



Bar Council response to the Civil Justice Council consultation on the use of artificial intelligence for preparing court documents

This is the response of the General Council of the Bar of England and Wales (the Bar Council) to the Civil Justice Council (CJC) consultation on the use of artificial intelligence (AI) for preparing court documents.¹

The Bar Council is the voice of the barrister profession in England and Wales. Our nearly 18,000 members – self-employed and employed barristers – make up a united Bar that aims to be strong, inclusive, independent and influential. As well as championing the rule of law and access to justice, we lead, represent and support the Bar in the public interest through:

- Providing advice, guidance, services, training and events for our members to support career development and help maintain the highest standards of ethics and conduct
- Inspiring and supporting the next generation of barristers from all backgrounds
- Working to enhance diversity and inclusion at the Bar
- Encouraging a positive culture where wellbeing is prioritised and people can thrive in their careers
- Drawing on our members' expertise to influence policy and legislation that relates to the justice system and the rule of law
- Sharing barristers' vital contributions to society with the public, media and policymakers
- Developing career and business opportunities for barristers at home and abroad through promoting the Bar of England and Wales
- Engaging with national Bars and international Bar associations to facilitate the exchange of knowledge and the development of legal links and legal business overseas

¹ <https://www.judiciary.uk/wp-content/uploads/2026/02/Interim-Report-and-Consultation-Use-of-AI-for-Preparing-Court-Documents-2.pdf>

To ensure joined-up support, we work within the wider ecosystem of the Bar alongside the Inns, circuits and specialist Bar associations, as well as with the Institute of Barristers' Clerks and the Legal Practice Management Association. As the General Council of the Bar, we are the approved regulator for all practising barristers in England and Wales. We delegate our statutory regulatory functions to the operationally independent Bar Standards Board (BSB) as required by the Legal Services Act 2007.

Introduction

The Bar Council has issued guidance to its members on the use of Generative Artificial Intelligence.²

At the core of our guidance is the requirement on practitioners to understand the technological process that is occurring in the particular tool being considered for use and to use that information to inform their decision as to if, and if so to what extent, the use of the tool is consistent with their professional obligations. Our guidance explains why that approach is necessary by reference to specific weaknesses inherent in the technological processes that underlie the relevant models and the known risks associated with those processes.

How AI interacts with the administration of justice is clearly going to be a significant question across the range of civil, criminal, family and other proceedings. In our view, it is very difficult to resolve in the abstract the myriad issues that may arise. It is likely that these issues will need to be resolved on a use by use or case by case basis within the specialist practice areas.

Consistent with the approach in our own guidance, we suggest that the resolution of these issues by the courts will require and depend on a close examination of the particular technological processes used, or to be used, and the inherent weaknesses and risks associated with them. We suggest that this background should inform the present consultation.

We broadly support the distinction made at the beginning of the consultation between AI being used to generate evidence on which the Court is being asked to

² [Bar Council – Considerations when using ChatGPT and generative artificial intelligence software based on large language models, \(Published in January 2024 & Updated in November 2025\).](#)

rely and administrative uses, although, as the consultation acknowledges, those categories do not exhaust the potential uses of AI in a litigation context.

We also agree that the question of AI use by litigants in person raises particular issues. We suggest that these questions do require a further specific consideration. While access to justice issues arise from the use by litigants in person of AI tools, nonetheless, the unreliability of AI outputs, especially where commercial general AI Large Language Model (LLM) tools (such as ChatGPT) are used, has been shown to raise significant administration of justice issues, and give rise to significant consumer protection issues.

In line with our understanding of the scope of the consultation, our responses to the consultation questions are confined to those cases where a person is legally represented. The CJC may want to consider whether, if there is a possibility that its view of a particular requirement might be different where a party is self-represented, it should modify its present approach whereby generally all litigants should be treated in the same way.

In regard to those areas where the consultation suggests that the use of AI should be prohibited, we do not see a reason why the prohibitions should not apply equally to litigants in person. The same public interest considerations apply.

Scope

1. The scope of this work has been concerned with rules relating to legal representatives, on the basis that guidance is a matter for their professional bodies. Do you agree with that approach to guidance? If not, please explain why not.

