The Mayflower Lecture

Rt Hon Sir Geoffrey Vos

The Roland Levinsky Building, University of Plymouth, Devon PL4 8AA

Thursday 12 November 2026

Human Rights in a Rapidly Changing World

Introduction

1. Many thanks to Chris Hoile, President of the Plymouth Law Society, for that kind introduction.

2. The Mayflower Lecture has now become a distinguished feature of the calendar here in Plymouth, having changed its name from the original title, namely the "Pilgrim Fathers Lecture". It commemorates, of course, the Mayflower's voyage from Plymouth departing in August 1620 and making landfall at Cape Cod near Plymouth, Massachusetts on 21 November 1620. It has traditionally, I am told, been associated with freedom and human rights. I was told, rather charmingly, in my briefing that "[t]he subject [of the lecture] is not restricted to the voyage itself".

3. That is really rather lucky, as whilst I have modest expertise in a few areas of history and the law, I cannot claim to be an expert on the voyage of the Mayflower.

4. I have spoken on several occasions about the risks that capable artificial intelligence (AI) poses to human judicial decision-making, when it becomes clear that, in many legal fields at least, machine-made decisions are as reliable, but much quicker and cheaper than human judicial decisions.

5. I have suggested that, in that situation, which I do not, as an opening aside, think is all that far away, economics might dictate the direction of travel. The economic pressures to which I am referring will operate on the parties to legal

disputes as much as on justice systems and those responsible for them. There will be economic pressure on parties to agree to an AI-model deciding their cases rather than waiting months or years for a human judge to do so. And there will be economic pressure on governments and justice systems to insist, whether by legislation or otherwise, on AI-assisted or AI-enabled judicial decision-making in order to clear case backlogs and to save judicial and court time.

6. In recent speeches to the Legal Geek Conference[1] and the ICC,[2] I explained why I foresaw difficulties in identifying a robust dividing line between: (a) cases that could quite sensibly and justly be decided by AI and (b) those cases which ought always, anyway in a democracy, to be the subject of human judicial decision-making.

7. On the one side of the line, we can all see how it might be thought to be perfectly reasonable for the parties to agree to: (a) small claims, (b) personal injury damages, or even (c) maintenance payments payable on divorce, being decided by a machine if those parties were able to appeal to a human judge. On the other side of the line, many would find it unimaginable for a machine to be allowed to decide: (a) whether children should be removed from their parents, (b) final appeals in place of an apex national court, or (c) on sentencing for serious criminal offences.

8. The reasons why many of us can understand how objectionable it might be for such decisions to be taken by a machine is not just about empathy, human emotions or the quality of mercy. It is also because judicial decision-making is the last opportunity that any individual or business has to assert their rights or to defend themselves from the power of the state.

---

1 https://www.judiciary.uk/speech-by-the-master-of-the-rolls-legal-geek-conference/

2 https://www.judiciary.uk/speech-by-the-master-of-the-rolls-the-future-of-cross-border-dispute-resolution/

9. The problem, as I have expressed it thus far, however, is that there is genuine disagreement about, the levels of human oversight that are needed if machines are to be used at all in the judicial process.

10. In England and Wales and the EU, there is a cautious approach. Our judges in England and Wales are provided with a closed LLM, from which data cannot escape, and may use it to assist them in their judicial work. Our judges are advised not to allow the LLM, to which they have access, to write their judgments, and they are cautioned that every word of their work product is their complete responsibility.[3]

11. In the EU, the AI Act has denominated law enforcement AI systems and those intended to be used in the administration of justice as "High Risk AI systems".[4] Article 14 of the AI Act requires such high-risk AI systems to be "designed and developed in such a way, including with appropriate human-machine interface tools, that they can be effectively overseen by natural persons during the period in which they are in use". These provisions, therefore, envisage that AI will only be used in judicial processes and decision-making if it is entirely overseen by human judges.

12. Article 22 of the GDPR (and of the UK GDPR) provides that: "The data subject shall have the right not to be subject to a decision based solely on automated processing … which produces legal effects concerning him or her …". Article 22 was considered by the CJEU in the *SCHUFA Holding* case.[5] Although *SCHUFA Holding* did not decide whether article 22 **prohibited** automated decision-making, paragraphs 42-44 of

---

[3] See Artificial Intelligence Guidance for Judicial Office Holders October 2025 at https://www.judiciary.uk/wp-content/uploads/2025/10/Artificial-Intelligence-AI-Guidance-for-Judicial-Office-Holders-2.pdf

[4] Article 6(2) and paragraph 9(a) of annex III to the EU's AI Act provides for the following precise definition: "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".

[5] Joined Cases C-26/22 and C-64/22 of 7 December 2023.

the CJEU's decision suggests that article 22 may give data subjects the right not to have solely automated decisions made about them, if they will result in legal consequences for them.

