# AI and Legal Reasoning

James Wood Lecture

Glasgow University 20 November 2026

*A version of this speech was given by Sir Colin Birss, Master of the Rolls, 20 November 2026.*

1. Good evening. It is a pleasure for me to be here in Glasgow University to give the James Wood public lecture. May I start by saying thank you to Martin Kretschmer of the CREATe centre here in Glasgow for inviting me. He has been trying to get me to come here for some years and we finally found a way, nearly a year ago. And it is a particular honour for me to lecture here, since my father lectured here in the 1950s after doing his PhD in what I think was still called the Natural Philosophy Department.

2. Now in this lecture I intend to examine the capabilities of artificial intelligence and its use in the law today, and address how this relates to legal reasoning. I should make clear that the context I am describing is primarily the civil justice system in England and Wales, since that is the context of my work, but I hope it will also be of interest more broadly.

*The development of AI*

3. Let me start with how AI has developed. Chat GPT suddenly emerged on the scene in November 2022. It was immediately obvious that something new and quite extraordinary had happened. But what was it? Now when I left university in the 1980s, before working in the law, I worked in computers, and one of the relics I have from that time is a wee folder with the title "Artificial Intelligence Training". So something called AI was around even then. And Professor Richard Susskind's original work on AI and law started in that time – the 1980s. Indeed, if you know your technology, you will know that the artificial neural networks on which these modern AI systems are built, date right back to at least the 1960s.

4. Now things have clearly moved on a bit since the 1960s. I gather that one of the key technical developments in artificial neural networks, which is behind these new and highly capable large language models, is something called the transformer architecture. Despite my interest in technology, I am not qualified to explain what that is. Nevertheless, in my experience it is helpful in a legal context to have some sense of what these modern AI systems are doing, or at least appear to be doing. I am going to give you an explanation as I see it. It is a lawyer's explanation which technologists will no doubt tell you is not right, but I find it is helpful for those in the law to look at it this way.

5. The first point to make is that one of the notable things which has changed is that these machines can "*understand*" human language in a manner which was

unheard of before. By this, I am not trying to suggest that they are the equivalent of human beings. What I mean is that they *appear* to understand and certainly operate *as if* they understand human language – but it gets in the way to repeatedly refer to these machines as operating "as if" they work that way. Nevertheless, it bears emphasising that in putting it that way I am not intending to anthropomorphise the machines.

6. This capability of understanding human language is striking, it is new, and it is a fundamental aspect of the reason why AI is important to the practice of law. After all, law is language. In law it is the words which have power. That is true whether these words are in the form of legislation or in the form of court judgments. The interpretation and grappling with all this language is what the practice of law is all about.

7. Of course computers have been used in legal practice for decades. I am old enough to remember when computers were only just coming into the law. We called them "*word processors*" and for good reason. The machines were not processing language: they were handling words. And that essentially was still true – for most people involved in law – until the day before ChatGPT appeared. Let me give you a simple example to make the point. Imagine I am looking for an old email in my computer which was about a dog. In October 2023 I would type "dog" into the search bar. And it would do a word search for the string of characters D O G. It might find other words like DOGE or DOGS, but what it would not do is find an email just because that message included the word "HOUND". The reason being that the machine treated the word essentially as a meaningless string of characters. Today, if I do the same thing using the AI enabled email search function, and I ask for emails about a DOG, it will offer emails in which the only relevant word is "HOUND". That is new. And that is what I mean by the machines understanding language.

8. This takes me to the second point: why is it that these machines now do this? The reason is because they build models. You have heard the expression "Large language model", or LLM. What that is referring to is a multidimensional mathematical model in which the relationships between things are mapped in a statistical way. To take the example of DOG and HOUND, if the things in the model are words, then these two words are close together – whereas, for example, the word CAT is further away. The distance between these words allows the machine to operate using probabilities. So when I say I am looking for DOG, the closeness of HOUND to DOG means there is a good chance something with HOUND in it is relevant. Whereas the larger distance between DOG and CAT means that the probabilities used by the machine are not likely to suggest that CAT is a sensible answer.

