

**BLACKSTONE LECTURE**

**Pembroke College, Oxford**

**Are rights sufficiently human in the age of the machine?**

**Sir Geoffrey Vos, Master of the Rolls and Head of Civil Justice  
in England and Wales**

**Wednesday 27 November 2024**

**Introduction**

1. I am grateful to Sir Ernest Ryder for inviting me to deliver this lecture.
2. Perhaps the most popular contemporary subjects for judicial articles and lectures are artificial intelligence and climate change. I have certainly not been backward in coming forward to make my contribution on the former, if not the latter, subject. Tonight, I want to take a step back from the usual arguments about whether AI can or should be used for various hitherto human tasks, whether AI is likely to make lawyers and/or judges redundant and whether either or both of public and private law need to be adjusted to reduce environmental damage caused by the continuing global use of fossil fuels.
3. I want to ask, even if I cannot answer, a more fundamental question that should, I think, be concerning the modern-day legal community. That question, put at its broadest, is whether the current international legal order, as it affects the rights of humans, is fit for purpose in the light of the changes and challenges of what I think we can now call, the machine age.
4. I was struck, when speaking last week at an international conference on Artificial Intelligence at the Swiss Institute of Comparative Law, by the different approaches that judges and lawyers from different parts of the world are taking to AI and to the changes to our legal landscape that it is causing. In essence, as it seems to me, most are more comfortable speaking about what is going on at the operational level, rather

than venturing to consider higher level principles. I will explain what I mean.

5. On the subject of AI, the Chinese speaker<sup>1</sup> explained how Chinese judges are using AI in their decision-making, and how that enables them to deal with their large volume of cases and to ensure that courts produce consistent judgments in comparable cases. There is obvious and appropriate awareness of the risks of hallucination and bias, but these problems are tackled rather than used as a reason for objecting to the deployment of capable new technologies.
6. Conversely, the Francophone African speaker<sup>2</sup> was less sanguine. He was concerned that big tech companies and the Western world more generally neither understood nor appreciated the problems that AI is causing and would cause in African judicial systems, and warned against re-colonisation resulting from an irresistible push towards using AI in a continent, where there is: (a) less digital literacy, (b) less internet coverage and power capacity, and (c) endemic corruption which makes AI tools a particular risk.
7. The US judge<sup>3</sup> thought AI was a useful aid to lawyers and judges, but that it could never replace human decision-making in trial courts. He was extremely concerned about cyber-fakes and the injustices potentially caused by their deployment in family and criminal cases. To make his point, he posted online an avatar of himself introducing the conference speaking fluent French and German, when in reality he speaks only English.
8. I was, however, particularly interested by the French administrative judge<sup>4</sup> who understood completely the problem that AI causes in government administration. He did not want

---

<sup>1</sup> Dr Zhiyu Li, Associate Professor in Law and Policy at Durham University.

<sup>2</sup> Judge Jean Aloise Ndiaye of the Supreme Court of Senegal.

<sup>3</sup> Judge Scott Schlegel of the 1st District of the Louisiana 5th Circuit Court of Appeal.

<sup>4</sup> Judge Marc Clément, Presiding judge at Administrative Tribunal of Lyon.

to stop its use, recognising that that was impossible, but wanted it to be understood that the European and French regulations that say that automated decisions cannot be taken where they affect a person's legal rights are easily circumvented by a machine making the calculations and 'advising' the human decision-maker as to what they need to decide. He recognised something very important, namely that, as machines become more and more capable, economic imperatives will prevent human decision-makers from questioning the products of artificial intelligence. There will simply not be the time or money to allow each of perhaps thousands or millions of administrative decisions concerning pensions, benefits or immigration, as examples, to be checked by a human. The regulation that says decisions must be taken by humans can be satisfied by the human signing off on a machine-prepared schedule of many decisions.

