horizon
DIGITAL ECONOMY RESEARCH

Response to ICO consultation on the draft AI auditing framework guidance for organisations

Submitted by

Prof. Derek McAuley	derek.mcauley@nottingham.ac.uk
Dr. Ansgar Koene	ansgar.koene@nottingham.ac.uk
Dr. Jiahong Chen	jiahong.chen@nottingham.ac.uk

Horizon Digital Economy Research, University of Nottingham
https://www.horizon.ac.uk

Horizon is a Research Institute centred at The University of Nottingham and a Research Hub within the UKRI Digital Economy programme. Horizon brings together researchers from a broad range of disciplines to investigate the opportunities and challenges arising from the increased use of digital technology in our everyday lives, specifically related to opportunities and challenges in personal data use.

1 May 2020

ico.
Information Commissioner's Office

Consultation on the ICO’s Guidance on the AI auditing framework

The ICO are consulting on our guidance on the AI auditing framework. This guidance aims to give organisations practical advice to help them create, use and audit AI systems that are compliant with data protection laws.

We are looking for a wide range of views from organisations across all sectors and sizes.

If you would like further information about the consultation, please email AIAuditingFramework@ico.org.uk.

Please send us your response by 17:00 on 01/04/2020 by completing the online version of this survey.

Privacy statement

Please note, your responses to this survey will be used to help us with our work on the AI auditing framework only. The information will not be used to consider any regulatory action, and you may respond anonymously should you wish. For more information about what we do with personal data see our privacy notice.

Please note that we are using the platform Snap Surveys to gather this information. Any data collected by Snap Surveys for ICO is stored on UK servers. You can read their Privacy Policy.

Q1 Is the draft guidance clear about what you should consider when creating and using AI-systems that are compliant with data protection law?

☐ Yes

☒ No

Please outline what parts, if any, you think could be improved:

There is a general disconnection between the title of the guidance, which suggests that it addresses auditing of AI system, and the actual content that focuses overwhelmingly on showing compliance with data protection law.

For specific suggestions of improvements for each section, see responses below to individual questions.

Q2 How well-pitched are the sections in the draft guidance?

a ‘About this guidance’

☒ Too detailed

☐ Just right

☐ Too vague

Please provide your suggestions on how we can improve on the level of detail:

The draft guidance would benefit from greater clarity about the specific audience that is being targeted. Some parts appear to be written for people involved in the development or deployment of AI systems while other parts seem to address those providing internal oversight. It would help if the intended core audience is indicated.

b 'What are the accountability and governance implications of AI?

☐ Too detailed

☐ Just right

☒ Too vague

Please provide your suggestions on how we can improve on the level of detail:

One point worth noting in the "controller/processor relationships" section is that many SMEs are relying on standardised AI services such as Microsoft Azure Machine Learning or Cognitive Services. Smaller organisations often make use of these AI solutions on a "plug and play" basis, with very little influence on the underlying algorithms or training datasets. The guidance could on the one hand provide further clarification on the role of such services as data controllers, data processors or simply "facilitators" in different scenarios. On the other hand, in order for these smaller organisations to fulfil their data protection duties, it would also be helpful to provide some general advice on what to look out for when choosing these solutions.

In the "AI-related trade-offs" section, the draft guidance discusses a number of common considerations in designing or selecting the appropriate AI systems. However, the guidance should also cover the general trade-offs between AI solutions and the less complex algorithmic or even manual ones. Many businesses choose to adopt AI systems on the assumption that they are more reliable, scalable and affordable – which is not necessarily true – without a thorough assessment of the longer-term and less tangible costs in terms of, for example, compliance, auditing and consumer trust. In fact, in several places across the guidance, it has been suggested that using AI systems is not always the most appropriate option – or sometimes even not an option if no solution fulfils the minimum requirements. It would be helpful to highlight in this section some of the common technical and commercial factors to take into account before making the decision to adopt an AI approach.

c 'What do we need to do to ensure lawfulness, fairness, and transparency in AI systems?'