9. Now the consequence of what I have described is that when you think about it, these models have become models of ideas. The location in the model space of DOG and HOUND represent a concept. DOG, HOUND and presumably CANINE

and CORGI and so on are all located close to one another. And it is more complicated still, because this model allows the machine to "understand" context. So in some contexts, a CANINE is a dog, but if we are at the dentist a CANINE is a tooth. Of course, these two ideas relate to each other, but this modelling approach allows these concept maps to take that context into account. So when I do a search for an email about a DOG the AI will not return a message about a dental appointment even if it refers to work to be done on my CANINE tooth.

10. The third thing to note is what happens when you give these AI machines some text, such as a judgment. The model I have just described means that the machine can read the text and in doing so it builds a model of that body of text of a similar kind to the one I have just described. In other words, it builds a model of the concepts or ideas found in the text. And with that model the machine can produce a summary of the ideas in the text.

11. Now if you think about it, what I have just described is something lawyers do every day. And until recently computers could not do it – at least not the computers available to us. Well now they can, *to some extent*.

12. The fourth thing to consider about AI is about the output. What these AI machines produce is an output based on probabilities and a degree of randomness. The machine predicts the most likely words needed to follow from the prompt. You might have heard the expression "*stochastic parrot*", which is a term used to make that point. AI is said to be a parrot because what it produces is based on what it was given in the first place, and the word *stochastic* refers to the use of randomness in the statistics used to create the answer. This randomness is an important feature because it explains why the machine will not necessarily give you the same answer when you ask it the same question twice. The randomness is built in by design.

13. The final aspect of AI I want to mention concerns "machine learning". This is important and refers to the way in which the machines are set up in the first place. They were not trained to reason. They were not trained to produce the right answer to a question. They were trained to produce an answer which looked like the right answer.

14. Let me explain the difference. In early 2023, the early days of ChatGPT, I was thinking about this, so I typed in 2 + 2 =?. And I got the answer 4. So far, so good. Then I bashed the keyboard and typed in a long random string of numbers, followed by a plus sign then another long string of numbers and "=". It produced an answer. The answer was wrong. But there was an interesting feature. It was a string of digits of about the right length. If you add 10 digit number to a 12 digit number, the answer will probably be 12 digits long, but it might be 13. It will never be less, nor will it be more.

15. It was clear what was going on. ChatGPT was not doing arithmetic. It knew that the character "4" came after "2+2=" not because it was reasoning but because

that string of characters is common enough for it to have learned what is likely to follow. It had learned what the answer looked like. But for longer numbers all it had learned in its model was what the answer looked like – which was likely to be a string of digits the same length as the longer one of the pair.

16. That is the difference I am talking about between being trained to produce the right answer to a question and being trained to produce an answer which looks like the right answer. I believe this remains fundamentally how LLM AI still works, at least from my experience as a user in the context of law. In other words, if it has seen something like the question and the answer before, it can reproduce it, but it will not give a new answer to an old question (I will come back to this) and when asked a new question, it does not do so well at all.

17. By the way, before you all turn on ChatGPT or Clause at home, this trick with numbers does not work anymore. I do not know why, but I am pretty sure the reason is because – speaking metaphorically – these LLMs now have their own little co-pilots sitting beside them, so that if the machine spots a maths question, it turns off the LLM, gets out a pocket calculator, does the sum and then gives it to the LLM. Even though the number trick in particular no longer works, I believe it reveals a fundamental aspect of how these machines work and at least partly explains why, in a legal context, the AI machines can produce what are called hallucinations – fake cases which look like real ones.[1 and 2]

18. Now at the risk of repeating myself, these points are just my way of explaining AI in a legal context, from the perspective of someone who uses it. That is all.

*Legal practice and legal reasoning*

19. Now let me turn to legal practice and legal reasoning.

20. What do lawyers and judges do? Well in my experience a lot of legal practice involves producing a summary or précis of something. I never liked doing this at school, and it was only after I had started studying law that it dawned on me that a very great deal of the work comes down to that process. As a practising lawyer, the summarising exercise applies to facts and it applies to law. When a barrister or advocate stands up at the start of a hearing, the first thing they say, usually, is what the case is about. In other words, they present a very brief summary. Then they will go into a bit more detail, still giving in effect a slightly more detailed summary of the case and of their opponent's case. If you think about it, a very great deal of legal practice involves a variety of summaries at different levels of detail in different circumstances.

