The Bar Council

Considerations when using ChatGPT and generative artificial intelligence software based on large language models

Purpose:	To provide barristers with a summary of considerations if using ChatGPT or any other generative AI software based on large language models (LLMs).
Scope of application:	All barristers and chambers
Issued by:	The Information Technology Panel
Issued on:	30 January 2024
Last reviewed:	25 November 2025
Status and effect:	**Please see the notice at the end of this document. This is not "guidance" for the purposes of the BSB Handbook I6.4, the principles and warnings it contains reflect current professional expectations, particularly in light of recent High Court judgments on professional responsibility.**

Introduction

1. In the rapidly evolving technological landscape, particularly with the evolution of generative artificial intelligence (**generative AI**) based on large language model (**LLM**) systems, like OpenAI's ChatGPT, Google's Gemini, Perplexity, Harvey and Microsoft Copilot (which is also based on Open AI technology) to name a few; generative AI is increasingly being used by legal professionals for efficiency and practice management. The Bar Council has updated this guidance to assist barristers in understanding the technological basis and inherent risks in the use of such generative LLM systems. Although this updated guidance may not be exhaustive, it

aims to underscore the heightened imperative for barristers who engage with these tools to do so responsibly, ensuring adherence to legal and ethical standards; safeguarding client confidentiality, and maintaining trust and confidence, privacy and compliance with applicable laws.

2. The purpose of this guidance is to provide a useful summary of considerations for barristers if they decide to use ChatGPT or any similar LLM software, as well as systems specifically aimed at lawyers, such as Lexis+ AI, Clio Duo or Thomson Reuters Co-Counsel. It should also be noted that generative LLM technologies are developing rapidly and as the field of generative AI continues to evolve, with new models and advances being introduced regularly, it is always good to understand the underlying model and acknowledge its limitations prior to using these technologies. It is important to note that the legal and regulatory landscape on the use of AI is subject to constant change, and therefore barristers will need to be vigilant and adapt accordingly.

What is large language model (LLM) software?

3. It is easier to begin by explaining what it is not. It is not a conventional research tool, it does not analyse the *content* of data, and it does not think for itself. It is, rather, a very sophisticated version of the sort of predictive text systems that people are familiar with from email and chat apps on smart phones, in which the algorithm predicts what the next word is likely to be. LLMs use machine learning algorithms, first to be 'trained' on text and, based on that 'training' (which involves the application of *inter alia* mathematical formulae), to generate sequential text[1]. These programmes are now sufficiently sophisticated that the text often appears as if it was written by a human being, or at least by a machine which thinks for itself. This is a key risk factor in their use.

4. LLMs have not been around long enough and have not been sufficiently tested for it to be clear what tasks they can or should be used for in legal practice. Some practitioners and judges have made positive comments about using them to arrange text; others have expressed frustration at their over-use. However, **it is important for barristers who choose to use LLMs to do so responsibly and think about what they are doing, by weighing the potential risks and challenges associated with such use in light of their professional responsibilities.**

5. Crucially, barristers must understand that LLMs, while sophisticated, are not infallible. They are predictive tools, prone to generating plausible but entirely false information – a phenomenon known as 'hallucinations'. LLMs are not a

[1] The latest versions also include image capabilities.

substitute for human legal expertise, critical judgement, or diligent verification. The ultimate responsibility for all legal work remains with the barrister.

What is ChatGPT?[2]

6. ChatGPT is an advanced LLM AI technology developed by OpenAI. It is based on GPT architecture, which stands for 'Generative Pre-Trained Transformer'. The latest iteration of ChatGPT at the time of this guidance is GPT-5. Transformer architecture uses mathematical matrices, supplemented by corrective procedures and technologies. The number of parameters used by GPT-5 is thought to be in the many billions.

7. In common with other LLMs (such as Google's Gemini), ChatGPT is trained on huge amounts of data, which is processed through a neural network made up of multiple nodes and layers. These networks continually adjust the way they interpret and make sense of data based on a host of factors, including the results of previous trial and error.

8. Certain consequences inevitably follow from the nature of the technological process that is being carried out. LLM AI systems are not concerned with concepts like 'truth' or accuracy.

