

Written Evidence Submitted by Connected by Data (GAI0052)

There is an urgent need for more effective regulation and governance of AI in the UK. We would draw attention to three criteria of particular importance:

1. It must recognise the **collective impacts** of AI systems on groups, communities, societies and economies, not just the impacts on individual subjects of automated decisions.
2. It must include a **right of participation** in the development and governance of AI systems by the people and communities affected by it.
3. It must enable **accountability and enforcement**, through transparency; rights granted to organisations representing affected people and communities to act on their behalf; and meaningful penalties and redress mechanisms.

The remainder of this response expands on these while addressing the questions posed in the [Call for Evidence](#).

How effective is current governance of AI in the UK?

The governance of AI in the UK is largely reliant on the self-regulation of both public and private sector organisations. While there are provisions to govern automated decision making under the UK GDPR and Data Protection Act, these are ineffective in that they:

- only apply to AI that uses personal data to make decisions, when AI systems that use non-personal data such as satellite imagery or sensor readings can also be used to make decisions that impact the lives of people and communities
- are focused on individual impacts and therefore do not effectively regulate the impact of AI on communities or societies as a whole¹
- provide controls for data subjects, but not *decision* subjects who are affected by AI systems; this distinction is particularly important because AI may be used to make decisions about people and communities about whom little or nothing is known²
- have an emphasis on keeping a “human-in-the-loop” as a mechanism for limiting negative impacts, when these have been proven to be ineffective³
- are not effectively enforced, partly because concerns can only be raised by data subjects rather than others affected by decisions or organisations representing their interests (the UK opted out of implementing [Article 80\(2\) of GDPR](#)⁴ which would have enabled organisations to raise complaints independently of the mandate of a data subject)
- are on track to be further limited in their scope through amendments within the [Data Protection and Digital Information Bill](#)⁵ which will permit algorithmic decision making except in cases determined to be risky, rather than limiting it to situations proven to be safe

¹ Smuha, Nathalie A., Beyond the Individual: Governing AI’s Societal Harm (September 2021). Internet Policy Review, 10(3). <https://doi.org/10.14763/2021.3.1574>, Available at SSRN: <https://ssrn.com/abstract=3941956>

² Viljoen, Salome, A Relational Theory of Data Governance (November 11, 2020). Yale Law Journal, Forthcoming, Available at SSRN: <https://ssrn.com/abstract=3727562> or <http://dx.doi.org/10.2139/ssrn.3727562>

³ Green B. The flaws of policies requiring human oversight of government algorithms. Computer Law & Security Review. 2022 Jul 1;45:105681. <https://arxiv.org/abs/2109.05067>

⁴ <https://gdpr-info.eu/art-80-gdpr/>

⁵ <https://bills.parliament.uk/bills/3322>

Frameworks for AI governance recognise that there are different risk levels and types associated with the use of AI in different contexts. For example, facial recognition technology used to unlock phones has a different risk profile to that used to identify criminal suspects. Small automated food delivery vehicles have a different risk profile to self-driving cars, and to military drones.

Closer governance is particularly required around the use of AI in:

- **Public sector and essential services** such as utilities, which people cannot avoid interacting with
- Services that impact people's **wellbeing**, including their mental and financial wellbeing
- Services that provide **access to opportunities** (such as jobs, training, housing and finance), including automated pricing and selection or vetting algorithms; search engines; and digital marketplaces

Effective regulation and governance of AI needs to be able to be context-aware and context-specific. While it is possible to set an overall high level framework for AI governance, the details of the deployment of particular AI systems in particular ways and particular contexts need both sector-level and case-by-case consideration, which must also include the voices and concerns of the people and communities who are affected by the use of those systems. For example, there should be effective general governance processes for the use of AI systems in recruitment, distinct from other purposes to which AI is put, created in consultation with the public. Equally, there should be governance over the specific use of algorithms in recruitment at an organisational level, given each organisation will have its own particular context, processes and adopt different AI systems, created in consultation with employees and prospective applicants.

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

Transparency and explainability around the use of AI should be oriented towards what they enable, namely:

- **confidence** in the proper functioning of the system, to facilitate adoption and acceptance of effective AI systems, and prevent adoption of ineffective ones
- **accountability** for the impacts of AI systems, and the impact of the collection of data that is required to support it, to discourage and provide redress for the development and use of harmful AI systems

The necessity for transparency and explainability is linked to the amount of evidence available about the effectiveness and impacts of a given AI system, and to the known potential for harms to arise from its use. The impacts of AI systems – particularly those at the level of communities, societies and economies – often emerge from their use over time, so greater levels of transparency should be required for novel systems, as well as those that pose higher risks as outlined above.

Considering different audiences for information about AI systems helps to identify relevant and useful areas for transparency and explainability, for example:

- **When data is gathered about people and communities** to act as training or input data for an AI system, they need transparency about what data collection is happening; how data is stored to protect their privacy and security; the purpose of the AI systems being created using data about them; and how to object or opt-out to data about them being used in that way.
- **When AI is used to make decisions that affect people and communities**, they need transparency about the fact that decisions are being made in automated ways; how this fits into the larger process to which they are subject; on what basis those decisions are being made; and how to challenge or appeal those decisions.
- **Prospective purchasers or users of AI systems, academic researchers, and civil society organisations** that champion human, civil and consumer rights need transparency about the design process of AI systems; risk, ethics and impact assessments carried out during that process, which may limit the contexts in which the system should be used; and actual use and measured impacts of the

system. This should include its effectiveness against its purpose, and any wider impacts, such as on equalities, as well as reporting on the number and nature of complaints or appeals received.

- **AI auditors and both sector-specific and data-specific regulators** need similar information but may also require in-depth access to the datasets and code underpinning the AI system.

