



Developing proposals for regulation of public sector use of AI: roundtable findings

Mia Leslie

November 2023





Public Law Project is an independent national legal charity.

We are researchers, lawyers, trainers, and public law policy experts.

Our aim is to make sure state decision-making is fair and lawful and that anyone can hold the state to account.

For over 30 years we have represented and supported people marginalised through poverty, discrimination, or disadvantage when they have been affected by unlawful state decision-making.

Public Law Project contributes and responds to consultations, policy proposals, and legislation to ensure public law remedies, access to justice, and the rule of law are not undermined.

We provide evidence to inquiries, reviews, statutory bodies, and parliamentary committees in relation to our areas of expertise, and we publish independent research and guides to increase understanding of public law.

Public Law Project's research and publications are available at:

www.publiclawproject.org.uk/resources-search/



Contents

Introduction	4
Transparency	5
Public consultation	7
Rights related to human involvement	9
Accountability and redress	9
Regulation	11



Introduction

In March 2023, the Department for Science, Innovation and Technology and Office for Artificial Intelligence published a White Paper ‘A pro-innovation approach to AI regulation’ setting out its policy proposals for regulating AI in the UK.¹

In responding to the consultation on the proposals presented, PLP and many others raised concerns about the White Paper’s failure to properly consider the need for regulation of the use of AI in the public sector. PLP convened a joint statement, signed by 30 civil society organisations and individuals, that warned that the Government’s approach does not properly protect individuals from the risk of unfairness or discrimination when automation is used to make decisions that affect them.²

On 11 September 2023 PLP convened a two-hour private roundtable discussion between 10 legal, policy and regulation researchers, academics, and legal practitioners – chaired by PLP’s Mia Leslie – to explore the further development of practical policy and legislative proposals around the use of AI by public authorities.

The roundtable focused on five key themes of public sector use of AI: transparency, public consultation, rights related to human involvement, accountability and redress, and regulation.

This report summarises this discussion, with thanks to the participants for contributing their time and thoughts during the roundtable.

¹ Department for Science, Innovation and Technology (DSIT) and Office for Artificial Intelligence, *A pro-innovation approach to AI regulation* (29 March 2023) <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper>.

² Key principles for an alternative AI White Paper (June 2023) <https://publiclawproject.org.uk/content/uploads/2023/06/AI-alternative-white-paper-in-template.pdf>.



1. Transparency

1.1. Is a statutory duty the best way to achieve certain elements of transparency?

Some roundtable participants expressed hesitancy around the introduction of a general statutory duty for ensuring transparency, due to the complexity of implementation for government, the need for exemptions and the risk that a general duty (by nature of its wide application) will water-down its impact.

Participants were interested in the idea previously floated by Government, that engagement with the Algorithmic Transparency Recording Standard (ATRS) could be put on a legislative basis. Questions were raised as to why (or if) this commitment has been dropped.

At present, public authorities are encouraged to submit reports to the ATRS for algorithmic tools that either: have a significant influence on a decision-making process with direct or indirect public effect, or directly interact with the general public. However, many participants considered that some of the reports submitted to the ATRS Hub to date are of low potential impact, whilst reports for tools that carry the potential for higher impact are not present. To remedy this, participants discussed the idea of raising the threshold of which algorithmic tools should be prioritised for the ATRS to those that carry the potential for 'significant impact' but were unsure on what that would mean in practice. Concerns were shared that by raising the threshold of reporting to 'significant impact' a loophole could be created that allows public authorities to circumvent the requirement to complete and publish reports regarding their algorithmic tools.

Some participants were of the view that the starting point should be to advocate for maximum transparency. When considering what more substantive disclosure via the ATRS could look like, participants referred to PLP's Tracking Automated Government (TAG) Register and queried the differences between the two algorithmic registers, and whether asking for TAG Register-level and type of disclosure would be a good starting point.

One participant raised that when considering what transparency requirements we might want to see it is useful to reflect on what transparency is trying to achieve. This was articulated through the lens of two separate, or possibly interconnecting audience



groups, that may seek different levels of transparency because of the difference in what they are trying to achieve with it. The first audience group is the general public and/or decision-subjects and the second, a more expert audience. The participant suggested that the first group are likely to seek transparency around the systems and processes that have had effects on their individual lives, whereas the second group probably need a process duty to “get into the weeds” and seek a different level of information. The participant raised the question of how you can make people feel like a decision was transparent or ensure they can understand how it was reached.

