

# The City of London Law Society

## Legal professional privilege in the Age of AI

22<sup>nd</sup> April 2026

1. Can I start by thanking the City of London Law Society for arranging this event this afternoon. Of course this was going to be hosted at Simmons & Simmons today but the current circumstances mean we have moved entirely online.
2. The reason why we were going to be at Simmons was that this all came about as a result of a conversation I had with Colin Passmore last year about AI. No doubt you all know Colin. He was the senior partner in that firm until 2021 and is the Chair of the City of London Law Society. However what you may not know is that Colin also sits as a fee paid Deputy District Judge. And that gives me the opportunity to thank him for his work as a judge, and also to draw attention to the point that there is, without question, a role for solicitors in our judiciary at all levels. The District Bench has always recruited solicitors to its ranks, but in addition I would like to take this opportunity to note today that in the Chancery Division, of which I am the head, there are there are more former solicitors as High Court judges than there have ever been. 20% of Chancery High Court judges have that background. The idea that solicitors cannot be appointed as High Court judges is demonstrably false. I would encourage anyone with interest in judicial work to speak to someone, and if you can't find anyone else to speak to, please speak to me.
3. But this afternoon's event is not concerned with judicial recruitment, it is about AI, so let me turn to that.
4. The potential use of computers in the law has been well understood for many years. There is a Society for Computers and the Law of longstanding. Lord Justice Brooke was an early proponent of it in the 1990s and today, although a surprising amount of paper can still be seen in courts and legal offices, in reality the bulk of legal work is done on, or at least involves, computers. However until recently the computers one encountered in the law were of the traditional kind, in other words they did not involve machine-learning or so called artificial intelligence or "AI".
5. But machine learning has always been part of computer science, right back at least into the 1960s. And in the 1980s Professor Richard Susskind was one of the pioneers, identifying the potential for machine learning systems to undertake legal work. And then, jumping ahead by a few more decades, in November 2022 ChatGPT burst onto an unsuspecting world. Moreover since that time, these new machine learning systems, the large language models, are continuing to become more and more capable before our eyes. These systems are what we mean when we talk today about AI.
6. It is worth reflecting on what makes this machine learning/AI so different from the computers we were used to. I think there are two features. The first is that we now have machines which at least appear to operate on the basis that they understand English.

That is new and significant. I say “operate on the basis that” because although it is convenient to say simply that the machine itself understands English, I do not want to get into a debate about what that means in terms of anthropomorphising. The second is that unlike traditional computers, the way these systems appear to operate is probabilistic. Again I am not trying to get into a debate about how they actually work. My point is that they will not answer the same question in the same way every time. There is a degree of variation. One cannot fully predict, in advance, what the machine will say or do in response to a given set of circumstances.

7. There are many other specific features of these systems but I believe these two features together are at the heart of both the potential of these systems and some of the challenges they throw up. Overall and bearing this in mind, the stance I believe we should take is one of cautious optimism. There are good reasons to take real care but also good reasons for optimism about the potential.
8. Now in the judiciary of England and Wales an important response to these developments was the Judicial Guidance on the use of AI. The original guidance came out within about a year of the launch of ChatGPT and the current form of the Guidance was published in October 2025. The Guidance has three aspects.
9. The first point is that judges are not prohibited from using AI. The decision to do so is a matter for the individual judge. However the second point, is to make clear that judges do and must take full personal responsibility for whatever goes out in their name. So, although it might seem like a statement of the obvious, the content of a judge’s judgement is their full and sole responsibility.
10. The third aspect is that a judge who wishes to use any kind of AI system should only use a system which they are sure is secure. When the Guidance was first introduced it was clear that there were many public systems, such as public ChatGPT in which one’s data was not secure. That was both in the sense that entering data into such a system risked violating data protection law, and also because the AI system itself would or could learn from what was entered into it. These were good reasons why it was not appropriate for court data to be entered into those systems.
11. However thanks to important work done internally involving collaboration between the judiciary, the Judicial Office, His Majesty’s Courts and Tribunals Service (HMCTS) and the Ministry of Justice’s AI unit, we now have two AI systems which we are satisfied are secure. One is a secure form of the Microsoft Copilot system and that is available to all judges in England and Wales. The other is a secure system based on in-house development with by HMCTS and MOJ.
12. We have online guidance and training material provided via the Judicial College, which includes among other things a Copilot prompt library. And we are running a series of webinar training courses, starting at the end of April, with a focus on deep fakes.
13. Let me mention four specific examples of what is going on:
  - a. First, working with HMCTS and MoJ, we are looking at using AI as an in house transcription tool. This is very exciting and has potential to make a big difference in all our courts and tribunals.

- b. Second, we have found that AI is useful for helping judges produce anonymised judgments. This is a particular issue in the family area (but not limited to that). The proper anonymisation of judgments is a good thing because it helps transparency in the justice system, by allowing details of cases to be published. The way this works is that the judge prepares their reserved judgment in the normal way. The AI is then used as a tool to help the judge produce an anonymised version, by making suggestions. The judge can use them or not, and of course it is the judge who is responsible for the final version. However the tool has been found to be valuable and, in particular, some judges have commented that the AI has identified pieces of information as candidates for anonymisation, which are not the obvious things to redact (like the names and so on). The AI identified information combinations which might risk a kind of jigsaw identification of the individuals concerned.
- c. A third example is something I only came across relatively recently. This is to use AI to identify internal inconsistencies in my own work. So what I have now done since then is thus. Once I have finished writing a judgment, I give it to the secure Copilot system on my computer and ask it to identify any internal inconsistencies. It is remarkably effective. What I choose to do with the proposals is up to me. I don't always agree with the AI but it has been helpful and I have clarified wording in draft judgments as a result.
- d. The fourth and final example is the use of AI for administrative tasks. With the form of Copilot available to leadership judges like me, the ability to find things in emails and files has been transformed. No longer do I need to do word searches on old emails.