Key risks with LLMs

9. **Anthropomorphism:** the first key risk inherent in LLMs is that they are designed and marketed in such a way as to give the impression that the user is interacting with something that has human characteristics. One of the mechanisms by which this is sought to be achieved is by the use of anthropomorphic language to describe what is happening. Perhaps the most obvious example of this is the use, by OpenAI, of the word 'Chat' in the name of its LLM products (ChatGPT) and the fact it may state it is 'thinking' when one is awaiting the result of a prompt. As set out above, LLMs (at least at the current stage in their development) do not have human characteristics in any relevant sense, they don't understand concepts, emotions or causality in the way humans do, or have social or emotional intelligence or conscience.

[2] The reason for highlighting ChatGPT is that it remains the most widely known LLM and also shares technology with Microsoft Copilot.

10. **Hallucinations:** it has been said that LLMs are prone to 'hallucinations', a term which is used to describe the phenomenon where the outputs generated by these LLMs may sound very plausible but are either factually incorrect or unrelated to the given context.[3] However, whilst the use of this term is helpful for illustrative purposes, it demonstrates the widespread tendency to anthropomorphise the technology. As we say, it is necessary when using LLMs to keep well in mind the actual technical process that is being carried out. The experience of the legal profession (see the Appendix to Dame Victoria Sharp's judgment in *Ayinde* below) and of other professions (such as medicine) is that general purpose LLMs are unreliable tools for 'source-based' research.

11. Although hallucinations may be much less frequent in sophisticated fine-tuned LLM-based legal research tools like Lexis+ AI, they still occur, as shown in a Stanford University study, 'Hallucination-free? Assessing the reliability of leading AI legal research tools'[4], which found that Lexis+ AI and Thomson Reuters (Westlaw AI-Assisted Research and Ask Practical Law AI) each relatively hallucinate between 17% and 33% of queries. Therefore, it is essential that barristers (and all legal practitioners) verify that any sources or authorities cited by such systems actually support the propositions asserted, and ensure that the citations are accurate and extant. Errors or complacency in this regard will most likely have serious professional consequences.

12. The *R (Ayinde) v The London Borough of Haringey* [2025] EWHC 1040 (Admin) judgment serves as a stark precedent. In this case, although the lawyers denied using AI, counsel submitted pleadings containing five fabricated legal citations, including a non-existent Court of Appeal case, and misrepresented a statutory provision (Section 188(3) of the Housing Act 1996) as mandatory ('must') rather than discretionary ('may').

13. Mr Justice Ritchie unequivocally found this conduct to be "improper, unreasonable, and negligent," constituting "professional misconduct." He explicitly stated that putting fabricated cases in pleadings is "wholly improper" and "misleading the Court," profoundly undermining the "integrity of the legal profession and the Bar." The judge said that AI use without proper checking would be negligent. The professional consequences were severe, including wasted costs orders against both counsel and the instructing solicitors, a significant reduction in the claimant's overall costs, and mandatory referrals of both barrister and solicitors to their respective regulatory bodies.

---

[3] (1) Survey of hallucination in natural language generation, Ji Z, Lee N, Frieske R, et al. ACM Comput Surv. 2022. (2) Abstracts written by ChatGPT fool scientists, Holly Else

[4] 'Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools' 30 May 2024 Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning, Daniel E. Ho

14. The *Ayinde* judgment makes it unequivocally clear that 'minor citation errors' or 'cosmetic errors' (which is how the lawyers sought to explain what had happened) are grossly unprofessional categorisations. Such conduct directly breaches Core Duty 1 (duty to the court) and Core Duty 3 (duty to act with honesty and integrity), and will lead to severe sanctions, including wasted costs, disciplinary proceedings, and professional negligence claims. Also note the case of *MS v Secretary of State for the Home Department (Professional Conduct: AI Generated Documents) Bangladesh* [2025] UKUT 305 (IAC).