13. My own experience has, however, shown that different judiciaries and legal systems have very different opinions about how far it is permissible to use AI at all in judicial decision-making.

14. On one side of the argument are hugely respected judges like Chief Justice Bell of New South Wales, who has described his approach to AI as "relatively conservative".[6] He points to all the downsides of AI including "rubbish in, rubbish out", embedded biases in algorithms, algorithmic failure, cyber-security risks, and the still unsolved issue of hallucinations. Alongside, Chief Justice Bell, I heard a Philippines' judge in September 2026 at the Conference of Chief Justices of Asia and the Pacific describe AI as "an irresponsible lawyer and a catastrophic judge". Justice, he said, could not be automated.

15. But on the other side of the line, automation is already in full swing in China, Korea, Brazil, Argentina and Kazakhstan. Shenzhen Intermediate Peoples' Court published its White Paper on *AI in the Judiciary – the Shenzhen Practice* in September 2026 as a blueprint for the roll out of its small claims resolution model across China. The White Paper explains how in Shenzhen Intermediate Peoples' Court: "Judges no longer leaf back and forth between documents and systems to establish the facts of a case. Review becomes a matter of checking, correcting, and supplementing, and work that took a full day now takes tens of minutes". It continues: "The System has now assisted in drafting 678,000 judicial documents". The Shenzhen system allows AI to assist the judge by distilling the important points to be derived from the evidence, listing the issues that the machine thinks need to be decided in order to determine the case, and providing drafts of

---

[6] See Chief Justice Bell's 2026 Harold Ford Memorial lecture delivered at the University of Melbourne Law School: *Corporate responsibility and directors' duties in the era of Artificial Intelligence* at [26].

the reasons and order that the judge might want to provide to explain and encapsulate their decision.

16. But, in every jurisdiction, I have encountered thus far, it is made crystal clear that judges still have the last word in decision-making whatever role the machine may play in summarising, advising or drafting reasons.

17. These points should provide an adequate introduction to what I really want to discuss this evening, namely the subject of *Human rights in a rapidly changing world*. Do we need more or different rights to protect our societies against algorithmic decision-making or will the problem resolve itself, like so many apparently insoluble problems have in the past?

18. What additional fundamental rights, if any, are needed in the context of fake images and fake news, the use of facial recognition technologies by state authorities and police forces, and violations of data privacy and data protection in a world where AI tools are trained on materials drawn from the world wide web.

19. I shall address these topics through the lens of the European Convention on Human Rights (the ECHR), but they could equally easily be addressed through the lens of other human rights charters and national constitutions.

What rights are we talking about?

20. Let me start by setting out some of the existing rights in the ECHR. Many of them, as I say, are reflected elsewhere internationally.

21. First, article 6 of the ECHR provides that "[i]n the determination of [their] civil rights and obligations or of any criminal charge against [them], everyone is entitled to a fair and public hearing within a reasonable time by an independent and impartial tribunal established by law". It has not yet been determined authoritatively whether a machine or an AI tool can properly be regarded as an "independent and impartial

tribunal established by law". But I have repeatedly expressed the view that a machine might well **not** be properly so regarded.

22. Article 8(1), of course, provides that: "[e]veryone has the right to respect for his private and family life, his home and his correspondence". Will article 8 protect humans from intrusions by AI agents and other cyber-actors?

23. Article 10 provides that: "[e]veryone has the right to freedom of expression. This right shall include freedom to hold opinions and to receive and impart information and ideas without interference by public authority and regardless of frontiers". Article 11 provides that "[e]veryone has the right to freedom of peaceful assembly and to freedom of association with others". Do these rights apply to protect the producers of LLMs from claims in respect of outputs from LLMs?

24. Article 14 provides that: "[t]he enjoyment of the rights and freedoms set forth in this Convention shall be secured without discrimination on any ground". Does this apply to protect humans from the bias and hallucinations of machines? Will the adoption of AI more broadly create new kinds of problematic discrimination?

25. These rights are entrenched in our law under the Human Rights Act 1998, but we cannot assume that will always be the case. But even on the assumption that the Human Rights Act 1998 survives in the long term, is it sufficient in the age of ever more capable AI and quantum computing?

<u>Does article 6 require amendment or addition?</u>

26. Many commentators have spent time considering whether a particular machine with particular qualities might or might not satisfy the requirements of article 6. For example, the requirement of independence has been held to mean that the tribunal must be both independent from government and the

executive and from the parties.[7] One might analyse the algorithmic process itself to see whether in a particular case, it could satisfy those requirements.

27. Likewise, "impartiality" implies a requirement that the tribunal is not actually biased and appears to be neutral.[8] One might spend time considering whether an AI programme did or not have the appearance of neutrality. I wonder, however, whether these are the most useful questions.