We agree that guidance is a matter for professional bodies. However, we think that the nature of the technological processes, the weaknesses and risks associated with them and the fact that the courts cannot necessarily have confidence in the levels of AI literacy (even among legal practitioners) it is necessary for there to be at least some Rules or Practice Directions or Guidance in this area.

Statements of Case

2. The CJC proposes that provided a statement of case bears the name of the legal representative who is taking professional responsibility for it, there is no

need for any (further) rules relating to statements of case produced with the assistance of AI. Do you agree? If not why not?

We agree. Statements of Case are generally not in a client's own words and most legal representatives will be subject to their ordinary professional obligations, when putting their name to it. All practising Barristers are regulated by the Bar Standards Board (BSB) and are subject to the duties set out in the Code of Conduct, including the duty not to mislead the court (rC6) and to act with honesty and integrity (CD3). These obligations require barristers to ensure that any statement of case is accurate, properly arguable (rC9.2.b), and supported by an appropriate evidential and legal basis, regardless of the tools used in its preparation.

In that context, the use of AI does not alter the fundamental position. Whether material is drafted by the barrister, by a member of their team, or with the assistance of technological tools. (e.g. researching via WestLaw /LexisNexis etc), the barrister remains personally responsible for verifying its accuracy and ensuring compliance with their professional duties. The existing regulatory framework therefore provides a clear and robust safeguard against the risks identified by the CJC, including the inclusion of inaccurate or hallucinated material. Further, we are aware that the BSB is aiming to provide specific rules on the use of AI tools by barristers.

Against that background, we do not consider that additional AI-specific rules are necessary in relation to Statements of Case. The imposition of further requirements risks duplicating existing obligations and introducing unnecessary complexity, without materially enhancing the protection afforded to the court or to other parties.

3. An alternative would be to require a specific declaration to make clear if the legal representative has used AI in the preparation of the statement of case. Do you prefer this alternative? If so, please explain why and consider which uses of AI ought to be covered.

If there were to be a requirement for a specific declaration, it should be confined to the use of AI to generate text, to set out evidence or assertions of facts.

Skeleton arguments and other advocacy documents

4. The CJC proposes that provided the skeleton argument or other advocacy document bears the name of the legal representative who is taking professional responsibility for it, there is no need for any (further) rules relating to these documents produced with the assistance of AI. Do you agree? If not why not?

We agree with this proposal. Again, any practising barrister will be subject to their professional obligations as already mentioned. Similar professional obligation will apply to solicitors and other regulated practitioners. As aforementioned the Bar Council has issued guidance and the BSB is very likely to do so in the near future. The court should be able to assume that this guidance has been followed. Further, as the primary purpose of skeleton arguments are produced by the advocate to assist the court by clearly and concisely identifying the issues, the relevant legal principles, and the submissions to be advanced. Introducing AI-specific rules in this context would give rise to further difficulty, not least because of the lack of a clear and workable definition of “AI”, and the increasingly blurred distinction between AI-assisted and conventional research or drafting tools. In those circumstances, the existing professional framework—focussed on responsibility for the accuracy and integrity of the material presented by regulated practitioners—provides a more coherent and principled basis for regulation than attempting to categorise or prescribe the means by which such documents are prepared.

In particular, there is no principled basis for requiring disclosure of AI-assisted research where no such obligation exists in relation to the use of established tools such as Westlaw or Practical Law, even where those tools increasingly incorporate AI functionality. Some barristers further observe that AI-assisted research using reputable, curated legal sources is, in many respects, a continuation of existing practice. Barristers have long relied on such platforms to trace authorities and develop arguments. The material distinction lies in the capacity of some AI tools to generate content, including hallucinated authorities, rather than merely to retrieve it. That risk is real, but it is one which is properly addressed through the barrister’s overriding duty to the court and obligation to verify sources, rather than through broad disclosure requirements tied to the means by which research is conducted.

5. An alternative would be to require a specific declaration to make clear if the legal representative has used AI in the preparation of these documents. Do you prefer this alternative? If so, please explain why and consider which uses of AI ought to be covered.

If a declaration were to be required, merely specifying the types of AI or uses would be unlikely to be helpful and most likely cause unnecessary burdens on practitioners and the court. However, if a declaration is required then such a declaration as mentioned in the CJC consultation being placed in New South Wales might be appropriate. It would state that in so far as AI has been used the practitioner has checked that the references exist, are accurate and are relevant. Such a statement sets out what the practitioner is required to do in any event.

