ARTIFICIAL INTELLIGENCE IN DISPUTE RESOLUTION:

POSSIBLE ROLES AND LEGAL LIMITATIONS

1. My subject is artificial intelligence in dispute resolution. It's a huge subject, which is generating an enormous amount of academic literature across a wide range of disciplines. Deservedly so, as over time it will radically transform the practice of law. In an attempt to narrow the scope of the topic, I am going to focus on how our existing legal process may be transformed over the coming years.

A little history

2. Our practice of law has been radically transformed before, but one needs to go back to the middle ages to see the change. In those days an accused was exposed to a trial by ordeal. For example, by fire or by water. The easiest to pass was probably trial by cake. Success was measured by avoiding choking on a dry lump of consecrated bread.

3. Another example of trial by ordeal, which I owe to Mr Justice Zacaroli, was in relation to land disputes. If you claimed title to a piece of land in the possession of another you would first have to provide some prima facie evidence of the claim. But, assuming that hurdle was overcome, each party would then present a "champion", someone willing to "fight it out" on their behalf. The court would direct the battle to commence at noon on the appointed date at the "lists" which was an area of open ground. The Judges would be sitting on a raised bench. The court would declare that the land belonged to the party whose champion killed his opponent or forced him into submission. If both champions remained standing by nightfall then a stalemate was declared and the claimant was held to have failed to prove his case. An early example of what one might call the burden of proof.

4. Trial by battle obtained its public legitimacy because it was believed that God, by protecting the champion of the righteous party, determined where truth lay.

5. An alternative justification was provided by Peter Leeson of George Mason University in an article called "Trial by Battle" in 2011. In his view, trial by battle produced an economically efficient allocation of disputed property rights, by handing the spoils to the highest bidder in an "all-pays" auction. In his view the person who wanted the land most, would be willing to pay more to hire the best champion and achieve victory.

The legal process today and its flaws

6. Trial by battle is long gone. But it was replaced by a process, the essential nature of which has remained almost entirely unchanged for hundreds of years. We are all familiar with it. Witnesses give evidence of fact. Advocates argue on their behalf. The investigation of the facts is forensic, the argument is largely rational in nature and the proceedings are usually in public. At its centre is the judge, who weighs the evidence, determines the law and announces the result, dispassionately and without fear or favour.

7. This process, although obviously a considerable improvement on trial by ordeal, is also not without its flaws.

8. Some of the problems are mundane. Until about 50 years ago trials would be conducted with little or no documentary evidence. Evidence would be given orally by witnesses. In 1977 Megarry VC spent a year hearing a dispute about the effect of phosphate mining on an island in the Pacific. It was reputedly the longest-ever trial. Despite the length of the trial, the documentary evidence barely filled one file. In contrast, last year Mr Justice Hildyard heard a trial concerning Hewlett Packard's acquisition of a company called Autonomy, in which the parties' written closing submissions alone ran to thousands of pages.

9. Other problems are more fundamental. It turns out that Judges, like the rest of us, are fallible.

10. Recent work on the theory of the mind by people like Daniel Kahneman and Amos Tversky has illustrated this. They discuss the now familiar dual process model of the brain. The idea is that we employ two fundamentally different modes of thought, System 1 and System 2. System 1 is fast, intuitive, impressionistic and cannot be turned off. System 2 is slow, deliberate and effortful. It also tires easily, a process called ego depletion. The result is that System 2 usually accepts what System 1 tells it.

11. Whilst System 1 is mostly pretty good at what it does, it is hopelessly bad at statistical thinking, jumps wildly to conclusions and is subject to a fantastic suite of irrational biases and interference effects. Kahneman and Tversky discuss these biases or mental short cuts to solving problems. The labels are now well known: such things as confirmation bias, anchoring bias, availability bias.

