



Master of the Rolls

THE RIGHT HON. SIR GEOFFREY VOS

KEYNOTE SPEECH

HUMAN RIGHTS, ALGORITHMIC JUSTICE, AND GLOBAL AI POLICY CONFERENCE
GLOBAL POLICY INSTITUTE, DURHAM UNIVERSITY
THURSDAY, 15TH MAY 2025

Introduction

1. It is a great pleasure to have been invited to speak here at Durham University for the first time. I am sorry that I have missed the earlier contributions to this conference. I hope I shall not be reprising ideas you have already discussed.
2. A series of engagements starting in Lausanne and culminating in the Blackstone lecture in Oxford last November led me to suggest that we should be considering whether and how the European Convention on Human Rights and Fundamental Freedoms (ECHR) might be amended to protect the rights of humans in the rapidly approaching machine age. I would loosely define the “machine age” as the age of hugely capable AI.
3. It may be worth mentioning at the outset that, when I discuss these ideas with many lawyers, they do not accept the premise that hugely capable AI, without significant hallucination, is ever likely to be available in the legal space. As will appear, I am willing to accept that premise, whilst acknowledging that we are not quite there yet.
4. I don’t want to repeat what I said in my previous lectures tonight. I want instead to try to take the thinking a little further forward.
5. I think the key to the potential human rights problem caused by AI is predicting what humans really care about and what humans will come to accept. In my lifetime, I have been constantly surprised by the cultural and societal changes that I have observed. I vividly remember debating at school whether it would ever be possible to enforce speed limits with cameras. The prevailing view was that people would never accept that kind of restriction on their liberty. Nowadays, of course, speed cameras are accepted and acceptable, as are street cameras and even facial recognition cameras. Much of what people accept in their everyday lives passes under the radar.



MASTER OF THE ROLLS

6. It is for that reason that I have had cause to review my earlier, perhaps too unequivocal, statement that people would never accept machines making peculiarly human, emotional, or empathetic decisions about their lives. I am still not convinced that such acceptance would occur easily. But I accept the arguments put to me by others that some humans, but certainly not all, may actually prefer machines to take the emotion out of, for example, a decision concerning the custody of children or the end of a life, by allowing it to be made by an AI decision-making tool.
7. Let me, therefore, base what I am about to say on the premise that there will be a valid and, perhaps hard fought, debate about what decisions machines should be allowed to take and what decisions machines should never be allowed to take. Before that debate begins in earnest, it is worth considering some regulatory and other parameters. In other words, where and how are the lines to be drawn?
8. It seems to me to be already clear that, in many parts of the world, some sort of red line is being drawn between legal advice being delivered by machines, on the one hand, and judicial decision-making being delivered by machines, on the other hand.
9. We already see AI tools giving simple legal advice online in the employment, property, and consumer fields, to name but a few. I have no doubt that there will be significant uptake as these tools become more readily and cheaply available. After all, a “client” can always seek human legal advice if they do not like or cannot properly understand what the machine predicts or advises.
10. AI is also already used to decide many administrative issues such as the amount of a pension or a social security benefit. AI decides how much we are all charged by major corporations, utilities, and many other consumer facing organisations. As a society, we generally do not object to the use of AI in these cases, notwithstanding that article 22 of the GDPR (and article 22 of the UK GDPR) gives data subjects “the right not to be subject to a decision based solely on automated processing”, which produces legal effects on them. I have recently given a lecture about the issues thrown up by article 22, but we do not need to take those questions further this afternoon.



MASTER OF
THE ROLLS

11. The EU has suggested in its AI Act that judicial decision-making and, even, decisions made in the context of the administration of justice, are to be treated as different and special. The AI Act makes AI systems concerned with the administration of justice into “High Risk AI systems”. The definition in article 6(2) and paragraph 9(a) of annex III to the EU’s AI Act is as follows: “AI systems intended to be used by a judicial authority or on their behalf to assist a judicial authority in researching and interpreting facts and the law and in applying the law to a concrete set of facts”. That definition may only include systems that advise judges, but it would be surprising if it did not also encompass the decisions themselves. I wonder whether that definition will prove too broad and how robust it will be when tested. The EU is, of course, now working on an AI Liability Directive to complement its AI Act.
12. In countries outside Europe, there is much less scepticism about the appropriateness of the potential, and even actual, use of AI, at least to assist, in judicial decision-making. China, in particular, already uses AI in several judicial decision-making contexts – something that I believe you have already heard quite a bit about today. And other countries are experimenting with how AI can expedite and streamline the work of judges. There is little doubt expressed, in such countries, that AI will be capable of reaching reasoned decisions. Some jurisdictions are already using machines to guide or direct judges towards their decision-making and to ensure consistency within their justice processes. Some will undoubtedly ask why Western Europe is so resistant to the speed and efficiency that automated judicial decision-making might be able to bring.
13. We can, I think, safely assume that, in time – probably not all that much time – machine made decisions will come to be as reliable, if not more, reliable than human decisions, in many areas at least. I have long suggested that litigants in personal injury claims might well prefer to have an AI decide the level of their damages on well-established mechanical principles, rather than wait for two years to achieve a similar, perhaps identical, result from a human judge. That would surely be even more attractive if an appeal to the human judge remained an available option.
14. So, I think, against that background, there are three questions that must now be asked in the context of highly capable AI tools.
15. First, is judicial decision-making different from other kinds of decision-making, and, if so, why?
16. Secondly, is judicial decision-making different from legal advice, and if so, why?
17. Thirdly, what rights should humans have to protect just and fair decision-making in addition to article 6 of the ECHR?



