

THE MANCHESTER LAW SOCIETY

AI Conference 2024: Transforming the Legal Landscape

**Keynote Speech by Sir Geoffrey Vos Master of the Rolls and
Head of Civil Justice**

AI – Transforming the work of lawyers and judges

Friday 8 March 2024 at 3.05pm

This speech was supported by AI generated images, which are referenced throughout.

Introduction

1. It is an honour to have been invited back to Manchester to talk at this important AI conference.
2. I guess I am speaking to the converted, but it is, in my view, incredibly important that lawyers and judges get to grips with new technologies in general and AI in particular. AI is changing the way things are done in every conceivable sector of the global economy and the legal sector is no exception. The problem is that some lawyers and judges are, even now, hoping that they will be able to retire before they have to, as you might say: “get with the program”.
3. I have done so. My first image today is Dall-E’s view of what it looks like when I sit in an AI-technology enabled court.
4. The first serious thing to say is that there is nothing scary about AI. It is just a technological tool that has, by the way, been around for years. You use it happily every time you pick up your smart phone.
5. What is scary, as always, is a very small number of ill-intentioned people. Such people might use AI inappropriately if we do not protect ourselves properly, and build in human controls. But that is not really any different to other

technological developments that history has produced. Cars, aeroplanes, industrial machinery, oil, mining and almost every other technological innovation can be very dangerous to people, and even to humanity itself, if misused.

6. The second thing to say is that everyone is talking about AI now because, in March 2023, generative AI was made widely available with the launch of ChatGPT and then a series of other Large Language Models or LLMs, such as Google Gemini and Microsoft CoPilot. I guess that all of you have now had a go at using these LLMs. If you have not, I recommend you do. These LLMs are only the start. ChatGPT answers text questions. Dall-E makes images from text prompts, and Open AI's new product, Sora, produces videos instantly based only on a text description. I am sure, when you use these AIs in the future, you will bear in mind some of the things I am going to say in a moment.
7. The third thing to say about AI is that its use may rapidly become **necessary** in order to perform workplace duties. One may ask rhetorically whether lawyers and others in a range of professional services will be able to show that they have used reasonable skill care and diligence to protect their clients' interests if they fail to use available AI programmes that would be better, quicker and cheaper. I will return to this in a few minutes.
8. This afternoon, I want to try to explain where I think the use of AI is taking lawyers and the legal system. But before I do so, I want to refer briefly to two sets of principles. First, two guiding principles as to technology that I have adumbrated to several conferences over the last year. Secondly, the Judicial Guidance for the use of AI that I and other senior judges have recently promulgated.

Guiding principles

9. The first principle takes us back immediately to what I have just said about the need, in the future, to use AI.

10. It is the principle that we all owe a duty to those we serve – namely citizens and businesses here in England and Wales – to make constructive use of whatever technology is available if it helps to provide a better, quicker and more cost effective service to clients and the public, if you are a lawyer, and to provide a better, quicker and more cost effective dispute resolution process if you are a judge.
11. The second principle is that it is an integral part of the adoption of new technologies that we need to do all we can to protect the very same citizens and businesses from their adverse effects. That means that, where appropriate, we need to promote effective regulation, rule-making, data protection, the protection of confidential material, and the minimisation of cyber-crime and cyber-fakes. All these present risks to the communities we serve to a greater or lesser extent.
12. But none of that means that we should forsake new technologies and the benefits they bring. Many fear, as I have said, that they pose threats to the way things have always been done. And they really do. But the simple fact is that we would not be properly serving either the interests of justice or access to justice if we did not embrace the use of new technologies for the benefit of those we serve.

Judicial Guidance for the use of AI

13. The messages contained in the Judicial Guidance are very simple. They apply just as much to lawyers as well.
14. They can be summarised as follows.
15. First, before using generative AI, you need to understand what it does and what it does not do. Generative AI does not generally provide completely reliable information, because the LLM is trained to predict the most likely combination of words from a mass of data. It does not check its responses by reference to an authoritative database. So, be aware that what you get out of an LLM may be inaccurate, incomplete,

misleading or biased. To cheer you up, I can show you an image of how ChatGPT thinks that an LLM works.

16. Secondly, lawyers and judges must not feed confidential information into public LLMs, because when they do, that information becomes theoretically available to all the world. Some LLMs claim to be confidential, and some can check their work output against accredited databases, but you always need to be absolutely sure that confidentiality is assured.
17. Thirdly, when you do use a LLM to summarise information or to draft something or for any other purpose, you must check the responses yourself before using them for any purpose. In a few words, **you** are responsible for your work product, not ChatGPT.

How then will new technologies affect legal practice?