12. Judges are just as susceptible to such biases as anyone else. In an experiment to test the "anchoring effect", highly experienced judges were given a description of a shoplifting offence. They were then "anchored" to different numbers by being asked to roll a pair of dice that had been secretly loaded to produce only two totals, three or nine. Finally, they were asked whether the prison sentence for the shoplifting offence should be greater or fewer, in months, than the total showing on the dice. Normally the judges would have

given sentences of very similar length. However, those who rolled nine proposed an average of eight months, while those who rolled three proposed an average of only five months. All were unaware of the anchoring effect.

13. Another example concerns the way in which sentencing length differs depending on when, during the court day, the judge rules. Sentences handed down at the start of the day tend to be lower and increase towards lunchtime. They then go down slightly after lunch, only to increase again towards the end of the day. What is the reason for this? Apparently, it is all down to the judges' blood sugar levels.

14. One important aspect of Kahneman's and Tversky's work was the persistence of such problems. System 2, even if engaged, is often unable to overcome System 1. Kahneman gives the example in his book "Thinking, Fast and Slow", of someone who is convinced that they are a good judge of people and able to tell whether a candidate for a job will be a success. You can provide them with evidence that indicates that they would make a better decision if they relied on data. But they will still trust their instinct.

15. One should not overstate these problems. My point is simply that it would be extraordinary if the legal process we have been using for hundreds of years, with a human being as its central decision maker, is not without its flaws or is incapable of improvement.

16. It is inevitable that there are some things done by a Judge that could, in one way or another, be better done with the assistance of AI, or indeed by AI on its own. After all, the legal process can in some sense be described as an algorithmic discipline. In other words, if X and Y exist, then the result should be Z. Given this, AI must be capable of being applied to one or more aspects of the process, as it is in other areas of life.

17. It is therefore inevitable that AI will change the legal process to a lesser or greater extent. The potential problem is one of ensuring that it is applied only where it assists and that it does not take over by default. Marshall McLuhan said that "once a new technology comes into a social milieu it cannot cease to permeate that milieu until every institution is saturated". And this applies to AI as much, if not more, than other technology.

18. But this then requires one to identify where AI will assist and where it will not. Which, in turn, requires one to have a clear-eyed view of what parts of the existing legal process need to be improved and where AI may assist, and what parts are valuable and should not be lost.

19. The main advantages and disadvantages of AI are obvious. The process is quick, cheap, certain and consistent. Its most obvious disadvantages include a lack of transparency and the danger of bias if the data set which was used to train the program was itself biased. More subtle issues arise from the fact that AI reduces or excludes human agency from the process.

20. When one discusses how and to what extent AI is going to affect the legal process, two different aspects can be identified. The first is that AI will increasingly become the subject matter of legal disputes. In other words, we will be arguing *about* AI. The second is that AI will be used to assist in *determining* legal disputes. I want to say something about each of these.

AI as the subject matter of disputes

21. It is inevitable that courts will increasingly be concerned with the legal acceptability of a decision made by AI. Indeed this is already happening, both in the context of administrative law and in the commercial sphere.

22. Algorithmic government is already here, and will be used increasingly in determining policy and taking decisions large and small. I could give thousands of examples. But one example at the more humdrum end of the scale concerns an algorithm used by Durham police to decide whether someone under arrest should be referred to a rehabilitation programme as an alternative to prosecution. The algorithm assesses the risk that the person will re-offend in future based on 34 variables, one of which is the person's postcode.

23. One question is whether it is permissible to include a person's postcode at all. There may be a strong correlation at the *population* level between certain postcodes and offending, but it does not necessarily follow that a public official should be entitled to use such a correlation to decide whether a particular *individual* is likely to reoffend.

24. A second question, perhaps a more fundamental question, concerns what part such a variable actually plays in the ultimate decision. Because of the algorithm's complexity, a person's postcode has no direct impact on the forecasted result. Instead it is combined with all the other predictors in thousands of different ways before a final forecasted conclusion is reached. Disentangling the role it played may be far from straightforward.

