



Policy Briefing

Learn fast and build things

Lessons from six years of studying AI in the public sector

[Imogen Parker](#), [Anna Studman](#), [Elliot Jones](#)

14 March 2025

Reading time: 53 minutes

Research domain [Society, Justice & Public Services](#)

Keywords Public trust AI in the public sector

Key insights

At the Ada Lovelace Institute, we have spent the past six years examining AI and data-driven technologies in the public sector. As governments seek to accelerate the use of AI in policymaking and public services, here are the lessons for success that we have drawn from more than 30 reports and research publications.

Contextualise AI

Lack of clear terminology about 'AI' is inhibiting learning and effective use.

Shared definitions for AI are needed for effective communication, evaluation and to make strategic decisions. Governments should build on existing work to develop a taxonomy of AI tools that accounts for the significant variation in their purpose, technology type and context of use.

AI is only as good as the data underpinning it.

AI systems rely on data, which is never neutral, often partial and can encode existing societal inequalities. For AI to be effective and for data-driven systems to work as expected, they need high-quality, relevant and complete enough data to solve a specific problem. Those using AI need to mitigate limitations in data collection, which range from the underrepresentation of digitally excluded and marginalised groups through to structural inconsistencies in how data is recorded, collected and understood.

AI systems are not deployed in a vacuum and context is important.

AI is 'sociotechnical', in that it influences and is influenced by the social contexts in which it is deployed, often with unintended ripple effects. The success and acceptance of AI tools depends on their interaction with existing social systems, values and trust. Focusing exclusively on technical criteria while failing to consider these factors can lead to scepticism, and ultimately hinder adoption and use.

Learn what works

The public sector does not have a comprehensive view of where AI is being deployed in government and public services.

The public sector's understanding of its own AI usage is severely lacking, which hinders both democratic accountability and internal knowledge sharing. More transparency is crucial for enabling scrutiny and ensuring effective AI implementation. This includes transparency in procurement processes, tool functionality and public impact, as well as addressing the 'slipstreaming' of AI into existing products.

There is not enough evidence on the effectiveness of AI tools.

There is a surprising lack of evidence on the effectiveness and impact of AI tools, even from a purely technical standpoint. Evaluating AI interventions in context is crucial to determining their performance and value compared to existing manual or traditional methods. The public sector needs structures to learn from achievements and failures, to identify conditions for success and to disseminate that learning; for example, an expansion of the existing 'What Works Centres' model.

Deliver on public expectations and public sector values

Successful use of AI requires public licence.

Moving out of step with public comfort can undermine the ability for the public sector to effectively use AI. Public backlash has led to people withdrawing from data sharing or hampered the use of existing tools. Some frontline professionals have been reluctant to use AI tools where they do not see them as defensible and legitimate.

Public procurement of AI is not fit for purpose.

Existing public sector procurement processes are not effective in ensuring transparency and fairness around AI products, or for avoiding vendor lock-in. Market concentration and knowledge asymmetries exacerbate these issues. Public sector procurers need consistent guidance and terminology to help them buy AI well.

Gaps in AI governance undermine the sector's ability to ensure tools are safe, effective and fair.

Public sector leaders need more assurance and governance mechanisms so they can procure, deploy and use AI tools with confidence, particularly given the increasing use of tools built on foundation models. This includes independent testing, post-market monitoring and redress mechanisms for affected individuals, as found in other regulated sectors such as pharmaceuticals. Improving governance will create a more competitive market for AI tools, allowing the public sector to access more reliable and cost-effective tools for its services.

Think beyond the technology

The adoption of AI will have wider societal consequences.

The public sector will inevitably have to deal with the wider societal consequences of AI adoption, regardless of direct use within public services. This includes potential effects on employment, trust in institutions and information, social inequalities and the environment. Governments should set up a future-facing

programme of economic and social policy development to anticipate changes and to support individuals and communities.

See AI not as an opportunity to automate the public sector, but to reimagine it.

We welcome a long-term vision for public service transformation where AI follows rather than leads, one that is grounded in public and professional legitimacy. Public sector leaders should see the rollout of AI as an opportunity to reimagine the state, rather than focusing solely on immediate efficiency gains or automating the status quo. AI should be viewed as a catalyst for fundamental service redesign, placing the citizen at the centre of public service delivery.

Introduction

A fresh wave of enthusiasm for the use of AI has swept across the public sector, beginning with the latest generation of chatbots like ChatGPT and now with the promise of general-purpose AI and increasingly agentic systems. This enthusiasm is often framed around expectations of increased efficiency, lower costs and better outcomes in areas like healthcare and education.

The implementation of this latest wave of AI also raises critical questions about transparency, accountability, equity and public trust as well as more practical questions about implementation, performance and value for money. AI is not a new phenomenon in the public sector: narrow forms like predictive analytics or image classification tools have been experimented with and used in public services for many years. General-purpose AI has also been adopted as part of formal tools, as well as through informal use.

We can learn from the existing examples of AI and data-driven technologies in the public sector to ensure these new tools genuinely serve the public interest. At the Ada Lovelace Institute (Ada), we have spent the past six years examining the use of data and AI across the public sector, in healthcare, education, social care and in cross-cutting work on transparency and foundation models.

This briefing sets out the 'lessons for success' from our research, including more than 30 reports and research publications. While not comprehensive, these lessons have consistently reoccurred in studies looking at different use cases of AI in the public sector. We believe these lessons can support governments aiming to accelerate the use of AI in policymaking and public services, to ensure AI works for the sector and for the public they seek to serve.

AI in the public sector: lessons for success

Contextualise AI

To enable the wide deployment of AI across the public sector, governments should adopt clear terminology, address data challenges and recognise that these technologies operate within complex social systems rather than in isolation.

1. Lack of clear terminology about 'AI' is inhibiting learning and effective use

There is an issue with definitions: no one can completely agree on the fundamental parameters for what AI actually *is*, or how much we should believe it is capable of. AI itself is also a moving target; the foundation models¹ that have spurred the last couple of years of excitement around AI are not the same tools as narrow forms like predictive analytics.

In our research on public sector procurement of AI, we found that definitions of AI varied within the government's own collection of guidance materials.² Even when focused on a certain type of AI – in this case foundation models – our evidence review found that 'terminology in this area is inconsistent and still evolving'.³

Without clear definitions, the public sector is left without actionable or cohesive frameworks for decision making about where to prioritise and pilot AI.