MASTER OF THE ROLLS

Article 6

18. Article 6 provides that: “[i]n the determination of his civil rights and obligations or of any criminal charge against him, everyone is entitled to a fair and public hearing within a reasonable time by an independent and impartial tribunal established by law”.
19. One may ask first whether an AI could ever be regarded as “an independent and impartial tribunal established by law”. That is itself a difficult question to which I shall come in a moment. If it is accepted that an AI **could be** independent and impartial, then one needs to ask whether other rights may be needed to protect the legitimate rights of humans in a world of AI decision-making. If a machine **cannot be** an independent and impartial tribunal, then article 6 may be breached every time any judicial decision is informed by machine-learning. In that event, article 6 might need to be softened or adjusted to take account of the assistance that AI can properly offer to the judicial decision-making process. Either way, we need to consider what other rights may be needed to protect humans in the age of hugely capable AI.
20. It is worth noting that the “established by law” qualification will depend on domestic rules and provides some threshold protection for individuals. But ultimately, I am assuming that some domestic legislatures will at some stage validate some form of machine-made intervention in the judicial decision-making process.

Is judicial decision-making different from other kinds of decision-making, and, if so, why?

21. I have already alluded to the question of trust and confidence. The key to the article 6 right is the requirement that, in every dispute, whether with the state or between citizens or businesses, a fair and open hearing is guaranteed before an independent and impartial judge. The transparency of the hearing is critical to ensure public scrutiny of the delivery of justice, and the confidence of the public in the process.
22. The problems, as I see them, with judicial decisions being taken by machines, so far as transparency and confidence is concerned, include the following.
23. First, whilst the outcomes of machine-made decisions can be tracked and evaluated, it will become increasingly difficult to check either the methodology or the process itself.
24. Secondly, it may not be clear how the AIs and algorithms that result in the decisions are being trained or what steps are taken to ensure that bias is reduced or eliminated.
25. Thirdly, it may be hard to assess the levels of human trust and confidence in the process when some 50% of people, by definition, win their cases, and 50 % of people lose.



MASTER OF THE ROLLS

26. It may be, as I have said, in previous lectures, that people will eventually, perhaps easily, come to accept routine judicial decisions being made by machines, particularly if an appeal to a human judge is available. But I think we need to be very clear that judicial decision-making is the last line of defence for the liberty of the individual.
27. Thus, to answer my own question, I am absolutely sure that judicial decisions are different from administrative decisions that governments or large corporations may make about our lives. In all those other cases, the individual can appeal to the court to review or challenge the decision on public or private law principles. Judicial decisions are, to put the matter bluntly, the end of the road.
28. The difficulty with this analysis is that, if there remains an appeal to a human judge, then many will ask why we should not take advantage of technology to streamline the process by utilising AI to inform judicial decisions or take routine decisions as part of the process, leaving the human judge as the decision-maker of last resort. I will return to those issues.

Is judicial decision-making different from legal advice, and, if so, why?

29. I also think that judicial decision-making is different from legal advice. In that respect, I agree with the EU's approach. Availability of independent legal advice is one element of the rule of law. Independent legal advice feeds into the independence of the judiciary because the work of the courts is made harder, if not impossible, without an independent and trustworthy legal profession.
30. It is entirely clear that any justice system is less effective without an independent legal profession. But ultimately legal advice does not automatically affect the rights of humans, whilst judicial decisions do.
31. In my view, therefore, the issues that affect the desirability of using machines to give legal advice are, at least, different in scale, if not quality, from those that affect the desirability of using machines to make judicial decisions.
32. Moreover, I do not think that human lawyers will ever be replaced by machines. As our lives become more complex, as regulation becomes more intense, and as technology affects every aspect of our society, citizens and businesses will find it increasingly hard to understand the law and the legal position. The explanations that lawyers provide will become of increasing importance.



MASTER OF
THE ROLLS

What rights should humans have to protect just and fair decision-making in addition to article 6?

