Case reference IC-479685-D4P5

Disclosure

Data Protection Impact Assessment (DPIA) – [Insert the title of this DPIA]

Document Name	Data Protection Impact Assessment – [insert the title of this DPIA]
Author or Owner (name and job title)	[Your Name and Role]
Directorate and Team	[Your directorate and team]
Document Status	Choose an item.
Version Number	v0.1
Release Date	
Approver (if applicable)	
Review Date	
Distribution	Choose an item.

Guidance for completing this template

- If you're unsure whether you need to complete a DPIA, use the DPIA screening assessment to help you decide.
- **Must** and **should** are used throughout the guidance notes in this template to help you understand legislative requirements that **must** be met and additional steps that the ICO considers **should** be done as best practice to comply effectively with the data protection laws.
- You **must** complete a DPIA, using this template, if your screening assessment indicates a DPIA is required.
- You **should** aim to complete your DPIA as early as possible as the outcome of the assessment could affect the viability of your plans. In extreme cases, you won't be able to continue with your plans without changing them, or at all.
- We recommend that you fill out each section of this template in order, as each subsequent section builds upon the last. You won't be able to complete later sections correctly if you skip ahead. You **should** read the guidance notes throughout this template to help you with each section and you can also find training videos on IRIS that will help you.
- If you are struggling to complete this template and need advice the Information Management and Compliance Service is available to provide support. Please keep in mind their service standards if you require assistance.

1. Data processing overview

1.1. Ownership of risk

Guidance notes:

- There **must** be a clear owner for any residual risk resulting from your data processing. At the ICO our Information Asset Owners (IAO), as our service directors, are our senior risk owners and **must** sign off on your plans.
- You should identify below the individuals with responsibility for competing this DPIA and signing off on the residual risk. They will also be responsible for the periodic review of this assessment to ensure risk is being monitored and managed appropriately.

Project title	
DPIA lead	
Information Asset Owner	Director of….

1.2. Describe your new service or process

Guidance notes:

- Provide a summary below of the service or process you want to implement. Include any relevant background information and your key aims and objectives.

1.3 Use of Artificial Intelligence

Guidance notes:

- Artificial Intelligence (AI) is a term for a range of technologies which attempt to mimic human thought to solve complex tasks.
- AI technologies may cause or exacerbate risks to individuals and specific controls may be needed to ensure an AI system is compliant with data protection law and provides sufficient safeguards for individuals' rights and freedoms.

Guidance links:
Guidance on AI and data protection
AI Governance and accountability toolkit
AI and data protection risk toolkit

Q.1 Please indicate whether your plans involve the use of Artificial Intelligence technologies.

a) I am developing a new AI technology. For example, you'll be selecting datasets to train your own model.

Choose an item.

b) I am deploying an existing AI technology. For example, you're intending to implement an already developed AI system from a third party, or a piece of software that has AI features.

Choose an item.

If you've answered **No** to both questions above, you can skip the remainder of this section and move to section 1.4 - Understanding our role.

Guidance notes:

- Provide a summary below of the AI technology you are developing and/or intend to deploy. You should explain below what this is and what it is intended to be used for and by who.
- If you are developing a new AI technology, you should explain how you plan to train and develop the model. For example, what data will be used, where this has been obtained from as well as what steps you intend to take to improve the statistical accuracy of your AI system.
- If you are procuring AI as a service or off-the-shelf models, you should ask the supplier for information that explains how the model has been developed and seek assurances about the product's statistical accuracy. You can link to any relevant product

documentation or similar within your explanation below where this is useful.

- You should outline any testing you have planned as part of the AI systems development or in advance of your intended deployment. Or if you've already conducted testing explain what was learned.

1.4 Understanding our role

Guidance notes:

- We **must** understand our role in relation to the personal data you will be processing. Our obligations will vary depending on whether we are a controller, joint controller or processor.
- Controllers are the main decision-makers who exercise overall control over the purposes and means of the processing of personal data. Where two of more controllers jointly determine this, they are joint controllers. Processors act on behalf of, and only on the instructions of, the relevant controller.
- If you are procuring a new product or service from a third party, you will typically find information about data protection roles and responsibilities within the service terms and conditions, or any contract being agreed between us and the third party.

Guidance links:

Controllers and processors

What are the accountability and governance implications of AI?

Controller(s):	
Data processor(s):	

1.5. Personal data inventory

Guidance notes:

- We **must** have a clear understanding of the personal data being processed. This is **essential** for identifying and managing risks.
- Use the table below to list each category of personal data being processing. Use a new row for each data category. You can add as many rows to the table as you need.
- Categories of data may not be obvious to you from the outset e.g. tracking data (IP or location) or data collated via cookies, and you need to take the time to fully understand the extent of the personal data you will process.
- Your data subjects are the individuals the personal data relates to. For example, these could be members of the public, ICO employees, our contractors etc. Try and provide an estimate of how many data subjects are involved.
- Recipients will be anyone who the data is shared with.
- UK GDPR restricts transfers of personal data outside of the UK so any overseas transfers **must** be identified.
- Personal data should be kept for no longer than is necessary. You **must** identify a retention period for the personal data you intend to process.

Guidance links:

Personal data - what is it?

Category of data	Data subjects	Recipients	Overseas transfers	Retention period
			Choose an item. If yes, list the countries the data will be transferred to:	**Choose an item.** If selecting other, please specify the length of time personal data will be retained:
			Choose an item. If yes, list the countries the data will be transferred to:	**Choose an item.** If selecting other, please specify the length of time

Category of data	Data subjects	Recipients	Overseas transfers	Retention period
				personal data will be retained:
			Choose an item. If yes, list the countries the data will be transferred to:	**Choose an item.** If selecting other, please specify the length of time personal data will be retained:

1.6. Lawful basis for ICO processing

Guidance notes:

- To process personal data, you **must** have a lawful basis. Select a lawful basis for processing the personal data in your inventory from the drop-down lists below.
- If you are planning to use AI personal data can be processed at different stages of the AI lifecycle, for different purposes, and you should identify a lawful basis for each stage.

Guidance links:

Lawful basis

Lawful basis interactive guidance tool

How do we ensure lawfulness in AI?

First, select a lawful basis from Article 6 of the UK GDPR.

Choose an item.

If more than one lawful basis applies to your processing, please list any additional basis here:

Guidance notes:

- If your personal data inventory includes any **special category data,** you **must** identify an additional condition for processing from Article 9 of the UK GDPR.

Guidance links:

Special category data

Next, if applicable, select an additional condition for processing from Article 9 of the UK GDPR:

Choose an item.

If you have selected conditions (b), (h), (i) or (j) above, you also need to meet the associated condition in UK law, set out in Part 1 of Schedule 1 of the DPA 2018. Please select from the following:

Choose an item.

If you are relying on the substantial public interest condition in Article 9(2)(g), you also need to meet one of the conditions set out in Part 2 of Schedule 1 of the DPA 2018. Please select from the following:

Choose an item.