25. At the moment, our legal framework for reviewing decisions of public officials is not designed to address such matters. Until now, any decision that was taken, was taken by a human being. Reviewing such a decision is straightforward. Case law requires you to investigate what matters he took into account, and check whether he took into account all relevant matters and omitted all irrelevant matters. You then analyse his reasons for reaching the conclusion and decide whether he acted rationally and in good faith. These tools, although capable of being applied to AI at a high level of abstraction, don't work in practice when a judge is confronted by a decision made by a black box. The operation of an algorithm will often be opaque. The data which it used may be sufficiently large that it is incapable of assimilation by a court. It may not be clear what factors the algorithm took into account in reaching its decision, or what role each factor played. AI does not give reasons for reaching the conclusion it did.

26. The courts are already having to grapple with such issues. There is an emerging concept of structural procedural review, under which policies and systems may be unlawful if they give rise to an unacceptable risk of unlawful or unfair decision making. But there is an issue, which I will come back to shortly, as to how a court should go about such a review.

27. Algorithmic determination is also present in the commercial sphere. Again, our existing legal toolbox is not really designed for such matters.

28. A recent decision in Singapore in B2C2 Ltd v Quoine Pte Ltd provides an example. A currency trader was using an algorithmic trading platform. A glitch between this and the program run by a currency trading platform resulted in the trader being able to purchase trades at 1/250th of their true value. The trading platform was not allowed to unravel the trades. Defences based on implied terms, mistake and unjust enrichment all failed. The judge had to make sense of the concept of mistake in contract when two computer platforms traded with each other. He did so by looking at the minds and expectations of the programmers, although they were not involved in the trades. But in future programs may operate sufficiently independently that looking back to the minds of those who created them will seem increasingly irrelevant.

29. There are other questions we will need to address. Is an AI system to be treated as a product or should it be given personhood like a company? Assuming that the creators of AI should not be able to shield themselves behind it, should we be applying contract or tort concepts of liability, or developing doctrines of vicarious liability or agency?

30. The courts are going to be increasingly concerned with algorithmic decision making, and will need to develop a new set of tools to deal with such cases. This may not be straightforward.

31. The first problem is the opacity or lack of transparency of such systems. There are a number of reasons for this. Lawyers may be ignorant about coding or the way such systems work. The algorithm may be confidential or secret. There may be a fundamental difficulty in understanding how the output was derived. If the decision of the model is based on an assessment of a volume of material that a judge would never be able to assimilate, how is he meant to decide if, in the light of that material, the output from the algorithm is correct or even rational?

32. Lord Sales, in his Henry Brooke Lecture in November 2019, suggested that to deal with this problem we need a new agency for the scrutiny of programs, which would constitute a public resource for government, parliament, the courts and the public generally. He envisaged an expert commission, staffed by coding technicians, with lawyers and ethicists to assist. He suggested that there would also need to be a pro-active mechanism for identifying cases which raised systematic issues. Lord Sales' lecture is an important one, and his suggestion deserves serious consideration, although I am not convinced that it would provide more than one part of the answer.

33. In the meantime, in my view, judges need to develop tools, both factual and legal, for understanding and interrogating AI derived decisions. This is not easy, as AI does not give reasons. But, in every case, judges should insist on an intelligible explanation for how the algorithm reached its conclusion. What forms might such an explanation take?

34. One kind would be to attempt to explain or replicate the model's reasoning. This is sometimes referred to as a decompositional approach. Another model-centric approach does not attempt to explain the inner-workings of the model, but instead attempts to explain how it works using extrinsic methods. For example, explaining the creator's intentions behind the modelling process, the type of model used, the parameters adopted, descriptions of inputs and how the model was tested for undesirable properties. A different type of approach might audit the outcomes of the machine learning system and scour the system's decisions for appearances of bias or error.

35. Attempts are apparently underway to develop systems that can provide explanations, so-called explainable AI. Unfortunately, it seems that the explain-ability function tends to undermine the accuracy of outcomes.

