

BBC Corporate Policy

Responsible AI Policy

Effective from: 1 December 2023

Policy owner: Director of Product Data

Owner's department: Product Data

Version control: v1.2

Please note that many of the links in this document are to BBC internal documents and processes, and as such may not be available to everyone.

Contents

1. Why this Policy matters	3
2. Who this Policy applies to	3
3. BBC AI Principles	3
We will act in the best interests of the public.....	3
We will prioritise talent and creativity.....	4
We will be open and transparent.....	4
4. The essential things you must do or not do.....	4
Risk awareness and internal transparency	4
Risk management	5
Using AI tools.....	5
Ongoing human oversight	6
Iteration and evolution	6
5. Where to go for help or an exception.....	6
Document Control.....	6
Appendix: Key AI Terms and Definitions.....	8

Responsible AI Policy

1. Why this Policy matters

Artificial Intelligence (AI) offers significant opportunities, but also poses risks to our values, audience trust, copyright, creators' rights, privacy, security, and compliance.

This Policy sets out our core principles and requirements to make sure that use, development and procurement of AI manages these risks and reflects our values

2. Who this Policy applies to

This Policy applies to you if:

1. You work for, or represent, the BBC or one of its subsidiaries, and
2. You use, develop or procure AI (or AI-powered tools, content or services) for BBC purposes, whether on a BBC device or a personal device.

This includes, but is not limited to:

- employees, full-time or part-time
- those on a fixed term or temporary contract
- freelancers, contractors or consultants
- volunteers or interns

If you work in Public Service Broadcasting outside the UK, local variations may apply in line with relevant laws and regulations. Check with your local management.

If you work in Studios, refer to the Studios Policy Hub in case of policy variations specific to Studios.

3. BBC AI Principles

The following principles guide us in using AI responsibly.

We will act in the best interests of the public

BBC Values: Our use of AI will reflect the public service mission and values of our organisation: upholding trust; respect and inclusivity; boosting creativity; putting audiences at the heart of everything we do; being accountable and delivering quality; and working as one BBC.

BBC Editorial Values: When we use AI to create, present or distribute content we will make sure that this complies with the BBC's editorial values, guidance and guidelines.

Fairness: Our use of AI will be fair, equitable and inclusive for our audiences and employees. We will make sure that everyone has the opportunity to access and benefit from the range of content and services we provide and use.

Security & Robustness: We will use AI in a way that is secure, robust and safe. We will assess its accuracy and reliability, monitor its performance, document its use, and maintain continuity of service.

We will prioritise talent and creativity

Respecting Rights: Creators, contributors and suppliers play a vital role in our industry. We will consider the rights of creators, artists, contributors and rights holders when using AI. We will make sure our use of AI respects the data protection and privacy rights of individuals.

Human Creativity: We will use AI to support and empower human insight, talent and creativity.

We will be open and transparent

Transparency & Clear Explanations: We will be clear with audiences and employees about where we use AI and what data we collect. We will make sure they can understand why we use it, how it works and how it affects them.

Accountability: We will make sure that there is proper supervision of and clear accountability for our use of AI, including our use of content or services provided by others that involve AI.

Human oversight: Our use of AI will be accompanied by effective and informed human oversight, including over the content or data involved as inputs and outputs.

4. The essential things you must do or not do

These high-level essential mandatory requirements are the most important for you to understand and follow to meet the Policy objectives. In addition to these, there are detailed requirements and best practices outlined in supporting procedures and guidance, linked below. Read this information when you need to understand the detail.

If you have any concerns or conflicting priorities that prevent you following essential requirements, please discuss them with one of the people listed in the section under 'Where to go for help and exceptions'. Without an approved exception, breaches may put the BBC at risk and may result in disciplinary action, up to and including dismissal, or termination of our relationship with you.

Risk awareness and internal transparency

1. For any new use of AI, you must follow the AI Checklist or relevant Tool Guide, which will help you determine next steps relative to risk level.