This leaves a gulf between the stated intentions of the government to 'radically improve our public services'⁴ with AI and individual decisions about specific tools. It also hinders the ability of the public sector to evaluate AI or discuss its capabilities. For example, stakeholders involved in the procurement of AI told us that this lack of clarity often forces procurers without expertise in AI to fill in the gaps on their own to understand which type of technology offers the best solution for the issues they are trying to address.⁵

There are many different ways to classify AI in public services, including by the purpose of the technology, the type of technology, the context of its use and its impact. The challenge facing policymakers is to see how these factors intersect, which is crucial for making the right decisions about AI in the public sector.

Some work has already been done on classifying AI according to broad public sector use cases, for example by the National Audit Office (NAO), the Incubator for Artificial

Intelligence (i.AI), the thinktank Reform and the European Commission. These are useful foundations but we see an opportunity to build upon them.

i.AI recently published its taxonomy for AI in government, which does link different types of challenges in the public sector with corresponding technology solutions.⁶ It groups five user challenges in government (public-facing services, fraud and error, matching and triage, casework management, and data infrastructure), with the type of technology proposed to meet them (from generative AI to databases and APIs).

For instance, in grouping 'matching and triage' with 'optimisation,' the i.AI taxonomy suggests that 'although no project proposal is the same, there are often common solutions to the problems they seek to tackle'. It says, 'for example, the same underlying optimisation techniques could be applied to prioritise what to connect to the national grid, solve timetabling problems such as scheduling appointments, or match people to services such as accommodation'.

However, there is a lot more to unpack with these disparate use cases. We know that prioritising what connects to the national grid comes with a completely different set of considerations than matching potentially vulnerable people to accommodation services. Beyond ethical considerations, some areas have datasets with a certain level of ground truth or at least completeness, whereas others have partial data generated from subjective and often cash-strapped processes.

Similarly, the thinktank Reform recently published a paper on scaling AI in public services, which outlined some categories of use cases that it says are well evidenced for further deployment. One use case category is 'assessment streamlining' to increase the speed at which decisions might be made, and the paper includes the examples of processing asylum claims and assessing the outputs of diagnostic tests like X-rays. Again, these are different contexts with different levels of underlying data quality and completeness, and where levels of subjectivity differ quite drastically.⁷

Categories like 'assessment streamlining' may be useful when thinking at a broader level about government strategy around AI. However, it is important to note that between those two use cases there may be very different organisational structures and cultures, differing levels of discretion on the part of caseworkers and radiologists, and potentially differing levels of opportunity or agency on the part of the affected people to challenge decisions. None of these are necessarily problems with the technology itself, although different contexts will affect the quality of data that feeds into a system and whether the technology is automating an existing process or generating new outputs.

We therefore see a missing 'middle layer' between government rhetoric which focuses on the generic function of AI, and a purely case-by-case analysis to describe AI in the public sector. A more nuanced taxonomy that accounts for significant variations in the purpose of AI tools, the type of technology and the context of its use will enable more effective decision-making across the public sector, as well as the sharing of actionable lessons.

Further reading

Spending wisely: Redesigning the landscape for the procurement of AI in local government

Buying AI: Is the public sector equipped to procure technology in the public interest?

Foundation models in the public sector: AI foundation models are integrated into commonly used applications and are used informally in the public sector

2. AI is only as good as the data underpinning it

AI systems require solid data foundations, and not all datasets are created equally. Data comes from the world we live in: datasets reflect historical and social biases, and decisions about collecting and coding are not neutral. AI interventions should therefore only be prioritised for areas where data is as comprehensive, clear and complete as possible.

Acknowledging limitations in data collection might include considering, for example:

the underrepresentation of digitally excluded and marginalised groups

inconsistencies in how data is recorded (e.g. around ethnicity)

a lack of clarity on what the data represents (for example, data on homelessness presentations being portrayed as representing overall homelessness figures)

the historical biases in existing datasets which may be inaccurate (for example, poor data for ethnic minorities in healthcare), or 'accurate' but bound up with systemic inequalities that should not be 'baked in' to future decision making (for example, stop and search data in policing).

In October 2024, the Secretary of State for Health and Social Care Wes Streeting said that the NHS is 'the best-placed healthcare system in the world to take advantage of rapid advances in data, genomics, and predictive and preventive medicine', in part due to its 'large and diverse set of data'.⁸

While NHS data is rich and detailed in many ways, it still often lacks local and historical context, which can lead to misinterpretation and poorly targeted interventions. This results in some demographics being overlooked or incorrectly categorised. It means that what health data says about someone does not fully reflect their lived reality and can be devoid of important nuance. In fact, there are always limitations to what data can tell us, and trusting too much in data as a proxy for reality could lead to poor quality insights.⁹ In our report 'A knotted pipeline', a product owner shared:

'Dashboards can be seductive, give you nice clean view of the world and you forget it's just a representation of the world that may be skewed. Part of our job is to tell you why it's skewed, and that's a constant challenge.'¹⁰

As one example, we have found that clinical information systems used in the NHS are not consistently designed to accommodate gender diversity, and that they conflate gender with physiological sex characteristics, undermining the provision of relevant health testing like cervical scans.¹¹

People who are digitally excluded may also be excluded from datasets. Members of the public who lack adequate access to a computer, broadband, an up-to-date smartphone or data allowance may be less able to benefit from data-driven technologies if they are less represented in the data.¹² Any technology built on top of data systems will also be affected by inaccurate or absent data.¹³ Even cutting-edge AI technologies like AI-powered genomic health prediction are affected by historical data. Evidence shows that, for example, historical phenotype data reflects the prejudices of those responsible for the labelling. Clinical notes recorded by psychiatrists reflect the historical tendency to make different treatment recommendations for minority ethnic groups and female patients. This contributes to the current lack of accuracy in individual polygenic risk scoring.¹⁴

Some data issues can be improved on by more complete data collection, like recording ethnicity data in health records consistently. However, while it is important for people to be represented in data insofar as it ensures equitable provision and experience of services, there is a balance to be struck. Over-datafication can raise ethical concerns and erode public trust and willingness to engage with services or technologies.¹⁵

While data-driven technologies are increasingly powerful tools for understanding the world we live in, these tools only present one view, based on what can and is captured in datasets, and who collects the data and how. Therefore, we need to complement these data-driven technologies with other forms of understanding, such as relational, qualitative and deliberative engagement.

Further reading

A knotted pipeline: Data-driven systems and inequalities in health and social care

The data divide: Public attitudes to tackling social and health inequalities in the COVID-19 pandemic and beyond

'The computer won't do that': Exploring the impact of clinical information systems in primary care on transgender and non-binary adults

3. AI systems are not deployed in a vacuum and context is important

Central to Ada's work on AI and data-driven technologies in public services is the importance of taking a sociotechnical approach.¹⁶ These technologies must be seen as shaping and being shaped by the complex social systems in which they exist. We have seen that success and public acceptability is contingent upon how AI is integrated into public services.

