

Committee on Standards in Public Life submission to the House of Lords Communications and Digital Committee Inquiry on Large Language Models

Background on Committee

1. The Committee on Standards in Public Life (CSPL or the Committee) is an independent, non-departmental public body that advises the Prime Minister on the arrangements for upholding standards of conduct across public life in England. We are not a regulator and do not have investigative powers or consider individual cases. The Committee's remit and current membership is attached at **Annex A**.
2. The Committee articulated the Seven Principles of Public Life – commonly referred to as the Nolan Principles – in its first report in 1995: honesty; objectivity; openness; selflessness; integrity; accountability; and leadership. These Principles apply to all public office holders, including those who are elected or appointed, and to private providers of public services.

Introduction

3. The Committee is pleased to give evidence to this timely inquiry by the House of Lords Communications and Digital Committee into large language models.
4. Our submission is based on the findings and recommendations of the Committee's 2020 report, *Artificial Intelligence and Public Standards*, which is summarised below.¹ The report used the term "AI" liberally to describe anything from automated data analysis to complex neural networks, including machine learning. The evidence we took during the review pre-dated the widespread use of large language models and therefore may no longer be comprehensive. However, we believe that the principles announced in the report remain relevant.
5. Henceforth, this paper sets out points thought to be pertinent from evidence taken during the Committee's 2019-20 AI review, which may be helpful to the Lords' Communications and Digital Committee.²
6. The CSPL regards the regulation of AI as an issue of immense and immediate importance. As such, the Committee is currently following up the recommendations made in the 2020 AI report, and has written to regulators³, government departments and public bodies⁴ – as well as the Minister for AI⁵ and Adviser to the Prime Minister on AI⁶ – with the intention of holding a seminar and publishing an update on progress in the new year. This further work may cause the Committee to add to or modify its earlier recommendations.

¹ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020

² Committee on Standards in Public Life, [evidence collection from AI review](#), 2019-20

³ Committee on Standards in Public Life, [AI and public standards: 2023 regulators survey](#), 2023

⁴ Committee on Standards in Public Life, [AI and public standards: 2023 public bodies survey](#), 2023

⁵ Committee on Standards in Public Life, [AI and public standards: letter from Lord Evans to AI Minister](#), 2023

⁶ Committee on Standards in Public Life, [AI and public standards: letter from Lord Evans to AI Adviser](#), 2023

Summary of 2020 report on Artificial Intelligence

7. In 2020, the Committee published a report, *Artificial Intelligence and Public Standards*, which looked at the risks and opportunities for public standards posed by AI, and examined whether the then regulatory and governance landscape for AI was fit for purpose.
8. The report found that while the Nolan Principles remain strong, relevant and do not need reformulating for AI, AI posed a particular challenge to three principles: openness, accountability, and objectivity. On openness, we found there was a lack of information about the government's use of AI; on accountability, we found AI could make it difficult to hold public organisations accountable for their decisions and to provide meaningful explanations for decisions; and on objectivity, we found that the prevalence of data bias in AI risks entrenching or amplifying discrimination in public sector practice.⁷
9. We found that a robust and coherent legal and regulatory framework for AI in the public sector is a work in progress. For instance, at the time of writing there had been good progress around the establishment of expert bodies and government departments, including the Office for AI and the Centre for Data Ethics and Innovation, and in establishing ethical principles and guidance for AI. However, the specific functions of some of those bodies and sets of ethical principles remain unclear. Similarly, we found that although AI is subject to the provisions of the GDPR⁸, the Equality Act, and sections of administrative law, there remains some uncertainty about how the law applies to automated decision-making in practice.⁹
10. At the time of writing (2020), the weight of evidence was that the UK did not need a specific AI regulator, but that all regulators should consider and respond to the regulatory requirements and impact of AI in the fields for which they have responsibility. However, given the complexity of AI, we felt regulators would require advice and guidance from a central body about issues associated with AI, and recommended – in line with the government's then published intention¹⁰ – that the Centre for Data Ethics and Innovation take on this responsibility.¹¹
11. It was clear from the evidence we received that there is nothing inherently new about the governance needed for AI, and that public standards can be upheld with a traditional risk management approach. Because decisions about implementing AI in the public sector lie with individual government departments and public bodies, we took the view that they should each be responsible for establishing appropriate

⁷ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.16-29

⁸ At the time of writing, the EU GDPR had direct application in UK law through the Data Protection Act 2018. The provisions of the EU GDPR have since been incorporated directly into UK law as the UK GDPR.