28. Instead, it seems to me that we should be debating where the line is to be drawn between where a state **can properly** in a democratic society require its citizens to accept machine-made decisions, and where it cannot. If the parties to a dispute consent to machine-made decision-making, this question may be easy. But they will not always consent. I think, therefore, the real question is how far can machine-made judicial decision-making be properly **imposed** on commercial parties, individuals with administrative claims against the state, criminal defendants or parents in family proceedings.

29. A serious and informed debate about where the boundaries are to be drawn has not yet really taken place. Though I have mentioned several times that I think it should. If there were to be consensus, one could imagine that such consensus could be incorporated into an international convention or an amendment to the ECHR and other human rights conventions.

30. Such a debate would also circumvent two further problems created by article 6. First, whether a judicial decision taken or informed by a machine could be considered to have been taken at "a fair and public hearing", even if the machine could be regarded as "an independent and impartial tribunal

---

[7] See *Beaumartin v. France,* No. 19951/92, 24 October 1994, at [38] and *Sramek v. Austria*, No. 8790/79, 22 October 1984 at [42].

[8] *Micallef v. Malta* [GC], No. 17056/06, 15 October 2009 at [93-101].

established by law". Secondly, whether a machine could ever comply with the duty to give reasons required by article 6.[9]

31. If there were to be a consideration of how article 6 might be amended to protect the rights of citizens to have their cases, or some cases, decided by human judges, two further things are, in my view clear.

32. First, the technique adopted by article 14 of the EU's AI Act will not ultimately be sufficient in itself. Article 14, as I have said, requires AI judicial systems to have "appropriate human-machine interface tools" so that "they can be effectively overseen by natural persons". The problem there is that machines are quicker and more efficient. They are, and certainly will be, capable of "advising" the human judge on the outcomes of tens or hundreds of cases at a time. The human judge would never have time to check, or even review, the outcomes that the machine was proposing, making it inevitable that, over time, the human oversight required by article 14 would risk becoming a tick-box exercise.

33. Secondly, it is not, in my view, a solution to identify parts of the judicial process that might properly be undertaken by machines and part that must not. As I have explained in relation to the Shenzhen system, it assists the judge by distilling the important points to be derived from the evidence, listing the issues that need to be decided, and providing draft reasons. If one accepted that these were reasonable "assistance" functions, the question would resolve itself, once again, into asking how much time should be devoted by the human judge to checking the recommendations of the machine. Moreover, even the seemingly unobjectionable process of summarising the relevant facts has a critical judicial element. Judges are trained to recognise facts that may crucially influence outcomes. It may be hard to train machines reliably to recognise such factual minutiae.

---

9 *H v. Belgium*, No. 127-B (Series A), 30 November 1987 at [53].

34. It may be, therefore, that attempts to lay down human rights as to what citizens are entitled to demand by way of human judicial decision-making will be better focused on *types* of case rather than parts of the decision-making process.

35. To conclude, therefore, on article 6, I do think that a further protection will need to be considered. There should be an attempt to reach international consensus as to the types of case that can never, on human rights grounds, be decided by machines, even where there are extensive backlogs. This debate is being undertaken already in the commercial field through the work of the Standing International Forum of Commercial Courts (SIFoCC) in developing Model Rules for the use of AI in commercial court and arbitral decision-making. The debate will be equally important in other areas of law including criminal cases, family cases and administrative challenges to decisions of the state.

Does article 10 require amendment or addition?

36. Freedom of speech under the First Amendment to the US constitution[10] has been invoked by US AI companies to protect them against liability for harms caused by their Large Language Models (LLMs). The claims by US Attorneys General against Meta and other companies in California and other states resulted in a $17 billion settlement of claims alleging that children had been harmed by deliberately addictive features on social media platforms.

37. The argument about how far the First Amendment protects material created by AI has not, however, been settled. A decision in Florida has rejected the First Amendment defence, whilst other US decisions take a different view.

---

[10] The First Amendment to the US Constitution provides that: "Congress shall make no law … abridging the freedom of speech, or of the press; or the right of the people peaceably to assemble, and to petition the Government for a redress of grievances".

38. Article 10 also protects freedom of speech, but it is qualified by additional wording in article 10(1) and by article 10(2) which provides exceptions. Article 10(1) says that the right: includes: "freedom to hold opinions and to receive and impart information and ideas without interference by public authority". Article 10(2) provides that the right "carries with it duties and responsibilities" and "may be subject to such formalities, conditions, restrictions or penalties as are prescribed by law and are necessary in a democratic society", "for the prevention of disorder or crime, for the protection of health or morals, [or] for the protection of the reputation or rights of others".