If you are an employee, you must complete the required AI training modules if you are using, developing or procuring AI; or overseeing others; or engaging others outside the BBC where AI use is likely (e.g. content creators).

Some [commonly used AI tools have specific guides](#), which include when you must seek advice and [what training is required](#). For all other uses of AI, the [AI Checklist](#) will tell you when you must seek advice, and [what training is required](#). You can find more detail on how to get advice on the [Responsible AI Hub](#). You may be asked to complete a [Use of AI form](#) to gather more details about your use case. This will help determine whether more detailed advice or further risk assessments are required.

Ownership and accountability

2. You must define and document who is accountable for the use of AI, owning and accepting the identified risks across the use case lifecycle. They must sign off on the risk prior to implementing your use of AI.

There must be a senior Editorial figure responsible for Editorial risks across the lifecycle if your use of AI has significant editorial impact.

If you are using an AI tool with a specific guide, in line with what that guide allows, accountability has already been defined at the Divisional level. In these cases, you do not need to define separate accountability.

For all other cases, define risk owners at appropriate seniority in line with the level of risks, advised by your risk assessment, if applicable. If you are a Team Lead or Manager, make sure you are aware of and effectively managing use of AI by your team. Consult with your [divisional AI representative](#) for advice on risk ownership where needed.

It is the responsibility of BBC employees (and freelancers) to make sure that any content, services or tools they commission or procure from others are delivered and used in line with this Policy.

When commissioning content from external parties or using external services, make sure you are aware of potential AI use in the generation of this content or provision of the service. [Further guidance](#) is available for those working with external parties who may be using AI for BBC purposes.

Risk management

3. You must apply appropriate risk mitigations to your use case. These should be based on the outcome of your discussions with Responsible AI, the BBC AI Handbook, and any risk assessments required.

You may find further information in the [BBC AI Handbook](#) on appropriate mitigations. These may be included in a formal risk assessment or guide you if a risk assessment is not required and you are self-managing risk. Where directed by the Responsible AI team, provide information on your use case to support a detailed [AI Impact Assessment](#) and Impact Report to help you identify and implement appropriate mitigations.

AI governance involves multiple risk areas and different use cases may require different risk assessments, tracked on the Responsible AI team's use case log. For example:

- If your use of AI is directly for BBC audiences or has a material impact on News, Current Affairs or Factual content for BBC audiences, consult with Editorial Policy to help apply the requirements in the [Editorial Guidance on AI](#) and the AI [Transparency Advice](#) for how to disclose your use of AI to audiences.
- If your use of AI processes personal data, or if automated decision-making could impact individuals, complete a [Data Protection Impact Assessment](#) in line with the Data Protection Corporate Policy.

Using AI tools

4. You must only use external AI tools, software or models, including free online tools, which have been authorised for use at the BBC. New AI features in existing tools must also be authorised.

Note that approval of the tool is NOT default approval for AI use cases.

Use the existing [software authorisation](#) or procurement process. Material investments in technology are approved through the [Technology Investment Committee](#).

You may read more about [authorised AI tools](#).

Ongoing human oversight

5. You must make sure there is effective human oversight of any use of AI. This can include direct oversight of the Inputs and Outputs involved, and/or measuring and monitoring appropriate data.

Regularly measure and monitor your risks and the quality of Inputs and Outputs, as well as the effectiveness of mitigations.

Iteration and evolution

6. You must regularly review your AI use case to account for new and emerging risks.

If you have a detailed AI Impact Assessment, review your assessment in line with the schedule agreed as part of the assessment process.

5. Where to go for help or an exception

For general advice on this Policy or to suggest improvements, contact the First Point of Contact listed below. Alternatively:

Responsible AI Team: dedicated to supporting the responsible use of AI. They can help you consider and manage your use of AI and advise you on guides, processes and governance. The [Responsible AI Hub](#) contains everything you need to use AI responsibly.