This is clearly evident when it comes to digital transformation in the NHS for example. In our research on access to digital healthcare, people routinely cited wider concerns about the NHS, including underfunding, workforce strain and the perceived incursion of private organisations in the health service as reasons for scepticism towards new digital health services or data sharing projects. These contextual factors mean that new technologies may be perceived as cost-cutting or efficiency initiatives rather than efforts to improve patient care.¹⁷

Our work on COVID-19 technologies demonstrated the need to be conscious about the values being built into any new technology. The introduction of new technologies may shift the sociopolitical fabric of society, and cannot be decoupled or isolated from the society it shapes. For example, members of the public were clear that any measures that might undermine solidarity – like individualised risk scoring or immunity certification – should be taken with extreme caution.¹⁸

Members of our Citizens' Biometrics Council implicitly acknowledged how technologies are not used in isolation from a social and organisational structure but are intertwined with it. One participant said: 'For me, I think it's about trust. Stop and search has been abused over the years and to add on top of that – to have technology that supports stop and search – it's not going to make young Black males trust the police any more than they already do.'¹⁹

Our three-year legal, public-facing and policy work on biometrics concluded with the need for a dual assessment of biometric technologies. We recommended that they were assessed to meet scientifically based and clearly established standards of

accuracy, reliability and validity, and that they should be assessed for proportionality *in their proposed contexts*. We have since made the case that two overarching questions should be applied to other applications of AI in the public sector:

1. Does it work? And does it work well enough for everyone?
2. Is it proportionate to use in this context?

In addition to appreciating how social context may influence AI technologies, we have also found evidence of AI technologies influencing existing social and professional norms.

Our research examining the use of predictive analytics in a London council showed that the development and deployment of the system had an impact not only on IT systems in the council but also on the day-to-day work and practice of frontline workers.²⁰ Interviewees highlighted how the system could affect the relationships between residents and council staff that are crucial for effective and trusted social work. When social workers did not trust the tools or had concerns about their legitimacy, they did not use them.

Sociotechnical factors should therefore be considered as part of any assessment of AI technologies. As well as questions about the AI system itself (which might include the methodology being used, the type of data being used, whether the tool is bespoke or off the shelf, or the role of the private sector), evaluations will need to consider the evidence and effectiveness of the AI system as an intervention. For example, the type of problem AI is being applied to, the potential public-facing impact, the level of evidence behind the proposed solution, and the relative risk of the AI system not working compared to the current system.

The algorithmic impact assessment template we created in partnership with the former NHS AI Lab encouraged reflection and reflexivity about the impacts of an algorithmic system on the part of those developing it. Some of the considerations included whether there are groups that might interact differently with a product, and which people or groups might be harmed when the system fails. Other questions included which stakeholders will use a system, how they would optimally interact or work together for the system to succeed, how information would be shared (and with who), and what social, technical and workflow dependencies may need to exist.²¹

Taking a sociotechnical approach also means acknowledging the position, power and possible biases of the individuals or teams who are making decisions about data collection, data quality, data curation and standardisation methods, as well as the design and deployment of interventions.²²

This requires a view of the whole system or service an AI tool might be deployed in. Successful adoption requires careful consideration of the financial, social, infrastructural and other pressures that the existing system or service faces, as well as identifying how AI might alleviate, worsen or be shaped by those same pressures.

Further reading

Predicting: The future of health?: Assessing the potential, risks and appropriate role of AI-powered genomic health prediction in the UK health system

Access denied?: Socioeconomic inequalities in digital health services

Confidence in a crisis?: Findings of a public online deliberation project on attitudes to the use of COVID-19 related technologies for transitioning out of lockdown

The Citizens' Biometrics Council: Report with recommendations and findings of a public deliberation on biometric technologies, policy and governance

Learn what works

Making informed decisions about AI requires transparency, rigorous evaluation and improved procurement processes. Governments should prioritise the establishment of structured approaches to assess AI effectiveness in applied public sector contexts and ensure these technologies deliver genuine value while maintaining accountability.

4. The public sector does not have a comprehensive view of where AI is being deployed in government and public services

There remains a persistent, systemic deficit in understanding where and how AI and data-driven systems are used in the public sector.²³ This hampers knowledge sharing across the public sector about what is working and what is not. It also reduces the ability for journalists, academics, civil society and ultimately the public to scrutinise and hold the government democratically accountable for its use of AI.

Across our work, from examining AI procurement to the deployment of foundation models in the public sector, we have seen and heard that there is a real problem with visibility across government about where AI projects are being piloted and deployed.^{24,25} Within local government, for example, this has left some teams in charge of procuring AI systems to rely on the claims of AI suppliers rather than on their own assessments.²⁶

Part of the problem is there are several mechanisms for transparency (including impact assessments, procurement documents, open data, FOIs and standardised disclosure of data), but they do not necessarily enable adequate scrutiny of automated decision making or AI systems in the public sector.²⁷ Many of these mechanisms are not mandatory, and those that are suffer from a lack of enforcement. There is also not enough support for public authorities in ensuring best practice.²⁸

Our research on transparency has highlighted several different 'targets' for transparency.²⁹ For example, focusing solely on the procurement of AI highlights different categories of transparency: transparency in relation to process of procurement itself, to ensure that suppliers have a clear understanding of the competition and what is expected; transparency in relation to a specific tool; and transparency in relation to the tool's impact on the public.³⁰

Following the welcome announcement that the Algorithmic Transparency Recording Standard (ATRS) would be made mandatory for central government departments, there was a disappointing lag in the publication of new records, with only nine public records by summer 2024. Since December 2024, we have seen multiple new records published each month, and in early March 2025, there were 56 published records. We hope this marks the start of the ATRS becoming a routine part of government operations.³¹

We see this as a basic and foundational requirement for the successful and trustworthy implementation of AI in the public sector. Peter Kyle, the Secretary of State for Science, Innovation and Technology, has admitted the public sector 'hasn't taken seriously enough the need to be transparent in the way that the government uses algorithms'.³²

Even when the ATRS is functioning well, it will only be a partial account of AI being used in the public sector. It covers only public-facing AI which significantly influences a decision. It allows for exemptions on sensitivity and security grounds, and does not currently include examples of informal usage or cross-departmental implementations.³³

From the perspective of democratic accountability, the ATRS is now being used more extensively and it would be valuable to evaluate how well it is working, and whether it covers what the public and their representatives expect to see. In terms of the ATRS as a tool for the public sector, it is significantly limited and cannot provide the full picture of where and how AI tools are being used, and to what effect.

