Executive summary

Artificial intelligence has the potential to revolutionise the delivery of public services, creating an opportunity for more innovative and efficient public service delivery. Machine learning in particular will transform the way decisions are made in areas as diverse as policing, health, welfare, transport, social care, and education.

This review found that the Nolan Principles are strong, relevant, and do not need reformulating for AI. The Committee heard that they are principles of good governance that have stood, and continue to stand, the test of time. All seven principles will remain relevant and valid as AI is increasingly used for public service delivery.

If correctly implemented, AI offers the possibility of improved public standards in some areas. However, AI poses a challenge to three Nolan Principles in particular: openness, accountability, and objectivity. This review examined how public officials and government departments can uphold these principles as AI is increasingly rolled out across our public services.

Our concerns here overlap with key themes from the field of AI ethics. Under the principle of openness, a current lack of information about government use of AI risks undermining transparency. Under the principle of accountability, there are three risks: AI may obscure the chain of organisational accountability; undermine the attribution of responsibility for key decisions made by public officials; and inhibit public officials from providing meaningful explanations for decisions reached by AI. Under the principle of objectivity, the prevalence of data bias risks embedding and amplifying discrimination in everyday public sector practice.

This review found that the government is failing on openness. Public sector organisations are not sufficiently transparent about their use of AI and it is too difficult to find out where machine learning is currently being used in government. It is too early to judge if public sector bodies are successfully upholding accountability. Fears over 'black box' AI, however, may be overstated, and the Committee believes that explainable AI is a realistic goal for the public sector. On objectivity, data bias is an issue of serious concern, and further work is needed on measuring and mitigating the impact of bias.

Governance and regulation

To uphold public standards, government and public sector organisations should set effective governance to mitigate the risks we have identified. In this sense, AI is a new challenge that can be solved with existing tools and established principles. Public standards can be upheld with a traditional risk management approach.

This is not a challenge that public sector organisations can tackle alone. Government needs to identify and embed authoritative ethical principles and issue accessible guidance on AI governance to those using it in the public sector. Government and regulators must also establish a coherent regulatory framework that sets clear legal boundaries on how AI should be used in the public sector.

Attempts to establish this governance and regulatory framework are emerging and developments are fast-moving. In the area of ethical principles and guidance, the Department for Culture, Media and Sport (DCMS), the Centre for Data Ethics and Innovation (CDEI) and the Office for AI have all published ethical principles for data-driven technology and AI. The Office for AI, the Government Digital Service (GDS), and the Alan Turing Institute have jointly issued A Guide to Using Artificial Intelligence in the Public Sector and draft guidelines on AI procurement. The Information Commissioner's Office (ICO) has also published its Auditing Framework for AI.

In the area of regulation, the use of AI is subject to the provisions of the GDPR, the Equality Act, and sections of administrative law. The government has also established the Centre for Data Ethics and Innovation to advise on regulation.

These developments are positive and are to be welcomed. However, at the time of writing, this review has found that the governance and regulatory framework for AI in the public sector is still a work in progress and one with significant deficiencies.

This is mostly because key documents have only recently been published and government AI institutions are very new. Multiple sets of ethical principles are confusing and the application of each is unclear. Public sector guidance is not yet widely used and public officials with no AI expertise may find it difficult to understand and comply with.

We conclude that a new AI regulator is not needed but existing regulators will need to adapt to face the challenges AI brings. They will need assistance from a central body to do so, but the CDEI does not yet have a clearly defined purpose and is not yet on a statutory footing. Two areas in particular – transparency and data bias – are in need of urgent attention in the form of new regulation and guidance.

Our recommendations

Our recommendations to government and regulators are intended to assist in the development of a stronger and more coherent regulatory and governance framework for AI in the public sector.

We recommend that government should establish consistent and authoritative ethical principles and issue easier to use guidance. Procurement processes should be reformed and the Digital Marketplace should offer greater assistance to public bodies seeking technologies that are compliant with public standards.

Though no new AI regulator is needed, the CDEI should advise regulators on how to adapt to new technologies and be set on an independent statutory footing. The application of anti-discrimination law to AI needs to be clarified and new transparency guidelines are needed. AI impact assessments should be mandatory, published, and set by the CDEI, and new guidelines are needed to enforce transparency.

We also provide recommendations to providers of public services, both public and private, to help them develop effective risk-based governance for AI. During project planning, our recommendations focus on legal and legitimate AI, system design, and diversity. During project implementation, our recommendations cover setting responsibility, internal and external oversight, monitoring and evaluation, appeal and redress, and training and education.