⁹ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.39-56

¹⁰ DCMS, [Centre for Data Ethics and Innovation: government response to consultation](#), 2018

¹¹ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.47-9

governance arrangements to manage the ethical risks associated with AI and address regulatory compliance. As such, we recommended a number of specific risk management mechanisms that public bodies using AI should put in place, including using AI in ways that are legal and legitimate; monitoring AI systems to ensure they operate as intended; and setting appropriate oversight and appeal mechanisms for AI systems and decisions.¹²

12. The report's recommendations are set out in full at **Annex B**.

Relevant points from evidence taken during CSPL AI review

Opportunities and risks of AI

13. Although we did not take evidence specifically on the trends and capabilities of large language models, we did hear during the course of our review that the impact of AI across the public and private sectors is potentially vast, with advances in AI capability having the potential to revolutionise the delivery of public services, creating an opportunity for more innovative and efficient public service delivery in areas as diverse as policing, health and education. We also heard that AI could potentially improve public standards in some areas. However, these opportunities can only be realised if AI is implemented in accordance with the Nolan Principles. Any change in how the government makes policy decisions or delivers public services must not undermine public standards or the public's confidence in its institutions. This is especially important because AI has the potential to change how decisions are made in sensitive policy areas, where the impact on individuals can be significant.¹³

14. As noted in paragraph 8, our report identified three areas where AI poses a risk to the public standards: openness, accountability and objectivity. These are summarised below.

Openness

15. Evidence received in 2020 suggests that the government and public bodies are not sufficiently transparent about their use of AI, with most information resulting from Freedom of Information requests and procurement data.¹⁴ We heard that transparency is further complicated by the public sector's use of private sector commercial organisations in the development and provision of AI. These organisations may use commercial confidentiality arrangements to avoid certain forms of disclosure.¹⁵ This can make it difficult for the public and others to scrutinise and hold organisations accountable for their decisions.

¹² Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.57-66

¹³ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, p.10

¹⁴ The Bureau of Investigative Journalism, [Government Data Systems](#), 2019

¹⁵ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, p.18

16. In our report, we considered whether the requirements for proactive disclosure under the Freedom of Information Act 2000 are sufficient to increase transparency around the use of AI by public organisations. We said that an expectation on public bodies to think about openness is not enough to change behaviour, and recommended that the government establish guidance for public bodies about what information to disclose about their AI systems. We did not recommend the establishment of a centralised AI transparency register, which could be an overwhelming bureaucratic challenge.¹⁶
17. We also heard that some more complex forms of AI are opaque (eg. you cannot see how decisions are made). AI systems of this form are often referred to as “black boxes”. Contributors to the review told us that most AI used in the public sector will be processing simple data meaning that less complex and less opaque AI systems can and should be used. Where more opaque AI systems are used, public bodies should justify why certain trade-offs have been made. The CSPL took the view in the 2020 report that the technical obstacles to providing explanations for automated decisions were small and that it would be possible for public bodies to provide meaningful explanations of AI-decisions. We said that public bodies would need to consider explainability in the early stages of AI design and development, and during the procurement process, where requirements for transparency could be stipulated in tenders and contracts, to achieve this.¹⁷

Accountability

18. Human oversight over an AI system and its decision-making process and outcomes is a standards imperative. Responsibility for AI will likely be shared by individuals across an organisation, which should be clearly allocated. Senior leadership should have oversight over the whole AI process, from making decisions about procuring AI systems to reviewing the impact of automated decisions. In high risk policy areas, such as health or policing, independent oversight bodies, such as ethics committees, are useful tools for ensuring that ethical challenges relating to AI are given proper consideration, and for providing independent scrutiny. Public bodies using AI should also establish processes to monitor and evaluate issues relating to the performance of the technology.¹⁸

Objectivity

19. The evidence we received in 2020 suggested that data bias could cause AI to produce decisions and policy outcomes that are discriminatory, which may breach the Equality Act 2010. Contributors to the review told us that the Public Sector Equality Duty, established under the Equality Act in 2011, is the “single best tool available” for dealing with data bias if used correctly.¹⁹ For example, many public bodies told us in

