Response – final version

Establishing a pro-innovation approach to regulating AI: BCS Response

September 2022

Table of Contents

BCS
The Chartered Institute for IT
3 Newbridge House,
Newbridge Square,
Swindon SN1 1BY
BCS is a registered charity: No 292786

This document

This is the BCS response to the UK government's policy paper 'Establishing a pro-innovation approach to regulating AI'[1]. The policy paper was published on July 18, 2022. It discusses how government can establish clear, innovation-friendly, and flexible approaches to regulating AI and how the regulatory regime will be able to keep pace with and respond to new and distinct challenges and opportunities posed by AI. Further details of the consultation can be found in this BCS briefing.

1   Executive Summary

The regulatory proposals are broadly welcome, but with caveats which are explained in the following sections. We agree that a light-touch, risk and context-based approach is sensible given that AI is still a set of emerging and rapidly evolving technologies.

We welcome the proposals to:
- extend the remit of existing regulators to deal with AI based on its use and likely impact
- focus on addressing issues where there is clear evidence of real risk or missed opportunities, provided regulators have suitable discretion to adopt a more precautionary approach with novel applications that are untried in real world settings
- use cross-sectoral principles tailored to the distinct characteristics of AI, and agree that those set out in the consultation are appropriate and should prove effective as a basis for future regulation (see Section 2), but there are gaps that need to be addressed, as explained below.

There are areas that need more consideration to ensure the proposals maximise the public benefit of AI:
1. The proposed cross sectorial principles are appropriate and useful, but should be extended further. There should be additional cross sector principles including:
   - AI systems must have appropriate safeguards to ensure they remain technically sound and are used ethically under reasonably foreseeable exceptional circumstances, as well as under normal circumstances
   - Organisations must show they have properly explored and mitigated against reasonably foreseeable unintended consequences of AI systems
   - AI systems should be standards compliant to enable effective use of digital analysis/auditing tools and techniques
   - Auditable data about AI systems should be generated in a standardised way that can be readily digitally processed and assimilated by regulators
   - Where necessary there should be recording of the outputs of AI systems to support analysis and demonstration that outcomes are appropriate and ethical. Note, such recording may include personal data, thereby adding an additional potential data protection challenge. Without such recording and analysis, the organisation would not be able to demonstrate the appropriateness of AI outputs nor demonstrate such appropriateness to regulators. In cases where external challenge arises about

---

[1] https://www.gov.uk/government/publications/establishing-a-pro-innovation-approach-to-regulating-ai/establishing-a-pro-innovation-approach-to-regulating-ai-policy-statement

> potential bias/unethical decision making, such recording and analysis will be an essential part of verifying or refuting any claims

- Organisational governance must be capable of dealing with complex software supply chains that are distributed across different legal jurisdictions (see Section 5)

In addition the proposals need to:

2. Discus how regulators will support organisations develop governance that enhances their freedom and autonomy to innovate responsibly (see Section 3).

3. Ensure there will be minimal divergence between the UK approach and that currently being developed by the EU, to enable UK companies to more easily compete in European markets.

4. Discuss how regulators, such as for example the Information Commissioners Office and Ofcom, will be able to handle the increased workload caused by the extensions in their remit. This could be significant given the frequency of changes that AI systems undergo, but also, for example, the possible impact of the Online Safety Bill on Ofcom's capacity.

5. Consider as a separate issue data quality, particularly of the input data to an AI system. There is a risk that an algorithm tested as acceptable based on 'good' data may deliver unacceptable outputs when using 'real world' data (e.g. such as data containing invalid/missing entries or that are not sufficiently accurate). Note, consideration of undesirable bias should be seen as a key aspect of assessing real world data quality.

6. Explain how regulatory overlap will be managed. I.e. when an AI system falls within the remit of multiple regulators, each with different interpretations of the cross-sectorial principles. For example, fairness may be interpreted differently due to different contexts.

7. Consider how to carefully phase in new AI regulations, given the change management challenges organisations will face in preparing for AI regulatory compliance (see Section 4 for further details).

8. Ensure transparency and appropriate checks and balances to address legitimate concerns over fundamental rights and freedoms that may occur if AI regulation is subject to legislative exceptions and exemptions (e.g. as in the 2018 Data Protection Act[2] where there are exceptions for Law Enforcement and Intelligence Service data processing).

9. Consider how regulation can foster the development of good professional practice for the design, development, and use of AI systems (see Section 3 for further details).

---

[2] https://www.legislation.gov.uk/ukpga/2018/12/part/3/enacted

2   AI Characteristics and Cross Sectoral Principles

We welcome the proposed AI Characteristics and the cross sectoral principles as likely to be appropriate and useful for dealing with a large class of problematical AI systems, as explained in more detail in the rest of this section.

Previous BCS studies highlighted an AI system should trigger alarm bells from an ethical perspective when it is:

- an automated system that must process data streams in real-time
- uses probabilistic self-learning algorithms to inform decisions that will have significant consequences for people
- is used in such a way it is difficult to uncover how decisions are derived
- is used where contestability of a decision is not deterministic and
- ultimately decisions rely on some form of best judgment that requires understanding of the broader context

We call an AI system problematic when it has the above attributes.