15. In *Ayinde v London Borough of Haringey and Al-Haroun* [2025] EWHC 1383 (Admin), Dame Victoria Sharp issued this warning:

> "9…There are serious implications for the administration of justice and public confidence in the justice system if artificial intelligence is misused. In those circumstances, practical and effective measures must now be taken by those within the legal profession with individual leadership responsibilities (such as heads of chambers and managing partners) and by those with the responsibility for regulating the provision of legal services. Those measures must ensure that every individual currently providing legal services within this jurisdiction (whenever and wherever they were qualified to do so) understands and complies with their professional and ethical obligations and their duties to the court if using artificial intelligence. For the future, in *Hamid* hearings such as these, the profession can expect the court to inquire whether those leadership responsibilities have been fulfilled."

16. She noted that a barrister can always access authoritative legal resources at their Inn of Court library. The four libraries provide a comprehensive range of legal research resources and training in how to carry out legal research effectively. The Inn libraries offer training at the qualifying session (QS), pupillage, and continuing professional development (CPD levels), both in-person and remotely. They also provide document supply services, to ensure that all members of the Bar have access to authoritative legal resources.

17. **Information disorder:** the ability for ChatGPT inadvertently to generate information disorder, including misinformation, is a serious issue of which to be aware.

18. The danger is illustrated by the affidavit that was filed by a New York lawyer to explain his conduct when he had included 6 fictitious cases suggested by

ChatGPT in his submissions.[5] Screenshots of what he had done were annexed to his affidavit:

> "Is varghese a real case"
>
> "Yes, Varghese v. China Southern Airlines Co Ltd, 925 F. 3d 1339 (11th Cir. 2019) is a real case."
>
> "What is your source"
>
> "I apologize for the confusion earlier. Upon double-checking, I found that the case Varghese v. China Southern Airlines Co. Ltd, 925 F.3d 1339 (11th Cir.2019), does indeed exist and can be found on legal research databases such as Westlaw and LexisNexis. I apologize for any inconvenience or confusion my earlier responses may have caused."
>
> "Are the other cases you provided fake."
>
> "No, the other cases I provided are real and can be found in reputable legal databases such as LexisNexis and Westlaw."

19. As can be seen, the lawyer thought that the LLM was engaging in the human process of reading and understanding the question, searching for the correct answer and then communicating the correct answer to the lawyer. In fact, all the LLM was doing was producing outputs (which just happened to be in the form of words) which its mathematical processes related to its inputs (which also just happened to be in the form of words). (The judgment and the Appendix to the Divisional Court's judgment in *Ayinde* contains a selection of cases from the UK and around the world concerning the misuse of AI in this area. There are also examples of LLMs being used to manufacture entirely fictitious allegations of misconduct against individuals.[6]

20. **Bias in training data:** another key risk is inherent in the manner in which an LLM is 'trained'. The fact that the training data in general purpose LLMs is trawled from the internet means that LLMs will inevitably contain biases or perpetuate stereotypes or world views that are found in the training data. There is now a growing body of research on how a range of AI-based tools contain inappropriate biases based on, for example, race and gender. Although the developers of ChatGPT and other LLMs have attempted to put safeguards in place to address these issues, it

[5] *Mata v. Avianca, Inc.* [Civil Action No: 22 Civ 1461]

[6] We do not identify the examples for obvious good reason, but they are serious and personally damaging.

is not yet clear how effective these safeguards are. Of course, it is also possible to game and manipulate the LLM in certain ways. Ensuring safe and appropriate behaviour from all users can be a significant challenge.

21. **Mistakes and confidential training data:** ChatGPT and other LLMs may use the inputs from users' prompts to continue to develop and refine the system. In consequence, anything that a user types into the system may be used to train the software and might find itself repeated verbatim in future results. This is plainly problematic not only if the material typed into the system is incorrect, but also if it is confidential or subject to legal professional privilege. It is important for barristers to fully understand how any model they are working with uses inputs and how to engage any relevant protective settings.

22. **Cyber security vulnerabilities**: the increasing integration of LLMs into legal tech platforms introduces new attack vectors for cyber criminals. Barristers and chambers must implement robust cyber security measures and be aware that AI can be used for more sophisticated phishing, business email compromise (BEC) scams, and other forms of fraud. Due diligence on the security protocols of AI tools is essential.