¹⁶ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.54-6

¹⁷ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.22-4

¹⁸ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.19-22 and pp.63-4

¹⁹ Equality and Human Rights Commission, [Public Sector Equality Duty](#), 2011

2020 that they undertook Equality Impact Assessments to consider the potential impact of policy decisions on protected characteristics. We said in the report that the same could be done for automated decisions. However, contributors also told us there was uncertainty about how the Equality Act 2010 applies to automated decision-making in practice. We recommended that The Equality and Human Rights Commission develop guidance on how public bodies using AI should comply with the Equality Act 2010.²⁰ We are glad to see that the EHRC has since done so.²¹

20. Other methods suggested by contributors to reduce the risk of data bias include: maximising diversity and taking into account a diverse range of behaviours, backgrounds and perspectives at all stages of the AI process, including in the workforce and in training and education; and periodically retesting AI models to observe whether any groups are being systematically advantaged or disadvantaged and updating AI systems accordingly.²²

Domestic AI regulation

21. As noted in paragraph 9, our 2020 report found that a coherent regulatory framework for AI in the public sector remains a work in progress. For example, healthcare practitioners told the Committee they were confident that AI could be implemented safely and ethically in medicine because it operates within a well regulated system, where there are professional standards in place for testing and implementing new technologies, and for reporting and research.²³ In contrast, the same established regulatory framework does not exist in policing, which led to the “unlawful and unethical” use of facial recognition technology by the police in late 2022, according to researchers at the University of Cambridge.²⁴ Evidence submitted to the 2020 review suggests that the use of AI in policing is more representative of the overall use of AI in the public sector.²⁵
22. As noted in paragraph 10, most contributors to the review said that the UK did not need a specific AI regulator, which would “inevitably overlap with existing regulations”, arguing that any system of ethical regulation for AI in the public sector will require sector based oversight to account for the specific risks and challenges of automated decision-making across sectors. As such, we recommended that all regulators should consider and respond to the regulatory requirements and impact of

²⁰ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.45-7

²¹ Equality and Human Rights Commission, [Artificial Intelligence in Public Services](#), 2022

²² Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.25-9

²³ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, p.39

²⁴ Minderoo Centre for Technology and Democracy (University of Cambridge), [A Sociotechnical Audit: Assessing Police use of Facial Recognition](#), 2022

²⁵ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, p.39

AI in the fields for which they have responsibility.²⁶ This is also the approach taken by the government in their AI White Paper, *A pro-innovation approach to AI regulation*.²⁷

23. However, three years on, given the pace of AI developments in that period, we are currently following up the recommendations made in our report, as noted in paragraph 6.
24. On the question of whether regulators have sufficient expertise and resources to respond to generative AI, we are of the view that the majority of regulators do not. We said in our report that regulators will need guidance from a central and independent body to help them navigate the challenges associated with AI and any unforeseen regulatory issues. As noted in paragraph 10, we recommended in our report that the Centre for Data Ethics and Innovation take on this responsibility, in line with the government's previously published intention that they oversee and advise on AI governance and regulation.²⁸ We also supported the government's intention to establish the centre on a statutory footing.²⁹ However, the specific functions of the Centre remain unclear. At present, it is described as a "government expert body enabling the trustworthy use of data and AI"³⁰, and makes no mention of its intended role in identifying and addressing "areas where guidelines or regulations" are needed. This means there may still be a significant gap in the regulatory landscape for AI and lack of clarity for the public.
25. This view is supported by further evidence we took in November 2020 from regulators about how they were adapting to the challenges posed by AI. Although some regulators were "quite prepared" for AI, most regulators told us that they were in the very early stages of thinking about AI, with some raising issues around limited expertise and lack of resources, and others noting that the regulation of AI is beyond their scope. Most regulators said they were limited in the extent to which they could place controls on the organisations they regulate, and that more could be done to improve and distribute guidance to regulators to help develop a more coordinated approach to AI regulation.³¹ As noted in paragraph 6, we are currently following up with regulators to see what progress has been made in this area.
26. On the question of other regulatory and non-regulatory options to address the risks associated with AI, we said in our report that procurement processes can act as a form of soft regulation, and that government should use its purchasing power in the market to set procurement requirements that ensure private companies developing AI for the public sector address public standards. For example, procurement processes should be designed so products and services that facilitate high standards are preferred and companies that prioritise ethical practices are rewarded. As part of

