




# **Artificial Intelligence (AI)**

Guidance for Judicial Office Holders

8 October 2025



# **Contents**

<b>1. Introduction</b>	<b>2</b>
<b>2. Common terms</b>	<b>2</b>
<b>3. Guidance for responsible use of AI in Courts and Tribunals</b>	<b>4</b>
<b>4. Examples: potential uses and risks of Generative AI in Courts and Tribunals</b>	<b>7</b>

# 1. Introduction

This refreshed guidance has been developed to assist judicial office holders in relation to the use of Artificial Intelligence (AI). It updates and replaces the guidance document issued in April 2025.

It sets out key risks and issues associated with using AI and some suggestions for minimising them. Examples of potential uses are also included.

Any use of AI by or on behalf of the judiciary must be consistent with the judiciary's overarching obligation to protect the integrity of the administration of justice.

This guidance applies to all judicial office holders under the Lady Chief Justice and Senior President of Tribunal's responsibility, their clerks, judicial assistants, legal advisers/officers and other support staff. This guidance is published online to promote transparency, open justice and public confidence.

## 2. Common Terms

Algorithm:	A set of instructions used to perform tasks, such as calculations and data analysis, usually using a computer or another smart device.
Alignment:	Any process by which Artificial Intelligence systems are aligned with organisational goals or ethical principles, such as in relation to 'de-biasing'.
Artificial Intelligence (AI):	Computer systems able to perform tasks normally requiring human intelligence.
AI Agent:	A software programme that uses AI to become aware of its environment, process information, and take actions to achieve its goals, based on the inputted information.
AI Prompt:	An input or instruction given to an AI system which will generate a specific response or result. Typically, a prompt is in the form of text, but many chatbots will now accept voice prompts.
Generative AI:	A form of AI which generates new content, which can include text, images, sounds and computer code. Some generative AI tools are designed to take actions.
Generative AI Chatbot:	A computer program which simulates an online human conversation using generative AI.

Publicly available examples are ChatGPT, Google Gemini and Meta AI.

Hallucination:

AI hallucinations are incorrect or misleading results that AI models generate. These errors can be caused by a variety of factors, including insufficient training data, the model's statistical nature, incorrect assumptions made by the model, or biases in the data used to train the model.

Large Language Model (LLM):

LLMs are AI models which learn to predict the next best word or part of a word in a sentence having been trained on enormous quantities of text.

Machine Learning (ML):

A branch of AI that uses data and algorithms to imitate the way that humans learn, gradually improving accuracy. Through the use of statistical methods, algorithms are trained to make classifications or predictions, and to uncover key insights in data mining projects.

Natural Language Processing:

Programming computer systems to understand and generate human speech and text. Algorithms look for linguistic patterns in how sentences and paragraphs are constructed and how words, context and structure work together to create meaning. Applications include speech-to-text converters, online tools that summarise text, chatbots, speech recognition and translations.

Responsible AI:

The practice of designing, developing, and deploying AI with certain values, such as being trustworthy, transparent, explainable, fair, robust and upholding privacy rights.

Technology Assisted Review (TAR):

AI tools used as part of the disclosure process to identify potentially relevant documents. In TAR, a machine learning system is trained on data created by lawyers who identify relevant documents manually, then the tool uses the learned criteria to identify other similar documents from very large disclosure data sets.

White Text:

Text formatted to be invisible to human readers (e.g. white font on a white background) but still detectable by computers. It can be used to manipulate

search engines or large language models by embedding hidden instructions or keywords.

## **3. Guidance for responsible use of AI in Courts and Tribunals**

### **I. Understand AI and its applications**

Before using any AI tools, ensure you have a basic understanding of their capabilities and potential limitations.

Some key limitations:

Public AI chatbots do not provide answers from authoritative databases. They generate new text using an algorithm based on the prompts they receive and the data they have been trained upon. This means the output which AI chatbots generate is what the model predicts to be the most likely combination of words (based on the documents and data that it holds as source information). It is not necessarily the most accurate answer.

As with any other information available on the internet in general, AI tools may be useful to find material you would recognise as correct but have not got to hand, but are a poor way of conducting research to find new information you cannot verify. They may be best seen as a way of obtaining non-definitive confirmation of something, rather than providing immediately correct facts.

The quality of any answers you receive will depend on how you engage with the relevant AI tool, including the nature of the prompts you enter, and the quality of the underlying datasets. These may include misinformation (whether deliberate or otherwise), selective data, or data that is not up to date. Even with the best prompts, the information provided may be inaccurate, incomplete, misleading, or biased. It must be borne in mind that “wrong” answers are not infrequent.

The currently available LLMs appear to have been trained on material published on the internet. Their “view” of the law is often based heavily on US and historic law, although some do purport to be able to distinguish between that and the law of England and Wales.

### **II. Uphold confidentiality and privacy**

Do not enter any information into a public AI chatbot that is not already in the public domain. Do not enter information which is private or confidential. Any information that you input into a public AI chatbot should be seen as being published to all the world.

The current publicly available AI chatbots remember every question that you ask them, as well as any other information you put into them. That information is then available to be used to respond to queries from other users. As a result, anything you type into it could become publicly known.

You should disable the chat history in public AI chatbots if this option is available, as it should prevent your data from being used to train the chatbot and after 30 days, the conversations will be permanently deleted. This option is currently available in ChatGPT and Google Gemini but not in some other chatbots. Even with history turned off, though, it should be assumed that data entered is being disclosed.