33. If I am right thus far, human judicial decision-making might need some protection. The level of protection required may, as I have said, depend on whether an AI could ever be regarded as an “independent and impartial tribunal” within article 6. But even if it could not, the question remains as to the appropriate use of AI in the run-up to human judicial decision-making.
34. I come back to the problem that I alluded to in the Blackstone lecture, namely what is to happen when it just becomes too expensive and time consuming to allow humans to sit as judges because the machine can decide everything far more quickly and cheaply. And even if we do legislate for human judicial decision-making, should the humans be allowed simply to endorse the AI’s advice as to what decisions to reach? It seems likely that it will be very expensive and time consuming for humans to check the AI’s advice in every case. That is, even if, as machines become cleverer, that would actually be possible.
35. In understanding why protection is necessary and what that protection might be, there are a number of factors that are worth identifying.
36. The first thing to understand is that there is not always, in judicial decision-making, a single right answer. Most legal systems give judges a significant latitude in terms of discretion. Sentencing often has an element of discretion. Relief from sanctions in the civil context does too. Appeals are not allowed simply because the appeal court thinks it would have reached a different conclusion on the facts, were it to have decided the case at first instance.
37. Secondly, I think there is a risk that the development of the law would be stultified if machines took all the decisions. As I have said, in many legal situations, there is more than one right answer. What is to happen when one machine reaches one answer, and another machine reaches the opposite answer? We already know from experiments in some jurisdictions that machines can happily give equally compelling reasons for a positive or negative decision in a whole range of case types.
38. Ultimately, though, the type of further protection required depends on the answer to the question I have now posed twice already. Can an AI be regarded as an “independent and impartial tribunal”?
39. I think there are strong arguments that an AI cannot, or at least should not, be regarded as an “independent and impartial tribunal”.



MASTER OF THE ROLLS

40. Machines are a function of their training. Machines can, of course, be trained to mimic independence, but they cannot actually be independent, because their response is always dictated by their programming and learning. The same may be said of humans, of course. Humans are also the product of their upbringing and environment and education. But we understand the words “independent and impartial” as applied to humans to encompass those demographic and societal differences.
41. But, as it seems to me, the crucial difference is that machines cannot actually have human emotions, even human empathy. They cannot actually show mercy or sympathy. They cannot laugh or cry. I still think, despite the push back in relation to my Blackstone lecture, that humans genuinely value the idiosyncratic qualities that define our humanity.
42. Drawing the threads together, as I see it, the essential distinguishing feature of judicial, as opposed to other kinds of, decision-making, is the fact that it is the last resort of every individual and business, and even the state. The ability to bring legal challenges to the actions of large corporations, other citizens or businesses or the state is fundamental. The judiciary is the third arm of state, because the decisions made by judges are final, whether they affect rights to life, liberty, or property.
43. Article 6 is, therefore, a critical protection. Since machines will arguably never be able to mimic the perhaps peculiarly idiosyncratic nature of independent and impartial human judicial decision-making, I would suggest that machines cannot, therefore, properly be regarded as an “independent and impartial tribunal”.
44. If that were to be accepted, the debate would move to the question of what protections were needed to prevent machines making, anyway final, judicial decisions by default. How can the economic and time pressures driving the path towards machine-made decisions be resisted? This also engages the need to work out what we actually want humans to do in the future, even in situations where machines can do whatever that is equally effectively, and certainly more quickly and cheaply.
45. Since, as I have already said, I accept that humans will differ as to what types of decision should be reserved for humans, this is a very difficult question to answer.
46. In human decision-making, the quality, even the reasoning, of the outcome is not necessarily what leads to maximum levels of party acceptance. In end-of-life or children cases, for example, the question may not be whether the right answer has been reached, but more about whether society as a whole accepts and empathises with the outcome.



MASTER OF THE ROLLS

47. The types of judicial decisions (if any) that are accepted as being appropriate for AI tools to take will affect the way we interact as humans, and also the operation of our economies and social structures.

Conclusions

48. I can summarise briefly. Article 6 is crucial to the debate I started with my Blackstone lecture. The critical question is whether an AI can ever be regarded as an “independent and impartial tribunal”. My current view is that an AI is probably not properly to be regarded as an independent or impartial tribunal within the meaning of article 6. I can, though, see both sides of that argument.
49. But, even if article 6 does protect individuals against automated judicial decisions, that only leads us back to the question I first posed, namely how can we protect ourselves against the economic and time pressures that will push justice systems towards using AI to advise how judicial decisions are taken, even if they are ultimately actually approved by humans?
50. The debate about what sort of judicial decisions humans will accept being taken by machines, and how the rights of individuals can and should be protected in this new environment of highly capable AI, must continue.

Sir Geoffrey Vos.