In terms of securing transparency for the first audience group, some participants noted that in developing policy and legislative proposals for AI regulation, they have drawn inspiration from the Canadian regime which imposes an obligation on the relevant institutions to provide notice that a decision will be undertaken in whole or in part by an automated decision system, and to provide a meaningful explanation to affected individuals of how and why the decision was made. They consider the requirements around engagement with the ATRS to help with the transparency required by both audience groups, but that the detail of reports submitted may be of more use to the second audience group.

It was also suggested that transparency is not necessarily about understanding the minutia, it is about seeing how a decision is reached, and therefore notice about the presence of automation or AI in a decision-making process should be prominent, clear and provide information on the possible avenues for redress.

It was suggested by another participant that perhaps a ‘middle layer’ audience group that requires explainability of AI be considered. This could include actors such as the Information Commissioner’s Office to enable them to understand what is going on in a system and frontline professionals who are using AI as ‘decision-support’ technologies, to enable them adequate understanding so as to be able to explain to their clients the totality of how decisions are made.

Reflecting on the overall question of whether a ‘transparency duty’ should be introduced through AI regulation, a participant raised the questions: when do you want the duty to have effect? What about systems which are not yet in use? What will happen when significant changes are made to systems while it is in use? The duty will be ‘continuing’, so it is necessary to consider what amount of change to the system, process or algorithm is enough to require change of the disclosure made by the public authority?

Some participants have been looking at safeguards for data subjects and ex-post



safeguards around data controllers' obligations, and expressed views that decision subjects as well as data subjects should get information through transparency mechanisms. Another raised the question of whether we should open up the definition of 'legal or similarly significant effect on individuals' to 'on communities'.

2. Public consultation

2.1. To what extent would participation in an AI (or automated) tool's design and deployment be expected or wanted?

Participants were interested in the proposal of public consultation but flagged that it is not clear what forms of consultation are held around these systems already. It is possible that small levels of engagement (e.g. user research) already exist.

Whilst public consultation would be positive in principle, concerns were raised about the scale at which AI and automated tools are being or will continue to be introduced, which could mean endless streams of consultation. Participants wondered how the Government might strike a balance and reflected that without examples of where public consultation is being conducted well, it is difficult to propose how it might be carried out in relation to public authority use of AI within ADM systems and processes. One participant suggested that Camden's Data Charter might be a good example of public consultation. The question of resourcing also came up as the facilitation of effective public consultation requires adequate funding.

Some felt that the professional background of those attending the roundtable can mean that we sometimes overstate the effect of public consultation. One participant raised concerns around the danger of public consultation, expressing that currently the majority of public opinion supports hostile policies towards certain groups within society, such as migrants. They wondered whether consultation with those who hold such views and support hostile policies could permit the development and deployment of harmful AI or automated systems that affect such groups. Another said that obligations for transparency would undermine use of public consultation in horizon scanning, but that consultation can be good for raising problems within implementation early on.

Another participant referenced the findings of the Ada Lovelace Institute and The Alan



Turing Institute's 'How do people feel about AI?' survey,³ which revealed that understanding of AI and ADM is much lower for public use than private or commercial.

Some raised that there are examples where consultation processes have been in place, but the process did not result in 'meaningful' consultation and the views collected weren't represented in final decisions of how to proceed with the suggested policy or legislative change.

2.2. Could effective public consultation be achieved through the creation of a group that has direct engagement with the relevant Government department(s) and regulator?

Some participants thought that a user/engagement/stakeholder group might be effective, but wondered what the right level of engagement might be. Another supported the idea in terms of transparency but felt such groups usage is limited from a human rights campaigning perspective, as they can be easily sidelined.

One participant shared that their experience of user groups in practice is that they are ineffective. For example, this participant felt that there are many HM Courts and Tribunals Service user/engagement groups, but only one with real power. They also shared concerns around the framing of user group consultation questions, that often take the form of 'how should we do X?', rather than 'should we do X?'

It was mentioned that the introduction of automation or AI systems is mostly advocated for in terms of savings, and it is those sorts of reasons and/or decisions that engagement groups must look at. In initial deployments there are lots of vague discussions of what the system or tool will do, but "money is where the decision-maker's mind is". One participant suggested that there may be a role for engagement groups here - in interrogating what the success criteria for automation and AI systems are and how this will be measured.

³ Ada Lovelace Institute and The Alan Turing Institute, How do people feel about AI? A nationally representative survey of public attitudes to artificial intelligence in Britain (2023). Available at: <https://adalovelaceinstitute.org/report/public-attitudes-ai>.



3. Rights related to human involvement

3.1. Should new rights be established for decisions which aren't solely automated, but where automation makes up a meaningful part of the decision?

One participant suggested that in discussions around prohibiting automated decision-making in relation to Article 22 UK GDPR, we should start by demanding more than the status quo and remove the distinction between sole and partial decision-making.

