BCS submission:
Science and Technology Select Committee call for evidence -
Governance of Artificial Intelligence

November 2022

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

Table of Contents

Introduction

This is the BCS submission to the Science and Technology Select Committee call for evidence on the governance of Artificial Intelligence[1]. The section headings are taken directly from the questions in the call for evidence.

In September 2021 the UK government published the National AI Strategy. Included in the strategy was an action to consult on proposals to regulate the use of AI under certain circumstances. Those proposals are published here: 'Establishing a pro-innovation approach to regulating AI proposals'. The BCS response to the consultation based on input from our professional membership can be found here. The key focus of the Science and Technology Select Committee is examining the government's proposals to regulate AI. Hence, the BCS evidence to the Committee is derived from our response to the government's consultation on those proposals.

1   How effective is current governance of AI in the UK?

It is important to appreciate that Artificial Intelligence (AI) is still a set of nascent technologies.

–   There are organisations in all UK regions struggling to build management and technical capability to successfully adopt AI.

Standards of AI governance in the UK are not uniformly high

Good governance of technology that impacts on people's lives, whether that is AI or some other digital technology, leads to high standards of ethical practice and high levels of public trust in the way the technology is used. Our evidence strongly indicates there is not a uniformly high level of ethical practice across information technology in general, and there

---

[1] https://committees.parliament.uk/work/6986/governance-of-artificial-intelligence-ai/

is a low level of public trust in the use of algorithms (including AI) to make decisions about people. Given governance of AI is often integrated within existing governance structures around data and digital, our evidence indicates:

– There is not a uniformly high standard of AI governance across the UK.

Further details of BCS surveys

Our most recent survey of ethical standards in information technology had responses from over 4,400 BCS members (see Table 1 for details) working in all sectors of the economy and at all levels of seniority. When asked to assess the general standard of ethical practice in the organisations they worked for 20% stated the standard of ethical practice was quite low and 4% that is was very low. 32% of members felt ethical standards were quite high or very high.

Current Ethical Standard	Overall
Very low	4%
Quite low	20%
Neither high nor low	36%
Quite high	28%
Very high	4%
Don't know	8%

Table 1: Results from survey of over 4,400 BCS members in 2018

Whilst it is encouraging that ethical standards were perceived to be high by almost a third of those responding, the responses also show a high level of variability overall, where a quarter of members reported low ethical standards.

Unweighted base	2076
Base: All UK adults	2076
The Government	10%
Social media companies (e.g. Facebook, Instagram etc.)	8%
'Big Tech' companies (e.g. Apple, Google etc.)	11%
Financial services (e.g. banks, insurance companies etc.)	16%
Health and social care (e.g. the NHS, private health care, the council etc.)	17%
Armed Forces	7%
The education sector	7%
The police	11%
Social Services	7%
National Security and Intelligence services	12%
Housing associations	6%
Other	1%
Don't know	16%
I do not trust any organisations to use algorithms to make decisions about me	**53%**

Table 2: Survey results to question - Which, if any, of the following organisations do you trust to use algorithms to make decisions about you personally

In 2020 BCS commissioned YouGov to survey a representative sample of 2,000 members of the public across the UK on trust in algorithms. The headline result from the survey was that:

– 53% responded 'I do not trust any organisations to use algorithms to make decisions about me'

The survey question in full was 'Which, if any, of the following organisations do you trust to use algorithms to make decisions about you personally' where the range of options to choose from is shown in Table 2 [NB - bold emphasis in the table was added only after the survey results were analysed].

2 What measures could make the use of AI more transparent and explainable to the public?

For the use of AI to be more transparent and explainable organisational governance should:

1. Identify reasonably foreseeable exceptional circumstances that may affect the operation of an AI system and show there are appropriate safeguards to ensure it remains technically sound and is used ethically under those circumstances, as well as under normal circumstances.

2. Evidence that organisations have properly explored and mitigated against reasonably foreseeable unintended consequences of AI systems.

3. Ensure AI systems are standards compliant to enable effective use of digital analysis/auditing tools and techniques.

4. Ensure auditable data about AI systems is generated in a standardised way that can be readily digitally processed and assimilated by regulators.

5. Where necessary record 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 potential bias/unethical decision making, such recording and analysis will be an essential part of verifying or refuting any claims

6. Treat data quality as a separate issue, particularly 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.

7. Be capable of dealing with complex software supply chains that are distributed across different legal jurisdictions.

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).

---

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

3   How should decisions involving AI be reviewed and scrutinised in both public and private sectors?

Properly reviewing AI decisions requires governance structures to follow the principles in Section 1. Decisions involving AI can then be properly reviewed as aspects of the governance structure. That is, the review of decisions made by an AI system, whether in the public or private sector, should focus on assuring that the proper governance structures are in place and governance processes are followed. That will mean decisions are based on data that is standards compliant and enables effective use of digital analysis/auditing tools and techniques to validate decisions.