9. These perspectives covered a good part of the world, albeit excluding much of the global South and India.
10. What they demonstrated to me is that we need to be careful to ensure that we do not sleep-walk into usages of extremely capable AI that change what humans do and how they do it for good, without our having even embarked on a debate about the fundamental rights of those humans in an age of AI and climate change.
11. Let me, then, make some introductory remarks about climate change, because the problems are related. First, the rights of humans are indeed being asserted in relation to climate protection. I can point to the recent decision of the Grand Chamber of the ECtHR in *Verein KlimaSeniorinnen Schweiz v. Switzerland*<sup>5</sup> where it was held that article 8 could be construed as providing certain environmental protection rights against the state. Likewise, in the recent case of *Milieudefensie v. Royal Dutch Shell plc* in The Hague Court of Appeal, where it was explained that the citizen claimants had not established that Shell had a "social standard of care"

---

<sup>5</sup> Case number 53600/20.

to reduce its emissions by 45% or any other amount. But the court did acknowledge an obligation on Shell owed to citizens to limit emissions.

12. Human rights have always been rights that vest in individuals against groups of people and the state. Historically, business was generally subservient to the state, but what we see now is that the rights of humans are being asserted in the environmental field against both states and international corporations. Moreover, even if a single small state, such as Switzerland, were forced to reduce its emissions to net zero, that would not necessarily solve the world's climate change problem. This too seems to me to indicate that we may need to take another look at the rights that humans ought to have against both states and big business for the coming generations. I am not sure that there is any real thinking going on about the adequacy of the legal protections for the rights of humans in the light of the dramatic effects of more frequent heat waves and floods affecting every part of the world, and certainly the global South.
13. To summarise this introduction, then, I want to examine in this lecture whether it is right to say that AI and climate change are creating completely new situations that necessitate a re-think about the fundamental rights of humans.
14. The last time such a re-think occurred was after the Second World War, which led to the drafting and ratification of the European Convention on Human Rights and Fundamental Freedoms (ECHR). The question is whether the changes to the current legal order is such as to merit revisiting that Convention and possibly others. I am not suggesting that, even if such a process were desirable, it would be easy to achieve in the current precarious global political situation. But that is not really the point. If the international legal community were to think that the fundamental rights of humans needed to be reformulated in the light of these changes, then the debate would need to begin many years before a satisfactory outcome could be expected to be achieved.

15. I will now address: (i) whether the current legal and regulatory approach to automated decision making and the use of AI is fit for purpose, (ii) whether the current legal and regulatory approach to environmental protection is fit for purpose, and (iii) if not, what can or should be done to remedy those lacunae.

Is the current legal and regulatory approach to AI fit for purpose?

16. There is, or will soon be, a spectrum of hitherto human tasks and decisions that machines will be able to take for us or help us to undertake.
17. That spectrum starts at the left hand extreme with broadly mechanical decisions, like those about the amount of a pension or benefits or the calculation of personal injury damages and loss of earnings. It may be that humans are, or at least will become, fairly relaxed about those kinds of decisions being taken by machines, since the algorithms will be fairly transparent. The use of machines will undoubtedly save a vast amount of time and money. An appeal to a human judge, if the machine gets the decision wrong, would probably add to the human confidence in using AI in such areas.
18. More complex advice or decision-making lies in the middle of the spectrum. We are told by big tech companies that the situation is likely soon to be reached whereby AIs will be able to give legal advice very quickly and cheaply. It seems inevitable that such advice will be available in circumstances in which a lawyer would take hours or days to do the same thing, or even to check the accuracy of what the AI has advised. The same problem applies to judges. If the AI will be able, in future, to write a judgment in minutes, even if the human judge is required by regulations to take the final decision, how is that judge to check if the AI's "advice" about the outcome and the written judgment of the machine are correct without doing work that may take that human days or even weeks. And what will happen in this or any other field if the human cannot check the answer given by the machine? I would suggest that the human decision-maker may find themselves obliged to accept what the AI or maybe two AIs

are telling them are the answers. In the legal sector, it seems unlikely that clients will want to pay for their lawyer to spend days checking an AI's legal view once they have established the reliability of that machine. It is perhaps equally unlikely, in that situation that litigating parties will want to wait months for a human judicial decision, when a machine assisted one can be taken for little money almost immediately.