Disclosure

6. The CJC proposes that there does not appear to be a pressing case to introduce a requirement that disclosure lists/statements have a section addressing the extent to which AI tools/software have been used. Do you agree that disclosure lists/statements do not need to contain such a statement? If not why not?

As the consultation acknowledges, the use of AI in disclosure is a developing area. The increasing size of datasets and the innovative attempts to deal with them mean that the appropriateness of using particular tools is likely to be the subject of litigation. In those circumstances there is a case for addressing the use of AI in disclosure lists/statements as a basis for any such issues to be resolved. However, we anticipate that members of the Bar (particularly in different practice areas) will have different views on this issue. It may therefore be that this should be dealt with in a separate process as suggested in the consultation.

Witness statements

7. The CJC makes different proposals for different kinds of witness statements, in particular drawing a distinction between trial witness statements and non-trial witness statements. Do you agree with that approach? What distinction if any would you propose?

If there is to be a distinction between trial and non-trial witness statements, such a distinction should be consistent with the CJC's approach to trial statements (which we support). Non-trial witness statements may deal with a range of issues. To be consistent with the CJC's approach to trial statements there should be the same prohibition on using AI in any part of a non-trial witness statement that deals with a disputed or potentially disputed issue of fact. Given this, any declaration should be applied equally to witness statements across the board.

8. In relation to non-trial witness statements, the proposal is that provided the statement bears the name (or firm name?) of the legal representative who is taking professional responsibility for its preparation, there is no need for any (further) rules relating to these documents produced with the assistance of AI. Do you agree? If not, why not?

See answer 7 above.

9. An alternative would be to require a specific declaration to make clear if the legal representative has used AI in the preparation of non-trial witness statements. Do you prefer this alternative? If so, please explain why and consider which uses of AI ought to be covered.

We do not believe there is any principled justification for applying a specific declaration only to trial witness statements. The concern that needs to be addressed is that if AI is allowed to 'generate the content' of trial witness statements, this could result in (1) the statement not being expressed in the witness's own words, and/or (2) the evidence containing factual assertions outside of those which the witness would have made of their own volition. This should apply equally to trial and non-trial witness statements. Given this, any declaration should either be applied to witness statements across the board.

10. In relation to witness statements covered by PD57AC and within the scope of this paper (i.e. trial witness statements prepared with the involvement of a legal representative), the proposal is that there be a rule requiring a declaration on such a statement that AI has not been used for the purposes of generating its content (including by way of altering, embellishing, strengthening, diluting or rephrasing the witness's evidence). Do you agree? If not, why not?

We agree with this approach. It would be inconsistent with the requirement that a witness statement be in a witness' own words and that the witness not be asked leading questions, for a witness to be permitted to use Generative/Agentic AI tools in its preparation. Most relevant AI tools are text enhancement tools and suggest what should be written. These features make their use inconsistent with the purpose of a trial witness statement.

11. In relation to witness statements under CPR Part 32, not covered by PD57AC but within the scope of this paper (i.e. trial witness statements prepared with the involvement of a legal representative), the proposal is that there be a rule

requiring a declaration or such a statement that AI has not been used for the purposes of generating its content (including by way of altering, embellishing, strengthening or diluting or rephrasing the witness's evidence). Do you agree? If not, why not?

We agree for the reasons given in response to the previous questions.

12. In relation to witness statements involving translation, one issue relates to use of AI by translators. Should there be a rule making provision for the use of AI by human translators? If a translator is prepared to sign a statement of accuracy, taking responsibility for it, is there any need to enquire further? A further proposal is to permit the use of publicly available machine translation, provided the tool used is identified, and provided (if necessary) that provision is made clarifying that other parties are entitled to check the translation themselves by using such a tool. Do you agree? If not, why not? Do you favour the alternative below? If so why?

There would seem to be no reason to make a rule dealing with the use of AI by translators. As the consultation says, if a human translator is prepared to sign a statement of accuracy, taking responsibility for it, there would seem to be no need to enquire further.