Be aware that some AI platforms, particularly if used as an App on a smartphone, may request various permissions which give them access to information on your device. In those circumstances you should refuse all such permissions.

In the event of unintentional disclosure of confidential or private information you should contact your leadership judge and the Judicial Office. If the disclosed information includes personal data, the disclosure should be reported as a data incident. Details of how to report a data incident to Judicial Office can be found at this link: [Data breach notification form judiciary](#)

You should treat all public AI tools as being capable of making public anything entered into them.

### **III. Ensure accountability and accuracy**

The accuracy of any information you have been provided by an AI tool must be checked before it is used or relied upon.

Information provided by AI tools may be inaccurate, incomplete, misleading or out of date. Even if it purports to represent the law of England and Wales, it may not do so. This includes cited source material which might also be hallucinated.

AI tools may “hallucinate”, which includes the following:

- make up fictitious cases, citations or quotes, or refer to legislation, articles or legal texts that do not exist,
- provide incorrect or misleading information regarding the law or how it might apply, and
- make factual errors.

### **IV. Be aware of bias**

AI tools based on LLMs generate responses based on the dataset they are trained upon. Information generated by AI will inevitably reflect errors and biases in its training data, perhaps mitigated by any alignment strategies that may operate.

You should always have regard to this possibility and the need to correct this. You may be particularly assisted by reference to the [Equal Treatment Bench Book](#).

## **V. Take Responsibility**

Judicial office holders are personally responsible for material which is produced in their name.

Judges must always read the underlying documents. AI tools may assist, but they cannot replace direct judicial engagement with evidence.

Judges are not generally obliged to describe the research or preparatory work which may have been done in order to produce a judgment. Provided these guidelines are appropriately followed, there is no reason why generative AI could not be a potentially useful secondary tool.

Follow best practices for maintaining your own and the court/tribunal's security. Use work devices (rather than personal devices) to access AI tools. Before using an AI tool, ensure that it is secure. If there has been a potential security breach, see (II) above.

If clerks, judicial assistants, legal officers/advisers, or other staff are using AI tools in the course of their work for you, you should discuss it with them to ensure they are using such tools appropriately and taking steps to mitigate any risks. If using a HMCTS or MoJ supplied laptop, you should also ensure that such use has HMCTS service manager approval.

## **VI. Be aware that court/tribunal users may have used AI tools**

Some kinds of AI tools have been used by legal professionals for a significant time without difficulty. For example, TAR is now part of the landscape of approaches to electronic disclosure. Leaving aside the law in particular, many aspects of AI are already in general use: for example, in search engines to auto-fill questions, in social media to select content to be delivered, and in image recognition and predictive text.

All legal representatives are responsible for the material they put before the court/tribunal and have a professional obligation to ensure it is accurate and appropriate. Provided AI is used responsibly, there is no reason why a legal representative ought to refer to its use, but this is dependent upon context.

Until the legal profession becomes familiar with these new technologies, however, it may be necessary at times to remind individual lawyers of their obligations and confirm that they have independently verified the accuracy of any research or case citations that have been generated with the assistance of an AI chatbot.

AI chatbots are now being used by unrepresented litigants. They may be the only source of advice or assistance some litigants receive. Litigants rarely have the skills independently to verify legal information provided by AI chatbots and may not be aware that they are prone to error. If it appears an AI chatbot may have been used to prepare submissions or other documents, it is appropriate to inquire about this, ask what checks for accuracy have been undertaken (if any), and inform the litigant that they are responsible for what they put to the court/tribunal. Examples of indications that text has been produced this way are shown below.

AI tools are now being used to produce fake material, including text, images and video. Courts and tribunals have always had to handle forgeries, and allegations of forgery, involving varying levels of sophistication. Judges should be aware of this new possibility and potential challenges posed by deepfake technology.

Another form of fake material of which you must be aware is so called “white text”, which consists of hidden prompts or concealed text inserted into a document so as to be visible to the computer or system but not to the human reader. This possibility underscores the importance of judicial office holders’ personal responsibility for anything produced in their name.

## **4. Examples: Potential uses and risks of Generative AI in Courts and Tribunals**

### **Potentially useful tasks**

- AI tools are capable of summarising large bodies of text. As with any summary, care needs to be taken to ensure the summary is accurate.
- AI tools can be used in writing presentations, e.g. to provide suggestions for topics to cover.
- Administrative tasks can be performed by AI, including composing, summarising and prioritising emails, transcribing and summarising meetings, and composing memoranda.

### **Tasks not recommended**

- Legal research: AI tools are a poor way of conducting research to find new information you cannot verify independently. They may be useful as a way to be reminded of material you would recognise as correct, although final material should always be checked against maintained authoritative legal sources.
- Legal analysis: the current public AI chatbots do not produce convincing analysis or reasoning.

### **Indications that work may have been produced by AI:**

- References to cases that do not sound familiar, or have unfamiliar citations (sometimes from the US),
- Parties citing different bodies of case law in relation to the same legal issues,
- Submissions that do not accord with your general understanding of the law in the area,
- Submissions that use American spelling or refer to overseas cases,
- Content that (superficially at least) appears to be highly persuasive and well written, but on closer inspection contains obvious substantive errors, and
- The accidental inclusion of an AI prompt, or the retention of a ‘prompt rejection’, such as “as an AI language model, I can’t ...”