Guidance notes:

- If you are processing **criminal offence data,** you **must** meet one of the 28 conditions for processing criminal offence data set out in paragraphs 1 to 37 Schedule 1 of the DPA 2018.

Guidance links:

Criminal offence data | ICO

Finally, if applicable select an additional condition for processing any criminal offence data:

Choose an item.

1.7. Necessity and proportionality

Guidance notes:

- You **must** assess whether your plans to process personal data are both necessary and proportionate to you achieving your purpose and outline any alternative options you have considered. For example, if you're intending to use Artificial Intelligence, you should also outline the non-AI alternatives have you considered.
- Your assessment **must** include consideration of whether your processing will be within the reasonable expectations of your data subjects. For example, where you intend to use Artificial Intelligence, you need to assess whether individuals would reasonably expect an AI system to conduct the processing.
- You **must** take steps to minimise the personal data you process; processing only what is adequate, relevant and necessary.
- You **should** think about any personal data you can remove without affecting your objective.

- You **should** consider if there's any opportunity to anonymise or pseudonymise the data you're using, for example can you anonymise data used for training AI systems.

1.8. Consulting with stakeholders

Guidance notes:

- You **should** consult with relevant stakeholders to help you identify any risks to your data subjects. This might include consulting internally with colleagues, for example in our Cyber Security or Legal Services teams. Or externally, with service suppliers whose products you might be purchasing.
- Where your data processing includes the use of new or complex technologies, such as AI, you **should** consider whether consultation with product developers or subject matter experts is required to help with your understanding of the technology and any risks. For example, our AI policy or AI tech teams.
- Briefly outline who you've consulted to inform your DPIA.
- Where appropriate you **should** seek the views of your data subjects, or their representatives, on your intended processing. Where this isn't possible, you should explain why below.

2. Personal data lifecycle

Guidance Note:

- You **must** provide a systematic description of your processing from the point that personal data is first collected through to its disposal.
- This **must** include the source of the data, how it is obtained, what technology is used to process it, who has access to it, where it is stored and how and when it is disposed of.
- If your plans involve the use of any new technology, for example a new piece of software, you **must** explain how this technology works and outline any 'privacy friendly' features that are available. You **should** obtain relevant information from the developer or supplier of the technology.
- If developing or deploying an AI system you **must** outline when AI processes and automated decisions may produce effects on individuals and outline the degree of any human involvement in the decision-making process. You should demonstrate an understanding of any limitations, relevant variation or margins of error in the performance of the system.
- If helpful you can use the headings provided below to help you construct your lifecycle. You can also include flow diagrams if these help with your explanation.

Guidance links:

What are the accountability and governance implications of AI?

What do we need to know about accuracy and statistical accuracy

Fairness in the AI lifecycle

Data source and collection:

Technology used for the processing (including any use of artificial intelligence):

Storage location:

Access controls:

Data sharing:

Disposal:

3. Key UK GDPR principles and requirements

Guidance notes:

- Answering the questions in this section will help you comply with the data protection principles.
- You may identify specific actions that are needed, and you should add these to your list of DPIA outcomes in section 6.0.

Guidance links:

A guide to the data protection principles

3.1. Purpose and Transparency

Guidance notes:

- In most cases you will need to communicate essential information about your data processing to your data subjects. A privacy notice is the most common way of doing this. This is to ensure we're transparent about what we're doing with personal data. Specific guidance covering AI is also linked to below.
- You **must** review the existing privacy notice on the ICO website. If your data processing involves the personal data of ICO staff, review our Staff Privacy Notice on Iris.
- You need to decide if our existing privacy notices sufficiently cover your plans. If not, you **must** get them updated or you **must** provide your data subjects with a separate, bespoke privacy notice.

Guidance links:

Principle (a): Lawfulness, fairness and transparency

The right to be informed

How do we ensure transparency in AI?

Explaining decisions made with AI

Q1. How will you provide your data subjects with information about your data processing?

Choose an item.

Guidance notes:

- If you identified consent as your lawful basis for processing in section 1.4 you **must** maintain appropriate records of the data subject's consent.

Guidance links:

Consent

Q2. Are you satisfied you're maintaining appropriate records of data subjects' consent?

Choose an item.

Guidance notes:

- If you identified legitimate interests as your lawful basis for processing in section 1.4 you **should** complete a Legitimate Interests Assessment (LIA). We have a template LIA available.

Guidance links:

Legitimate interests

Q3. If legitimate interests is your lawful basis for processing have you completed a legitimate interest assessment?

Choose an item.

If applicable, please provide a link to your completed assessment.

3.2. Accuracy

Guidance notes:

- All reasonable steps should be taken to ensure personal data is kept accurate and up to date. Steps **must** be taken to ensure that personal data that are inaccurate are erased or rectified without delay.
- If using an AI system, you need to consider the accuracy of its outputs. You should provide assurances below about the statistical accuracy of the AI system and consider any measures you can take to mitigate accuracy risks as part of your risk assessment, for example, ensuring you have sufficient human oversight of outputs.

Guidance links:

Principle (d): Accuracy

AI - what do we need to know about accuracy and statistical accuracy?

Q4. Are you satisfied the personal data you're processing is accurate?

Choose an item.

Q5. How will you ensure the personal data remains accurate for the duration of your processing?

3.3. Minimisation, Retention and Deletion

Guidance notes:

- You should only collect and hold the minimum amount of personal data you need to fulfil your purpose. Data should be retained for no longer than is needed for that purpose and then deleted without delay.

Guidance links:

Principle (c): Data minimisation

How should we assess security and data minimisation in AI?

Q6. Have you done everything you can to minimise the personal data you're processing?

Choose an item.

Q7. How will you ensure the personal data are deleted at the end of the retention period?

Q8. Will you need to update the ICO retention and disposal schedule?

Choose an item.

3.4. Security: Confidentiality, integrity and availability

Guidance notes:

- Personal data **must** be processed in a way that ensures it is appropriately secure and protected from unauthorised access, accidental loss, destruction or damage. For AI systems, you must consider if extra security measures are needed to protect the system.
- You **must** make sure access to the personal data is limited to the appropriate people and ensure you're confident the processing system being used is secure.

Guidance links:

Principle (f): Integrity and confidentiality (security)

A guide to data security

How should we assess security and data minimisation in AI?

Q9. Where will the personal data be stored and what measures will you put in place to maintain confidentiality, integrity and availability?

Q10. Have you confirmed there are appropriate access controls to keep the personal data secure?

Choose an item.

Q11. Has the cyber security team completed a security assessment of your plans?

Choose an item.

Q12. If yes what was the outcome of their assessment?

Q13. Please explain the policies, training or other instructions you intend to put in place to enable staff to operate the new system or process securely.

3.5. Accountability and governance

Guidance notes:

- The accountability principle makes us responsible for demonstrating our compliance with the UK GDPR. We do this by clearly assigning responsibilities for compliance tasks, and by maintaining relevant records relating to our processing activities and decision making.
- Your Information Asset Owner is the risk owner for any residual risk associated with your data processing and **must** sign off this DPIA.

