



MASTER OF
THE ROLLS

THE RIGHT HON. SIR GEOFFREY VOS

LAW SOCIETY OF SCOTLAND

LAW AND TECHNOLOGY CONFERENCE

ONLINE LECTURE

WEDNESDAY 14TH JUNE 2023

Introduction

1. I am honoured and delighted to be able to deliver this talk to the Law Society of Scotland's Law and Technology Conference. I am sorry that I have not, on this occasion, been able to come to Edinburgh, but I promise you I was there only a few weeks ago.
2. It is three short months today since GPT-4 was released by OpenAI on 14 March 2023. I doubt that even the developers realised the transformational effect that it would have, even if it is a multimodal large language model called "Generative Pre-trained **Transformer 4**".
3. This morning, I want to explore a little the effect that generative AI is already having, and is likely to have, on legal services and on dispute resolution methods, including the courts.

The story of Mr Schwarz

4. I was struck a couple of weeks ago by an article in the Times about a lawyer in New York, one Steven Schwartz. He was an early adopter of ChatGPT and used it to prepare his submissions to Judge P Kevin Castel in New York. The case concerned Mr Schwartz's client's injury when, in 2019, a metal trolley had struck his leg on a flight in the USA. Mr Schwartz's submissions included 6 cases which the judge described as "bogus decisions with bogus quotes and bogus citations". One of these cases was called *Varghese v. China Southern Airlines*. It was allegedly heard in 2019 before the 11th circuit of the US Court of Appeals, but it was in fact almost entirely fictional. The reference number for the case was a reference for an extradition decision, and no case involving anyone called Varghese had appeared in the preceding decade. What was perhaps more alarming was that Mr Schwartz had asked three follow-up questions. He asked ChatGPT first if *Varghese* was a real case and it had replied in the affirmative. When asked for its source, it said "Upon double-checking, I found that the *Varghese* case does indeed exist". When asked if the other cases it had provided were fake, it said "No, the other cases I provided are real and can be found in reputable legal databases". Mr Schwartz had later to apologise and to answer at a sanctions hearing for the "citation of non-existent cases".



MASTER OF
THE ROLLS

5. Let that be a lesson to you all! But there is a very serious point here. Mr Schwartz was not uncovered because of the language of his brief, but because the judge in that case took the trouble to look up the cases cited. No doubt that does not always happen. The risks of litigants in person using ChatGPT to create plausible submissions must be even more palpable. And indeed such an event was reported as having happened in Manchester only a few days ago.

Should the development of generative AI be paused?

6. There is, of course, an active debate raging at the moment about whether advanced machine learning is actually proceeding too fast for humanity to cope with. For my part, I somehow doubt that there will ever be sufficient international cooperation for there to be an effective pause in the research and development of generative AI. We may hope instead for effective regulation, but even that may take some time and may be somewhat left behind by the speed of the developments themselves.

What mechanisms will the legal community need to deal with the use of generative AI?

7. In my view, we are going to have to develop mechanisms to deal with the use of generative AI within the legal system. We may even hopefully turn it to the advantage of access to justice and effective and economical legal advice and dispute resolution.
8. The story of Mr Schwartz demonstrates that one thing generative AI cannot do effectively for lawyers is to allow them simply to cut corners. I suspect that non-specialised AI tools will not help professional lawyers as much as they may think, though I have no doubt that specialised legal AIs will be a different story. Spellbook is already claiming to have adapted “GPT-4 to review and suggest language for your contracts and legal documents”.
9. ChatGPT itself identified its own most valuable uses in the dispute resolution field in an article published on 4 May 2023 online by Enyo Law. IT said that its most valuable uses were to assist lawyers with (a) drafting, (b) document review, (c) predicting case outcomes to inform strategy, and (d) settlement negotiations.
10. As I said in a recent lecture, and I think is very important for lawyers to understand, clients are unlikely to pay for things they can get for free. Mr Schwartz would have done well to read Enyo Law’s article, which emphasises that the large language model is there to “assist” the lawyers and needs to be carefully checked. Nonetheless, if briefs can be written by ChatGPT and Spellbook, checked by lawyers, clients will presumably apply pressure for that to happen if it is cheaper, and saves some of an expensive fee earners’ time.



MASTER OF
THE ROLLS

11. In the end, as Mr Schwartz's case also shows, in litigation at least, the limiting factor may be the court or tribunal adjudicating on the dispute. One can envisage a rule or a professional code of conduct regulating whether and in what circumstances and for what purposes lawyers can: (i) use large language models to assist in their preparation of court documents, and (b) be properly held responsible for their use in such circumstances. Those will be things that the existing rules committees, regulators, and the new Online Procedure Rules Committee (announced in England and Wales this week) will need to be considering as a matter of urgency.
12. I will return in a minute to the question of robo-judges that has so taken the imagination of legal journalists in recent months. First, let me say something on the question of alignment.