Problematic AI systems describe a significant class of systems that would be very challenging to ethically deploy, and hence should be below the necessary threshold for regulatory approval. Problematic AI systems like these easily match against the proposed AI characteristics of the consultation that would make them of interest to regulators. It is likely problematic AI systems as described above would fail to meet the proposed cross-sectoral principles and would not gain regulatory approval without significant design changes (technically and at a governance level). Hence, the proposed characteristics and principles are a welcome contribution to ensuring the ethical and competent development and use of AI.

However, the cross-sector principles should be further enhanced, as outlined in the executive summary, since there are a range of other cross-sectoral principles that are significant to assuring the ethical and technically sound development and use of AI.

3   Responsible Innovation and the role of Professional Practice

The National Innovation Strategy[3] makes it clear good regulation should enable and promote responsible innovation. Innovation thrives on freedom and autonomy in the pursuit of a clear purpose and vision.

---

[3] https://www.gov.uk/government/publications/the-government-technology-innovation-strategy/the-government-technology-innovation-strategy
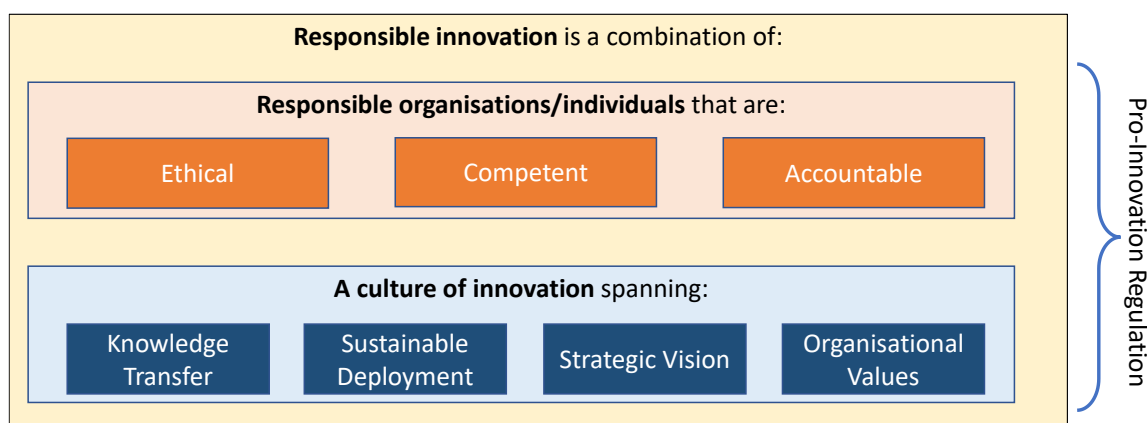
Response – final version

Figure 1: The role of pro-innovation regulation

The BCS view is that regulation should allow organisations as much freedom and autonomy as possible to innovate, provided those organisations can demonstrate they are ethical, competent and accountable when measured against standards that are relevant to the area of innovation. Pro-innovation regulation should enable effective knowledge transfer, the sustainable deployment of new technologies, as well as stimulate organisations to embrace innovative thinking as core to their strategic vision and values, as illustrated in Figure 1.

- A primary purpose of professional practice is to enable exactly the kind of responsible innovation as outlined in Figure 1.
- Which means, regulators should see supporting the development and adoption of good professional practice as a key enabler of responsible innovation.

The DCMS consultation[4] on 'App security and privacy interventions' is proposing introducing a (currently voluntary) code of practice for App developers and App stores to ensure they meet high cyber security standards. Note that consultation regards an App as any software product or service that can be downloaded from an online App store. E.g. the Microsoft App store falls into this category, which for example includes entire integrated software development environments as downloadable Apps. That consultation will have far reaching consequences across a large range of technology companies if its proposals are implemented.

The App store consultation provides a precedence for other regulators to follow. Working with support from professional bodies and other stakeholders with appropriate experience, regulators should support the development of codes of professional practice that facilitates effective regulation and responsible innovation. This is especially important since a significant amount of good professional practice will be cross-sectoral and relevant to all regulators. E.g. the Goldacre review[5] points out that much established software engineering practice can readily be adapted to the needs of the NHS and doing so should be a priority.