Others were less supportive of introducing 'new rights', expressing that more rights might not be the best approach as it adds complexity to the existing framework. Instead, they supported the proper enforcement of the current status quo.

3.2. Should individuals have a right to make representations to the human reviewer/decision-maker, and to have the decision re-made by a human without the involvement of automation?

There can be scepticism about the extent to which further representations to decision makers will have any effect, so one participant suggested that perhaps higher quality understanding of what an automated system's recommendations are would be more effective. They suggested this could be brought to effect by incorporating clear confidence measures in systems.

Another participant shared favourable views of how the Data Protection and Digital Information (DPDI) Bill clarifies what is meant by 'meaningful' human involvement and solely automated decision making, and emphasised that effective enforcement of provisions is key.



4. Accountability and redress

4.1. What is the best vehicle for ensuring clarity and effectiveness of our legal framework and securing proper redress for individuals affected by public authority use of AI? Is a legislative solution appropriate or would guidance be sufficient?

One participant suggested that proper redress for AI and automated-related harms may be the creation of a 'digital bill of rights', that utilises the Human Rights Act 1998, and the Data Protection Act 2018 as example frameworks. They envision it to be a new piece of legislation that would ban the most harmful uses of AI, such as live facial recognition technology.

Another flagged the need to consider 'collective redress', similar to that allowed under Article 80(2)EU GDPR which allows unions and representative bodies to be involved. Within the UK, GDPR provisions do not permit this.

Another participant reminded us that there can be no meaningful avenues for redress without transparency, prompting the group to consider whether in decisions based on statistical likelihood, would disclosure of statistics alone be enough?

It was raised that general questions on what would be required to allow for 'redress' are challenging as it differs from sector to sector – the NHS and immigration are worlds apart.

In many sectors, there are already inspector-type mechanisms in place (such as the Independent Chief Inspector of Borders and Immigration for the Home Office, or the National Audit Office for the Department for Work and Pensions) which given specific powers/obligations, could play a role in holding government use of AI and automation to account. This was preferred by one participant, because they perceive Government as struggling to implement general obligations, but that more tailored provisions offer more protection in practice. Another participant supported this suggestion as it would spread responsibility for securing accountability.



4.2. Would the publication of assessments of the impact of policies and processes that involve AI and/or automation go far enough to allow routes to seek redress?

Some participants felt that this would be necessary but not sufficient. One recognised that the impact assessment frameworks that exist do not cover the full range of impacts of systems and wondered how to bring in participation of systems' users, builders and subjects.

5. Regulation

5.1. Do we need a specific AI Regulator?

In relation to individual decisions, a participant mentioned that their research shows that different decisions follow different processes, but they have all been designed in a pre-ADM world. So those processes might not work to articulate complaints. What's needed is greater flexibility and application of those rules.

It was also raised that at present, the Parliamentary and Health Service Ombudsman does not have powers to initiate their own investigations, but an expansion to that effect would be a good step forward. However, they went on to express that there is a need for more ombudsmen to deal with AI and ADM issues, as there is a lack of effective forum to bring parties together when complaints are levelled.

Another participant shared that following the AI white paper, they sensed consensus building towards the idea of empowering existing regulators, rather than creating a single AI regulator. But the question remains of what new powers will be awarded to deal with this.

Some posited that smaller sectoral regulators will not retain sufficient in-house expertise to look into use of AI and ADM, and that the different powers, skills and expertise in different existing regulators presents a challenge.

Some felt that the Information Commissioner's Office (ICO) would be best placed to act as an 'AI regulator', as there is a risk that decoupling AI and data protection would jeopardise a cohesive approach. However, others articulated that the ICO is not currently well-resourced or sufficiently bold, and that there is a mismatch between



what the ICO could be doing, as per the Supreme Court interpretation of their powers and obligations, and the same approach taken in practice.

Others added that they were sceptical about the effectiveness of wide-reaching regulators, and whilst they would want to give a regulator the power to bring judicial reviews, they would also have to enforce existing regimes.

One concern held if the ICO were to be the 'AI regulator' is that AI and automation-related harms require the consideration of a wider range of rights, beyond those within the data protection framework.

A participant made a two-pronged recommendation, that first there must be an increase in the resources of existing regulators, and second that there must be an appropriate ombudsperson to provide first point of contact for those seeking redress and to provide guidance to different regulators.



Public
Law
Project

The Public Law Project (PLP) is an independent national legal charity. Our mission is to improve public decision making and facilitate access to justice. We work through a combination of research and policy work, training and conferences, and providing second-tier support and legal casework including public interest litigation.

www.publiclawproject.org.uk