



MASTER OF
THE ROLLS

THE RIGHT HON. SIR GEOFFREY VOS

LAWTECHUK GENERATIVE AI EVENT

LONDON

WEDNESDAY 5 FEBRUARY 2025

1. I am delighted to have been given a slot at the Lawtech UK's Generative AI event.
2. As many of you will know, I am often asked to speak about AI and the Law. Only last week, I spoke at the launch of Justice's "AI in Our Justice System" Report. I was struck by the reactions of some of the lawyers in the audience: nodding vigorously when the risks of AI are mentioned and freezing when it was suggested that even lawyers might have to find ways to use AI to expedite and reduce the cost of both legal advice and dispute resolution.
3. I think it is imperative to build bridges in the legal community between the AI sceptics and the AI enthusiasts. There is no real choice about whether lawyers and judges embrace AI – they will have to – and there are very good reasons why they should do so – albeit cautiously and responsibly, taking the time that lawyers always like to take before they accept any radical change.
4. I want to touch on some of those reasons this morning.
5. First and foremost, the legal system and lawyers themselves serve all other industrial, financial and consumer sectors. All those sectors will be using AI at every level – there is simply no way that lawyers can set themselves apart and say that GenAI is too dangerous or the work of lawyers is too precise to use it.
6. Secondly, one of the biggest fields of legal activity in years to come is likely to be the claims that will be brought in respect of the negligent or inappropriate use of AI, and also the negligent or inappropriate failure to use AI. Lawyers will, as I have frequently said, be at the forefront of these AI liability disputes. That is why the UKJT is embarking on preparing a legal statement, similar to the ones it has prepared in relation to legal questions concerning digital assets, asking and answering questions like: "In what circumstances, and on what legal basis, will English common law impose liability for physical and economic loss caused by the use of AI? How does vicarious liability apply to loss caused by AI? When can a professional be liable for using or failing to use AI in the provision of their services?"



MASTER OF
THE ROLLS

7. If lawyers are not adept at the understanding the capabilities and weaknesses of generative AI, they will not be able to advise their clients properly about the issues that will undoubtedly arise from its applications.
8. The third reason why lawyers and judges must embrace generative AI, is that it will save time and money and allow advice to be given and disputes to be resolved far more quickly and efficiently.
9. That is why I am so committed, through the OPRC, to the creation of the Digital Justice System, which will allow millions of disputes to be resolved online, using AI where appropriate, without the need for those disputes to enter the more expensive and time-consuming court process. This will operate across civil, family and tribunals disputes. It will bring together all the existing online providers of ombuds services, mediation providers, arbitration providers, third sector online legal advice and information platforms such as Advice Now and the CAB, and, I very much hope, legal aid online.
10. Whenever I say that generative AI will save lawyers time and money, someone pipes up with the example of a lawyer who used GenAI to write submissions which included a fictitious case reference. The first and best example of that was the hapless Steven Schwartz in New York, who got his comeuppance from Judge P. Kevin Castel (who I have recently met). But that is what I mean. We should not be using silly examples of bad practice as a reason to shun the entirety of a new technology.
11. AI tools are not inherently problematic, so long as we understand what they are doing, and use them appropriately. For that reason, we published our *Judicial Guidance for the use of AI* last year. There are 3 simple messages in that guidance that apply as much to lawyers as to judges.
12. First, before using generative AI, you need to understand what it does and what it does not do. Large Language Models are trained to predict the most likely combination of words from a mass of data. Basic GenAI does **not** check its responses by reference to an authoritative database.
13. Secondly, you must avoid inputting confidential information into public LLMs, because doing so makes the information available to the world. Some LLMs claim to be confidential, and some can check their work output against accredited databases, but it is paramount that that confidentiality is always guaranteed.
14. Thirdly, when you do use a LLM to summarise information, draft a document or for any other purpose, you must carefully review its responses before using them elsewhere. In a few words, you are responsible for your work product, not ChatGPT.
15. Last week the Supreme Court of New South Wales published its practice note on the subject. Its rules include the following:



MASTER OF
THE ROLLS

- (1) Gen AI must not be used to generate the content of affidavits, witness statements, or other material that is to be used in evidence or cross examination.
 - (2) Where Gen AI has been used in the preparation of written submissions or skeleton arguments, the author must verify the accuracy of all citations and authorities.
 - (3) Gen AI must not be used to draft or prepare experts' reports without prior leave of the Court.
16. It will be interesting to see how that more restrictive approach in New South Wales works out as compared to our approach. I would comment, though, that AI is already being used in many jurisdictions for some of the purposes that the NSW guidance says it should not be. I doubt we will be able to turn back the tide. Our guidance is within the grain of current usage, making clear that the lawyers are 100% responsible for all their output, AI generated or not.
 17. So, to summarise, there are three excellent reasons why all lawyers and judges should embrace AI: those we serve are using it. It will make what we do available to more people, more cheaply, and allow us to do necessary things more quickly, and it will be at the centre of the future work of lawyers, when claims are all about when AI has been used for the wrong things, and AI ought to have been used but was not used.
 18. Before finishing, I want to mention something that, I think, is very important in relation to the future uses of generative AI.
 19. I recently gave the Blackstone Lecture, which concluded by suggesting that the use of AI within the justice system was creating a completely new situation that necessitated a re-think about the fundamental rights of humans. I suggested that current regulatory tools such as the EU's AI Act might not be sufficient to protect people from decisions being made by machines. It might not be enough, in the age of hugely capable machines, to legislate, as does the AI Act, that decisions must be taken by humans.
 20. In a world in which machines are **so much** more capable than humans, it may become simply too time-consuming and expensive for anyone to check the integrity of every decision that machines recommend humans to make. Even now, the levels of your benefits and pensions are likely calculated by an algorithm. I suggested that we, as humans, would want to decide which types of decision were, in the future, genuinely a human prerogative, and which types of decision we were content to be taken by machines. I was certainly not suggesting that machine-made decisions were inappropriate in many cases, but I suggested an early debate about the detail.
 21. The legal community, internationally, not just here in the UK, needs to consider what kinds of advice and decision-making should and should not be undertaken by a



MASTER OF
THE ROLLS

machine. I suggested in my Blackstone lecture that it was fairly obvious that people would never have the requisite confidence in peculiarly human decisions, like whether children should be removed from their parents, being made by machines. But some of the distinguished Oxford academics present questioned my assumption. They thought that emotive decisions of that kind would be just the type of decision-making that parents would really prefer to be taken out of human hands. I don't know who is right. But what the disagreement shows is that we need to start an urgent and wide-ranging discussion about what we want machines to do, and more importantly what we feel that machines should not be allowed to do.

22. Reduced to the essentials, the reason why we might decide that particular advice or decision-making should not be undertaken by a machine must be in order to protect fundamental human rights. The question I posed in Oxford was whether the existing European Convention on Human Rights needs itself to be revisited in the world of machines that will undoubtedly be more capable than humans.
23. Human rights and the rule of law remain fundamental to the justice system. It must be uncontroversial that we should always use generative AI with an eye to promoting and improving access to justice and the quality of decision-making. AI provides massive opportunities in that regard.