Guidance links:

Accountability principle

What are the accountability and governance implications of AI?

Q14. Is your Information Asset Owner aware of your plans?

Choose an item.

Q15. Will you need to update our article 30 record of processing activities?

Choose an item.

Q16. If you are using a data processor, have you agreed, or will you be agreeing, a written contract with them?

Choose an item.

3.6. Individual Rights

Guidance Note:

- UK GDPR provides several rights to data subjects when their personal data is being processed.

- As some rights are not absolute, and only apply in limited circumstances, we may have grounds to refuse a specific request from an individual data subject. But you **must** be sure your new service or process can facilitate the exercise of these rights, and it should be technically feasible for us to action a request if required.
- The UK GDPR has additional rules to protect individuals if you are carrying out solely automated decision-making that has legal or similarly significant effects on them.

Guidance links:

A guide to individual rights

What is the impact of Article 22 of the UK GDPR on fairness?

Q17. Is there a means of providing the data subjects with access to the personal data being processed?

Choose an item.

Q18. Can inaccurate or incomplete personal data be updated on receipt of a request from a data subject?

Choose an item.

Q19. Can we restrict our processing of the personal data on receipt of a request from a data subject?

Choose an item.

Q20. Can we stop our processing of the personal data on receipt of a request from a data subject?

Choose an item.

Q21. Can we extract and transmit the personal data in a structured, commonly used and machine-readable format if requested by the data subject?

Choose an item.

Q22. Can we erase the personal data on receipt of a request from the data subject?

Choose an item.

Q23. Does your data processing involve any automated decision making that could have legal or similarly significant effects on your data subjects?

Choose an item.

If you've answered **yes**, please explain below the decision being made and outline how you will ensure individuals can request human intervention or challenge the decision.

4. Risk assessment

Guidance Note:

- You **must** use the table below to identify and assess risks to individuals. You can add as many rows to the table as you need.
- **Remember:** we have an **Averse** risk appetite towards compliance risks (see our Risk Management Policy and Appetite Statement for more information).
- You **must** identify measures to reduce the level of risk where possible.
- In the risk description column, you can select from common risks to individuals in the drop-down list provided. Alternatively, you can enter your own risk descriptions if preferred.
- **The drop-down list is not exhaustive**, and you must identify and assess risks within the context of your planned data processing. The AI and data protection risk toolkit may be useful to you if you're using AI.
- Mitigation measures can be existing, i.e. they're already in place and reduce the risk without any further action being needed. Or they're expected i.e. these are additional measures you intend to take before the data processing begins to further reduce risk.
- Use the risk scoring criteria in Appendix 1 to score your risks. You **must** score both the impact (I) and probability (P). The expected risk score total is the result of I multiplied by P.
- When considering probability, you should score based on all your mitigation measures having been implemented to get an *expected* risk score.

	Risk description	Response to Risk	Risk Mitigation	Expected Risk Score		
				Impact	Probability	Total
	Example: Access controls are not implemented correctly, and personal data is accessible to an unauthorised party.	Treat: this risk is being reduced by management action such as implementing controls or tackling the cause	Existing mitigation: We have checked that the system we intend to procure allows us to set access permissions for different users. Expected mitigation: We will appoint and train a system administrator who will be responsible for implementing access controls and monitoring access. The system administrator will also	3	1	3 - low

	Risk description	Response to Risk	Risk Mitigation	Expected Risk Score		
			audit the system periodically to review access permissions.			
1.	Choose an item.	**Choose an item.**	Existing mitigation: Expected mitigation:	**Choose an item.**	**Choose an item.**	**Choose an item.**
2.	Choose an item.	**Choose an item.**	Existing mitigation: Expected mitigation:	**Choose an item.**	**Choose an item.**	**Choose an item.**
3.	Choose an item.	**Choose an item.**	Existing mitigation: Expected mitigation:	**Choose an item.**	**Choose an item.**	**Choose an item.**

	Risk description	Response to Risk	Risk Mitigation	Expected Risk Score		
4.	Choose an item.	**Choose an item.**	Existing mitigation: Expected mitigation:	**Choose an item.**	**Choose an item.**	**Choose an item.**
5.	Choose an item.	**Choose an item.**	Existing mitigation: Expected mitigation:	**Choose an item.**	**Choose an item.**	**Choose an item.**

5. Consult the DPO

Guidance Note:

- Once you have completed all the sections above you **must** submit your DPIA for consideration by the DPIA Forum who will provide you with recommendations on behalf of our Data Protection Officer (DPO). The DPIA process outlines the next steps.
- Any recommendations from the DPOs team will be recorded below and your DPIA will then be returned to you. You **must** then record your response to each recommendation and then proceed with completing the rest of this template.

	Recommendation	Date	Project Team Response
1			**Choose an item.** Any comments: If rejecting DPO recommendations explain why:
2			**Choose an item.** Any comments: If rejecting DPO recommendations explain why:

	Recommendation	Date	Project Team Response
3			**Choose an item.** Any comments: If rejecting DPO recommendations explain why:
4			**Choose an item.** Any comments: If rejecting DPO recommendations explain why:

6. Integrate the DPIA outcomes

Guidance Note:

- Completing sections 1 to 5 of your DPIA will have helped you identify several key actions that you now **must** take to meet UK GDPR requirements and minimise risks to your data subjects. For example, you may now need to draft a privacy notice for your data subjects; or you could have risk mitigations that you need to go and implement.

- You **should** also consider whether any additional actions are required because of any recommendations you received from the DPOs team.

- Use the table below to list the actions you need to take and track your progress with implementation. Most actions will typically need to be completed **before** you can start your processing.

Action	Date for completion	Responsibility for Action	Completed Date

Action	Date for completion	Responsibility for Action	Completed Date

7. Expected residual risk and sign off by the IAO

Guidance notes:

- Summarise the expected residual risk below for the benefit of your IAO. This is any remaining risk **after** you implement all your mitigation measures and complete all actions. It is never possible to remove all risk so this section shouldn't be omitted or blank.

- If the expected residual risk remains high (i.e. red on the traffic light scoring in the Appendix) then you **must** consult the ICO as the regulator by following the process used by external organisations.

7.1 IAO sign off

Guidance Note:

- Your IAO owns the risks associated with your processing and they have final sign off on your plans. You **must** get your IAO to review the expected residual risk and confirm their acceptance of this risk before you proceed.

- Once your DPIA has been signed off it is complete. You should review it periodically or when there are any changes to your data processing.

IAO (name and role)	Date of sign off

8. DPIA change history

Guidance notes:

- You should track all significant changes to your DPIA by updating the table below.

Version	Date	Author	Change description
V0.1			First Draft

Appendix 1: Risk Assessment Criteria

The following criteria are aligned with our corporate risk assessment criteria.

Impact

Impact is the consequence of a risk to the rights and freedoms of individuals being realised. Factors to consider include the financial harm or emotional distress that can be expected to occur.