21. Now that takes me to another thing I did in 2023, which attracted a fair bit of attention. I used an AI generated summary in a judgment. Let me say

[1] Incidentally these sorts of fakes are quite easy for IT systems to spot based on access to online legal databases of authentic cases.
[2] And also, incidentally, it is not a good idea to set out hallucinated cases in a real judgment, since that will be picked up by AI and risks compounding and perpetuating the hallucination.

something about that. A piece of legal reasoning is commonly structured in three parts. You set out the facts (including fact finding), then you set out the law, and then in the third section you apply the law to the facts. Conceptually this structure is always a good starting point, and it works nicely on many appeals.

22. So in 2023, when I was writing an appeal judgment, the facts were summarised and I was starting the section on the law. The appeal was on a point of law. If you look at many judgments where the issue is the law, their legal sections also often have a common structure. A topic is introduced with an uncontroversial high level – even basic – summary of the law in question, to orient yourself and the reader. Then the judgment delves into the particular legal issue which needs to be considered in detail. In a copyright case you might start by saying that "*Under UK law copyright subsists in original artistic works (s1, Copyright, Designs and Patents Act 1988) and the term artistic work is defined in such a way that a photograph is an artistic work regardless of artistic quality (s4(1), 1988 Act*". You can then go on to address the particular issue, which could be for example about the nature of originality in the context of a photograph.

23. So in the case I was dealing with, I was about to write the basic summary. I wondered what ChatGPT would produce. I forget the prompt exactly, but it was something like - please give me a 400-word summary of UK law on X. And it produced it. And it was accurate. I knew it was right because I was about to write it anyway and I done the preparation. I think I changed one word and used it. I used it and I took – and take – full personal responsibility for every word of my judgments – including that passage.

*Trust and effort*

24. A critical point is that one does not need to trust the machine to find it useful in this way. I have heard people say that they do not use AI because they do not trust the machine. However, while I agree with the premise – that you should not trust the machine – I do not agree with the conclusion. The insight, I suggest, is to see that the machines are still useful without having to trust them. The AI is like a tireless but not brilliant assistant.

25. I would explain this with two analogies. The first is an analogy between hunting and fishing. When hunting you are looking for a specific target which you will recognise when you see it. Whereas with fishing, your hook is dropped over the side of the boat, and you will bring up whatever happens to bite. The distinction is that AI is a good tool if you can check or recognise that the answer is what you want (hunting) but a bad tool if you take whatever you will get (fishing).

26. The second is an analogy with mathematics. In mathematics there are a number of problems which are easy to check but hard to solve. An example is the fact that finding the prime factors of a given large number N is much much harder than checking whether a given pair of numbers are indeed the factors of N. Let's

say A and B are prime numbers and A x B = N. The only way to work out what A and B are, is to work through the possible options, testing each one.

27. Say you are given N as 899 and say your assistant proposes the factors are 11 x 83. It is easy to check. 11 x 83 is actually 913 and not 899. So your assistant was wrong.

28. But given just 899, in order to find the factors yourself, you have to start by dividing by 3 (299.66..), then 5 (no) then 7 (128.4..) then 11, 13, 17, 19, and 23 before getting to 29 and finding the answer – which is 29x 31. You do not want to spend time working through all that, but your tireless assistant can. And you do not have to trust them when they give you the answer – because you can check and identify that the answer is right or wrong without anything like as much effort.

29. And by the way, this kind of mathematics is at the heart of internet encryption (at least today).

30. So focussing now on how we might use such a tool in legal practice, the point is that once the judge has read the case material and listened to the evidence and argument, it still requires significant effort for them to produce a suitable summary. However, the judge in that position will be able to use their judgment to recognise, and if necessary correct, a summary produced by someone or something else. It is *judgment*, which still resides with the human judge, that enables us to check and recognise a good answer as opposed to a bad one.