☐ Too detailed

☒ Just right

☐ Too vague

Please provide your suggestions on how we can improve on the level of detail:

In the “purpose and lawful basis” section, it should be emphasised that the legitimate interest of the data controller would not form a valid legal basis when sensitive data or solely automated individual decision-making is involved.

In the “bias and discrimination” section, the guidance could provide a brief summary of existing approaches in detecting discrimination based on protected categories of characteristics without (or with a minimum level of) collecting sensitive data.[1]

Many of the issues discussed in this part are generically true for any IT system. It would be helpful if those parts that are specific to AI systems were highlighted. Perhaps some formatting design could be used to separate the AI specific concerns from the general IT system concerns.

d 'How should we assess security and data minimisation in AI?'

☐ Too detailed

☒ Just right

☐ Too vague

Please provide your suggestions on how we can improve on the level of detail:

This part should also cover how certification mechanisms could help developers and users of AI systems demonstrate compliance with data protection principles. While the ICO has not yet accredited any certification schemes, relevant work is already underway.[2] Once such schemes are up and running, data controllers may consider having their products or services certified. It is therefore important to explain to organisations what can be certified and to what extent this is encouraged as part of their risk assessment of the use of AI systems.

e 'How do we enable individual rights in our AI systems?'

☐ Too detailed

☐ Just right

☒ Too vague

[1] See for example, Veale, M. and Binns, R., 2017. Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. Big Data & Society.

[2] https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2019/12/statement-on-ico-approved-certification-schemes/

Please provide your suggestions on how we can improve on the level of detail:

In the "enable individual rights" section, the draft guidance has rightly underlined that in general processing of personal data for automated individual decision-making, the "safeguards differ to those in the GDPR if the lawful basis for such processing is a requirement or authorisation by law." Considering the way the GDPR has explicitly laid down the safeguards for such processing based on consent or contract but left that matter for Member States to decide when it comes to processing based on authorisation by law, it is reasonable to believe that the level of protection in the first two categories should not be lower than the third category. For the third category, Section 14(4)(b) DPA 2018 allows the data subject to request the data controller to "(i) reconsider the decision, or (ii) take a new decision that is not based solely on automated processing." We believe this should be considered a baseline safeguard equally applicable to the "right to obtain human intervention" when such decision-making is based on consent or contract.

3 Is it easy to find information in the draft guidance?

☐ Yes

☒ No

Please provide your suggestions, if any, on how the structure could be improved:

Generally the draft guidance lacks a formal structure to highlight the key components of an AI audit framework. As highlighted above, the guidance could be restructured into a number of parts addressing, for example, the roles of different groups of target audience, the types of risk involved in generic IT or specific AI systems, the major phases of the development and deployment of AI systems where auditing can be put in place, or the key elements of the data protection framework against which an effective assessment can be conducted.

Also, the guidance covers the related issues in the form of questions, making the sub-headings too long and hard to locate the information from the table of contents. Using more concise key words for as section and sub-section headings would significantly improve the usability of the guidance.

4 Are the risk statements and the examples of controls useful?

☒ Yes

☐ No

Please provide any suggestions, if any, on how these could be improved:

5 Do you have any examples of using the draft guidance in practice that you think would be useful for us to know?

☐ Yes

☒ No

If yes, please provide further details:

6 What industry is your organisation in?

Higher education.

7 Do you develop AI in house, or provide/procure it to/from others?

Multiple options allowed

☐ We procure AI from a third party

☐ We create and use AI in-house

☐ We provide AI to a third party/parties

☒ N/A

If yes, please provide further details:

8 Where did you hear about the consultation?

- ☐ ICO Twitter
- ☐ ICO LinkedIn
- ☐ ICO enewsletter
- ☐ ICO website
- ☐ Twitter
- ☐ LinkedIn
- ☐ Other organisation’s enewsletter
- ☐ Other website (please specify)
- ☐ Media, blog or podcast
- ☐ ICO staff member
- ☐ Colleague

Thank you for completing our survey