From a wide-ranging evidence review of quantitative and qualitative research on public attitudes to AI, it is clear that transparency is expected by the public.³⁴ In fact,

the public prioritise having an explanation of how AI-driven decisions are made, above a more accurate decision.³⁵

Further reading

Transparency mechanisms for UK public-sector algorithmic decision-making systems: Existing UK mechanisms for transparency and their relation to the implementation of algorithmic decision-making systems

What forms of mandatory reporting can help achieve public-sector algorithmic accountability?: A look at transparency mechanisms that should be in place to enable us to scrutinise and challenge algorithmic decision-making systems

Meaningful transparency and (in)visible algorithms: Can transparency bring accountability to public-sector algorithmic decision-making (ADM) systems?

A window into the black box: New entries on the UK government's AI register

AI in the public sector: From black boxes to meaningful transparency

Spending wisely: Redesigning the landscape for the procurement of AI in local government

Buying AI: Is the public sector equipped to procure technology in the public interest?

Foundation models in the public sector: AI foundation models are integrated into commonly used applications and are used informally in the public sector

5. There is not enough evidence on the effectiveness of AI tools

Despite optimism for the potential of AI to deliver benefits to the public sector, from productivity to performance, there is not yet systematic or comprehensive evaluation of AI tools in the public sector.

Across our work, we have seen the importance of evaluating AI systems to understand whether a given AI application is effective at delivering an intended outcome, and how it performs relative to other interventions, e.g. investing in staff training, better digital infrastructure or more staff. Evaluation which considers performance and value for money, and compares these to alternative approaches, should allow governments to target resources at scaling-up tools that are net-positive and discontinuing harmful or ineffective interventions. Critically, these evaluations need to understand the broader service and social context these systems are being deployed in, and to take an iterative approach to evaluation as the models and systems continue to develop.

Our research on the early adoption of data analytics for local authority service delivery recommended that local authorities should include the development of clear, actionable success criteria and plans for how they will be evaluated when they procure and implement analytics systems, including in pilot deployments.³⁶ In addition, local authorities should develop success criteria and evaluation methods for the system as a whole, with the participation of those who will be most affected by the use of the system. These recommendations are equally applicable across the public sector.

Our work on the early deployment of foundation models in the public sector found a need for iterative monitoring and evaluation, beyond the initial development and procurement of foundation model applications, given how frequently models are updated by developers. Local and central government representatives in our roundtables felt that this ongoing monitoring and evaluation by public services, of both public and private applications, was needed to ensure foundation model systems operate as intended and to discover when AI systems are failing to deliver results.³⁷

Particularly when it comes to using foundation models like GPT-4 or Gemini, our work on evaluations has shown the importance of contextual evaluations. While model evaluations can provide useful information about the overall capabilities and general risks, the existing model evaluations have a number of limitations. These evaluations often lack external validity (where results fail to generalise to real-world conditions) and the results can vary significantly with small changes to the model or context. These evaluations also often overlook the sociotechnical context: how users interact with the model, the design of its interface and the broader societal impacts, especially regarding embedding systematic biases.³⁸

Deployers of AI in the public sector should take both a systematic and iterative lens on evaluating deployments of AI, understanding both their immediate impact and how users, workers and wider society may change and adapt to their use over time. This includes how staff adapt working practices, as well as reflecting the ongoing and iterative nature of AI product development, and involving those affected by the systems, including, where appropriate, both service users and staff. Successful evaluations should include establishing clear success criteria, procedures for continuous monitoring and contextual evaluations that account for impacts on society, user interaction and interface design.

In December 2024, the UK government published 'Guidance on the Impact Evaluation of AI Interventions'.³⁹ This guidance builds on the government's existing evaluation guidance, the Magenta Book, and has a welcome focus 'on the systematic assessment of the outcomes of an intervention with the aim of establishing whether, to what extent,

how and why an intervention has resulted in its intended impacts'. This guidance is a good start but, as the guidance itself notes, robust evaluation will also require drawing on the knowledge of evaluation experts in a given domain like social care or education.

Further reading

Foundation models in the public sector: AI foundation models are integrated into commonly used applications and are used informally in the public sector

Critical analytics?: Learning from the early adoption of data analytics for local authority service delivery

Under the radar?: Examining the evaluation of foundation models

Deliver on public expectations and public sector values

For the public sector to successfully benefit from AI adoption, it must earn and maintain public trust. This requires developing systems that are not only technically sound but also ethically designed, properly governed and deployed with consideration for their broader societal impacts and alignment with long-term public service goals.

6. Successful use of AI requires public licence

Moving out of step with public attitudes can undermine the ability for the public sector to effectively deploy AI. There have been high-profile examples of people withdrawing their data because of concerns about how it will be used. Public backlash to the General Practice Data for Planning and Research (GPDPR) programme has led to 3.6 million people opting out of data sharing.⁴⁰

But we have also seen examples where concerns about public legitimacy can hamper use of existing tools. In a study with one London local authority, some social workers using predictive tools to identify at-risk residents shared their reluctance to use them because of their perceived lack of legitimacy:

'One of the areas that we were a little bit reluctant with was, if OneView identifies the family, and we pick up the phone, what do we actually say to them? [...] Nothing has actually happened to this family. There's been no incident.'⁴¹

In another example, when teenagers protested the 'A-level algorithm' used during the pandemic to assign results, public servants told us that a number of other AI projects were shelved as a consequence of the backlash.⁴²

Getting AI 'wrong' in the eyes of the public may be a big risk to harnessing its potential. Therefore public attitudes around data sharing and use must be taken into account for technologies to gain public legitimacy.

In new initiatives like the National Data Library, the government needs to learn from unpopular data sharing initiatives like care.data and the GDPR programme.⁴³ NHS England and the Department of Health and Social Care are working on a large-scale public consultation to inform future data policy.⁴⁴ It is important that this also draws from the rich bank of evidence that already exists on how people feel about data sharing.

In our work with people who experience health inequalities, it became clear that people do not understand how their data is being used or protected by health and care professionals. Participants said they did not know who exactly has access to their health data, and whether partnerships between the NHS and private companies affect the security of that data.⁴⁵ In our 'Access denied?' report, one interviewee with lived experience of poverty said:

'I think we all have to be just a little careful and make sure as services change that, you know, making sure that the public are on board about how things are working, and making sure that [data controllers are] not misusing data, let's say, or sharing without agreements and so on.'⁴⁶

Our Citizens' Biometrics Council emphasised that improving data security would be crucial before any further rollout of biometric technologies and that there should be clear parameters in place for biometric data that are communicated to the public.⁴⁷

Aims and outputs of systems must also be clearly understood by users and seen as having a legitimate purpose. For example, when we studied the deployment of a predictive analytics system in one local council, we found that one of its functions (COVID-19 case management) was more widely accepted and used than the others (case summaries for social workers and predictive risk alerts in children's social care).⁴⁸

For the latter, staff felt that they were unable to identify how the system created its outputs from data held about individuals. A common view among both social workers and management staff was that there was not an adequately transparent and clear explanation for the predictions. This was especially the case for knowing what factors contributed to case summaries and predictive alerts, and the rationale for recommendations and outputs. Some frontline staff were also unconvinced that the analytics were as objective, neutral or accurate as had been described to them.