Impact	Scoring criteria
Very low (1)	No discernible impact on individuals.

Impact	Scoring criteria
Low (2)	Individuals may encounter a few minor inconveniences, which they will overcome without any problem (time spent re-entering information, annoyances, irritations, etc).
Medium (3)	Individuals may encounter significant inconveniences, which they will overcome despite a few difficulties (extra costs, denial of access to business services, fear, lack of understanding, stress, minor physical ailments, etc)
High (4)	Individuals may encounter significant consequences, which they should be able to overcome albeit with serious difficulties (misappropriation of funds, blacklisting by financial institutions, property damage, loss of employment, subpoena, worsening of health, etc).
Very high (5)	Individuals which may encounter significant, or even irreversible consequences, which they may not overcome (inability to work, long-term psychological or physical ailments, death, etc.).

Probability

Probability is the likelihood of a risk to the rights and freedoms of individuals being realised. Factors to consider include the expected frequency of occurrence, and the motivation and capability of threat

sources (e.g. does the threat require insider knowledge or significant technical resources to exploit any vulnerability).

Probability	Scoring criteria
Very low (1)	0-5% - extremely unlikely or improbable For example, the risk has not occurred before or is not expected to occur within the next three years.
Low (2)	6-20% - low but not improbable For example, the risk is expected to occur once a year.
Medium (3)	21-50% - fairly likely to occur For example, the risk is expected to occur several times a year.
High (4)	51-80% - more likely to occur than not For example, the risk is expected to occur once a month.
Very high (5)	81-100% - almost certainly will occur For example, the risk is expected to occur once a week.

Risk level

Risk level is a function of impact and probability and is represented by a RAG rating.

Probability / Impact	Very low (1)	Low (2)	Medium (3)	High (4)	Very high (5)
Very high (5)	Amber (5)	Amber (10)	Red (15)	Red (20)	Red (25)
High (4)	Green (4)	Amber (8)	Amber (12)	Red (16)	Red (20)
Medium (3)	Green (3)	Amber (6)	Amber (9)	Amber (12)	Red (15)
Low (2)	Green (2)	Green (4)	Amber (6)	Amber (8)	Amber (10)
Very low (1)	Green (1)	Green (2)	Green (3)	Green (4)	Amber (5)

Risk acceptance criteria

These criteria are guidelines only, and any risk treatment decisions should be made on a case-by-case basis. For example, it may be prudent to reduce a low risk because of legal and regulatory requirements.

Risk level	Acceptance criteria
Low (Green)	Within this range risks can be routinely accepted.
Medium (Amber)	Within this range risks can occasionally be accepted but shall be kept under regular review.

Risk level	Acceptance criteria
High (Red)	Within this range risks shall not be accepted, and immediate action is required to reduce, avoid or transfer the risk.

Template document control and change history (for Information Management Service use only)

Document Name	Data Protection Impact Assessment - template
Author or Owner (name and job title)	[redacted] Team Manager Information Management and Compliance Service
Department or Team	Information Management and Compliance Service
Document Status (draft, published or superseded)	Published
Version Number	V4.1
Release Date	07/10/2020
Approver (if applicable)	N/A
Review Date	31/01/2026
Distribution (internal or external)	Internal

Version	Date	Author	Change description
v0.1	01/06/2020	[redacted]	First draft
v1.0	07/10/2020	[redacted]	First release

Version	Date	Author	Change description
v1.1	07/01/2021	[redacted]	Amendment to guidance note page 2.
v1.2	18/03/2021	[redacted]	Addition of Privacy by design at the ICO (pages 2 and 3)
v1.3	24/06/2021	[redacted]	Section 3.0 Q13 amended. Removed request for link to security assessment.
v2.0	07/03/2022	[redacted]	Full document review. Simplified privacy by design explanation on page 3 and made minor format changes throughout. Guidance notes for 2.0 was updated and flow headings inserted to the text box. Next review date set to 31/1/2023.
v2.1	11/05/2022	[redacted]	Amended title of section 2 from 'data flows' to 'personal data lifecycle'
v2.2	26/10/2022	[redacted]	Guidance notes updated throughout following feedback from Project Management Office.
V3.0	16/01/2023	[redacted]	Annual review. Inclusion of further guidance notes to reflect feedback received from colleagues. Introduction of drop-down lists in sections 1.3, 1.4, 3.0

Version	Date	Author	Change description
			and 4.0. Addition of Q2 and Q14 in section 3.0. Removal of Appendix 2.
V3.1	31/05/2023	[redacted]	Change to reminder about risk appetite in section 4 – amended to averse after update to Risk Management Policy and Appetite statement.
V3.2	17/10/2023	[redacted]	Formatting changes to meet accessibility requirements
V3.3	03/01/2025	[redacted]	Links checked and amended where needed.
V4.0	22/01/2025	[redacted]	Annual review, no changes.
V4.1	27/06/2025	[redacted]	Template revised to incorporate AI focused section and additional AI guidance notes throughout.

Internal AI Use Policy

Foreword

It is imperative that we embrace the positive benefits that Artificial Intelligence (AI) can bring to our work at the ICO. I am delighted to see the launch of this policy to promote and support our responsible adoption and use of AI. I firmly believe that by using AI responsibly, we can enhance our decision-making processes, streamline operations, provide customer experiences that better meet the diverse needs of those we are here to serve and set a positive example to those we regulate.

I'm already an enthusiastic daily user of the initial AI capabilities we have made available. I'm routinely impressed by the time they have saved, the additional knowledge and insight they enable me to explore and, perhaps more importantly, the clear potential for the technology to improve and become ever more reliable at an astonishing pace.

As the UK's data protection regulator, it is vital that we are able to give those we regulate confidence that we are able to responsibly deploy the same technology they are also striving to use. This policy underscores our commitment to harnessing the power of AI to drive our organisation forward, while ensuring that we do so in a manner that is ethical, transparent, and aligned with our core values as well as our responsible position as a regulator.

As part of our adoption of AI we must acknowledge that our use of it comes with its own set of risks. It is crucial that we approach these risks with a proactive mindset, using our growing adoption of AI as an opportunity to best understand and mitigate them effectively. One of these risks is that we don't have the confidence to use the AI capabilities available to us because we aren't clear how to do that responsibly. This policy is intended to provide us with the practical guidance to do this.

Our journey towards AI adoption will not be without its hurdles, but it is a journey we must complete successfully. I am confident that by integrating the robust safeguards set out in this policy and continuously refining our approach, we can navigate the complexities of AI and unlock its full potential together.

I am excited about the possibilities that lie ahead and look forward to the positive impact that our responsible use of AI will have on our work together and for our relevance and impact as a regulator.

Paul Arnold MBE

Chief Executive

August 2025

Key messages

The main objective of this policy is to ensure any Artificial Intelligence (AI) used at the ICO is governed in a way that maximises benefits and minimises or mitigates related risks. The policy focuses on:

- how AI should and should not be used in the ICO;
- how internal AI use at the ICO should be considered, managed and governed; and
- what is meant by AI and how it might manifest/appear in the ICO's organisation.