*The Judicial Guidance*

31. I want to use today as an opportunity to set out how AI is being used by the judiciary in England and Wales. I am the judicial lead on AI in England and Wales, and in December 2023 we published Judicial Guidance which has been updated on a rolling basis ever since. There are three aspects.

32. ***First***, data security is vital. You must not use AI with private data such as case data, unless you know the AI system you are using is secure.

33. In late 2023, that meant there was no AI you could use with case data. But in the courts of England and Wales we now have AI systems which are secure, in the sense both that we are satisfied with the data governance and that the model is not learning from us. And such systems are much more common now than they were then.

34. ***Second***, judges take full personal responsibility for what goes out in their name. That may be obvious, but it is important. Some people call this a "human in the loop" – which is true as far as it goes, but it is often stated as if that is a problem, which I do not accept.

35. ***Third***, judges can use AI in their work provided it is secure and they remember they are taking full personal responsibility.

*Practical examples*

*Use of AI as a judicial assistant*

36. So that is the guidance. How is it applied in practice? I will give some examples in the context of different aspects of legal reasoning.

37. Reasoning in a legal context, which is what I mean by legal reasoning, involves a number of different activities. No doubt there are academic studies on this, but I will adopt a practical approach to classifying the activities based on my own experience in the law.

38. I mentioned fact finding a moment ago. It bears emphasis because in the civil justice system this is the task which the majority of judges do, most of the time. The reason why is because most judges sit at "first instance" rather than on appeal, and most of their work involves facts much more than law.

39. The exercise requires the civil judge to grapple with the oral and written evidence, listen to cross-examination and make findings about the reliability of witnesses. Overall, the exercise involves two tasks: (i) fact management and (ii) resolving disputed issues of primary fact applying the civil standard of proof. In effect asking: what happened?

40. Now taking these aspects of fact finding – fact management and issue resolution – in my experience, AI is most capable in the first one. It can produce useful summaries including *dramatis personae* and chronologies. These are good examples of work which takes more effort to perform than it does to check. And judges are already entirely familiar with being given the same sort of material by the parties and deciding if they accept it as it is or amend to their satisfaction.

41. It has been suggested that this kind of use of AI risks injecting bias into judicial reasoning. But this objection is based on an unduly narrow conception of the judicial role in our common law system. The judge's role is to receive material produced by the parties to a case, to decide the case and give reasons for that decision. The parties are entitled to present the material to the court in a manner which suits their case. They will summarise their case in a manner favourable to them and they are entitled to do so. Judges are used to dealing with that and asking themselves whether that way of presenting the argument or the facts is one that they will use in their reasoned judgment or not. In a sense the judicial task is to operate and make a decision in a noisy environment, with assertions coming at the court from different perspectives.

42. In any event this concern also underplays the vital importance of the judge's stated reasons. As any judge will tell you, the judicial role is not simply to make decisions. Hard though that is, it is not a complete statement of the task. The role is to make reasoned decisions, and the reasons for the decision are every bit as important as the decision itself. Those reasons are contestable. They can be

scrutinised, criticised and taken to an appeal court. What matters is that these are the reasons the judge is prepared to put their name to and stand by.

43. On the task of resolving contested issues of fact, in my view that is a judge's job. It is not a job for the AI. I will come back below to the question of AI making decisions. At this stage I am talking about AI as a tool.

44. Having done the facts, I now turn to the law. The first step is for the law to be identified. In many cases this is – at least conceptually – a relatively straightforward exercise. But there is one aspect I would highlight here, given that the context of this discussion is AI. A judge seeking to identify the law normally does so by identifying the primary source material: legislation or the decided legal authorities. Commentaries, such as textbooks, might be used to orient yourself in an unfamiliar area, but mostly you would use a commentary to find the right source material and then read that. I have not found AI particularly useful as a tool in this context, and I know other judges have had similar experiences. This is because while AI is quite good at providing a tailored exposition of the law, that is not what we need. We need to work on the source material.

45. Our legal training and experience already equip us to seek out and identify the relevant source material. Therefore I do not find AI of great assistance as a tool in that context.