On the other hand, the COVID-19 case management function was seen to have a clear purpose, transparent and visible risk factors, and higher explainability. Staff were more confident in its benefits and how to use the outputs in their work.⁴⁹

Public legitimacy can be bolstered by the technology having a clear and narrow purpose, and, conversely, it can be compromised by 'scope creep'. This is where technologies introduced to address a certain problem leak into other areas of application, without scrutiny of their appropriateness or broader impact.

Our citizens' juries during COVID-19 said that use of data must balance public health needs with risks to individuals and society, and that the pandemic response measures must not extend into post-pandemic data futures. For example, by repurposing vaccine passports to create digital identities, or by continuing the use of a particular risk scoring algorithm beyond the time-limited purpose of vaccine distribution.⁵⁰

In our work on public attitudes to biometric technologies, we found that people fear the normalisation of surveillance and only support facial recognition technology when there is a demonstrable public benefit *and* appropriate safeguards in place.⁵¹ We believe this warrants greater investment in testing and articulating the potential public benefits not just of biometric technologies but of AI technologies in the public sector more broadly.

Using vaccine passports as a case study for the introduction of novel technologies in the public domain, we highlighted the importance of public legitimacy, especially around the sensitivities of using personal health data in combination with biometric data. We recommended that government should undertake public deliberation, especially with groups who may be particularly affected by the technology.⁵² This could again apply to future novel technologies deployed in the public sphere.

Further reading

[The Citizens' Biometrics Council: Report with recommendations and findings of a public deliberation on biometric technologies, policy and governance](#)

[Access denied?: Socioeconomic inequalities in digital health services](#)

[Critical analytics?: Learning from the early adoption of data analytics for local authority service delivery](#)

[The rule of trust: Findings from citizens' juries on the good governance of data in pandemics](#)

[Beyond face value: public attitudes to facial recognition technology: First survey of public opinion on the use of facial recognition technology reveals the majority of](#)

people in the UK want restrictions on its use

Checkpoints for vaccine passports: Requirements that governments and developers will need to deliver in order for any vaccine passport system to deliver societal benefit

7. Public procurement of AI is not fit for purpose

The use of AI in the public sector is likely to increase partnerships with the private sector and procurement is an essential first step in the process. Getting the procurement of AI right is vital for ensuring data and AI work effectively and in the public interest – and for mitigating harms.

The enthusiasm for the potential of AI to improve public services contrasts with the reality of limited human and financial resources in the public sector, and particularly within local government. This rapidly evolving technology presents potential opportunities for the public sector to do more with those limited resources, but may require upfront resource and, if adopted uncritically, may damage public trust and cause harm.

Recent prominent examples of data and AI systems not working as intended – such as the Post Office's Horizon software⁵³ and the Home Office's visa application streamlining algorithm⁵⁴ – have raised important questions about whether the public sector is equipped to procure and oversee complex technologies across different stages of their lifecycle.

The public sector is subject to higher levels of scrutiny and accountability than the private sector around issues of legitimacy, trust, fairness and equality, for example by the Public Sector Equality Duty.⁵⁵ Higher levels of transparency and explainability are required regarding important decisions about public services such as welfare, healthcare and education.

Pursuing a high standard of transparency around procurement of AI can contribute to public trust and fairness. If both procurers and companies know they will be scrutinised on decisions, there is more incentive to 'build in' fairness and anticipate or flag damaging outcomes.⁵⁶

An important part of scrutinising procurement process is protecting against vendor lock-in. Market failures have led to an excess of power that rests in the hands of a few large suppliers, pricing out smaller and medium-sized vendors –limiting the choice available to procurement teams and creating vendor lock-in.⁵⁷ This both stems from and further entrenches power imbalances between public sector procurers and private sector suppliers. It leaves the public sector beholden to potentially unverifiable

claims from larger tech companies, and less able to demand contract conditions that would help ensure societal impact is captured.⁵⁸

Multiple organisations within central and local government are working on the knotty issue of procuring AI in the public sector. The Department for Science, Innovation and Technology has announced 'AI Management Essentials' to help public sector buyers make better and more informed decisions on the use of AI, with a long-term aim to embed this within government procurement policy.⁵⁹ The IEEE has developed standards for the procurement of AI which aim to address sociotechnical questions and public interest.⁶⁰ We are also aware of the emerging collaborative work among regulators and local government to embed the Public Sector Equality Duty into processes for procuring AI.

Our in-depth research into the procurement of AI in local government has led us to the conclusion that there is real potential in taking a more joined-up approach with dedicated resource.⁶¹ There should be ways of working that bring local government into discussions and collaboration from the outset. This means empowering local stakeholders to implement tools that help set criteria for success while providing the necessary support and tools to reshape markets and support skill development.

Further reading

Spending wisely: Redesigning the landscape for the procurement of AI in local government

Buying AI: Is the public sector equipped to procure technology in the public interest?

8. Gaps in AI governance undermine the sector's ability to ensure tools are safe, effective and fair

Our research suggests that to achieve the best results for public sector deployment of AI, the rollout should be accompanied by strengthened AI governance to ensure that AI tools being sold into the public sector are safe and effective.

The lack of monitoring and evaluation of AI within the public sector is compounded by the opaque market environment that exists outside of it. Significant information asymmetries exist between the large companies selling AI systems and those procuring, using or building on top of them, including in the public sector. Too often, important information about AI systems – how they work, and their safety and effectiveness when deployed in particular contexts – either does not exist or is kept private.

The ability of public sector decision-makers – from procurers in schools and NHS foundation trusts to senior leaders in the civil service – to make effective decisions about AI is limited by this one-sided market. They are left overly reliant on sales pitches from private providers, supplemented by the limited investigative research by the media and civil society.⁶²

All of this means that while the regulation of AI companies in the private sector is often seen as a separate policy agenda from rolling out AI across public services, the two are in fact deeply linked. Bringing in ‘binding regulation’ on foundation model developers to ensure these systems are tested before coming to market, as committed to in the Labour Party’s 2024 general election manifesto, will be an urgent first step towards providing the assurance public sector leaders need to procure, deploy and use these models with confidence.