As to the use of AI translation tools without human input, our view is that this should be approached with appropriate caution. The fundamental consideration is the need to ensure that the court is not misled. Consistent with the approach to the use of AI we set out at the beginning of this document, our view is that there should not be a rule change without a proper technical analysis of the reliability of translation tools and any inherent weaknesses they have. The sorts of issues that will need to be considered include: whether there is a difference in reliability amongst tools; whether there is a difference in reliability depending on the language that is being translated; whether there is a difference in reliability depending on the subject matter and what any difference might be; whether the tool can cope with nuance; and what the risks of bias in translation outcomes may be.

The existing rule in PD 32 paragraph 23.2(b) requires the human translator to certify the accuracy of the translation, thereby assuming responsibility for the output placed before the court. That correctly reflects the principle that accountability should attach to the individual certifying the translation, not to the tool or method used in producing it. With a human translator a court can also assess the qualifications and competence

of the individual. There is no reason why a similar process should not be available for AI translators used where no human translator has vetted the content produced.

13. An alternative to the previous proposal would be only to permit such use by a legal representative and to require that the legal representative involved in the preparation of the translation should identify what tool has been used. Do you favour this alternative?

It is not clear why there should be a different rule for legally represented litigants in this respect. The reason for caution is the need to ensure that a court is not misled. A legal representative will be in the same position as a litigant in person in assessing a specific piece of translation if they do not speak the relevant language.

Experts

14. The proposal is that the specific provisions for statements of truth used by experts should be amended to add a further requirement confirming that the expert's report identifies and explains any AI which has been used, other than for administrative uses such as transcription. Do you agree? If not why not?

We broadly agree that further guidance as to the use of AI in experts' reports is necessary. We anticipate that members of the Bar (even within practice areas) will have different views on this issue.

However, there is some concern by barristers that introducing further wording into the statement of truth requiring the expert to explain where AI has been used would not be beneficial (and may be counterproductive), for four reasons:

First, doing so risks introducing uncertainty as to the purpose of the statement of truth. The point of the statement of truth is to ensure the expert has independently confirmed (and will attest to) (i) the views expressed in his report, and (ii) the materials and calculations which they have used to produce that report.

AI can be used in the process of forming those views or gathering those facts, but the declaration adds an extra step – the expert must declare (essentially) that their conclusions are their own. To put it another way – the problem is not the use of AI in

itself, but experts' "*abdicat[ion] [of] their independent judgment and critical thinking skills in favour of ready-made AI generated answers*"³.

Introducing a further requirement into the statement of truth to declare AI use risks muddying the purpose as set out above. An expert might (for example) say that they have complied with their obligation to state their own views and what matters are/are not within their own knowledge by simply identifying where AI has been used, and without independently confirming the data gathered by the AI / the conclusions that it reached. That is unhelpful, because the two exercises are distinct – declaring that you have used AI does not mean that you have used AI in a manner which is consonant with your duty as an expert.

Second, experts are probably already required to explain where AI has been used in processes– (1) (gathering evidence) consists of facts and assumptions on which the expert's view is based, which the expert must declare; (2) (summarising evidence); and (3) (drawing certain conclusions from that evidence) involve obtaining assistance from others.

In *Kohls et al v. Ellison*, the motion to uphold the exclusion of the affected expert report was refused based on existing US law as to documentary evidence. It seems that declarations as to the use of generative AI are more naturally viewed as matters to be dealt with under the rules of civil evidence and procedure under the CPR, rather than something which ought to be dealt with in a statement of truth.

Third, any declaration in a statement of truth would have to be broad (because experts are asked to comment on a very broad range of fields) but that broadness would cause problems in compliance. The proposed wording asks experts to "*identify and explain any AI which has been used, other than for administrative purposes*", with no guidance as to the level of detail required in the identification/explanation.

Some experts would consider those duties discharged with a general statement or proforma wording (along the lines of "AI has been used in the preparation of this document", which is something that sometimes appears in expert reports seen by barristers). Others may be reluctant to sign such a declaration at all - as has been noted, it is often unclear whether a particular feature qualifies as AI, or AI for administrative purposes especially where it is integrated seamlessly into search

³ *Kohls et al v. Ellison*

engines and other software which might have been used. Such a proposal may also encourage tactical challenges, on the basis that minor use of AI as (say) a dictionary or thesaurus has not been declared, and cause further unnecessary time in cross-examination

Fourth, additional guidance on when use of AI should be declared would be very useful. The fact that we have an adversarial system means that most experts' reports are subject to some degree of review by the opposition's own experts. However, it is difficult for the court itself (particularly where a single joint expert report is provided) to verify details in the expert reports (as opposed to simply looking up a case-reference to determine whether or not it exists). The additional guidance can be tailored to particular disciplines on which experts are commonly consulted (e.g.: medical evidence in personal injury cases), or it can take the form of a guidance note or practice direction in the CPR (the latter would require careful thought).