19. Cases that are peculiarly human lie at the farthest end of the spectrum. To take some examples, we can consider: (a) the question of whether a child should be removed from its parents, (b) the question of whether a life support machine should be switched off, (c) the question of whether a murder should be treated less seriously as a result of provocation causing severe emotional distress, and (d) sentencing questions that arise where emotional mitigation is promulgated to the court. Most would probably say now that it is inconceivable that our societies would accept such decisions being made, or perhaps even assisted, by a machine. They depend on human empathy, which is something machines do not have and which, at the moment at least, machines cannot easily be trained to replicate.
20. Against this background, I think that there are three factors that should determine where, in this spectrum of decisions, it may be acceptable to use AI either to advise or decide: first, human confidence in the technology, secondly, the pressures created by the higher cost of using humans instead of machines, and thirdly ethics or fundamental human rights. Let me address each of these in turn.

#### *Confidence in technology*

21. We know that AI is already used very extensively indeed by public and private entities. Apple, Google, Facebook and Amazon and Open AI use it all day long to “assist” everyday tasks performed by all of us all the time. Governments already use AI to calculate the pensions, benefits and taxes of individual citizens and much more. Ordinary people are often not even aware of the many algorithms that are in use to affect their lives. It may be, though, that they do not care. Humans

seem generally to have a measure of confidence in the algorithms that are used in social media, and even by Government departments. They have a level of confidence, perhaps not complete confidence, that they will generally be treated fairly. And they believe that, if not, they will have the right to bring a claim before an independent court or judge in their home jurisdiction. But ordinary citizens probably do not yet have such a level of confidence in machines deciding matters that turn essentially on human empathy.

#### *Economic pressure to use AI*

22. It is easy for techno-sceptics to say that economics will never affect our ability, as humans, to decide what machines should do and what they should not do. In fact, though, the harsh realities of economic power make that far more difficult in practice. The big tech companies seek to reassure us that we will always have a choice, that machines will only be assisting us, and the final decisions will always be made by humans. That is the principle underlying the EU's AI Act and even the GDPR. But, as I have already said, I doubt whether such regulation is, or certainly in the future, will be effective.
23. First, even now, with the routine uses of AI that I have mentioned, there is only modest push-back. Secondly, AI will be relatively cheap at the point of use, even if its consumption of power can be alarming. So, once an AI can read the lawyers' case papers, retrieve the relevant legal materials, and advise as to the answers, the lawyers (and even the judges) may find it hard to maintain that the same work needs to be done again, or even checked, by humans with all the expenditure of time and money that that would involve.
24. In short, once the AI is good enough, the human adviser or decision-maker may be left with no choice but to accept the solution proposed by the AI, because that human will not have the resources to challenge it.

#### *Ethics and fundamental human rights*

25. If that is right, it seems to me that our society has a growing ethical problem. Generally, most humans would probably

prefer to decide what is suitable to be decided by a machine and what ought still, despite the AI's undoubted capability, to be decided by a human – even in the machine age. I cannot say what decisions are in which category, but I doubt that any of us would want all decisions, even those involving human empathy at the far end of my spectrum, to be made by a machine. If that were to happen, there would arguably be an existential challenge to our humanity, and to the democratic rights of our citizens. The law is ultimately all about how we relate to one another. It is not an end itself.

26. The problem is how we decide, as technology grows exponentially in capability what decisions **must** still be taken by humans, and how we stop the inevitable pathway towards humans formally taking these decisions, but being forced by economic pressures, in fact, to accept the advice or suggestions of ever-more capable machines.
27. This ethical problem is said by some to be easily resolved by regulation. I should say at once that I do not agree.

