on only one chosen industry sector)	Strongly prevent innovation, Somewhat prevent innovation, No impact on innovation, Somewhat support innovation, Strongly support innovation, Don't know
Multiple industry sectors, single regulator (support AI innovators in collaboration with a single regulator that is capable of working across multiple industry sectors).	Strongly prevent innovation, Somewhat prevent innovation, No impact on innovation, Somewhat support innovation, Strongly support innovation, Don't know
Single sector, multiple regulator (establish a sandbox that operates in only one industry sector, but is capable of supporting AI innovators whose path to market requires interaction with one or more regulators operating in that sector)	Strongly prevent innovation, Somewhat prevent innovation, No impact on innovation, Somewhat support innovation, Strongly support innovation, Don't know
Multiple sectors, multiple regulators (a sandbox capable of operating with one or more regulators in one or more industry sectors to	Strongly prevent innovation, Somewhat prevent innovation, No impact on innovation, Somewhat support

help AI innovators reach their target market.
The DRCF is piloting a version of this model),,

innovation, ~~Strongly support~~
~~innovation, Don't know~~

S2. What could government do to maximise the benefit of sandboxes to AI innovators?

Please limit your response to 2-3 sentences. [text box]

- Provide investment and make it easy to ensure sandbox developments comply with regulations.
- Facilitate participation of a wider range of stakeholders such as AI innovators, regulators, data rights holders, industry stakeholders, and relevant experts, including researchers who are considering the ethical and societal implications of AI.
- Allow a safe, low-risk and fail-fast environment and promote learning and data-sharing in a secure way.

S3. What could government do to facilitate participation in an AI regulatory sandbox?

Please limit your response to 1-2 sentences. [text box]

Financial support, access to datasets, and a guarantee that participation would mean the applications developed are aligned to regulatory requirements, including compliance with UK laws, such as copyright.

S4. Which industry sectors or classes of product would most benefit from an AI sandbox?

Please select from this list the sectors your organisation works in or interacts with that would most benefit from a sandbox.

~~Primary sectors (extraction of raw materials, farming, fishing)~~

~~Secondary sector (utilities, construction, manufacturing)~~

~~Financial services & insurance~~

~~Communications~~

~~Hospitality and leisure~~

~~Real estate~~

~~IT~~

~~Legal services~~

~~Retail~~

~~Transportation~~

~~Healthcare~~

~~Education~~

~~Public sector~~

~~Research and development~~

~~Arts and entertainment~~

~~AI, digital, and technology~~

~~Regulation~~

~~Other [text box]~~

Publishers would welcome conversations with government about a range of projects that could be explored via sandboxes across publishing subsectors.

UK AI regulation impact assessment

Question 1: Do you agree that the rationale for intervention comprehensively covers and evidences current and future harms? The rationale for intervention argues that intervention is required in AI regulation. It outlines that government is best placed to put forward a suitable cross-sectoral regulatory regime due to the large benefits of AI that need to be harnessed and the need to mitigate the new and amplified risks AI poses. Further consideration must be given to the harm occasioned by prolific unauthorised and unlicensed use of copyright works to train AI models.

~~Yes~~

No (please expand, text box)

Current and future harms are not adequately covered as it excludes harms relating to copyright, such as the harm to rightsholders and creators caused by infringement of their copyright by AI developers.

~~Don't know~~

Question 2: Do you agree that increased trust is a significant driver of demand for AI systems?

Please provide your evidence.

Yes (text box)

It is imperative that AI systems are trained on a sustainable supply of the latest, high-quality data inputs. This is the only way to ensure AI outputs can be trusted and potentially harmful 'hallucinations', mis/disinformation, and biases avoided.

~~No (text box)~~

~~Unsure~~

Question 3: Do you have any additional evidence to support the following estimates and assumptions across the framework? These statements refer to all three options proposed in the impact assessment.

If you have evidence specific to a single option then please make this clear in your answer.

~~[tick box] The proposals will impact an estimated 431,671 businesses who adopt/consume AI products and services significantly less than the estimated 3,170 businesses who produce/supply AI products and services [text box]~~

~~[tick box] Those who adopt/consume AI products and services will face lower costs than those who produce and/or supply AI solutions products and services [text box]~~

~~[tick box] Familiarisation costs (here referring to the cost of businesses upskilling employees in new regulation) will land in the range of £2.7m to £33.7m [text box]~~

~~[tick box] Compliance costs (here reflecting the cost of businesses adjusting business elements to comply with new standards) will land in the range of £107m to £6.7bn [text box]~~

**Question 4: Do you agree with the estimates associated with the central functions?
If no, please suggest alternative estimate and explain reasoning.**

~~[tick box] The average FTE cost for a regulator is estimated to be £106k [text box]~~

~~[tick box] A central AI regulatory coordination function would require 50 full time workers [text box]~~

[tick box] A central AI regulator would require 300 full time workers [text box]

Given the rapid and pervasive nature of these new technologies, and the gaps in the framework (e.g. an intellectual property regulator), it's important to rapidly resource central functions sufficiently.

~~[tick box] The average number of AI systems developed per small business is 2 [text box]~~

~~[tick box] The average number of AI systems developed per medium business is 5 [text box]~~

~~[tick box] The average number of AI systems developed per large business is 10 [text box]~~

~~[tick box] The proposals will impact an estimated 431,671 businesses who have adopt/consume AI products and services, and an estimated 3,170 businesses who produce/supply AI products and services [text box]~~

Question 5: Are you aware of any alternative metrics to measure the policy objectives?

~~Yes (please expand) [text box]~~

~~No~~

~~Don't know~~

Question 6: Do you believe that some AI systems would be prohibited in Options 1 and 2, due to increased regulatory scrutiny? Please provide evidence to support your conclusion.

The impact assessment evaluated three different options proposed for AI regulation in the UK.

- **Option 1: Delegate to existing regulators, guided by non-statutory advisory principles**
- **Option 2 (preferred): Delegate to existing regulators with a duty to regard the principles, supported by central AI regulatory functions**
- **Option 3: Central AI regulator, with mandatory requirements for businesses aligned to the EU AI Act**

~~Yes (please provide evidence) [text box]~~

No (please provide evidence) [text box]

There is little evidence tech companies will be deterred from breaking laws (such as copyright), rules and norms simply due to 'regulatory scrutiny'. Many current AI industry practices are not transparent e.g. the sources of input data. To be effective the framework requires the central functions and existing regulators to have real teeth, of the kind the largest tech firms pay attention to (akin to the fines of global turnover proposed for digital markets laws in the Digital Markets, Competition and Consumers Bill in the UK, and those already in place in the EU).

~~Don't know~~

7: Do you agree with our assessment of each policy option against the objectives? See Table 9W in the impact assessment for details. [tick one]

~~Strongly disagree, Somewhat disagree, Neither agree nor disagree, Somewhat agree, Strongly agree, Don't know~~

8: Do you have any additional evidence that proves or disproves our analysis in the impact assessment?

No