Matters covered might include what sorts of common AI use cases ought to be declared (e.g.: generation of chronologies, generation of references, collating facts from large number of documents). The issue can be returned to after the guidance note has been in operation to see if further measures are necessary.

General Issues

15. Is the term artificial intelligence sufficiently clear to be used in these proposed rules? If not do you have an alternative proposal?

We agree that there is significant debate as to what lies within the proper definition of "AI". The term 'AI' (standing alone) is extremely difficult to define in a way which is precise, as recognised globally, particularly with the use rate and complexity of which it is evolving. The OECD definition adopted by 47 governments and used as the basis for several national frameworks states:

"An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate output such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment."

The definition has been incorporated by the European Union, the Council of Europe, the United States and the United Nations into their respective legislative and regulatory frameworks. The EU AI Act's definition is largely derived from the OECD framework, and both the US (via NIST) and Australia reference it in their respective policy instruments too

This technology is emerging and it has not reached a stage in its development where clear lines can be drawn between reliable and unreliable tools or uses. However, in the context of the questions asked in this present consultation the term is sufficiently clear to be used, mainly considering generative AI tools.

A more material consideration than niceties of definition is whether the tools in question (however defined) are, or are not, reliable. Since this technology is emerging, but has not yet reached a stage in its development (if such a stage can ever be reached) where clear lines can be drawn between reliable and unreliable tools.

16. One of the distinctions drawn between different uses of AI is between activity defined in the report as administrative uses, which merely corrects spelling or grammar, provides transcription, operates as accessibility software, or assists with formatting and otherwise does not generate substantive content on the one hand, and activity which generates substantive text, images or videos on the other. Another distinction drawn is between fact evidence and the product of legal research. Do you agree with the distinctions drawn in these proposals? If not what alternatives do you propose?

As set out at the beginning of this document, we support the distinction between administrative uses and use to generate substantive content. Tools for the sorts of administrative uses that are identified have been around for some considerable time. The technological development that calls for separate careful consideration is the ability to generate substantive content. A particular risk for the administration of justice is the tendency of some AI tools to appear to be "thinking" as a human might think (a feature often programmed into general LLM's). This can give the appearance that such tools engage in human-like reasoning processes, whereas the reality is that the output is generated solely by technical means based on the processing of text or images. The tool does not understand meaning in a way that a human would.

We also accept that there is a distinction between fact evidence and the product of legal research and that the distinction justifies their different treatment in the way set out in this consultation.

However, we would not wish to be seen to be expressing any view on the reliability of AI legal research tools. Our guidance refers to a study that cast doubt on their reliability as at the date of that study. Individual barristers and chambers are carrying out their own work testing these tools. There is currently no consensus view as to how reliable individual tools are.

17. Should the endorsements proposed always identify the AI tool used? If so, to what end?

There may be an argument that in the case of skeleton arguments and similar documents, if the endorsement is of the type required in New South Wales (that any AI output has been checked by a lawyer) there is no reason for requiring the specific AI tool used to be identified. In all other cases in which an endorsement is required, the AI tool should be identified. The reason for the requirement for an endorsement is to ensure that the use of AI has not undermined the reliability of the relevant document. This can only be accurately assessed by knowing which AI tool has been used. AI tools are simply electronic processes, in so far as their use may need to be interrogated, it can only reliably be done by reference to the specific product that is engaging in that process.

18. Should there be a rule providing for a power to give a party permission to use AI for some specific purpose? If such a rule should be introduced, should it be general or confined to specific uses?

It is not possible to answer this question in a comprehensive or reliable way in response to this consultation. The better approach in this field of emerging technology is for the courts to consider each use case against specific relevant technologies on a case by case basis.

Bar Council

April 2026

For further information please contact:

Nell Ferrall, Policy Analyst: Legal Practice and Remuneration

The General Council of the Bar of England and Wales

289-293 High Holborn, London WC1V 7HZ

Email: NFerrall@barcouncil.org.uk