In the long term, the increasing use of AI in public services will necessitate robust governance mechanisms for AI within and beyond the public sector, comparable to other regulated sectors like pharmaceuticals that predominantly sell to public entities.⁶³ Ada’s research suggests that robust governance requires independent institutions with sufficient resources and powers to effectively oversee AI development and deployment, including pre-market assessments, post-market monitoring and redress mechanisms for affected individuals.

This will create a more competitive and higher quality market for AI tools. In turn, the public sector will then be able to buy and procure more reliable and cost-effective tools for its services.

Further reading

Countermeasures: The need for new legislation to govern biometric technologies in the UK

Algorithmic impact assessment: a case study in healthcare: This report sets out the first-known detailed proposal for the use of an algorithmic impact assessment for data access in a healthcare context

Algorithmic accountability for the public sector: Learning from the first wave of policy implementation

New rules?: Lessons for AI regulation from the governance of other high-tech sectors

Conclusion: think beyond the technology

While we have seen enthusiasm, hype and hope regarding AI's potential and power in the public sector, we have not yet seen enough truly radical thinking from governments about the structures that will be required if AI is to have the profound impact across society and public services that some anticipate.

Working on specific uses of AI in the public sector over the last six years has brought to our attention two 'big picture' considerations for any government or public sector organisation that is looking to fully reap the potential benefits of AI.

The adoption of AI will have wider societal consequences. The public sector will inevitably have to deal with the intended and unintended consequences of AI tools, regardless of their direct use within public services. This requires the state to proactively anticipate and address the impacts of AI beyond their immediate application. Therefore it is essential to prepare for and mitigate these impacts. This includes considering potential effects on employment, trust in institutions and information, social inequalities and the environment. Governments should set up a future-facing programme of economic and social policy development to anticipate changes and to support individuals and communities.

See AI not as an opportunity to automate the public sector, but to reimagine it.

We welcome work to establish a long-term vision for public service transformation where AI follows rather than leads, one that is grounded in public and professional legitimacy. Rather than 'ask for faster horses', public sector leaders should see the rollout of AI as an opportunity to reimagine the state with the citizen at its heart. AI should be viewed as a potential catalyst for fundamental service redesign, placing the citizen at the centre of public service delivery rather than focusing solely on immediate efficiency gains or automating the status quo. Through meaningful engagement with the public and relevant professions, governments can develop a shared understanding between citizens, staff and wider society of where AI has the potential to help reimagine more relational, effective and legitimate public services.

Ada's new strategy will place this agenda at its heart, by asking what a positive vision is for public services. It will consider how and where AI can be deployed in ways that benefit public benefit, while balancing the needs of services, the public and the state.

How did we approach this policy briefing?

This policy briefing is based on thematic analysis of more than 30 reports and research publications by the Ada Lovelace Institute since 2019 (see the full list below), which covered specific areas of the public sector and cross-cutting public sector topics:

Health

Education

Local government

Responses to the COVID-19 pandemic

Transparency

Accountability

Procurement

Use of foundation models

Biometric technologies

These reports span independent legal reviews, futures thinking, deliberative exercises and surveys of public opinions, landscape reviews, technical analysis, ethnographic case studies, and syntheses of expert views.

Reports

Health

- [‘The computer won’t do that’: Exploring the impact of clinical information systems in primary care on transgender and non-binary adults](#)
- [Policy briefing: Access denied?: Inequalities in data-driven health systems and digital health services](#)
- [Report: Access denied?: Socioeconomic inequalities in digital health services](#)
- [I.:](#) Early findings and emerging questions on the use of AI in genomics
- [Predicting: The future of health?: Assessing the potential, risks and appropriate role of AI-powered genomic health prediction in the UK health system](#)

- A knotted pipeline: Data-driven systems and inequalities in health and social care
- Algorithmic impact assessment: a case study in healthcare: This report sets out the first-known detailed proposal for the use of an algorithmic impact assessment for data access in a healthcare context
- The data will see you now: Exploring the datafication of health: what it is, how it occurs, and its impacts on individual and social wellbeing

Education

- A learning curve?: A landscape review of AI and education in the UK
- Can algorithms ever make the grade?: The failure of the A-level algorithm highlights the need for a more transparent, accountable and inclusive process in the deployment of algorithms

Local government

- Critical analytics?: Learning from the early adoption of data analytics for local authority service delivery

Responses to the COVID-19 pandemic

- The rule of trust: Findings from citizens' juries on the good governance of data in pandemics
- Checkpoints for vaccine passports: Requirements that governments and developers will need to deliver in order for any vaccine passport system to deliver societal benefit
- The data divide: Public attitudes to tackling social and health inequalities in the COVID-19 pandemic and beyond
- What place should COVID-19 vaccine passports have in society?: Findings from a rapid expert deliberation to consider the risks and benefits of the potential rollout of digital vaccine passports
- No green lights, no red lines: Lessons to assist Government and policymakers navigating difficult dilemmas when deploying data-driven technologies to manage the pandemic
- Confidence in a crisis?: Findings of a public online deliberation project on attitudes to the use of COVID-19 related technologies for transitioning out of lockdown

- Exit through the App Store?: A rapid evidence review of the technical considerations and societal implications of using technology to transition from the first COVID-19 lockdown
- Provisos for a contact tracing app: The route to trustworthy digital contact tracing
- Answers in the App Store?: Lessons from COVID-19 technologies
- Lessons from the App Store: Insights and learnings from COVID-19 technologies

Transparency

- Transparency mechanisms for UK public-sector algorithmic decision-making systems: Existing UK mechanisms for transparency and their relation to the implementation of algorithmic decision-making systems
- What forms of mandatory reporting can help achieve public-sector algorithmic accountability?: A look at transparency mechanisms that should be in place to enable us to scrutinise and challenge algorithmic decision-making systems
- Meaningful transparency and (in)visible algorithms: Can transparency bring accountability to public-sector algorithmic decision-making (ADM) systems?

Accountability

- Algorithmic accountability for the public sector: Learning from the first wave of policy implementation

Procurement

- Spending wisely: Redesigning the landscape for the procurement of AI in local government
- Buying AI: Is the public sector equipped to procure technology in the public interest?

Use of foundation models

- Foundation models in the public sector: AI foundation models are integrated into commonly used applications and are used informally in the public sector

Biometric technologies

- The Ryder Review: Independent legal review of the governance of biometric data in England and Wales

- Countermeasures: The need for new legislation to govern biometric technologies in the UK
- The Citizens' Biometrics Council: Report with recommendations and findings of a public deliberation on biometric technologies, policy and governance
- Beyond face value: public attitudes to facial recognition technology: First survey of public opinion on the use of facial recognition technology reveals the majority of people in the UK want restrictions on its use

Our analysis of these reports and research publications identified recurring themes and lessons for public sector use of AI, which are synthesised in this policy briefing.