14. But let me turn to legal professional privilege.

15. Thinking back to the months after ChatGPT was launched, when you started trying it out one of the striking things about using it was the slightly eerie sense that you were communicating with a person. You would type in a question (or even speak to it) and the system would respond in a manner which computers really never had before. The response created a facsimile of an interaction with another person. And it was quite clear you could ask it for what we can all identify was legal advice – and it would give it to you.

16. And like courts all over the world, we are seeing an increase in use of AI by unrepresented litigants in drafting the material they present to the court. Now the volume of this material can present a new challenge for us but it is worth emphasising that there is also a sense in which this use is pro-access to justice. The material can be very long and not right, it is not the case that all of it is wrong or of poor quality. My own experience, and the experience of other judges I speak to, is that quite often the litigant's case is presented more clearly and coherently than I would have expected in similar circumstances in the past.

17. Just as a test I tried using public ChatGPT last week. I typed in "*my builder broke my window when fixing my kitchen. Can I sue him?*". It responded in this way:

***"Short answer: yes, potentially—but it depends on the circumstances and evidence.***

*In the UK, this kind of situation usually falls under **negligence** or **breach of contract**. Here's how it breaks down:"*

18. The system then dealt five topics – negligence, contract, remedies, writing to the builder and finding out if they have insurance, and then how to go to court, particularly the small claims track. It ended:

*"If you want, tell me:*

- *roughly how it happened*
- *whether you had anything in writing*
- *and whether the builder is accepting responsibility*

*I can give a more precise view on how strong your case is and what to do next"*

19. Now the obvious point is simply that this interaction is the same as an exchange between a new client and their legal professional. And let us assume too that the "client" and the "adviser" then exchange more information – so that the "adviser" can provide a more precise view. If the adviser was a legal professional, the whole exchange would be covered by legal professional privilege.

20. Now this privilege is a fundamental constitutional right in England and Wales. In *Three Rivers* [2004] UKHL 48 Lord Scott at [34] famously credited Professor Adrian Zuckerman (in his textbook) with a rationale for this privilege as one based on the rule of law itself. One aspect of that rationale which I think has particular force is the point that it can only be fair to say that ignorance of the law is no excuse if everyone has the facility to find out what the law is. So the role of the lawyer is to level the playing field between those who are legally sophisticated and those who are not, ensuring that the presumption of knowledge of the law is not devoid of reality.<sup>1</sup>

21. The common law originally attributed the privilege to people interacting with legal professionals. And for a long time there have been attempts to expand the scope of the common law doctrine by reference to the nature of the adviser. For example in the 1980s there was an attempt to expand that common law concept to include legal advice from patent agents (*Wilden Pumps v Fusfield* [1985] FSR 159) but that failed<sup>2</sup> and it was done by legislation in s280 of the Copyright Designs and Patents Act 1988.

22. An important case was and remains *R (Prudential) v Special Commissioners of Income Tax* [2013] UKSC 1. Here the Supreme Court decided not to extend the common law privilege beyond the legal profession to include other advisers such as tax accountants. The judgments of Lord Neuberger (who gave the main judgment of the majority) and Lord Sumption (dissenting) demonstrate how difficult this area is and set out the pros and cons of extending the common law.

23. So if one wanted to extend the common law concept of this privilege to a situation like the one I have described above with ChatGPT (and I am not suggesting either way that

---

<sup>1</sup> See Zuckerman on Civil Procedure 5<sup>th</sup> Ed (2026) at 16.10.

<sup>2</sup> This was based on the common law because the Patents Act 1977 had only extended the privilege to cover patent agents advising in pending or contemplated patent proceedings.

one should or should not), the attempt to do so would need to confront the issues identified in *Prudential*.

24. However there is another dimension to consider here, confidentiality. What if – as public ChatGPT is – the AI service used by the individual is a public system? In the recent decision (November 2025) of *UK v Secretary of State for the Home Department* [2026] UKUT 81 (IAC) the judges in the Upper Tribunal had to deal with two cases in which the legal representatives of a party had used AI for research and to draft documents for use in court, with the result that the material provided to the court contained fake case citations. Most of the case is about the hallucinations but the case also addresses a consequence of what took place. At [60] the judges held as follows:

*“Uploading confidential documents into an open-source AI tool, such as ChatGPT, is to place this information on the internet in the public domain, and thus to breach client confidentiality and waive legal privilege, and any such conduct might itself warrant referral to the regulatory body and should, in any event, be referred to the Information Commissioner’s Office.”*

25. Now confidentiality has always been a pre-requisite for the attraction of the privilege, so even if one did extend the concept of the privilege to include advice from AI, on the current approach, it would not seem to attach to the interactions with these public AI systems because they do not appear to be confidential. Now there may be more to be said about the precise factual position of information entered into these public systems, but on any view one cannot assume that confidentiality is preserved.
26. But before I finish I would like to mention a further question – which I suspect is close to many lawyers’ hearts. One can imagine a common scenario in which a client comes to a lawyer for advice and it is the lawyer who uses AI to help them provide the advice to the client? Does that use of AI impact on the client’s privilege? (because of course the privilege we are talking about is the client’s not the lawyers). Now assuming the lawyer has used a secure system, I would suggest that it is hard to see how that could have an impact on privilege. Legal professionals are entitled to consult other sources of legal advice - text books (paper or online) and so one might think that ought not to undermine the privilege.
27. Now overall my purpose this afternoon is to identify live issues and to frame the debate we are going to have this evening on this panel. I am looking forward to an interesting discussion.
28. Thank you.