23. In short, while generative AI LLM systems have shown impressive capabilities in various natural language processing tasks, they also come with significant limitations.

Some considerations when using generative AI LLM systems

Practitioners should recognise the constraints and challenges presently embedded in the generative AI LLM software, including:

(1) Mandatory verification of outputs and human oversight:

24. **Barristers retain ultimate and complete responsibility for all advice to clients drafting and submissions made to the court or to clients, regardless of whether AI tools have been used in their preparation. AI tools *may* be aids for efficiency. They are not substitutes for a barrister's independent legal research verification, analysis, and judgment (see Ritchie J in *Ayinde*, above).**

25. The ability of LLMs to generate convincing but false content raises ethical concerns. Do not therefore take such systems' outputs on trust and certainly not at face value. It matters not that a misleading of the court may have been inadvertent, as it would still be considered to show incompetence and gross negligence on the part of the barrister. Such conduct brings the profession into disrepute (a breach of Core Duty 5), which may well lead to disciplinary proceedings (Ritchie J said that

the barrister in *Ayinde* should have self-reported to the BSB and the solicitor to the SRA). Barristers may also face professional negligence, defamation and/or data protection claims through careless or inappropriate use of these systems. As set out above, the data used to 'train' generative LLMs may not be up to date and can sometimes produce responses that are ambiguous, inaccurate or contaminated with inherent biases. Inherent bias may be invisible as it arises not only in the processing or training, but prior to that in the assembling of the training materials. LLMs may also generate responses which are out of context. For these reasons it is important for barristers to verify the output of AI LLM software and maintain proper procedures for checking the generative outputs.

26. The issues highlighted above are not necessarily limited to general purpose LLMs but may also arise, albeit to a lesser extent in relation to specialised legal LLMs, with the consequence that the user will be held to an equal level of professional responsibility as when using general purpose LLMs.

(2) 'Black box syndrome': lack of explain-ability

27. Like a number of AI tools, generative deep learning AI LLMs are often considered 'heavy black box' models, because it is difficult to understand the internal decision-making processes or provide clear explanations for the output. Some of the software remains 'proprietary' and therefore confidential. It can sometimes be difficult to interpret the results, due to the multilayer nonlinear model structures and the billions of parameters used. LLMs with attention mechanisms[7] may give some ability to see on which parts of the input text the model focuses when generating a response, thereby providing some insights into the decision making. But such insights cannot be a substitute for the exercise of professional judgement, quality legal analysis and the expertise which clients, courts and society expect from barristers.

(3) Respect legal professional privilege (LPP), confidential information and data protection compliance

28. Be extremely vigilant about sharing with a generative LLM system any legally privileged or confidential information (including trade secrets), or any personal data, as the input information provided may be used to generate future outputs and could therefore be publicly shared with other users. Any such sharing of confidential information is likely to be a breach of Core Duty 6 and rule rC15.5 of the Code of Conduct, which could also result in disciplinary proceedings and/or legal liability.

[7] Attention mechanisms allow the model to 'pay attention' to certain elements of the data to give them more weight.

As set out above, barristers need fully to understand how the tool they are using operates in this respect, including any relevant protective settings.

29. Barristers will also need to comply with relevant data protection laws. You should never input any personal data in response to prompts from the system. Note that in December 2024, the Italian Data Protection Authority fined ChatGPT based on its use of personal data. Italy, France and Spain have also investigated OpenAI's processing of personal data. Using only synthetic data (that is data that is artificially created) on prompts to the LLM represents one possible way to avoid the risk of falling into breach of the General Data Protection Regulation (EU 2016/679) as retained in English law (**UK GDPR**).

30. Barristers should check the terms and conditions of any generative LLM system to satisfy themselves that the LLM is compliant with Core Duty 6, rule rC15.5, and any relevant data protection laws.

31. As practitioners will be aware, the regulatory landscape in this area is in a state of flux and it is difficult to predict exactly what the UK position will be. Under the European Union (EU) AI Act,[8] certain uses of AI tools in legal practice are categorised as 'high-risk' which triggers heightened regulatory obligations. The UK Government's white paper, 'A pro-innovation approach to AI regulation'[9], published in March 2023, suggested that existing regulators should act in accordance with 5 principles (similar to the OECD principles on AI[10] although with different wording), namely:

(i) safety, security and robustness
(ii) appropriate transparency and explain-ability
(iii) fairness
(iv) accountability and governance
(v) contestability and redress.