18. Having said all that, there are many things that AI in general, and generative AI in particular, can do for lawyers that is likely to save time and money and to be of great value.
19. Last weekend, I asked ChatGPT how it could help a solicitor practising in Manchester. In a second, it told me that it could help with (i) legal research, (ii) drafting legal documents and contracts, (iii) factual and legal analysis of cases, (iv) continuing legal education, (v) drafting letters to clients, (vi) writing memos, briefs and opinions, (vii) discussing ethical considerations, (viii) translating, (ix) giving guidance on LawTech, and (x) practice management. I have been tracking answers to similar questions over some months, and I would say that the answers are slicker and more focused than they were.
20. Google Gemini answers the same question in a similar way, but emphasises that it can summarise complex information and large datasets, and can provide a “legal Q&A chatbot” and help with market research on competitors, legal demographics and legal trends in Manchester.

21. LLMs are very good at suggesting **draft contracts**. I have been truly amazed at how quickly they can be produced. Of course, they need checking and amending, but that process takes a fraction of the time it would take a lawyer to draft a contract from scratch.
22. There is no doubt that AI can save a vast amount of time in creating legal contracts for employment, company takeovers or sale and purchase agreements, just as examples. It will take time to get used to, and, as I have said, requires careful checking.
23. Generative AI will also be useful **in the creation of legal advice and court submissions**? But it was here that early adopters came a cropper, probably because they failed to understand that generative AI is prone to hallucination – it can make up facts, because of the way it operates. It is not in any way similar to Wikipedia or a data base searching system.
24. Even that problem does not, however, mean that generative AI will be of no use in these technical areas, but it means that serious caution is required. This image shows how Dall-E sees judges making decisions in a court room using advanced AI.
25. But there are in fact two points to be made about using generative AI in drafting advice and submissions.
26. First, it is very likely that specialist large language models trained on specialist legal data will be more accurate for lawyers. Such things already exist. One is called Spellbook. But they are not yet commonplace. There is, however, no reason why an AI could not be trained only on, for example, the 6,000 pages of the CPR or on the National Archives case law database, BAILLI, Westlaw, or Lexis Nexis, but unable to scrape the bulk of the internet. Such a tool would be likely to give answers that would be more accurate and useful than a public LLM.
27. Secondly, questions to generative AI programmes must, as I have already mentioned, take into account the way those programmes work. They will often give unhelpful literalist answers. The user needs to be skilful to get the best from AI

- that is particularly so in its valuable programming applications.
28. Just as an example of that last point, I asked Microsoft CoPilot who was the best AI lawyer in Manchester. It gave me the names of three firms that it would be invidious for me to share with you, and then warned me that the word “best” was subjective and depended on my specific needs and circumstances. I felt duly ticked off.
 29. LLMs can also, of course, also **predict case outcomes**. I would have thought that any litigation client would want to know, if they could, what an AI thought as to their prospects of success. That opinion could be compared with the opinion of their human lawyers. Since the AI has access to more and different data than the humans, its opinion would at least be worth taking into consideration.

What other effects will the inevitable adoption of AI have on the legal landscape?

30. Here, I think, there is a need for lawyers to start thinking out of the box.
31. Using AI is unlikely to be optional. First, clients will not want to pay for what they can get more cheaply elsewhere. If generative AI can draft a perfectly serviceable contract that can be quickly amended, checked and used, clients will not want to pay a lawyer to draft one instead.
32. Secondly, in a similar vein, if AI can summarise the salient points contained in thousands of pages of documents in seconds, clients will not want to pay for lawyers to do so manually. Having said that, ChatGPT tells me that it limits text input to some 450 words (or tokens), but Google Gemini (after a bit of an argument about whether the information I wanted was publicly disclosed) told me that it can accept more than 10,000 tokens (or words). CoPilot says it can accept 4096 tokens.

33. Thirdly, and perhaps more importantly, AI is not only quicker, but may do some tasks more comprehensively than a human adviser or operator can do. The consequence of this reality is that we may need to reconsider the way in which the common law applies to a vast range of activities.
34. I have already spoken quite extensively about the changes that automated decision making, whether in self-driving vehicles or in any other context will have on liability issues. It seems to me that, there are, perhaps, some even more fundamental questions for the common law.
35. As I have already said, the current law of negligence is based on the proposition that human beings will exercise reasonable care, when undertaking a range of activities, such as driving or operating machinery and actually just doing their jobs. Lawyers and other professionals are made liable for failing to exercise all reasonable professional skill and care in giving legal advice, in designing buildings, in providing auditing and accountancy services, in providing actuarial services and in every other imaginable field.
36. I have just been talking about the prospect of using AI for legal research. The same discussions are going on in all the professions. Just imagine, for example, that an AI tool was available to accountants, which was capable of quickly and easily identifying the warning signs of fraud. Would a firm that relied on old fashioned auditing methods, and shunned the use of AI, be exercising all reasonable professional skill and care to protect its client? You can ask the same question about lawyers, architects, computer programmers, health and safety executives and everyone in fact.
37. It seems to me that we may need to consider how the rapid advance of AI may affect the foundational principles of our common law. There may, for example, need to be reconsideration of the implication of terms, the regulations concerning unfair contract terms, and a range of other legal and regulatory provisions. There may even need to be a re-evaluation of the nature of the duty of care.