Previous BCS studies highlighted that the use of an AI system should trigger alarm bells from a governance 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 scrutinise or review decisions made by such systems. The overarching issue should be to prevent problematic AI systems being used to make decisions about people in the first place, which is best done by ensuring the governance principles described in Section 1 are always followed.

4   How should the use of AI be regulated, and which body or bodies should provide regulatory oversight?

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.
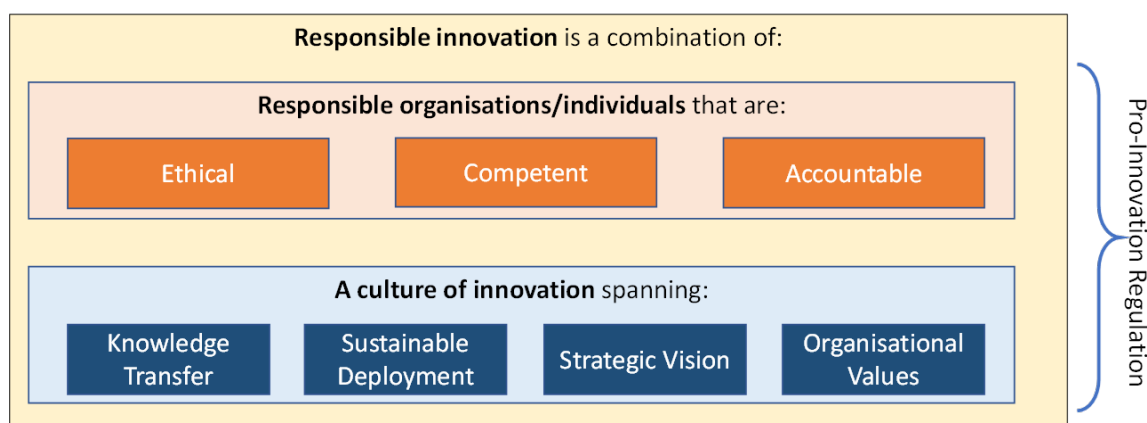
Figure 1: The role of pro-innovation regulation

In our response to the government proposals[3] for regulating AI we said 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, but there are gaps that need to be addressed, as explained in Section 1.

We believe that a light-touch, risk and context-based approach is sensible given that AI is still a set of emerging and rapidly evolving technologies.

4.1 Regulation of AI in the Health and Care Sector

The Health and Care sector is an area where there is an existing regulatory framework that applies to certain types of AI system. Currently regulation of software as a medical device is overseen by the Medicines and Healthcare products Regulatory Agency (MHRA). Evidence from BCS heath and care informaticians is that the MHRA will need to develop additional capabilities for effective regulation of AI systems in future.

Our evidence highlights informaticians are uncertain of the effectiveness of MHRA oversight of software systems, and therefore of AI systems in future. MHRA do not presently distinguish between an algorithm as a logic specification, versus an algorithm once it is embedded in an executable piece of software. Clinical knowledge developers like the National Institute for Health and Care Excellence (NICE) want to move from producing solely human-readable narrative to publishing digital artefacts into curated libraries. The libraries would contain clinical logic specifications expressed using open standards that can be implemented in multiple software products. Current MHRA practice does not seem to fit that model.

---

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

We note that the government's intention is to extend the remit of MHRA to include regulation of AI systems in the health and care sector in more general settings, rather than solely as embedded code within a medical device.

5 To what extent is the legal framework for the use of AI, especially in making decisions, fit for purpose?

Human review of AI decisions needs legal protection. Currently one of the key components of legislation that is applicable to AI is Article 22 of the GDPR, which focuses specifically on the right to review fully automated decisions.

Article 22 is not an easy provision to interpret and there is danger in interpreting it in isolation. There is need for greater clarity on the rights someone has in the scenario where there is fully automated decision making, which could have significant impact on that individual. Further consideration is needed that protection of human review of fully automated decisions is currently in a piece of legislation dealing with personal data. If no personal data is involved the protection does not apply, but the decision could still have a life-changing impact on someone.

We would also welcome clarity on whether Article 22(1) should be interpreted as a blanket prohibition of all automated data processing that fits the criteria or a more limited right to challenge a decision resulting from such processing.

6 What lessons, if any, can the UK learn from other countries on AI governance?

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 international supply chains. The Committee should consider how we

- ensure UK requirements of AI governance do not create tensions with governance requirements from other countries covering different parts of an international software supply chain
- minimise divergence between the UK approach and that currently being developed by the EU, to enable UK companies to more easily compete in European markets.

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.