46. Finally, we apply the law to the facts. Once the facts have been found and the law has been identified, this is a relatively straightforward task. However, a common situation – sometimes described as a mixed question of fact and law – is one in which the judge is required to reach an evaluative conclusion, perhaps by applying a legal standard of some kind. In negligence, this could be whether something was reasonably foreseeable. In copyright, this could be whether an alleged copy of an original work reproduces a substantial part of the original. (I will note that this is similar to, but not same as, an exercise of the court's discretion – but in this lecture I will not dwell on the differences.)

47. Like the task of resolving contested issues of fact, in my view this application of an evaluative standard (or exercise of a discretion) is a judge's job and as I said before, at this stage I am talking about AI as a tool.

*Further use of AI as a judicial assistant*

48. An example of what can be done is this. Once a judgment has been written, the judge can ask the AI (which in our case is Copilot) to identify internal inconsistences. When I do it, I do not always agree with the machine, but it is helpful see what it thinks. In this respect the AI is playing a similar role to that of our existing human judicial assistants.

49. The point about inconsistencies being internal one is that the AI does not have to go out into external websites (e.g. to look for caselaw) to see if it agrees with the

judge's reasoning. That would not be data secure, and we prevent our system from doing that.

*Use by legal professionals*

50. On 29 October 2026, the Civil Justice Council of England and Wales published its report on the use of AI in preparing court documents.

51. For those who do not know, the CJC is an Advisory Public Body. The Council's statutory duty is to review the civil justice system and make recommendations to the Lord Chancellor, the judiciary and the Civil Procedure Rule Committee on the development of the civil justice system to make it more accessible, fair and efficient. The chair is the Master of the Rolls.

52. The CJC's final report on AI came out a month or so ago. I recommend reading it. Let me give you a brief summary.

53. The report is about "court documents", which means documents created for the purpose of being provided to the court. In terms of legal professionals, the report proposes that starting point for the regulation of AI use in the preparation of court documents should remain with legal representatives and their professional regulators. For example, in England that would include the Solicitors Regulation Authority, the Bar Standards Board and the oversight regulator the Legal Standards Board.

54. The point is that when a document like a skeleton argument is prepared by a lawyer and presented to the court, the lawyer is taking professional responsibility for it. That is why it ought mainly to be a matter for the regulator.

55. Nevertheless, as the report recognises, these regulators have issued guidance. In doing this they have drawn attention to: "*the potential risks of using AI including (i) harm to a lawyer's clients through, for instance, the use of inaccurate information or potential breaches of client confidentiality and privilege; (ii) harm to the proper administration of justice through the submission of hallucinated authorities; (iii) potential breaches of data protection law and intellectual property rights; and, consequently (iv) harm to the public interest through trust and confidence in the provision of legal services being undermined.*"[3]

56. The report also notes that the court can and will in some cases issue sanctions of its own arising from the improper use of AI by legal professionals, the well-known example being the case of *(R) Ayinde v London Borough of Haringey* [2025] EWHC 1383 (Admin).[4]

57. A distinction is drawn in the report between the use of AI to produce documents setting out arguments, such as statements of case and skeleton arguments, and

[3] Civil Justice Council, *Use of AI for Preparing Court Documents: Final Report* (29 October 2026), https://www.judiciary.uk/wp-content/uploads/2026/10/Use-of-AI-for-the-Preparation-of-Court-Document.pdf, paragraph 4.1.
[4] See more cases at ibid, p. 12 footnote 8.

those setting out evidence, such as witness statements. What I have referred to so far relates to documents setting out arguments.

58. For witness statements, the CJC recommends a differential approach. Where a witness statement is that of a witness of fact, the witness should make a declaration that the statement has not been prepared through the use of an AI tool that has generated or influenced its substantive content. The purpose of this is to preserve the integrity of the witness's evidence, in their own words.

59. The point is that this approach does not prohibit the use of any specific AI tool, or type of tool. It is intended to ensure that AI use does not affect the substantive content of such statements. It is a declaration that focuses on the result of such use, not on the use itself or the tool that produces the result.