The Nolan Principles remain a valid guide for public sector practice in the age of AI. However, this new technology is a fast-moving field, so government and regulators will need to act swiftly to keep up with the pace of innovation. Our recommendations set out what we believe is needed to ensure the Seven Principles of Public Life are upheld as the public sector transitions into a new AI-enabled age.

List of recommendations

Recommendations to government, national bodies and regulators

The Committee makes eight recommendations to government, national bodies and regulators to help create a strong and coherent governance and regulatory framework for AI in the public sector.

Recommendation 1:
Ethical principles and guidance

There are currently three different sets of ethical principles intended to guide the use of AI in the public sector – the FAST SUM Principles, the OECD AI Principles, and the Data Ethics Framework. It is unclear how these work together and public bodies may be uncertain over which principles to follow.

   a. The public needs to understand the high level ethical principles that govern the use of AI in the public sector. The government should identify, endorse and promote these principles and outline the purpose, scope of application and respective standing of each of the three sets currently in use.

   b. The guidance by the Office for AI, the Government Digital Service and the Alan Turing Institute on using AI in the public sector should be made easier to use and understand, and promoted extensively.

Recommendation 2:
Articulating a clear legal basis for AI

All public sector organisations should publish a statement on how their use of AI complies with relevant laws and regulations before they are deployed in public service delivery.

Recommendation 3:
Data bias and anti-discrimination law

The Equality and Human Rights Commission should develop guidance in partnership with both the Alan Turing Institute and the CDEI on how public bodies should best comply with the Equality Act 2010.

Recommendation 4:
Regulatory assurance body

Given the speed of development and implementation of AI, we recommend that there is a regulatory assurance body, which identifies gaps in the regulatory landscape and provides advice to individual regulators and government on the issues associated with AI.

We do not recommend the creation of a specific AI regulator, and recommend that all existing regulators should consider and respond to the regulatory requirements and impact of the growing use of AI in the fields for which they have responsibility.

The Committee endorses the government's intention for CDEI to perform a regulatory assurance role. The government should act swiftly to clarify the overall purpose of CDEI before setting it on an independent statutory footing.

Recommendation 5:
Procurement rules and processes

Government should use its purchasing power in the market to set procurement requirements that ensure that private companies developing AI solutions for the public sector appropriately address public standards.

This should be achieved by ensuring provisions for ethical standards are considered early in the procurement process and explicitly written into tenders and contractual arrangements.

Recommendation 6:
The Crown Commercial Service's Digital Marketplace

The Crown Commercial Service should introduce practical tools as part of its new AI framework that help public bodies, and those delivering services to the public, find AI products and services that meet their ethical requirements.

Recommendation 7:
Impact assessment

Government should consider how an AI impact assessment requirement could be integrated into existing processes to evaluate the potential effects of AI on public standards. Such assessments should be mandatory and should be published.

Recommendation 8:
Transparency and disclosure

Government should establish guidelines for public bodies about the declaration and disclosure of their AI systems.

Recommendations to front-line providers, both public and private, of public services

The Committee makes seven recommendations to front-line providers of public services to help establish effective risk-based governance for the use of AI.

Recommendation 9:
Evaluating risks to public standards

Providers of public services, both public and private, should assess the potential impact of a proposed AI system on public standards at project design stage, and ensure that the design of the system mitigates any standards risks identified.

Standards review will need to occur every time a substantial change to the design of an AI system is made.

Recommendation 10:
Diversity

Providers of public services, both public and private, must consciously tackle issues of bias and discrimination by ensuring they have taken into account a diverse range of behaviours, backgrounds and points of view. They must take into account the full range of diversity of the population and provide a fair and effective service.

Recommendation 11:
Upholding responsibility

Providers of public services, both public and private, should ensure that responsibility for AI systems is clearly allocated and documented, and that operators of AI systems are able to exercise their responsibility in a meaningful way.

Recommendation 12:
Monitoring and evaluation

Providers of public services, both public and private, should monitor and evaluate their AI systems to ensure they always operate as intended.

Recommendation 13:
Establishing oversight

Providers of public services, both public and private, should set oversight mechanisms that allow for their AI systems to be properly scrutinised.

Recommendation 14:
Appeal and redress

Providers of public services, both public and private, must always inform citizens of their right and method of appeal against automated and AI-assisted decisions.

Recommendation 15:
Training and education

Providers of public services, both public and private, should ensure their employees working with AI systems undergo continuous training and education.