*The role of regulation*

28. There are a number of relevant existing regulations that attempt to stop automated decision making.
29. Article 47 of the French Data Protection Act (Law 78-17) of 6 January 1978<sup>6</sup> seems to have been the prescient origin of article 22 of the GDPR. It provides that a data subject has “the right not to be subject to a decision based solely on automated processing of personal data ... which produces legal effects concerning him or significantly affects him”.<sup>7</sup>
30. Article 22 of the GDPR (and of the UK GDPR) also gives data subjects “the right not to be subject to a decision based solely

---

<sup>6</sup> Loi n° 78-17 du 6 janvier 1978 relative à l'informatique, aux fichiers et aux libertés.

<sup>7</sup> In the original French: Aucune décision produisant des effets juridiques à l'égard d'une personne ou l'affectant de manière significative ne peut être prise sur le seul fondement d'un traitement automatisé de données à caractère personnel, y compris le profilage, à l'exception.



on automated processing”, which produces legal effects on them. The CJEU’s decision in the *SCHUFA Holding* case<sup>8</sup> (*SCHUFA Holding*) indicates that article 22 may really prohibit automated decision making. But, of course, it does not prohibit human decision-making assisted by a machine.

31. Moreover, the EU’s AI Act makes AI systems concerned with the administration of justice into “High Risk AI systems”.<sup>9</sup> Such AIs are those “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”. High Risk AI systems are not banned, but they must be closely monitored.
32. The Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy, and the Rule of Law was adopted in Vilnius on 5 September 2024.<sup>10</sup> Its aspirational contents are epitomised by article 4 that provides that states shall adopt measures “to ensure that the activities within the lifecycle of artificial intelligence systems are consistent with obligations to protect human rights, as enshrined in applicable international law and in its domestic law”. This does not, of course, extend the rights of humans in relation to automated decision making.
33. The central point raised by all current regulatory techniques is whether it can really be sufficient, if, for most applications, it is sufficient for the machine to **advise** and for the human notionally to **decide**, based on the machine’s advice.
34. I have considered whether the existing ECHR and its sister, the American Convention on Human Rights of 22 November 1969 (ACHR) are sufficient to provide the relevant protections for human decision-making in areas where that is necessary and appropriate. For my part, I am not sure that

---

<sup>8</sup> Joined Cases C-26/22 and C-64/22 of 7 December 2023.

<sup>9</sup> Article 6(2) and paragraph 9(a) of annex III to the EU’s AI Act.

<sup>10</sup> It has already been signed by the EU, the USA, the UK and 8 other countries.

article 6 of the ECHR<sup>11</sup> covers the situation, because, even accepting that AI can introduce bias, we may well reach a stage where it will be demonstrable that AI decision-making is fair and impartial. Moreover, if the human makes the decision informed by AI, can it really be said that the tribunal is not independent and impartial. Secondly, I am not sure that article 8 of the ECHR<sup>12</sup> covers the question, even if it can be construed imaginatively as we saw in *Verein KlimaSeniorinnen Schweiz v. Switzerland*. Other articles in the ECHR, ACHR and the Charter probably come nowhere close.

35. It seems to me that the answer to the question I have posed in this part of my talk is that the current legal and regulatory approach to AI may well not be fit for purpose. If, as I think we should be, we are concerned as humans to be the ones deciding what decisions are to be taken and advised upon by machines and what decisions should not be, we will need to consider how that is to be achieved both nationally and internationally. What may be required is to ask the question of what additional rights, if any, humans should have to require business and governments to make transparent choices as to which decisions can and should be decided and advised upon by machines, and which by humans.

36. I will return to this problem in a moment.

Is the current legal and regulatory approach to environmental protection fit for purpose?

37. It may be that the *Switzerland* and the *Shell* cases that I mentioned earlier demonstrate that, in the current legal order, citizens can, in theory at least call both government and big business to account in respect of actual or threatened environmental damage.

---

<sup>11</sup> Article 8 of the AHRC.

<sup>12</sup> Article 11 of the AHRC.