36. Judge's are entitled to insist on an explanation. This is particularly so in the context of administrative law. In R (Lumba) v Secretary of State, Lord Dyson said that "*it is in general inconsistent with the constitutional imperative that statute law be made known, for the government to withhold information about its policy relating to the exercise of a power conferred by statute*". A policy is, in substance, a set of rules for the decision maker to follow when exercising the relevant power. The government needs to be able to explain the policy embodied in an AI system.

37. In short, the subject matter of legal disputes is increasingly going to involve AI determinations, and judges need to develop tools for ensuring that they can interrogate and understand such systems. They are also going to need to re-shape existing legal concepts to address the issues that arise. The incremental approach of the common law, if sufficiently questioning, may be a useful process for this.

The intrusion of AI into the legal process

38. The second matter that I wanted to discuss is how AI is likely to intrude on the legal process itself. In other words, using AI to assist in deciding a case, or even to decide a case itself. This development is likely to be accelerated by the first. In other words, the more familiar courts become with AI, the more willing they will be to allow AI to play a part in the decision making process itself, or the less able they will be to resist it.

39. Some examples of the use of AI in legal proceedings are now commonplace. For example, in a large case disclosure of documents is now invariably done by e-disclosure, an algorithm that identifies relevant documents. The algorithm is trained on a limited data set, with its results being reviewed by a lawyer, and when the accuracy rate is sufficient is left to go on and identify the relevant documents in the overall data set.

40. At the other end of the spectrum is predictive AI. AI systems have been trained to predict the outcome of court proceedings. Impressive results have been achieved in predicting decisions of the ECHR and of the US Supreme Court. An AI model for the ECHR predicted 584 decisions with an average of 79% accuracy (and as high as 84% for violations of Article 6). An AI model for the US Supreme Court predicted 28,000 case outcomes with 70.2% accuracy and 240,000 judicial votes with 71.9% accuracy. Both studies used the available databases of judgments for each court to construct their model and applied natural language processing and machine learning algorithms to text-based material. Such programmes are now sufficiently sophisticated that some law firms are now apparently using predictive AI to decide whether or not to take on cases.

41. AI systems are going to start to intrude on dispute resolution. At first this is most likely to be by way of alternative dispute resolution, because here the parties can agree how their dispute is to be resolved. There are various papers on the scope for AI in arbitration. At some point commercial parties might be happy to use an AI model as the agreed dispute resolution procedure in their contracts. Rather than agreeing that any dispute should be decided by the court, an arbitrator or an expert, they agree to be bound by the outcome of an AI determination. The process is cheap and quick. It would also provide certainty, which is what commercial parties often say is their overriding need. And, as one pointed out to me recently, 70% accuracy is probably more he usually receives from his lawyers. Obviously, the other development in this area is smart contracts.

42. Small claims are another obvious example, where the value of the claim does not permit or justify expensive traditional court proceedings. Automated adjudication is effectively already happening through online dispute resolution platforms used by companies like eBay.

43. I was recently instructed in relation to a pay-day lender which was in financial difficulties due to the thousands of complaints that it was receiving. Unless a solution was found, it would have had to go into liquidation in which event its customers would not have received anything. In response, the company developed an algorithm which assessed claims. It was taught in the same way as an e-disclosure program. The model was assessed as having a better than 90% accuracy rate, which in practice was probably better than a human reviewer. Normally you cannot force such a process on parties. People are entitled to their day in court. The idea, however, was that a judge would have approved that process, as part of a scheme of arrangement, which would have bound customers.

44. Such programs could easily have a role even in more mainstream commercial cases. The courts are, for example, keen to encourage parties to mediate and have power to stay a case and require them to do so. There is software which applies game theory to the mediation process to try and encourage settlement. One input might be predicative AI as to the likely outcome. We have always had something called "early neutral evaluation" where a tribunal gives an early view on the merits, although it has rarely been used. This would merely involve AI providing such view.