*Use by Expert witnesses*

60. The CJC's report[5] also referred briefly to the question of the use of AI by expert witnesses. The recommendation is that when AI has been used to generate or influence the substantive content of the expert's report that ought to be declared. This declaration may include, where necessary, identification of the AI tool used, the methodology used, and the appropriateness of AI use for that purpose. Not surprisingly, this is similar to the position relating to fact witnesses.

*Use by litigants in person*

61. That takes me to the most notable trend we are seeing with AI in our courts and tribunals. This is its use by unrepresented litigants. This has a variety of dimensions.

62. First, it is striking that the public now has easy access to a source of free legal advice. Whether their providers admit it or not, that is what these AI models are providing. What is more, while that advice is nothing like as good as the legal advice a trained professional lawyer would be able to give the individual, by and large it is pretty good basic advice, and it can handle specific areas like intellectual property law.

63. The knock-on effect of this is that litigants are filing much longer briefs than they used to. But the difficult aspect in policy terms is that while a lot of this material is wrong and not helpful, some of it certainly is helpful. We are seeing individuals present their case in a significant better way than before – more clearly and more appropriately directed. Furthermore, we are seeing increases in cases (and appeals) filed by unrepresented parties. I am sure this is driven by advice from public AI systems like ChatGPT. I will come back to this briefly at the end.

[5] *Ibid,* paragraphs 8.1-8.17.

64. The CJC's report identified the absence of any existing guidance for litigants-in-person and stressed that this was an issue that needed urgent consideration[6]. It highlighted the fact that while access to justice is plainly facilitated by the use of AI by litigants-in-person, on the other hand there are problems caused by it.

65. There are a number of problems arising here: issues of confidentiality, possibly privilege, a serious problem caused by AI being a people pleaser and giving far too optimistic advice. Nevertheless, I find it hard to believe that we will be preventing people from using free tools available to them to help them. What would be worthwhile would be to find a way to make better tools available at a similar scale, and to find ways to direct people to those.

*Using AI as the decision maker*

66. Now let me turn away from the question of AI as an assistant and consider using AI as the decision maker. The first thing to observe is that to do that would, so it seems to me, require a proper legislative basis. In our current legal framework, a judge has no authority to use AI to decide a case.

67. Now let us assume for this purpose that you might provide a legislative basis for doing it. What are the factors to consider? There are two questions: could it be done, and should it be done. I will start with could.

68. We know, I suggest, that AI can perform pretty well on the fact management side of things. This is important because I think it may go a bit further than you might think. AI is quite capable of grappling with very large amounts of information.

69. Examining this material is a burdensome exercise. We may not be there yet, but it is foreseeable that an AI system could examine, classify and evaluate large bodies of evidence in a fraction of the time and cost of a court hearing. If we get that far, there may well be pressure to use it, subject to a form of human review or appeal.

70. We also know that it can identify the law, at least in many common areas. We know that because we see the effect of its use by unrepresented litigants.

71. In relation to other aspects of reasoning in legal practice, as far as I can tell AI is capable of identifying the factual issues to be resolved. It is not clear to me how it could resolve some primary factual disputes, but given a set of facts I have seen that an AI may well be able to give what looks like a reasoned decision on the application of a legal standard. There are also many cases in which the primary facts are not really in dispute.

72. So it does appear that AI is – or at least soon could be – capable of being given a body of evidence and submissions and providing a conclusion, with reasons.

---

[6] *Ibid,* paragraph 4.5.

73. However, I wonder if it is that simple. Let me consider an aspect of legal reasoning which we have not covered yet.

74. Once the right legal source material has been identified, the next task for a decision maker is construe it, and that applies whether the source material is legislation or decided cases, or both. "Construction" is the expression used in the law, rather than simple interpretation, because although it involves identifying the meaning of the words in the context in which they are used, the exercise also involves taking that meaning and determining its legal effect, particularly in the factual circumstances of a given case. When dealing with case law in our common law system of precedent, this also involves working out which principles in a given case are binding (so called *ratio decidendi*) and which are more or less persuasive (so called "*obiter dicta*"). This is the kind of reasoning which higher courts perform much more than the trial courts, and a feature of it which I would like to highlight is the element of creativity which arises when cases are distinguished from one another. A previous judgment is identified as a binding authority for a principle which appears to cover the case before the court, however the judge or judges may decide that there is a distinction between the present case and the previous binding case which was not identified in the previous case. This is how the common law develops and it a crucial aspect of our legal reasoning. Another way of saying this is that the common law is essentially a defeasible system: there is no such thing as a legal rule which permits of no exceptions.