AI Risk Advisory Group: co-ordinates advisory, compliance, and legal experts in all types of AI. They can support you in identifying and managing risks in your use of AI. This group contains representatives from Commercial Rights and Business Affairs, Legal, Information Security, Data Protection Office, Technology, Accessibility, and Editorial Policy.

For specific questions or personal concerns, you can also speak with your line manager or HR Business Partner. If you prefer to raise anonymous concerns about the Policy, or compliance, you may use the [Whistleblowing](#) channel.

For confidential personal support, you may access the [Employee Assistance Programme](#).

If you believe you need a business-critical exception to this Policy, discuss the need with the First Point of Contact listed below. If appropriate, they will facilitate exception approval from the Policy Owner.

Note that exceptions are rare and must be thoroughly justified. No exceptions can be made to legal requirements.

For employees outside the UK, please consult your local management for guidance on regional variations or additional support options.

Document Control

Date of change	Version	Changes/Comments
24/10/2023	0.1	Initial Draft
06/11/2023	0.2	Second Draft
14/11/2023	0.3	Third Draft; Approved by RAI Sub-Group
15/11/2023	0.4	Minor updates to align with policy template
01/12/2023	1.0	Approved by Data Governance Committee

06/02/2024	1.1	Minor updates
09/04/2025	1.2	Updates for annual review

Appendix: Key AI Terms and Definitions

Term	Definition
Artificial Intelligence (AI)	A machine-based system that can perform tasks commonly associated with human intelligence, including making content, predictions and/or decisions.
Machine Learning	A technique often used in Artificial Intelligence, which gives computers the ability to learn and improve at a task over time. This is done by using mathematical and statistical techniques to identify patterns, optimise systems and make predictions from data. This can be applied to solving different types of problems, such as identifying objects in pictures, recommending shows for a viewer, or targeting online adverts.
Generative AI	A type of artificial intelligence capable of generating text, images, speech, music, video and code in response to prompts from a user. Examples include ChatGPT and Copilot Chat (text), Adobe Firefly (images) Gen-2 (video) and GitHub Copilot (code).
AI Agent / Agentic AI	A type of AI-powered system that can take actions to achieve goals or outcomes set by the user. A key feature of agents is that they can perform multiple functions and can combine multiple tasks autonomously without direct human intervention or oversight. For example, an AI workflow that can book a trip for you, including flights and hotels, all from a single chat interface.
Input	Data or information that is provided to an AI system, whether by sending text prompts or uploading information, images, videos, or other types of data.
Output	Data or information, including text, images, and videos, that is produced by an AI system. This can include decisions or actions depending on the interface.
Training Data	Data or information that is used within the AI system's back-end to 'train' the algorithm. The AI system uses similarities between Inputs and the training data to determine the best Output, based on the requirements of the algorithm.
Bias	Bias has a number of meanings, and can be positive or negative. In Responsible AI, bias generally means a systematic skew in decision-making that results in unfair outcomes. Where this is related to protected characteristics such as race, gender, disability or sexuality, bias can constitute illegal discrimination. There are a number of causes of bias ranging from issues with data and algorithmic design to how humans interact with AI systems. Perhaps the most prominent cause is that AI systems trained to make decisions based on past data often replicate the historic biases in that data.
Personal Data	Personal data is any information that relates to a known person, or a person that could be identified (e.g. through a combination of the data held - such as a person's initials and postcode). The information is still personal data if you remove or replace indicators (such using a case number instead of a person's name). If the data is fully anonymised, and can no longer be identified, this is not considered personal data.
Special Category Data	Special category data is personal data that needs more protection because it is sensitive. The UK GDPR defines special category data as: <ul style="list-style-type: none"> • Personal data revealing racial or ethnic origin • Personal data revealing political opinions • Personal data revealing religious or philosophical beliefs • Personal data revealing trade union membership • Genetic data • Biometric data (where used for identification purposes) • Data concerning health (physical and mental) • Data concerning a person's sex life • Data concerning a person's sexual orientation