45. As AI begins to spread its reach, it is inevitable that it will begin to affect the legal process more directly. Indeed, it is already doing so. The criminal justice system in the US has already seen a widespread and growing use of predictive algorithms in the context of bail applications, sentencing and parole. Whilst such use raises serious questions about lack of transparency and potential bias, the evidence suggests that, for better or worse, their predictions are more accurate than those of judges. If they are more accurate than judges, it will be difficult to resist their use.

46. One can easily imagine ways in which over time AI is increasingly used to assist a judge with more substantial commercial trials.

47. A program could check witness evidence for consistency with the underlying documents, or even say whether a witness is telling the truth. Another cognitive failing we suffer from is that more detailed descriptions of events can give rise to higher judged probabilities. It's called the conjunction fallacy. There is already evidence suggesting that machines are better than humans in detecting lying.

48. Alternatively, a program could be used to provide the judge with an automated list of comparable cases which he needs to consider, or even perform deviation analysis on his draft judgment. He could also cross-check his decision with predictive models to see if the outcome was one that you would expect.

49. None of this will necessarily be resisted. A report by the Online Dispute Resolution Advisory Group in 2015 for the Courts and Tribunals Judiciary in the UK states that it envisages AI carrying out various tasks in the future, such as legal diagnosis, facilitating settlement without human intervention and acting as "intelligent assistants" to judges. Although the Report does say that, at the moment, they do not propose replacing judges with AI.

50. There are obvious risks with such an approach. One risk concerns what's called "automation bias". In other words, the tendency of human beings to defer to an automated decision. It's another example of System 1 taking the easy life, and System 2 finding it too tiring to override it. We are probably all familiar with continuing to follow our GPS even where the suggested route does not seem sensible, in some well publicised cases even as far as driving into a river. If a judge is confronted with an AI program that says that his judgment does not accord with the mainstream, he may be unduly influenced by that.

51. A connected risk is the withering of judicial expertise. It has been said that, in this AI world, judges will increasingly adopt a "managerial stance" which will develop into merely umpiring algorithmic functions in court.

52. One could get to a stage in which one begins to question the rationale for even having a judge. If the algorithm's decision is based on data that is too large for an individual judge to assimilate or it is impossible to explain how it was reached in a way that would allow a judge to review it, what actually is the judge's role? If the algorithm is more accurate than a judge and cannot sensibly be reviewed by a judge, why do we need the judge? Alternatively, if you prefer to view it this way, over time the algorithm would become the judge. A longer term but logical consequence of this may ultimately be the loss of the traditional court process entirely.

53. In the shorter term the steady infiltration of AI may have an increasing impact on the status and legitimacy of judges. The risk is that their authority would decline, as they would no longer be seen as the real deciders.

54. Much of this is no doubt a long way into the future. However, even our assessment of the time frame may be biased. Another of our cognitive failings is that we tend to assume that a problem will have a bigger short time impact than it will in fact have and, when the anticipated disaster does not immediately occur, assume that it will have little or no impact at all. Richard Susskind, who is one of the leading lights in this area, says that,

whilst it won't be tomorrow, AI heralds an end-game for lawyers. Whether or not he is right, studies of the impact of technology on jobs suggest that lawyers are only half-way down the list of those most likely to be replaced.

55. In short, it is inevitable that AI will increasingly affect the way that disputes are decided even in court. There is no reason to think that judges and the legal process will be immune to developments which are already affecting doctors, accountants and almost all other professionals.

The baby and the bathwater

56. This brings me on to my final point. We need to be conscious about what we might lose in a process which has already started and which will inevitably continue.

57. Some law is relatively mechanistic and susceptible to an entirely automated approach. I read recently that, in major Chinese cities, cameras and facial-recognition systems are being used at pedestrian crossings to identify those breaking traffic laws, and that systems developers are working with telecom providers to ensure that, as soon as pedestrians cross the road against a red light, they get a fine on their mobile phone. But, even in such cases, there is a rigidity and automaticity in the process that makes me uncomfortable. Something is lost, if we cannot voluntarily choose to comply and occasionally to ignore.