75. A hitherto well-established authority, which has stood perhaps for many years without qualification, is then qualified by a new distinction. This is not all that unusual, even though when it happens it can feel quite revolutionary in character. Even more strongly, although of course comparatively rare, the Supreme Court can go further still because it has the authority to depart from its own previous decisions.

76. These revolutions provide, I suggest, some similarity to the process of reasoning which involves different kinds of laws altogether. I am referring to scientific laws and the process identified by Thomas Kuhn in his work *The Structure of Scientific Revolutions* (1962). It is an analogy that I have been struck by for some years, and I thought I would take this opportunity to explain it. Thomas Kuhn's work challenged the traditional view of the development of scientific laws as a steady run of progress moving forwards to a more and more accurate set of laws which apply to explain the world we live in. He suggested that what really happens is that the development of scientific laws takes place in a context of scientific paradigms and happens in phases. A paradigm is a shared framework of laws, explanations and experiments. The first phase is called "normal science". In this phase the existing paradigm functions to frame the work being done. That work essentially involves puzzle solving aimed at fleshing out the existing paradigm. Sometimes anomalies are found, which the existing paradigm cannot convincingly explain, but they are often explained away and not regarded as

significant. Eventually the anomalies begin to pile up and at some point a second phase takes place, in which there is a scientific revolution and a so-called paradigm shift. This is an essentially creative process. As a result, a new paradigm emerges which fundamentally alters how scientists view the world, and at the same time explains the anomalies.

77. A famous example is the shift from a geocentric model of the universe with the Earth at the centre and the Sun, planets and all the stars in orbit in circles around the Earth, which had prevailed since Ptolemy's Almagest in the 2nd century CE; to the heliocentric model in which the Earth and the planets were all in elliptical orbits around the Sun, and the stars were not moving at all, which came from the work in the 16th-17th centuries of Copernicus, Kepler, Galileo and then Newton.

78. The point was that for many centuries those working in the old geocentric based paradigm were well aware of what, looking back, could be seen as anomalies in a geocentric model. For example, if you watch Mars from the Earth, it sometimes moves backwards in the sky. This is referred to as retrograde motion of Mars. It happens because both Mars and the Earth are actually in orbit around the Sun, moving at different speeds. However, Ptolemy fitted it into the geocentric paradigm by adding extra circles called epicycles to the orbit to cater for it. Once you have realised that the problem was with the paradigm itself, then the revolution occurs and a new paradigm is settled.

79. As it happens, when I did some research for this lecture, I found that I was not the first person to draw this parallel. I refer to a paper by Trueblood and Hatfield (2023) entitled *Precedent and Paradigm: Thomas Kuhn on Science and the Common Law.*

80. If you had trained your AI based on all the available material before Copernicus, the model would have been based on a resolutely geocentric system, with epicycles and all, with nothing to shift it. It would answer the questions posed from inside that framework because it had never seen a question answered in a different way. There would be no new answers.

81. In other words, even if AI could decide a case on the basis of "normal law" (like normal science), I have real doubt that it could generate the equivalent of the paradigm shifts necessary to allow the law to develop. The real practice of law is not just about *exploiting* (applying) existing rules. It is about *exploring* (discovering) exceptions to those rules – or putting it another way, creating new rules. Now, it is fair to note that these machines seem to be becoming more capable almost every day, so I cannot predict where they will end up, but I have not seen anything today which suggests that they will be able to do this.

*Should it?*

82. So even assuming AI could perform the task of deciding cases, that still leaves the question whether it should. The core issue is a societal one. In a democratic

society, the question is by who and how we want decisions to be made. Speaking for myself, I cannot imagine we would like a society in which the sorts of decisions our courts are called upon to make every day would be better decided by a computer, even if that decision was subject to review by a human judge or judges on appeal. These are some of the most consequential decisions in the lives of the parties to court proceedings.