---

[4] https://www.gov.uk/government/consultations/app-security-and-privacy-interventions
[5] https://www.gov.uk/government/publications/better-broader-safer-using-health-data-for-research-and-analysis

4   Capacity and capability building

The DCMS commissioned report[6] 'AI Activity in UK Businesses' outlined the following barriers to adoption of AI:

Internal barriers These are barriers that are within a firm's control.	External barriers These barriers are outside a firm's direct control.
Cost of AI adoption High procurement and operation costs of AI as well as an uncertain return on AI investments impact on a firm's **demand** for AI solutions.	**Labour supply** Lack of sufficient skilled personnel within the firm and in the overall labour market constrain the **supply** of AI solutions.
Data Legacy infrastructure and insufficient data sophistication to leverage data's potential reduce a firm's **demand** for AI solutions.	**Regulation** Regulatory costs, restrictive privacy laws, and uncertainties around the liability for damages caused by AI limit **demand** for AI solutions and **supply** of AI technologies.
The human factor An incomplete understanding of the benefits of AI technologies and a lack of desire to move off the well-trodden path limit a firm's **demand** for AI solutions.	**Ethics** Opacity and lack of explicability and accountability inhibit the creation of trust in AI technologies, slowing **demand** for AI solutions.

Table 1: Barriers to adopting AI - from DCMS commissioned report

To a significant extent the box in Table 1 relating to regulatory barriers is linked to the other boxes in the table and especially with the lack of capability and capacity across all parts and all levels of seniority of an organisation to adopt AI. The Royal Society report[7] 'Regional absorptive capacity: the skills dimension' points out that organisations are held back by a lack of managerial expertise needed to generate growth through technology adoption, as well as a significant lack of technical capability in the workforce at Levels 4 and 5 (technician levels), and that this is true in all regions.

This means it will be challenging for a wide range of organisations to develop the necessary capabilities to comply with new AI regulations. We anticipate there will be many organisations in primary and secondary education, NHS trusts, construction, logistics, and farming, for example, that will initially be very challenged to develop the capabilities needed within their governance and management structures to comply with new regulations. This has been the case with the introduction of GDPR, for example, where many smaller organisations decided it was better, as far as possible, to just not keep data as the simplest

---

[6]https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/104538 1/AI_Activity_in_UK_Businesses_Report__Capital_Economics_and_DCMS__January_2022__Web_accessible_. pdf

[7] https://royalsociety.org/topics-policy/publications/2022/absorptive-capacity/

way to comply. Which meant those organisations lost opportunities to innovatively use their data lawfully to drive growth, which is most definitely not what was intended.

Good regulation should lead to good governance, which enables senior mangers to identify where resources need to be directed to build capacity and capability in their business. Conversely, regulatory non-compliance should help organisations identify where governance needs improvement and take appropriate action.

Hypothetical scenario

Take as a hypothetical example a national logistics company, where an understaffed, underfunded IT Services Department is told it is responsible for ensuring a new AI system used by the HR department to make decisions about staff performance is compliant with new AI regulations.

This is inappropriate and unlikely to succeed, since:
- There is no direct line of visibility to the Board of Directors, as the IT Services Department is funded through the Finance Office. The Finance Director (who is also the Data Protection Officer) has no direct oversight of the way AI is used in the HR department, and so cannot provide effective assurance to the Board of Directors concerning AI regulatory compliance.
- The IT Services Department do not have the data science expertise needed to assure the way AI is used, and do not have appropriate stakeholder relationships with managers to instigate any necessary changes of behaviour.

In this hypothetical example the company fails AI regulatory compliance. This leads to a much needed shake up of their data and AI governance. The company is supported by regulators to understand the necessary changes to their governance, such as appointing a CTO reporting to the Board of Directors taking over the role of DPO as well as AI regulatory oversight. The company also identifies the need to further enhance professionalism across the HR and IT departments, and supports those teams become registered with appropriate professional bodies.

All of which leads to improved staff retention and satisfaction, as well as regulatory approval for the way AI is used by the company.

In summary, AI is still a set of nascent technologies and throughout all UK regions organisations are struggling to build management and technical capability to successfully adopt AI. Regulators should support the necessary change management that organisations, such as those mentioned above, will need to comply with new regulations.

5   The software supply chain

Today most digital systems are the result of complex software supply chains, integrating products and services from businesses based in different legal jurisdictions and developed by disparate teams whose members constantly change. Software components from third party suppliers within the chain are frequently updated and patched or sometimes

completely replaced by a component from a different third party, resulting in the need for constant maintenance of digital systems. Every additional component in the software supply chain significantly increases the effort to maintain the final service/product to appropriate quality standards (including ethical standards) that are specified by service level agreements. All of which creates significant challenges for businesses to have the proper governance to guarantee products and services do what they are intended to do (including ethically) now and in the future.

AI systems are of course digital systems, and so face the same issues as outlined above of good governance for complex supply chains. That is why the issue of governance of software supply chains should be an important area of focus for AI regulators, as highlighted in the Executive Summary.

Who we are

BCS is the UK's Chartered Institute for Information Technology. The purpose of BCS as defined by its Royal Charter is to promote and advance the education and practice of computing for the benefit of the public.

We bring together industry, academics, practitioners, and government to share knowledge, promote new thinking, inform the design of new curricula, shape public policy and inform the public.

As the professional membership and accreditation body for Information Technology we serve over 60,000 members including practitioners, businesses, academics, and students, in the UK and internationally.

We also accredit the computing degree courses in over ninety universities around the UK. As a leading information technology qualification body, we offer a range of widely recognised professional and end-user qualifications.