32. In the UK, the Information Commissioner has published guidance in relation to the development and use of technologies such as ChatGPT, 'Generative AI: eight questions that developers and users need to ask.'[11]

(4) Intellectual property (IP) infringement and brand association

[8] https://www.consilium.europa.eu/en/press/press-releases/2023/12/09/artificial-intelligence-act-council-and-parliament-strike-a-deal-on-the-first-worldwide-rules-for-ai/
[9] http://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1146950/a_pro-innovation_approach_to_AI_regulation_print_ready_version.pdf
[10] https://oecd.ai/en/ai-principles
[11] https://ico.org.uk/about-the-ico/media-centre/blog-generative-ai-eight-questions-that-developers-and-users-need-to-ask/

33. The interaction between intellectual property law and LLMs is rapidly evolving, particularly as the use of copyrighted and proprietary data to train AI systems becomes a focal point of legislative and judicial scrutiny worldwide. Recent developments, including high-profile litigation in the United States, the United Kingdom, and the European Union, highlight growing concerns over whether AI training datasets and outputs infringe third-party copyright, trade secrets, design rights, or confidentiality obligations. Consequently, barristers and legal practitioners must critically evaluate both the data provided to LLMs and the content generated by them to ensure compliance with intellectual property and confidentiality laws. As a sizable amount of text data, such as books, papers, and other written materials were used to train ChatGPT and other LLMs, it is clearly possible that output content produced may violate copyright or other IP rights in previously published materials. Several IP claims against generative AI owners have been lodged for allegedly unlawful copying and processing of millions of copyright-protected images, and associated metadata.[12]

34. Further, one should be careful not to use, in response to system prompts, words which may breach trademarks or give rise to a passing-off claim. Again, barristers must make sure they understand the terms of service of the LLM they are using in this respect.

Professional considerations

35. Irresponsible use of LLMs can lead to harsh and embarrassing consequences, including claims for professional negligence, breach of contract, breach of confidence, defamation, data protection infringements, infringement of IP rights (including passing off claims), and damage to reputation; as well as breaches of professional rules and duties, leading to disciplinary action and sanctions.

36. There is a growing body of material in which practitioners and others discuss their use of LLMs in the course of legal practice. This guidance is concerned only to explain some of the pitfalls. It is for barristers themselves to work out how and in what context a LLM might assist them in providing legal services. This process is likely to be a changing one as the technology itself develops as it is doing and with increasing speed.

[12] Cases such as (1) Getty Images against Stability AI Inc. for copyright infringement in AI training data; (2) Class actions in US against OpenAI challenging ChatGPT by Paul Tremblay and Mona Awad, and Sarah Silverman, Christopher Golden and Richard Kadrey and others suing OpenAI and Meta.

37. Barristers should also keep abreast of relevant Civil Procedure Rules, which in the future may implement rules/practice directions on the use of LLMs; for example, requiring parties to disclose to the court when they have used generative AI in the preparation of materials. This approach has already been adopted by the Court of the King's Bench in Manitoba[13] and the Civil Justice Council has setup a working group to consider specific rules for the use of AI in civil court proceedings.[14]

Conclusion

38. In conclusion, technical progress and the pressures of competition may lead to the increasing adoption of AI, including LLMs. The best-placed barristers will be those that make the effort to understand these systems and, if appropriate, use them as tools in their practice, while maintaining control and integrity in their use. There is nothing inherently improper about using reliable AI tools for augmenting legal services; but they must be properly understood by the individual practitioner and used responsibly, ensuring accuracy and compliance with applicable laws, rules and professional codes of conduct.

[13] https://www.lawgazette.co.uk/news/canadian-judges-demand-to-know-if-ai-used-insubmissions/5116452.article

[14] https://www.judiciary.uk/related-offices-and-bodies/advisory-bodies/cjc/current-work/use-of-ai-in-preparing-court-documents/