58. Much law is not, however, capable of being mechanistically applied in this way. In some cases, the problem is that we are dealing with questions of syntax, semantics and pragmatics to which there is no straightforward answer. In other cases, it is because the law involves standards which are open textured, like concepts of reasonableness, fairness or justice, or because the case requires one to resolve competing and incommensurable principles. Judges often talk about an iterative approach, checking their reasoning against their intuitions about the result, and vice-versa. In arguing cases, I am often struck by how relatively small changes in the facts can alter one's sense as to what the result should be.

59. This is not to say that AI will be incapable of dealing with such cases. It may be able to discover patterns that we cannot clearly discern ourselves, in the same way that the IBM computer can play the Chinese game GO, with a skill and, apparently beauty, that a human being cannot match. Perhaps we will look back many years in the future, on an AI synthesis of an area of law in the way that we look with admiration at Peter Birks' analysis of the structure of unjust enrichment.

60. But, even assuming this, there is something else that is embodied in such laws that we should not lose sight of, and that is flexibility and the scope for change. One danger with AI is that we would lose such flexibility. One simple example that Richard Susskind gives concerns playing the card game patience. If you play the game against a computer you are constrained by the rules. In contrast, a human being can always modify the rule to try and make the game better. As anyone familiar with the offside rule in football knows, we do it all the time, not always for the best.

61. One common question about judges is whether they make law. The short answer is that of course they do. Indeed, they have a constitutional role to make law as part of the development of the common law. In my view, it is idle to say that they do not also make law when they interpret a statute which is unclear or ambiguous. The issue, as it seems to me, is not about whether they make law, but about judges having an appropriate sense of the need for judicial restraint, given their constitutional position. That does not mean that judges cannot develop the law, rather that they need to be sensitive about when and how they do so.

62. Law develops over time, and needs to develop over time as societal views change. Lord Hoffmann used to explain how some areas of English commercial law in Victorian times reflected Britain's position in the world. Steven Pinker, in his book Enlightenment Now, has a series of graphs intended to show that, however you measure it, the world has got better, one of which shows changes in social attitudes. I was struck by the fact that social attitudes in some countries, which are often depicted as unenlightened, are apparently now little different from what they were in England in the 1950s.

63. The problem with AI is that, if we are not careful, we may lose some of the flexibility that the legal system presently provides, and its ability to develop, slowly and incrementally.

64. Another issue is the loss of agency. Various people have written eloquently about the social aspect of law, and the benefit of being able to present one's case to a judge, who listens, engages and gives it serious consideration. Adrian Zuckerman says that, in this way, the legal process recognises litigants as autonomous agents and as persons deserving of respect and concern. As a practicing barrister, who over the years has seen the occasional judge whose bedside manner is less than ideal, I think we should not exaggerate this.

65. But the social aspect of the process is important. I learned last year about a piece of legal software called Rhubarb which is running in the United States. The idea is that two parties submit a case for resolution. Lawyers would pay to participate and say what they thought the right answer was. The view that got the most votes would be treated as the right answer. The lawyers who had that view would share the pool, with the other lawyers receiving nothing. Most lawyers in the room were horrified at the idea of having to "pay to play". But it was another point that got my attention. The parties who submitted the case were not bound by the result. Over 90% nevertheless accepted it, simply because they had agreed to the process and because it was the majority view. We are social animals and find it easier to accept a result if our peers think it fair or right.

Conclusion

66. I will end with this thought. I wonder whether the legal process is not, in the long run, better for requiring the involvement of people rather than an algorithm. But perhaps that's just another social attitude that may in due course change, and that in years to come we will prefer the certainty and consistency of AI determinations. Or perhaps, for some of us, that change has already happened.

RD QC

21 January 2021