Our analysis was complemented by several interviews with staff across central and local government to sense check and contextualise these lessons. However, this synthesis briefing only draws on research conducted by the Ada Lovelace Institute. We do not suggest this briefing offers a complete picture of AI in public services, but we hope it offers useful insights to those considering the use of AI in public sector contexts. We will continue to undertake both primary and secondary research to expand the evidence base on AI in the public sector.

Acknowledgements

This briefing was co-authored by Imogen Parker, Anna Studman and Elliot Jones.

Footnotes

1. Foundation models are AI models designed to produce a wide and general variety of outputs. They are capable of a range of possible tasks and applications, such as text, image or audio generation. They can be standalone systems or can be used as a 'base' for many other applications. For more information, see our explainer on foundation models:
<https://www.adalovelaceinstitute.org/resource/foundation-models-explainer/>
2. Anna Studman and others, 'Buying AI' (Ada Lovelace Institute, 2024) <<https://www.adalovelaceinstitute.org/report/buying-ai-procurement/>> accessed 4 March 2025.
3. Elliot Jones and others, 'Foundation Models in the Public Sector' (Ada Lovelace Institute, 2023) <<https://www.adalovelaceinstitute.org/evidence-review/foundation-models-public-sector/>> accessed 4 March 2025.
4. 'PM Speech on AI Opportunities Action Plan: 13 January 2025' (GOV.UK, 13 January 2025) <<https://www.gov.uk/government/speeches/pm-speech-on-ai-opportunities-action-plan-13-january-2025>> accessed 7 March 2025.
5. Mavis Machirori and Anna Studman, 'Spending Wisely' (Ada Lovelace Institute, 2024) <<https://www.adalovelaceinstitute.org/report/spending-wisely-procurement/>> accessed 4 March 2025.
6. Michael Padfield, 'The i.AI Taxonomy' (Incubator for Artificial Intelligence – GOV.UK, 29 July 2024) <<https://ai.gov.uk/blogs/the-i-ai-taxonomy>> accessed 4 March 2025.
7. Joe Hill and Sean Eke, 'Getting the Machine Learning: Scaling AI in Public Services' (Reform, 2024) <<https://reform.uk/publications/getting-the-machine-learning-scaling-ai-in-public-services/>> accessed 4 March 2025.
8. Wes Streeting, 'I Love the NHS: It Saved My Life, but the Operation to Rescue It Must Be Led by the People and Its Staff' *The Guardian* (21 October 2024) <<https://www.theguardian.com/commentisfree/2024/oct/21/nhs-saved-my-life-rescue-health-service-wes-streeting>> accessed 4 March 2025.
9. *ibid.*

10. Mavis Machirori and others, 'A Knotted Pipeline' (Ada Lovelace Institute, 2022) <<https://www.adalovelaceinstitute.org/report/knotted-pipeline-health-data-inequalities/>> accessed 4 March 2025.
11. Kavya Kartik, 'The Computer Won't Do That' (Ada Lovelace Institute, 2024) <<https://www.adalovelaceinstitute.org/report/the-computer-wont-do-that/>> accessed 4 March 2025.
12. Reema Patel, Elliot Jones and Aidan Peppin, 'The Data Divide' (Ada Lovelace Institute, 2021) <<https://www.adalovelaceinstitute.org/report/the-data-divide/>> accessed 11 August 2023.
13. Kartik (n 11).
14. Harry Farmer, Maili Raven-Adams and Andrew Strait, 'Predicting: The Future of Health?' (Ada Lovelace Institute, 2024) <<https://www.adalovelaceinstitute.org/report/predicting-the-future-of-health/>> accessed 4 March 2025.
15. Valentina Pavel, 'Rethinking Data and Rebalancing Digital Power' (Ada Lovelace Institute, 2022) <<https://www.adalovelaceinstitute.org/project/rethinking-data/>> accessed 3 April 2023.
16. A perspective that recognises how the performance, effectiveness and downstream consequences of technologies derive neither from technical design nor from social dynamics in the abstract, but from the real-world interplay between the two. For more information, see Data & Society's explainer on 'A Sociotechnical Approach to AI Policy': https://datasociety.net/wp-content/uploads/2024/05/DS_Sociotechnical-Approach_to_AI_Policy.pdf
17. Anna Studman, 'Access Denied? Socioeconomic Inequalities in Digital Health Services' (Ada Lovelace Institute, 2023) <<https://www.adalovelaceinstitute.org/report/healthcare-access-denied/>> accessed 21 November 2024.
18. Ada Lovelace Institute and Traverse, 'Confidence in a Crisis? Building Public Trust in a Contact Tracing App' (Ada Lovelace Institute, 17 August 2020) <<https://www.adalovelaceinstitute.org/our-work/covid-19/confidence-in-a-crisis/>> accessed 18 August 2024.
19. Aidan Peppin, Reema Patel and Imogen Parker, 'The Citizens' Biometrics Council' (Ada Lovelace Institute, 2021) <<https://www.adalovelaceinstitute.org/report/citizens-biometrics-council/>> accessed 5 March 2025.

20. Laura Carter, Imogen Parker, Renate Samson and Octavia Reeve, 'Critical Analytics? Learning from the Early Adoption of Data Analytics for Local Authority Service Delivery' (Ada Lovelace Institute, 2023)
<<https://www.adalovelaceinstitute.org/report/local-authority-data-analytics/>>
21. Lara Groves and others, 'Algorithmic Impact Assessment: A Case Study in Healthcare' (Ada Lovelace Institute, 2022)
<<https://www.adalovelaceinstitute.org/report/algorithmic-impact-assessment-case-study-healthcare/>> accessed 19 April 2022.
22. Mavis Machirori and others, 'A Knotted Pipeline' (Ada Lovelace Institute, 2022)
<<https://www.adalovelaceinstitute.org/report/knotted-pipeline-health-data-inequalities/>> accessed 4 March 2025.
23. Cansu Safak and Imogen Parker, 'Meaningful Transparency and (in)Visible Algorithms' (15 October, 2020)
<<https://www.adalovelaceinstitute.org/blog/meaningful-transparency-and-invisible-algorithms/>> accessed 4 March 2025.
24. Mavis Machirori and Anna Studman (n 5).
25. Jones and others (n 3).
26. Mavis Machirori and Anna Studman (n 5).
27. Ada Lovelace Institute, 'Transparency Mechanisms for UK Public-Sector Algorithmic Decision-Making Systems' (Ada Lovelace Institute, 2020)
<<https://www.adalovelaceinstitute.org/report/transparency-mechanisms-for-uk-public-sector-algorithmic-decision-making-systems/>>.
28. Ada Lovelace Institute, 'Transparency Mechanisms for UK Public-Sector Algorithmic Decision-Making Systems' (n 27).
29. Safak and Parker (n 23).
30. Studman and others (n 2).
31. Cabinet Office, Department for Science, Innovation and Technology, and Government Digital Service, 'Find out How Algorithmic Tools Are Used in Public Organisations' (GOV.UK, 2 March 2025) <<https://www.gov.uk/algorithmic-transparency-records>> accessed 4 March 2025.
32. Robert Booth, 'UK Government Failing to List Use of AI on Mandatory Register' *The Guardian* (28 November 2024)
<<https://www.theguardian.com/technology/2024/nov/28/uk-government-failing-to-list-use-of-ai-on-mandatory-register>> accessed 4 March 2025.