Does this policy relate to me?

This policy should be read and understood by all ICO employees, as well as colleagues in a temporary role, on secondment. This also includes third parties working with ICO on a contractual basis (unless otherwise specified in the contract) or otherwise (each be referred to individually as a "ICO AI User", and, altogether, the "ICO AI Users").

Section 4 of this policy is of particular importance, as it details the requirements which apply to all ICO AI Users when using AI tools.

Contents

1. Introduction

1.1. Digital technologies such as AI can serve to increase the impact derived from the resources the ICO invests in and can improve the value offered to stakeholders. It is expected that the adoption of AI across the public sector, including the ICO will increase. As we explore further technological advancements, we will leverage the power of artificial intelligence to enhance regulatory compliance and streamline operations, driving cost down and increasing productivity.

1.2. As stated in the ICO's Enterprise Data Strategy:

> *"Whilst we will seek to fully understand any ethical, security or legal implications before commencing any development work, we will remain curious and daring with our stance. We don't want to lose the opportunity to fully embrace the capabilities offered by these new emerging technologies that could benefit how we serve our customers."*

This internal AI use policy aims to provide a way to find the appropriate balance between being curious and daring, and the goal of being responsible about what we do with technology and how we govern it, as well as being prepared for potential future developments in both technology and regulation.

1.3. This policy applies to AI in all its forms including bespoke applications and solutions, or where it is embedded in software-as-a-service platforms and services, or pilot AI projects, or those deployed in production. This policy applies to AI developed by a third party and to AI solutions developed in-house.

1.4. This policy aligns with the following documents and will be updated as guidance changes:

- Artificial intelligence | ICO
- ICO guidance on AI and Data Protection
- Information Commissioner's Office response to the consultation series on generative AI | ICO
- Automated decision-making and profiling | ICO
- AI Playbook for the UK Government
- ISO/IEC 42001:2023 - AI Management System
- Knowledge builder internal information (*internal link*)

Back to the top

2. How we define AI for the purpose of this policy

2.1. For the purpose of this policy, AI is an umbrella term for a range of technologies and approaches used to mimic human intelligence to solve complex tasks. For example:

- planning and optimisation (e.g. scheduling tasks to minimise downtime of finite resources)
- classification and prediction (e.g. filtering emails and content)
- interpreting and generating information and content such as video, imagery, audio and text (e.g. summarising the contents of documents and using chatbots to assist with research).

2.2. AI can perform these tasks by modelling and recognising patterns in data. Data can be internal to the organisation, taken from external sources or used in combination.

Back to the top

3. Internal AI use principles

3.1. ICO AI Users ("we") use AI collaboratively, thoughtfully and transparently, in line with current regulatory requirements and policy considerations, and the requirements outlined in section 4 of this policy.

3.2. For instance, we explore opportunities and needs for the internal use of specific AI tools that are appropriate, proportionate and sustainable. We include a range of perspectives in our decision-making around internal use of AI. We empower and dignify our people with the knowledge and skills needed to drive, support and challenge internal AI use.

3.3. In line with applicable regulatory requirements, we are transparent about our use of AI, including documenting information about our AI use and necessary risk assessments.

3.4. As with any technology, AI presents opportunities and risks. AI can power and accelerate automation and decision support but can present risks, for example, around inaccurate information, or bias and discrimination.

Back to the top

4. Policy requirements

4.1. Requirements when using AI at the ICO

This section covers the requirements that all users of AI at the ICO should follow.

4.1.1. You should only use AI that has been approved by the ICO for internal use (following the appropriate ICO governance process). You should only use ICO approved devices to access AI tools and systems for corporate work.

4.1.2. You should only use AI tools in a way which is consistent with published guidance and training on the use of the tool deployed by the ICO.

4.1.3. You should be transparent about your use of AI as appropriate and proportionate. If you have substantially edited or thoroughly reviewed your output, you do not need to state where you have used AI to draft, summarise or structure your documents, for example, when you summarise a guidance document or edit a report or an email. Where you use content that is substantially or wholly AI-generated, this should be marked as generated by AI (for example an image or video of an individual or song / poem). All outputs from Chatbots and similar tools, where individual outputs cannot be checked during operation, should be annotated as generated by AI.

4.1.4. You should ensure that all AI outputs are reviewed by a human reviewer, unless agreed otherwise by the appropriate approval body (Data, AI and Automation programme Board, DAA), taking into account relevant risks and impacts on users, data subjects or the quality of ICO's own outputs. For example, manual review of all outputs may not be necessary, where the impact is low, or other safeguards have been identified, documented and adopted. When reviewing AI generated outputs, you should ensure that the output is accurate, amending where necessary.

4.1.5. You should only use AI in activity involving the use or processing of any personal, sensitive or confidential information when permitted by the ICO (in line with ICO's appropriate internal governance processes and policies e.g. privacy by design procedures, security assessments and Architecture Design Authority approval). Where the ICO has approved use of AI which includes the processing of personal or sensitive information, data protection compliance is paramount. You should ensure that any processing meets all of the ICO's obligations. This includes the requirement to comply with the principle of data minimisation (i.e. only using personal data that is limited to what is necessary for your purposes). Please contact the Information Management team, Cyber Security team and the Architecture Design Authority if you require further information.

4.1.6. Advice that is typically given by licensed professionals, like accountants or lawyers, should only be obtained by the appropriate licensed professional and not AI. (This should not

prevent AI being used to assist professionals in their work provided that the work is validated before sharing).

4.1.7. You should only use AI for solely automated decision-making (ADM) when this is appropriate, i.e. when there would be no legal effect or similarly significant effect on an individual or group. If you think you may use AI for making such a decision, you must consult ICO's guidance on ADM and Legal Services colleagues as necessary, to ensure you are meeting regulatory requirements.

4.1.8. You should check AI-generated contents before publishing to make sure you are not infringing the intellectual property rights of a third party – if in doubt you should speak to the Contracts and Compliance team in Legal Services.

4.1.9. You should consider the appropriate security classification of inputs into and outputs from AI, taking account of the ICO's legal obligations as a public sector body.

4.1.10. You should not knowingly use AI in a way that causes, or may cause, significant risk of harm to individuals, groups or the reputation of the ICO, or breach any regulatory requirements.

4.1.11. You should flag any concerns, incidents or questions relating to internal AI use at the earliest opportunity to your manager and/or the system owner.

4.2. AI training and awareness

General AI literacy

4.2.1. The ICO's senior leadership team should ensure all staff have access to high-quality general AI literacy training sufficient to enable them to identify and specify potential AI opportunities, to appreciate the potential risks of AI use and to understand the importance of transparency and explainability in AI-enabled systems.

Specific training on an AI system

4.2.2. All users of a planned AI deployment should be given relevant training, prior to product deployment, covering its scope

(appropriate use), limitations including legal requirements, how to use it effectively and efficiently, and how to provide feedback and report incidents.