38. The problem here is that, as the *Shell* case shows, there is no clear delineation of what **level** of climate damage is permissible. That was also made clear in my judgment in the Court of Appeal in *Regina (on the application of Friends of the Earth) v. the Secretary of State for International Trade and the Chancellor of the Exchequer*.<sup>13</sup> There we decided that the Government was only required to reach a tenable view as to whether its \$1.15 billion investment in a liquified natural gas project in Mozambique was aligned with the UK's obligations under the Paris Agreement of 12 December 2015. The Government was not required to obtain a **quantification** of the indirect "Scope 3" emissions caused by the project before making its investment decision. Moreover, the UK's obligations under the Paris Agreement, whose terms were not incorporated into English law, did not give rise to domestic legal obligations.
39. Thus, whilst article 2 of the Paris Agreement provided, in effect, that its 197 state parties should "pursue efforts to limit the [global] temperature increase to 1.5°C above pre-industrial levels", it is hard to see how citizens can, by themselves, legally require big business and governments to adhere to that international obligation.
40. Like the use of AI, it may be that there is not an easy regulatory solution. What may be required here too is, at least, to ask the question of what additional rights, if any, humans should have to require business and governments to preserve the environment in which they live.

What can or should be done to remedy these regulatory lacunae?

41. As I see it, the two problems that I have identified are connected. They are about the future roles of humans and human relations. They are also about the future relationships between humans on the one hand and businesses and their governments on the other for the future.

---

<sup>13</sup> [2023] EWCA Civ 14.

*Summary of the two identified risks*

42. As I have tried to explain, I see two **genuinely new** risks affecting humanity.
43. To summarise, the first is the fact that AI will undoubtedly, in a relatively short time frame, be able to undertake complex tasks, including those involving varying elements of human empathy, with as much or more reliability than humans themselves. There is no unanimity about what rights individual humans have or should have against international businesses and governments to control the utilisation of AI to advise upon or make decisions that were previously always advised upon and made by humans. The risk is that, if nothing is done, AI will eventually decide, or advise on the outcomes, of even the most empathetic of human decisions, affecting human lives.
44. The second risk is that, if nothing is done, it will not be clear what rights humans have to require businesses and governments to preserve the environment in which humans live. There is, of course, some unanimity, though not complete unanimity, that the lifestyle of the Global North and its use of fossil fuels has caused the world environment to warm and, therefore, change. There is no unanimity about what rights individual humans have or should have against international businesses and governments to control the nature and speed of that change.

*The need for international debate*

45. Both the questions I have posed are essentially existential, and they are not questions that can be resolved by any one country individually. As I mentioned earlier, Switzerland alone achieving net zero will not solve climate change. Moreover, AI is generally used in cross-border applications by multinational corporations. A solution adopted in one country is unlikely to work without some international agreement.
46. There are, in addition, many sceptics who believe the problems are less acute than I have suggested and many who think they are more easily resolved. I do not claim to have the

answers. My view is that it falls to the international legal community to consider these two questions carefully and to attempt, at least, to propose solutions that can be considered by national governments and international organisations.

47. I return to what was said by the impressive African speaker at the conference in Lausanne. He told us to beware of assuming, in the field of AI, that the problems in Africa were the same as we diagnosed them to be in the first world. His concern about recolonisation by AI may be something that could be said also by other parts of the Global South. It indicates to me at least that we must be careful not to assume unanimity of thinking without discussing the issues widely and carefully. It is even clearer that there are very different approaches to climate change as between the Global North and the Global South. Again, these differences should not be left out of account in the debate I am suggesting.

48. If I am right that existing regulations may not be sufficient to protect the rights of humans in the face of the growth of ever more capable machines and an ever more challenged environment: **what are we to do?**

*Some more concrete ideas as to what might be done*

49. The first thing for lawyers to do, I think, is to identify the roles, responsibilities and rights of humans that might be fairly generally agreed to need protection.