²⁶ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.47-9

²⁷ DSIT, [A pro-innovation approach to AI regulation](#), 2023

²⁸ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.47-9

²⁹ DCMS, [Centre for Data Ethics and Innovation: government response to consultation](#), 2018

³⁰ CDEI, [About us page](#) (gov.uk)

³¹ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards: regulators survey and responses](#), 2020

the commissioning process, the government should set out the ethical principles expected of companies providing AI services to the public sector. Adherence to ethical standards should be given an appropriate weighting as part of the evaluation process, and companies that show a commitment to them should be scored more highly than those that do not.³²

27. As set out in paragraph 11 of this response, our report also stressed the importance of individual government departments and public bodies taking responsibility for establishing suitable governance arrangements to manage the ethical risks associated with AI.³³

International AI context

28. We have not taken sufficient evidence to compare the UK's regulatory approach to generative AI with a wide range of other jurisdictions. However, several contributors to our 2020 review told us that there were lessons to be learned from both Singapore and Canada on AI governance, as set out below.

29. In 2019, Singapore's Personal Data Commission published a model framework for AI governance (since updated), which states that the risks associated with AI can be managed by adapting existing or setting up internal governance structures to incorporate values, risks, and responsibilities relevant to AI decision-making.³⁴ We heard this approach was "a useful starting point for thinking about the mechanisms that public sector organisations in the UK should adopt when using AI technology".³⁵

30. Several contributors to the review argued that the absence of a mandatory AI impact assessment was a gap in the UK's regulatory framework for AI. Some contributors spoke favourably about the Canadian Algorithmic Impact Assessment, which is a compulsory risk-assessment tool intended to determine the social, environmental, and human rights impact of an automated decision-system.³⁶ We recommended in 2020 that the government consider how an AI impact assessment requirement could be integrated into the existing regulatory framework for AI. There may be lessons to be learned from Canada in doing so.³⁷

September 2023

³² Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.49-52

³³ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.57-66

³⁴ Personal Data Protection Commission (Singapore), [Model AI governance framework](#), 2020

³⁵ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, p.57

³⁶ Government of Canada, [Algorithmic Impact Assessment tool](#)

³⁷ Committee on Standards in Public Life, [Artificial Intelligence and Public Standards Report](#), 2020, pp.52-4

Annex A – Committee on Standards in Public Life remit and membership

The Committee on Standards in Public Life is an independent, advisory Non-Departmental Public Body (NDPB). The Committee was established in October 1994, by the then Prime Minister, with the following terms of reference:

To examine current concerns about standards of conduct of all holders of public office, including arrangements relating to financial and commercial activities, and make recommendations as to any changes in present arrangements which might be required to ensure the highest standards of propriety in public life.

The Principles of Selflessness, Objectivity, Integrity, Accountability, Openness, Honesty and Leadership remain the basis of the ethical standards expected of public office holders and continue as key criteria for assessing the quality of public life.

The remit of the Committee excludes investigation of individual allegations of misconduct.

On 12 November 1997, the terms of reference were extended by the then Prime Minister:

To review issues in relation to the funding of political parties, and to make recommendations as to any changes in present arrangements.

The Committee's terms of reference were further clarified following the Triennial Review of the Committee in 2013. The then Minister of the Cabinet Office confirmed that the Committee:

Should not inquire into matters relating to the devolved legislatures and Governments except with the agreement of those bodies. Secondly the Government understands the Committee's remit to examine "standards of conduct of all holders of public office" as encompassing all those involved in the delivery of public services, not solely, those appointed or elected to public office.

Committee membership:

- Lord Evans of Weardale KCB DL, Chair
- Rt Hon Lady Arden of Heswall DBE
- Rt Hon Dame Margaret Beckett DBE MP
- Rt Hon Ian Blackford MP
- Ewen Fergusson
- Baroness Finn
- Professor Gillian Peele
- [One independent member vacancy]

Annex B – 2020 AI report recommendations

Recommendations to government, national bodies and regulators

Recommendation 1

There are currently three different sets of ethical principles intended to guide the use of AI in the public sector. 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 and respective standing of the three sets currently in use.
- b. The guidance by the Office for AI, 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

All public 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

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

Recommendation 4

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

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 (CCS) 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

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

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

Recommendations to public bodies and private providers of public services

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.