Note, for example, that the [Algorithmic Transparency Standard](#) developed by the Central Digital and Data Office⁶, which is being piloted for the publication of information about public sector algorithms, includes much of this information and takes an approach where different levels of information are provided for different audiences.

Organisations should be required to openly publish most of this information – the exception being the inner workings that only auditors and regulators require access to. Open publication is necessary to enable the impacts of AI systems to be researched and understood, and to enable their developers and users to be held to account for any harms that arise.

We believe it is important for the people and communities whose lives are affected by the collection and use of data, including through AI, to have a powerful voice in their development and governance. For this purpose, it is less important for there to be transparency about the inner workings of AI systems than for there to be transparency about the process through which they are developed, and their effectiveness and impacts. Organisations (particularly those developing and using novel and high-risk systems) should be required to publish:

- risk, ethics and impact assessments generated during the design and development of the AI system, including a description of how these assessments were carried out so that it is possible to see whether the people and communities affected by the AI system were consulted, and their concerns understood and addressed
- the effectiveness and impacts of the AI system, particularly equalities impacts, including a description of how these are measured

Transparency is also only effective when information is communicated clearly for the intended audience, and there are groups able and equipped to receive and engage with the information that is published. Transparency measures should be paired with support for active engagement of groups affected by AI systems, both to improve the quality of transparency disclosures and to better close the loop between transparency, scrutiny and accountability.

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

Review and scrutiny of decisions involving AI should be enabled in four ways:

1. **People and communities who are affected by particular AI decisions must be able to review and scrutinise** those decisions, and the data used as an input to them, and if necessary to appeal those decisions.
2. **Company directors, charity trustees, and other members of governance bodies, should have accountability** for the impacts of decisions made by AI systems within their organisations, in much the same way as they hold financial accountability. Boards should make use of internal reporting and AI auditing services to support their governance and scrutiny.
3. **Independent third parties with formal powers**, such as the National Audit Office, regulators and consortia of regulators such as the Digital Regulation Cooperation Forum, should have both the power to and duties to review the use and impact of AI systems across relevant sectors and within the individual organisations that they regulate.
4. **Civil society, academia, and media organisations should have sufficient access** to information about the aggregate use and impact of AI systems such that they are able to provide review and scrutiny on behalf of those affected by AI decisions.

⁶ <https://www.gov.uk/government/collections/algorithmic-transparency-standard>

It is not sufficient to rely on individual objections and complaints as the trigger for review and scrutiny. Biases and unfairness in algorithms is seldom detectable in any individual decision. For example, racial biases in insurance pricing are not apparent to an individual being given a quote; it is only by comparing quotes across all those who are provided with them that the bias can be understood. This is why organisations using AI systems must publish data about their aggregate impact, and civil society organisations must have the right to challenge and raise complaints on behalf of the communities they represent.

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

The regulation of the use of AI should be **proportionate**, as described above, based on the novelty and risks associated with the AI system; **responsive**, given that AI harms can be unpredictable; and **contextual**, given that the types of risks and harms vary depending on the kind of AI and environment in which it is used.

High level regulation should be put in place that defines general requirements:

- on the processes involved in developing and using AI systems, including the involvement of the people and communities affected by the AI system in its development
- on transparency, as described above, particularly of risk and impact assessments and measurement
- on routes of appeal and redress, such as the right to human review of AI decisions, the right for organisations to complain on behalf of communities, and mechanisms for the resolution of complaints through regulators and tribunals

Further regulation should be put in place on a case-by-case basis, for example around the use of facial recognition systems, or the use of AI within recruitment. Given the need for responsiveness, high-level regulation should empower regulators to create binding codes of conduct on the use of AI.

Where possible, regulatory oversight should be provided by existing regulators, who understand the sector context in which AI systems are being deployed, with support from the ICO. For example, the use of AI in the energy sector should be regulated by Ofgem and in the legal sector by the Legal Services Board. Where there are no organisations with regulatory powers, the ICO should provide this oversight with support from relevant representative or professional bodies.

Regulatory oversight of the public sector's use of AI should be carried out by an independent body reporting directly to Parliament; one possibility is the National Audit Office.

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

Please see our response to the above question about the effectiveness of AI governance.

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

We would draw particular attention to the [US Blueprint for an AI Bill of Rights](https://www.whitehouse.gov/ostp/ai-bill-of-rights/)⁷, which describes many of the rights that people and communities should have around the development and deployment of AI systems. The first of these is a right to safe and effective systems, which includes the statement that “Automated systems should be developed with consultation from diverse communities, stakeholders, and domain experts to identify concerns, risks, and potential impacts of the system”.

This reflects a growing consensus amongst AI governance experts that effectively predicting, detecting and mitigating the potential risks and harms of AI systems requires the participation of those people and communities affected by AI systems in their governance. As a regulatory measure, requiring public participation in the governance of AI systems has the additional benefits of providing context-aware governance that is

⁷ <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>

responsive to changing social norms and expectations, supports greater public data and AI literacy, and promotes the adoption of innovative technologies by increasing public understanding and trust.

About CONNECTED BY DATA

CONNECTED BY DATA is a campaign to give communities a powerful say in decisions about data to create a just, equitable and sustainable world. We want to put community at the centre of data narratives, practices and policies through collective, democratic and open data governance. We are a non-profit company limited by guarantee founded in March 2022 with funding from the Shuttleworth Foundation, and a staff team of three.

Our Executive Director, Dr Jeni Tennison OBE, is an internationally recognised data expert from her years of leadership at the Open Data Institute (ODI) and her role as co-chair of the Data Governance Working Group of the Global Partnership on AI. She is an associate researcher at the Bennett Institute for Public Policy at the University of Cambridge and an adjunct Professor at the Web Science Institute at the University of Southampton.

(November 2022)