50. For centuries, the UK has made a significant contribution to the identification of the rights of humans. One can point to Magna Carta 1215, the Bill of Rights 1688-9, the Representation of the People Act 1832, and the UK's role in the formation of the ECHR. We have always had a clear understanding of the importance of individual freedoms and the rule of law.

51. This history might suggest that the UK could play a role in considering what legal protections are needed now to ensure that humans can still take essentially human decisions in the context of both AI and environmental change. It will be

important in this not connection not to make perfect solutions into the enemies of good solutions.

52. One could, perhaps, envisage a more far-reaching rule along the lines of article 22 of the GDPR's right not to be subject to a decision "based solely on automated processing, which produces legal effects". That can, however, still be circumvented by using the machine to advise, allowing the human to take a long list of scheduled decisions informed by the machine.
53. To be effective, it seems to me that a new foundational right would need to be created. It would have to address what truly needed to be protected, rather than the fringes of the problem.
54. Let me take some examples. It would serve no practical purpose to require humans to consent to the use of automated decision-making, because that consent could be obtained routinely as it is already for cookies and the like. Likewise, a requirement to be informed in advance about automated decision-making does not overcome the problem.
55. It seems to me also that it is hard to draw a simple line on the spectrum that I mentioned earlier. It is hard to say when a decision will involve human empathy, and whether human empathy is the only area where we, as humans, would prefer to keep machines out of the loop.
56. It might, in theory, be possible to have a fundamental right to have material human consideration of decisions requiring empathy or emotional intelligence. That would immediately raise the question of which decisions did require human empathy, but it might be a start for the debate. I accept that even mechanical or mathematical decisions can involve, or at least cause, human emotions, if only emotional reactions. For example, the injured pedestrians who are awarded less in damages by the machine than they thought they should receive might suffer annoyance and emotional trauma. Certainly, many of the complex decisions that I identified in the middle of my spectrum might raise or involve some element of human empathy.

57. The real difficulty is to identify areas where even AI-assisted decision-making infringes the fundamental rights of citizens, as we may wish them to be. That may involve, as a starting point, trying to identify, as I have said, where on the spectrum I have mentioned, the line is to be drawn. But even then, there are, of course, many spectrums in legal decision-making and many more in every other conceivable sector of consumer, financial and industrial activity.
58. There are many possible examples, but I can see that many of the decisions to the left of my spectrum might reasonably not be said to involve human empathy – the calculation of benefits and pensions are examples. I can envisage, though, that neither tech corporations nor governments would welcome a restriction on their ability to streamline decision-making by the use of machines just because it involved an element of human empathy.
59. The environmental problem is even more difficult. There are already some 2,500 pieces of climate litigation going on in the world today. Claimants struggle to identify what precisely they are entitled to require governments and business to do to protect the environment. The solution here could be an additional individual right, but it could also be a more granular international treaty. Both seem to me to be very difficult to achieve, particularly so in the light of the events at the recent Cop 29 summit.
60. But, as I said at the start, the fact that these two new problems are hard to solve is not, in my view at least, a reason for not trying.

### Conclusions

61. Let me try to draw the threads together.
62. My first conclusion is that existing forms of domestic, EU or international regulation are probably not competent to prevent AI being used inappropriately to make decisions that ought, for the benefit of humanity, to be taken by humans.

63. Secondly, the existing treaties and international conventions seem to provide individuals in different parts of the world with insufficient rights to impose limits on businesses and governments that are probably doing too little to prevent global environmental damage. It is obviously in the interests of humans to be able to limit environmental damage, just as it is in their interests to preserve appropriate human decision-making.
64. The international legal community is uniquely well-qualified to suggest solutions to both these problems. Such solutions will not, however, emerge until the problems are identified as existential ones that affect essential human rights and the rule of law. Once that is achieved, the hard work will need to begin. It will be difficult to gain any form of consensus, but that is, as I say, no excuse for not trying.
65. To conclude, as I started by asking, I do think it is right to say that AI and climate change are creating completely new situations that necessitate a re-think about the fundamental rights of humans.
66. I hope that I have provided some food for thought.

GV