Back to the top

5. Considerations you should give when using, procuring or developing AI-powered tools:

This section covers further requirements for anyone introducing AI powered functionality into the ICO.

5.1. Accountability, decision-making and governance, particularly in the procurement context

AI specifications for approval

5.1.1. All AI solutions should only be developed, procured or deployed at the ICO after:

- Proper consideration of alternative options, through-life costs and the benefits of data-driven, flexible, adaptive approaches (using the use case specification templates provided - see Annex A and Annex B); and
- approval by the appropriate board or function(s).
- Until the formal approval process is in place, staff should engage with the relevant product owner to ensure new AI use cases have been appropriately evaluated, risk assessed and documented.

5.1.2. To allow for consistency in prioritisation, and effective challenge and support, ICO staff involved in the procurement of AI tools should build appropriate time into the planning and preparatory stages of any procurement where AI comprises all or part of the solution. ICO staff should also be mindful of the possibility that AI tools may form ancillary parts of other products or services used by the ICO and should consider and account for this as part of their procurement planning.

5.2. Proportionality

5.2.1. To enable rapid approval and deployment of AI that presents minimal or low risk, a fast-track approval route can be used under the discretion of the relevant board or function. This should only be used where:

- The necessary Data Protection and Equality Impact assessments have been completed and have documented that the risks are assessed to be low/minimal; and
- The benefits and costs of a similar use case have been assessed and demonstrated at the ICO previously.

5.3. Logging decisions around AI

5.3.1. All decisions around AI governance should be logged by the appropriate governance body (currently DAA) for continuous improvement and auditing purposes.

Risk assessment

5.3.2. Every AI development/deployment should have a risk assessment e.g. part of the project risk register that references the internal AI risk framework.

5.4. Accountability for AI governance

5.4.1. Accountability for AI governance across the ICO should be assigned to the appropriate role or body eg Data AI and Automation board, along with ownership of the identified AI risks, and documented in relevant decision logs.

5.5. Logging AI available to staff

5.5.1. An AI inventory should be maintained as part of the ICO's service catalogue to log the AI functionality available to the ICO AI Users to support auditing and productivity.

5.6. Impact assessment, fairness and explainability

Data protection

5.6.1. You should consider data protection as paramount when considering the use of AI at the ICO.

5.6.2. This internal AI use policy does not replace or overwrite regulatory requirements. Any AI initiative should comply with applicable law, including data protection legislation, guidance and other applicable ICO policies. Specifically, every AI initiative involving personal data should be subject to the ICO's, data protection by design and default (*internal link*) approach. Some forms of AI may transform data in ways not initially anticipated, if in any doubt, ICO AI users should consult Information Management and/or Legal Service as necessary.

Early consideration of fairness and explainability

5.6.3. Use case specifications (see Annex B) should include consideration of those individuals affected by the internal use of AI, of ensuring fairness and of providing the necessary/appropriate explainability in the context of the AI's functionality and potential impact. Every AI initiative should be subject to an equality impact assessment (EQIA).

Impact on other ICO staff members

5.6.4. Any concerns about potential impacts on performance, productivity and equal access on the roles of the ICO AI Users should be discussed with the People Services business partner representative for the directorate at the first available opportunity.

5.7. Safety, security and robustness

5.7.1. The development, procurement and deployment of any AI-powered solution must conform with the relevant Architectural principles, Enterprise Architecture strategy and Technical Reference Model (form available on IT Self-Service Portal – *internal link*) and policies for cyber security and data protection (procedures and forms on IRIS – *internal link*).

5.8. AI verification and validation

5.8.1. No AI solution or functionality should be deployed and made available to users without it having passed a documented verification and validation phase and has been shown to work within specified safety and performance requirements.

5.9. AI performance monitoring

5.9.1. The deployment of any AI internally should be monitored in terms of compliance, task performance, fairness and usage, and cost-effectiveness. Monitoring of the data is the responsibility of the Product Owner (colleagues responsible for the major products deployed at the ICO e.g. Microsoft products, Workday etc)

5.10. Transparency and documentation

Technical documentation:

5.10.1. All AI functionality developed for and/or adopted by the ICO must have corresponding technical documentation available to the ICO including the following elements:

- Architectural description, including interfacing systems.
- Detailed description of the AI elements of any system or product including references to pre-trained or third-party elements, general logic of the AI functionality, and significant design choices/assumptions.
- Detailed description of the data sets used for training and testing including labelling procedures where relevant.
- Location of data, both where it is stored or moved, including for cloud-based solutions.
- Information on performance metrics.
- Measures for human oversight, cyber security and validation.
- Outcomes from previous validation.

- Specific risks identified and mitigations.
- Intended use (and any limitations or restrictions).

Transparency of AI outputs

Internal transparency:

5.10.2. All ICO staff should be able to access information about AI-powered products, functionality and developments, including pilots.

5.10.3. ICO AI users should have the chance to raise opportunities for AI use in a timely manner through the appropriate channel.

External transparency:

5.10.4. All AI products or use which either interact directly with the public or have a significant influence on a decision-making process should be logged using the Algorithmic Transparency Recording Standard ('ATRS'). For any ATRS-related issues or questions, please contact the Data and Information Management team.

5.11. Security classification of AI inputs and outputs

5.11.1. Where appropriate, an Information Asset Owner should be assigned to the AI tool/system in accordance with existing ICO policies and Information Governance Roles and Responsibilities Guidance.

5.12. Feedback loop and change management

Feedback and reporting mechanisms

5.12.1. For any AI deployment, Product Owners should ensure there are mechanisms for ICO AI Users to:

- provide feedback on AI functionality e.g. on what works well and not so well, and level of human intervention required.

- report any incident, issue or concern related to an AI deployment (urgent issues should be escalated by the responsible officer to the appropriate board or function (DAA) to be actioned promptly).
- be informed of any changes to the functionality or new or emerging risks.
- All reports should be regularly reviewed/audited by the relevant responsible officer.

5.13. Contestability and redress

5.13.1. For any AI deployment that is involved in decision making there must be a mechanism to enable a person to contest AI outcomes or decisions, and for redress if required. This applies whether the AI is public or ICO AI User-facing. The deployers must ensure that the mechanism is articulated and communicated as part of the transparency information.

Mechanism to stop or pause:

5.13.2. For any AI deployment, the Product Owner must ensure that there is a mechanism to pause or stop the AI elements of functionality in the event of an issue as well as a way to inform those impacted in a timely way. This mechanism must be specified in the deployment documentation.

5.14. Change management

5.14.1. Any material changes to an AI initiative or solution (such as change in scope, impact, user population, data source) must function as a new use case.

5.15. Re-evaluation and retirement

5.15.1. Any AI deployment should be reviewed on a regular basis to evaluate the costs and benefits to the ICO. Any AI that is no longer to be used should be formally retired and withdrawn from use. The Stakeholder/Sponsor is responsible for overseeing this process. The Sponsor is likely a director or an SLT member who

through their position as sponsor, play a crucial role in ensuring the success of a project or programme and its alignment with the ICO's strategic objective.

