ico.
Information Commissioner's Office

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

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 to identify an approach to internal use of specific AI tools, which must be 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.

Back to the top

4. Policy requirements

4.1. Requirements when using AI at the ICO

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. This includes being clear where your work includes AI generated outputs, marking them to be clear that is the case (e.g. providing a clear statement linked to a specific portion of text or image using a footnote to specify 'This text/image was generated using 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. Other considerations you should give when using, procuring or developing AI tools:

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).

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 AI Users 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 a data protection impact assessment (DPIA). 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) and policies for cyber security and data protection (procedures and forms on IRIS).

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 information management.

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 to provide it.

Back to the top

Document Control

Version number: 1.1

Status: Published **Department/Team:**

DDaT/EDS **Relevant policies:**

Acceptable Use Policy Data

Protection Policy

Information Management Policy Personal

Data Sharing Policy

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

Back to the top

Annex A – AI Screener

Introduction

This template is designed to elicit some high-level information that would be 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 had 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 needs 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 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) Harm

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 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 that need to be paying the data protection fee but are not).

The 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 in (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 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