38. Cast your mind over the commonly taught legal subjects, such as company and insolvency law, the law of intellectual property, contract and tort, the law of property, and even criminal law. The rapid adoption of AI tools will, I think, potentially affect every one of them.
39. I asked my friendly LLMs how they thought AI would affect legal liability. CoPilot told me that AI creators might be liable for any injuries if the AI products were defective when made. Self-consciously, it told me that use of generative AI could create technology and data risks that may not be fully understood as the technology is developing.
40. Gemini had rather greater insight. It told me that algorithms are increasingly used for loan approvals, risk assessments and even sentencing recommendations; it suggested liability might arise if the outcomes were biased or discriminatory. It told me, as is the case, that intellectual property liability may arise since LLMs are trained on large datasets, so that their training may involve potential copyright or patent infringement. Cases are already in progress in London on this question. Sorry – that is me talking, rather than Gemini. ChatGPT added contractual and professional liability and regulatory compliance to the list.
41. Now is not the time to expand on these thoughts, but I expect you will be able to see for yourselves how we will all need to rethink the way we do things, and the law may need to rethink how we allocate liability for things that go wrong in a world of capable AI tools.

Judicial use of AI

42. I expect you will not want me to finish this talk without mentioning how judges can use AI, and whether judicial decision-making is likely any time soon to be driven by AI.
43. The senior judiciary would not have issued the somewhat cautious judicial AI guidance that I have mentioned if we had

not thought that judges were as likely as any other group to be assisted by AI tools.

44. Once again, I asked my three AI friends what they thought. Gemini thought that it could not directly assist judges in making legal decisions, but that it could enhance their efficiency and research capabilities by analysing “vast amounts of legal documents and case law”, doing legal and factual research and drafting legal opinions. CoPilot thought much the same, save that it added unconscious “bias awareness” to the list of things it could help with. When asked what unconscious biases judges should be most aware of, it listed confirmation biases, contextual bias, racial and gender bias, and cognitive biases as the main headings. Interestingly, ChatGPT gave much the same answer, when asked how it could help judges as it had given when asked how it could help a Manchester solicitor!
45. More seriously, I think judges will need to become just as familiar with the use of AI as any lawyer. First, many cases will concern liability for the use or non-use of AI as I have been explaining. Secondly, if AI can help, without breaching confidentiality, in summarising complex material, there is no reason, in theory at least, why it should not be used for that purpose. Thirdly, AI is likely to be a valuable tool in the context of the digital justice system that is now being created both for court systems and for the pre-court ecosystem. As you may know, I chair the new online procedure rules committee, which is going to make rules and provide data standards for both the online court processes and the pre-action online dispute resolution processes within what I have previously described as the “funnel” of civil justice. AI will be needed to make these systems smart and to ensure that they operate as the parties, witnesses, experts and judges participating in them would expect.
46. I will leave over the question of whether AI is likely to be used for any kind of judicial decision-making. All I would say is that, when automated decision-making is being used in many other fields, it may not be long before parties will be asking why routine decisions cannot be made more quickly, and

subject to a right of appeal to a human judge, by a machine.
We shall see.

Conclusions

47. I hope that you do not think I am being alarmist. I asked Dall-E to create an image of lawyers alarmed by AI and I got this. I was myself shocked at the fact that it portrayed only men, and asked it to show me another image of female lawyers alarmed at AI.
48. For my part, I think it is important to take the adoption of AI one step at a time. But it is equally important not to assume that capable tools employing AI are likely to go away. We can and should make appropriate rules and regulations, but we will not be able to stop something that has beneficial uses for our society.
49. In my view, the task of understanding how the private law and regulatory backdrop needs to be adjusted to cater for the mainstream adoption of AI, cannot be started soon enough.
50. AI has great potential within the digital justice system which promises to provide quicker, cheaper and more efficient ways to resolve the millions of disputes that arise in British society every year.
51. Many thanks for listening. I look forward to answering your questions.