Back to the top

6. Policy compliance

6.1. The DAA Programme Board should verify compliance with this policy via various methods including, but not limited to, business reporting and internal audits of controls and processes.

6.2. In the event of a conflict, any applicable law takes precedence over the ICO’s internal policies.

6.3. Any exceptions to this policy should be approved by the DAA Programme Board in advance.

6.4 Improper use may expose you to civil or criminal liability under the applicable law. In addition, any ICO AI User, who is also an employee (or to whom the employee legal regime applies), found to have violated this policy may be subject to disciplinary action, up to and including termination of employment.

Back to the top

Feedback on this document

If you have any feedback on this document, please click this link (*internal link*) to provide it.

Back to the top

Document Control

Version number: 1.2

Status: Published

Department/Team: DDaT/EDS

Relevant policies:

Personal data sharing policy (*internal link*)

Data Protection Policy

Information Management Policy

Acceptable Use Policy (*internal link*)

Distribution: Internal

Owner: Director of Data

Consultees: Digital, Data and Technology (DDaT), Data, AI and Automation Programme Board (DAA), Information Management and Compliance, Legal Services Team, AI Policy Team

Approved by: Data, AI and Automation Board

Application date: August 2025

Review date: August 2026

Security classification: OFFICIAL – ORG USE

Version history

Version	Changes Made	Date	Made by
1.0	First published	05/08/2025	Iman Elmehdawy
1.1	Minor updates to tone and clarity	20/8/2025	Iman Elmehdawy

1.2	Clarification and addition of examples	18/11/2025	Iman Elmehdawy

Back to the top

Annex A – AI Screener

Introduction

This template is designed to elicit some high-level information that would help provide an understanding of the main features of an AI use case so it can be challenged, supported and prioritised as appropriate before more work is done on more detailed specification.

Title

[Provide a short, useful working title]

Product Owner

[The individual responsible for development, deployment and maintenance of the AI tool]

Source of use case

[Please indicate if this use case stems from: a newly identified business opportunity or need, availability of a new technology or tool, availability of a new functionality within an existing system or tool used in the organisation]

Problem statement/context

[Provide a short description of the current situation (eg manual processing, need for insight from data) and its impact]

Potential role for AI

[Be clear on the job the AI would be doing ie:

- *Summarisation*
- *Speech or audio analysis (including transcription)*
- *Image or video analysis*
- *Optimisation*
- *Prediction/Classification/Filtering*
- *Information/Data Retrieval*
- *Question and Answering/Chatbot*
- *Other*

If the AI will need to perform multiple tasks, list them.]

Potential business value (including baseline) – please quantify where possible

[State which of the following value drivers/types of benefit apply to this use case:

- *Increased revenue/income*
- *Increased productivity or efficiency, or decreased costs*
- *Improved stakeholder engagement/sentiment*
- *Reduction in risk/improvement in safeguarding for vulnerable people*

Availability of relevant data sources

[State the types of data likely to be needed for this use case eg first-party data (corporate documents and records), third-party data (such as firmographics, data from Companies House or other public sector body). State whether or not the ICO already has access to the data and permission/consent to use it, or if it needs to be obtained.]

Potential ethical issues

[Consider if there are any known ethical issues (e.g. equality)]

Legal considerations

[Consider legal issues (eg data protection law, Freedom of Information law, Equality Act 2010 and equality duties) that would need to be considered and addressed]

Stakeholder/sponsor

[Who has a significant stake in this use case being successful ie who is accountable or responsible for the process or system being impacted?]

Annex B – Full AI use case specification

Introduction

This template is designed to elicit a range of information that would be helpful in understanding the goals of an AI use case, the benefits it could realise, its potential connections and dependencies with other projects or technologies, its limitations and risks.

This should be completed after support has been given to an earlier AI Screener. Relevant information can be copied from the AI Screener.

It may not be possible to respond to all the questions early on so please submit all you can so the uncertainties and potential options can be assessed.

Title

[Provide a short, useful working title – ideally the same as the corresponding AI Screener]

Product Owner

[The individual responsible for development, deployment and maintenance of AI tool]

Problem statement/context

[Provide a short description of the current situation (eg manual processing, need for insight from data) and its impact]

Those impacted by the current problem

[State the roles or groups impacted by the current situation eg fee payers, colleagues in specific roles]

Options and proposed approach

[Briefly describe the overall approach eg 'The proposal is to automate parts of the X process', or 'The proposal is to use AI to filter out…'. Describe the likely elements of manual intervention. Please state what has already been tried in addressing this opportunity.]

Scope of proposed approach

[Here, describe how the approach will be bounded eg by process, by stakeholder type, by task]

Hypothesis/role of AI

[Be clear on the job the AI would be doing ie:

- *Summarisation*
- *Transcription*
- *Optimisation*
- *Prediction/Classification/Filtering*
- *Information/Data Retrieval*
- *Image analysis*
- *Question and Answering/chatbot*
- *Other*

If the AI will need to perform multiple tasks, list them.]

AI-specific requirements

Speed of response *[Consider how quickly the AI would need to respond to requests or incoming data eg in real-time, near-real-time, within x minutes]*	
Explainability *[Consider obligations to data subjects for explainability in the case of automated decision making Consider what others (including but not limited to end users) will need to understand about the AI in the wider solution]*	
Accuracy *[Consider whether the AI outputs need to be correct/accurate (vs more creative) and the impact of inaccurate outputs]*	
Need for labelled data *[Will labelled data be required for model training and testing?*	
Adaption mode *[Will the model(s) need to adapt over time and, if so, with what frequency will this be required eg daily, weekly, after every user interaction?]*	
Incorporation of existing knowledge *[Consider what existing knowledge or expertise should be incorporated into the solution eg market*	

segments, company structure, report structures]	

Deliverables

[State what the desired outcomes of the work are eg working AI models, findings of how users responded to a pilot]

Acceptance criteria/definition of done

[Be clear on the types of metrics, outcomes or outputs that would need to be provided to judge the project complete eg complete data and analysis on pilot results, working model shown to work with real data]

Out-of-scope

[Be clear on what would not be covered by the proposed project eg tasks that would not be covered by the proposed AI]

Related legal or compliance requirement

[State if the approach is to address a specific legal or compliance requirement]

Potential business value (including baseline) – please quantify as much as possible

[State which of the following value drivers/types of benefit apply to this use case:

- *Increased revenue/income*
- *Increased productivity or efficiency, or decreased costs*
- *Improved stakeholder engagement/sentiment*
- *Reduction in risk/improvement in safeguarding for vulnerable people*

State what baseline will be used to assess the benefit(s) of the AI and how any benefits will be measured.]

Data requirements

[State the types of data likely to be needed for this use case eg first-party data (corporate documents and records), third-party data (such as firmographics, data from Companies House or other public sector body). State whether or not the ICO already has access to the data and permission/consent to use it. What is known about its quality ie completeness and correctness?]