39. These caveats make article 10 a very different beast from the First Amendment. Indeed, it is more often used in the UK and the EU to protect citizens from risks to their freedom of expression such as social media's content moderation, rather than being used to protect big tech companies from claims against them.

40. It may be too early to say whether articles 10 and 11 of the ECHR require any amendment in the light of AI. What is already absolutely clear, however, is that private law liabilities for harms caused by AI will differ quite significantly across jurisdictions. Searching for universal international agreement as to the appropriate guardrails for AI within social media and the internet will take some time – if it can ever be achieved.

Does article 8 require amendment or addition?

41. Article 8 probably presents an even greater predicament in the context of cyber fakes and social media.

42. I asked ChatGPT how article 8 should be amended to protect citizens against cyber-fakes and was presented with an extraordinary essay, including the detailed drafting for three additional sub-clauses to protect against digital impersonation and fabrication.

43. It suggested first that everyone should have "the right to protection against the unauthorised creation, manipulation or dissemination, by digital or artificial means, of representations of their identity, likeness, voice, conduct or private life which falsely purport to be authentic and which cause, or create a serious risk of causing, substantial harm to their dignity, reputation, personal autonomy, security or private life". The remaining draft clauses required states to implement measures proactively to protect the right, and to balance it against article 10 rights. ChatGPT thought there were constitutional advantages in amending article 8, not least of which were: (a) the clarity that digital identity was an independent aspect of private life, (b) common international minimum standards, and (c) the ability to create effective remedies.

44. It may be that ChatGPT protests too much. Article 8 may be one area where "the right to respect for his private and family life" is sufficient to protect against cyber-fakes and even improper use of police facial recognition technology. Article 8(2), of course, allows public authorities to act in accordance with law where it is "necessary in a democratic society in the interests of national security, public safety or the economic well-being of the country, for the prevention of disorder or crime, for the protection of health or morals, or for the protection of the rights and freedoms of others".

45. So, article 8 may be the classic example of where it is appropriate to allow the courts to work out how the right to privacy, private life and an inviolable digital identity can be applied in the age of AI. This is a problem that may resolve itself.

Does article 14 of the ECHR require amendment or addition?

46. Article 14 provides, as I have said, that: "[t]he enjoyment of the rights and freedoms set forth in this Convention shall be secured without discrimination on any ground".

47. Whatever people say about the long history of AI, I believe that we are still at the very beginning of our journey into the world of highly capable AI. But we are already seeing how it is being used very differently in different parts of the world. The highly capable English language and Chinese LLMs originating in the United States and China are more effective tools in those countries than in the less developed world and the Global South. Countries with multiple languages suffer in the race for reliability. A lack of data centres, power and internet coverage in many countries in Africa and elsewhere put them at a disadvantage.

48. In these circumstances, there is a real question as to whether AI will create a new form of discrimination, which may not be covered by article 14, or perhaps a new form of digital disadvantage which may lead to a widening opportunity gap even within Europe.

49. There are also risks that unequal access to AI will create inequalities in terms of privacy and freedom of expression.

50. Once again, we are probably at too early a stage to make any concrete suggestions for Human Rights reform. But we should not assume that new protections will not be needed for the sections of our societies less able to take advantage of the opportunities that AI creates in some places and for only some people.

Conclusions

51. So, let me draw some threads together.

52. Lawyers and judges are undoubtedly worried about how AI will affect legal practice and dispute resolution. They are right to be thinking about those questions, but wrong, in my view, to think that they are likely to be replaced by machines. Humans value human decision-making and are unlikely to accept the most sensitive decisions about their lives being taken, at least ultimately, by machines.

53. It is obvious that any democratic society would want to protect its citizens from online dishonesty and cyber-fakes as much as it would want to protect them from offline dishonesty and fakery. The principles applicable to the detection of cyber-dishonesty are, though, much the same as those historically in place. Cyber-fakes are a new feature of an age-old problem. Judges are good at identifying those who lie in court, and will likely become adept also at distinguishing between reality and cyber-fakery.

54. The problem of judicial decision-making in the age of AI is the one that I find most difficult and troubling. It raises serious issues for citizens and businesses alike. It will, I think, become rapidly apparent in many jurisdictions that the only way to clear case backlogs will be by using AI to assist the judicial decision-making process. Younger people, brought up on Claude and Gemini, may accept that approach. Others will not. Either way, our national society and international societies will be faced with a real modern-day human rights problem that will require consideration, at least, of an amendment to the ECHR and other similar provisions.

55. Access to AI is not spread evenly across Europe and the world. There is a risk that this inequality will reduce privacy and freedom of expression for those with less privileged access to LLMs. Again, our societies will need to consider whether existing human rights protections are sufficient.

56. There is, I think, no cause for alarm, but a serious international debate about the issues I have mentioned, cannot, I think, come soon enough.

57. Many thanks for listening.