Alignment of the use of AI with human values

13. The issue of alignment of the use of AI with human values is a large one. Ought there to be more effort made to align the use of generative AI with human priorities, moralities, and values? Mr Schwartz once again illustrates this problem. According to the report in the Times, Mr Schwartz asked ChatGPT three confirmatory questions, no doubt because he suspected that his new-found arguments were too good to be true. He asked whether *Varghese* was a real case, and he was told that it was. That was true in part in that the reference was a reference to a real case but was certainly not a true answer to the question that Mr Schwartz had intended to ask.
14. This raises two problems: first, ChatGPT and AIs more generally need to be programmed to understand the full import of a human question. Secondly, humans using generative AI need to be savvier in checking their facts. Mr Schwartz could have asked: can you confirm that the case of *Varghese* is reported at such and such a reference, and that it decided such and such a point. Had he done so, he might have received a more accurate answer. Likewise, the programmers need to explain to the AIs they are creating what humans mean when they ask something as open textured as "is this a real case". What they mean, of course, is to ask whether everything ChatGPT has said derives from that case, does indeed derive from it: its reference, its name, the words it has reported as used by the judge etc. This requires careful and detailed programming beyond what might be required in other fields of activity.
15. There is, I am sure, much more work that could and should be done to align generative AI with human values and principles, but that would be costly and time-consuming – at the moment, perhaps, more energy is being expended on making ChatGPT's and other AI's work products look and feel as if it emanates from humans as opposed to making them absolutely accurate and aligned with human morality.



MASTER OF
THE ROLLS

16. If GPT-4 (and its subsequent iterations) is going to realise its full potential for lawyers in providing accurate legal advice, accurate predictions of legal outcomes and accurate assistance with dispute resolution processes, it is going to have to be trained to understand the principles upon which lawyers, courts and judges operate. As Mr Schwartz found to his cost, the present version of ChatGPT does not have a sufficiently reliable moral compass. In the meantime, court rules may have to fill the gap.
17. So, what then will be the effect of the future use of AI on the work of lawyers and judges? Will lawyers and judges soon be redundant? I don't think so.

Robo-judges and Robo-lawyers

18. I have already said that I have no doubt that lawyers will not be able to stand aside from the uses of generative AI. Clients will insist that all tools available are at least considered for application within the delivery of legal services.
19. But will judicial decisions be taken by machines rather than judges? As many of you will know, we are introducing in England and Wales a digital justice system that will allow citizens and businesses to go online to be directed to the most appropriate online pre-action portal or dispute resolution forum. That digital justice system will ultimately culminate at the end of what I regard as a "funnel" in the online court process that is already being developed for pretty well all civil, family and tribunal disputes.
20. I think that AI will be used within digital justice systems of the kind we are creating in England and Wales. It stands to reason that it should be used to enable everyone to be fully informed of the process that is being undertaken. It can also be used to help people read and understand complex sets of rules and instructions, by limiting the material from which the answers to questions can be taken.
21. I believe that it may also, at some stage, be used to take some (at first, very minor) decisions. The controls that will be required are (a) for the parties to know what decisions are taken by judges and what by machines, and (b) for there always to be the option of an appeal to a human judge.
22. The limiting feature for machine-made decisions is likely to be the requirement that the citizens and businesses that any justice system serves have confidence in that system. There are some decisions – like for example intensely personal decisions relating to the welfare of children - that humans are unlikely ever to accept being decided by machines. But in other kinds of less intensely personal disputes, such as commercial and compensation disputes, parties may come to have confidence in machine made decisions more quickly than many might expect.



MASTER OF
THE ROLLS

23. After all, there are many other areas of professional expertise where artificial intelligence is already more reliable than human advisers – take the old chestnut of the diagnosis of melanomas. The machine has seen many more skin cancers than any doctor and its advice is already, I believe, a valuable adjunct to the tools available to medical professionals.
24. If generative AI can predict outcomes of litigation, as ChatGPT claims, clients will want to know what it thinks, even if ultimately, they will also, for the moment anyway, want to know what a real-life lawyer thinks of the predictions so obtained.
25. Dr Irina Durnova, who prepared (rather than wrote) the article for Enyo law, to which I have already referred, records ChatGPT's own advice:
 - ChatGPT can analyse large amounts of legal data, including case law and case summaries, to predict how a particular case will likely be decided. This can help lawyers and clients make more informed decisions about settling a case or taking it to trial. **However, it is essential to note that ChatGPT is not infallible, and there is always a degree of uncertainty involved in legal decision-making.**
26. I might interpose: "you can say that again". ChatGPT's advice continues:
 - Additionally, the accuracy of ChatGPT's predictions will depend on the quality of the training data used to develop the model and the complexity of the legal concepts involved in the case. **The current version of ChatGPT has limited knowledge of the events (including cases) that took place after 2021, and it might be missing crucial legal developments.**
27. To interpose again, that is a critical limitation for lawyers to understand. And perhaps ChatGPT's own conclusion bears emphasis too:
 - **... while ChatGPT has the potential to be a valuable tool for predicting case outcomes, it should be used in conjunction with human judgment and expertise.** Ultimately, legal decision-making involves a range of factors beyond just predicting the outcome of a case, **including strategic and ethical considerations and client goals.**
28. That sounds like good advice for us all. And is perhaps a lifeline for the future of the legal profession.
29. I look forward to being able to ask ChatGPT for the answers to your questions!