83. That is not to say that AI has no utility in other ways – as an assistant, but also perhaps as a mediator. There is also evidence that commercial arbitration organisations are exploring the potential use of AI as an arbitrator, necessarily with consent because that is the nature of the arbitral process.

84. However, as ever, things are not so simple. There are countries today seriously contemplating using AI to make decisions in court. An important question, which could arise quite soon, would be our view on the enforcement of a foreign judgment made by AI. The grounds for challenge are quite narrow, but one of them is a public policy test (e.g. Dicey Rule 5). So it may not be all that long before a court here is asked whether an AI decision from somewhere else passes that test.

Conclusions

85. I would like to come back to the use of AI by unrepresented litigants. As the CJC's report highlights, this is a major current issue. The questions we need to ask are about how to respond to that.

86. With readily accessible AI chatbots today, we are now seeing an outline of how to address the problem of providing legal advice and assistance at scale to people in our society for whom professional help from a lawyer is unreachable. In my view, there is no substitute for the advice and assistance of a legal professional. But we have to recognise that with limited resources, the state has to prioritise and there will always be people who need that advice and assistance, but for whom a legal professional is not an option.

87. Another way of characterising the significant increase in cases being filed today by unrepresented litigants assisted by AI, could be to see that more people have access to justice than before. The assistance of AI is helping unrepresented litigants to present their case more clearly than would have happened without any assistance. However, a serious drawback is that AI usage often produces unnecessarily lengthy and unfocused documents, leading to unnecessary delay and time spent by the court. Another drawback is that AI will often give over-optimistic advice, and a further problem is that litigants may be putting confidential information into these machines.

88. Now as I have said in other contexts, I believe one simple response is that there may be a role for guidance and rules relating to word or page limits. We have them in our higher courts and the time may have come to use them much more widely.

89. However, there is a bigger picture here, in which three issues converge and present us with which I believe are real opportunities. The issues are AI sovereignty, legal regulation, and embedding the common law into the operational use and regulation of AI.

90. Imagine an AI system which was UK based and regulated, so that people in our society could use it without being concerned about data security or confidentiality, and in which the way we wanted it to work was deeply embedded in the architecture – including the large language model, perhaps some deterministic AI systems, the guardrails and no doubt other things too. Doing this would mean that we would be able to embed our own approach, based on the common law and our own views on professional regulation, into its ways of working. We would need to bring together AI technologists, law people and regulators. It would involve the public sector, the private sector and academics. I do not believe it is too late to work on AI sovereignty. Nor do I believe that we lack the technological expertise in this country to do it. We certainly do not lack the legal or regulatory expertise for this. There is huge potential.

91. I believe the upsides would include a significant improvement in access to justice without imposing such an extra burden on our courts and tribunals. But it would be much wider than that, because with a proper common law and UK regulatory framework embedded in the AI, there are so many other things we could do to allow everyone to benefit from this technology, but in a manner our society retains the ability to regulate.

92. I believe the legal sector, including the judiciary, has an important role to play in being open to the benefits AI can bring whilst being realistic about the risks, and upholding the integrity of our justice system. As an illustration, on 9 October 2026 the judiciary of England and Wales issued a position statement. If I may, I would like to close by giving you the core principles it identified. They are these:

> 1. The use of AI must remain consistent with the judiciary's overarching obligation to protect the integrity of the administration of justice and to uphold the rule of law.
>
> 2. Judicial decisions will be taken by a judicial office holder. AI may assist, but it must not displace reasoning, evaluation of evidence or the exercise of judicial discretion.
>
> 3. Judicial independence, impartiality, fairness, accountability, confidentiality and security must be preserved at all times.
>
> 4. Human oversight is essential. Every AI output used in judicial work must be reviewed, verified and approved by the responsible judicial office holder. The judicial office holder remains fully accountable.

5. The judiciary is taking a cautious but positive approach to AI adoption, supported by pilots, evaluation, training, governance and transparency, to maintain public confidence and uphold open justice.

93. That cautious but positive approach is where I am, and I suggest it is the place we all can be.

94. Thank you.

Sir Colin Birss

Master of the Rolls