- 33.** Department for Science, Innovation & Technology, 'Algorithmic Transparency Recording Standard (ATRS) Mandatory Scope and Exemptions Policy' (GOV.UK, 17 December 2024) <<https://www.gov.uk/government/publications/algorithmic-transparency-recording-standard-mandatory-scope-and-exemptions-policy>> accessed 10 March 2025.
- 34.** Octavia Field Reid, Anna Colom and Roshni Modhvadia, 'What Do the Public Think about AI?' (Ada Lovelace Institute, 2023) <<https://www.adalovelaceinstitute.org/evidence-review/what-do-the-public-think-about-ai/>> accessed 4 March 2025.
- 35.** Ada Lovelace Institute and Alan Turing Institute, 'How Do People Feel about AI? A Nationally Representative Survey of Public Attitudes to Artificial Intelligence in Britain' (2023) <<https://attitudestoai.uk/>> accessed 6 June 2023.
- 36.** Laura Carter, Imogen Parker, Renate Samson and Octavia Reeve, 'Critical Analytics? Learning from the Early Adoption of Data Analytics for Local Authority Service Delivery' (Ada Lovelace Institute, 2023) <<https://www.adalovelaceinstitute.org/report/local-authority-data-analytics/>> accessed 4 March 2025.
- 37.** Elliot Jones and others, 'Foundation Models in the Public Sector' (Ada Lovelace Institute, 2023) <<https://www.adalovelaceinstitute.org/evidence-review/foundation-models-public-sector/>> accessed 4 March 2025.
- 38.** Elliot Jones, Mahi Hardalupas and William Agnew, 'Under the Radar?' (Ada Lovelace Institute 2024) <<https://www.adalovelaceinstitute.org/report/under-the-radar/>> accessed 4 March 2025.
- 39.** HM Treasury and Evaluation Task Force, 'Guidance on the Impact Evaluation of AI Interventions' (GOV.UK, 17 December 2024) <<https://www.gov.uk/government/publications/the-magenta-book/guidance-on-the-impact-evaluation-of-ai-interventions-html>> accessed 4 March 2025.
- 40.** 'National Data Opt-Out' (NHS England Digital, 6 March 2025) <<https://digital.nhs.uk/services/national-data-opt-out>> accessed 6 March 2025.
- 41.** Laura Carter, Imogen Parker, Renate Samson and Octavia Reeve, 'Critical Analytics?' (Ada Lovelace Institute, 2023) <<https://www.adalovelaceinstitute.org/report/local-authority-data-analytics/>> accessed 6 March 2025.
- 42.** Elliot Jones and Cansu Safak, 'Can Algorithms Ever Make the Grade?' (Ada Lovelace Institute, 18 August 2020) <<https://www.adalovelaceinstitute.org/blog/can-algorithms-ever-make-the-grade/>> accessed 1 August 2023.

43. Mavis Machirori and Reema Patel, 'Turning Distrust in Data Sharing into "Engage, Deliberate, Decide"' (Ada Lovelace Institute, 11 June 2021)
<<https://www.adalovelaceinstitute.org/blog/distrust-data-sharing-engage-deliberate-decide/>> accessed 6 March 2025.
44. 'National Data Guardian 2023-2024 Report' (National Data Guardian, 2024)
<<https://www.gov.uk/government/publications/national-data-guardian-2023-2024-report/national-data-guardian-2023-2024-report>> accessed 4 March 2025.
45. Studman (n 17).
46. *ibid.*
47. Peppin, Patel and Parker (n 19).
48. Carter, Parker, Samson, and others (n 20).
49. *ibid.*
50. Ada Lovelace Institute, 'The Rule of Trust: Findings from Citizens' Juries on the Good Governance of Data in Pandemics' (Ada Lovelace Institute, 2022)
<<https://www.adalovelaceinstitute.org/wp-content/uploads/2022/07/The-rule-of-trust-Ada-Lovelace-Institute-July-2022.pdf>> accessed 4 March 2025.
51. Ada Lovelace Institute, 'Beyond Face Value: Public Attitudes to Facial Recognition Technology' (Ada Lovelace Institute, 2019) <
<https://www.adalovelaceinstitute.org/report/beyond-face-value-public-attitudes-to-facial-recognition-technology/>> accessed 4 March 2025.
52. Elliot Jones, Imogen Parker and Gavin Freeguard, 'Checkpoints for Vaccine Passports' (Ada Lovelace Institute, 2021)
<<https://www.adalovelaceinstitute.org/report/checkpoints-for-vaccine-passports/>> accessed 4 March 2025.
53. 'Post Office Horizon Scandal: Why Hundreds Were Wrongly Prosecuted' *BBC News* (21 April 2021) <<https://www.bbc.com/news/business-56718036>>
accessed 6 March 2025.
54. 'Home Office Drops "racist" Algorithm from Visa Decisions' *BBC News* (4 August 2020) <<https://www.bbc.com/news/technology-53650758>> accessed 25 February 2023.
55. Equality and Human Rights Commission, 'Public Sector Equality Duty'
<<https://www.equalityhumanrights.com/en/advice-and-guidance/public-sector-equality-duty>> accessed 28 March 2021.
56. Studman and others (n 2).

57. Mavis Machirori and Anna Studman (n 5).
58. *ibid.*
59. Department for Science, Innovation & Technology, 'AI Management Essentials Tool' (GOV.UK, 6 November 2024)
<<https://www.gov.uk/government/consultations/ai-management-essentials-tool>> accessed 6 March 2025.
60. 'IEEE Draft Standard for the Procurement of Artificial Intelligence and Automated Decision Systems' <<https://standards.ieee.org/ieee/3119/10729/>> accessed 4 March 2025.
61. Studman and others (n 2).
62. *Ibid.*
63. Julia Smakman and others, 'New Rules?' (Ada Lovelace Institute, 2024)
<<https://www.adalovelaceinstitute.org/report/new-rules-ai-regulation/>> accessed 4 March 2025.

Image credit: William Barton