Team requirements (expertise, skills, knowledge)

[Consider if specific expertise, skills or knowledge is required eg data domain expertise to ensure the data is interpreted correctly, data science/AI, knowledge of particular end users and their context. Provide any detail on how that expertise, skills and knowledge can be accessed by the organisation.]

Technology/tooling requirements

[Consider what technologies or tooling might be required to access the data, and to develop and run any AI]

Known/potential main risks, dependencies and constraints

Costs to consider and any estimates

[Consider source of through-life costs such as licenses to specific technology or data sources, investment into data labelling, number of roles to validate AI outputs, additional data storage, etc]

Ethical considerations

[Consider where the impacts may be different across Protected Characteristics, the impacts of outputs that may cause distress for users, the involvement of vulnerable users, and any possible unintended consequences]

Legal considerations

[Consider legal impacts (including the data protection law, Equality Act 2010 and equality duties, the Freedom of Information Act) eg what should be the classification for inputs into and outputs from the AI in the event of a FOI request?]

Stakeholder/sponsor

[Who has a significant stake in this use case being successful ie who is accountable or responsible for the process or system being impacted].

Annex C - What do you need to know about AI?

As a staff member you will engage with AI either as a user or as the person responsible for developing, deploying and maintaining AI systems. Below is some useful information and issues you may want to know about.

Machine learning and adaptability

Many forms of AI learn over time – that is, their performance improves with access to more data which updates the AI model (eg as new services become available, or as cyber threats evolve). This is known as 'machine learning'.

From automation to autonomy and the human-in-the-loop

Automation and autonomy exist on a spectrum from manual control in simple environments at one end to systems that can make a range of decisions operating in complex and dynamic environments with no human intervention at the other. AI functionality, such as language understanding, image analysis, machine learning and reasoning, can enable autonomy.

Automated Decision Making (ADM) sits on this spectrum. The role of a human-in-the-loop should be carefully considered to mitigate some of the risks associated with AI outputs used in decision making. The specific role of the human-in-the-loop should take into account the impact of AI outcomes on any individual, and the complexity of the task.

Profiling and automated decision making can be very useful for organisations and also benefit individuals in many sectors, including healthcare, education, financial services and marketing. They can lead to quicker and more consistent decisions, particularly in cases where a very large volume of data needs to be analysed and decisions made very quickly.

Any automated decision-making and profiling should be carried on in accordance with the ICO's published guidance on automated decision-making and profiling. If you have a question in relation to ADM that you cannot answer by referring to the ICO's current guidance please reach out to the AI policy team or the AI compliance team.

Agentic AI (or AI Agents)

Agentic AI generally refers to autonomous forms of AI that can perform tasks or communicate with each other to solve problems. It is often focused on agents performing customer service-related tasks, such as chatbots, but agentic AI is more autonomous and can have much broader applications.

Typically, agents are designed to specialise in a particular task, but future developments may lead to agents that can undertake more generalised and complex tasks requiring collaboration or negotiation with others – agent or human.

Technology vendors are increasingly marketing their AI functionality as agentic, so it is important to clarify the level of control and intervention needed or available to humans working with 'agents'. Please consult the AI policy team or the AI compliance team for relevant questions.

Generative AI

Generative AI is a branch of AI focused on generating content on demand, such as video, images, audio, text and software code, in response to natural language prompts. Often generative AI is powered by adaptive 'foundation models' (which is currently one of ICO's focus areas). Foundation models are vast in scale, complex in structure and trained using huge data sets.

Generative AI is powerful and general purpose but there are risks associated with it including hallucinations, bias, intellectual property rights infringement, data protection non-compliance and the potential for security breaches. In addition, foundational models are resource-hungry to train and use at scale.

Opportunities

AI can operate at high speed and at large scale in a way that is beyond human capacity. This makes it useful for data analysis and can support human decisions in a range of settings.

There are many opportunities for AI to provide benefits to the ICO. AI can be used to enable efficiencies through automation (eg template document

generation and email classification) and to generate insights from data that can be used to make strategies and actions more effective (eg analysing which organisations need to pay the data protection fee but do not).

The ICO encourages all staff to bring forward ideas for applying AI to ICO needs and opportunities.

Risks

AI can present significant risks to an organisation and to individuals due to its dependence on high quality data, its complexity and its adaptability together with the potential impact on individuals of data analysis and decision-making/profiling using AI.

AI may derive patterns that, if not sense-checked, may lead to AI outputs that are inaccurate (due to poor implementation or because it cannot account for the complete picture), or potentially harmful (eg biased, discriminatory, insensitive).

AI may also process personal data in ways that are difficult for data subjects to understand, make it challenging to exercise their individual rights, and may require their personal data to be transferred to other countries for processing. Indeed, AI has been an area of focus for the ICO for a number of years and the ICO has a range of guidance products and resources to address the broader risks and data protection implications of AI.

Additionally, AI tools procured without a robust cost benefit analysis and use-case are unlikely to deliver a positive impact and may expose the ICO to commercial and legal risk. These sorts of considerations are best addressed at the outset of a project or procurement, so it is essential to ensure that your requirements and objectives are well-understood prior to procuring any AI tools, and that you consult with colleagues as set out in our internal AI use governance process and structure.

Assessing opportunity and risk

It is important to assess the potential and actual benefits of introducing AI into the ICO’s processes to ensure any additional cost and complexity

is outweighed by gains, for example in terms of increased efficiency, productivity, revenue or customer satisfaction.

There may be specific reputational and enforcement risks to the ICO related to the use of AI, given its role as a digital technologies' regulator. These need to be accounted for in decision-making in relation to internal AI use in relevant use cases.

Forms of AI

AI can take various forms, including:

- Standalone applications, or embedded in a broader system or solution.
- Bespoke to the organisation's needs, generic for multiple organisations, or generic with the ability to tune parameters.
- Owned and maintained by the organisation (or a third party on its behalf) within the organisation's IT infrastructure or hosted and maintained elsewhere as a cloud-based service.

When selecting an AI solution, you should weigh the benefits, costs, and risks. Bespoke AI suits unique needs and offers transparency but requires significant investment. Cloud-hosted AI is ideal for common requirements, reducing IT costs and leveraging large-scale models. It is important to note that solutions may integrate AI and non-AI components.

Lifecycle

Like any software project or capability, AI has a lifecycle including the following phases: initiation; design and development; verification and validation; deployment; operation and monitoring; re-evaluation and retirement.

Specifically for AI, design and development include obtaining data, at the necessary quality and quantity, to train and/or test AI models. The quality and quantity of the data used is a major determinant of AI performance so must be considered for any AI initiative.

It is important to understand how the data has been derived and what it represents - you should consider consulting the relevant experts at the

ICO, the Artificial intelligence resources for the public sector - GOV.UK and the Guidance on AI and data protection | ICO. These considerations are also relevant in relation to fairness (in terms of how the data is obtained and used) and statistical accuracy of the data, in relation to which you should pay regard to official ICO guidance.

Back to the top