

Legal Geek Conference

Main Conference Keynote Speech

Wednesday 14 October 2026 at 10am

Sir Geoffrey Vos, Master of the Rolls

Introduction

1. I am delighted to return to this conference for the third time. On this occasion as has been said, I am talking to you in my last three weeks as Master of the Rolls. I hope, though, that my retirement will not dent my enthusiasm for, and my involvement in, the adoption of new technologies by the legal sector and judiciaries in the UK and beyond.
2. I have looked back at some of the things I have said to you before. In 2024, I said that:

Citizens will surely use these AI tools in practice to get legal advice and will surely want justice delivered more speedily and efficiently by the use of them. If lawyers and judges refuse to adopt them, they will be left behind. Once the human lawyer or judge is faced with something that may take days to check – or simply be uncheckable – what choice do they have but to accept the advice or verdict of the machine?

This is a serious problem for the future. It will ... arise from an inevitable commercial and perhaps human imperative to allow citizens and business to achieve what they want and need – namely effective legal solutions – more quickly and efficiently than humans can achieve ... This is the problem that requires our urgent attention.

3. In two years since I said all that, it has become a complete reality.

4. As I speak: (i) almost everyone with a legal problem asks an LLM for the answer before going anywhere near a lawyer, (ii) many, if not most, lawyers routinely ask an LLM or *Harvey* or *Legora* to assist with their legal work, (iii) judges in this jurisdiction are provided with a ring-fenced version of Microsoft CoPilot on their computers, and (iv) judges in other jurisdictions, such as South Korea, Kazakhstan and China routinely use AI-enabled programmes to assist them in deciding their cases.
5. So, this morning, I want to consider how these developments will affect the way legal services and dispute resolution are delivered in the next two or three years.
6. As I see it, the problem is essentially a conflict between principle and economics. The principle is that human judges read the documents in their cases and then decide them. Economics will or may dictate that the assistance of reliable machines that save judges time, perhaps in circumstances where judges do not have time to check what the machine is telling them, will clear backlogs and provide cheaper justice.

The Chinese model

7. I recently attended the conference of Asian and Pacific Chief Justices held in Seoul. Almost every speaker, even those from jurisdictions already using AI to help resolve court disputes adopted the principle I have stated. They said that decision-making **must** always be made by human judges, and lawyers must remain responsible for their work product, whether or not assisted by AI.
8. I also visited the Shenzhen Intermediate Peoples' Court in September 2026. There, an AI-powered platform has been developed that 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 the reasons and order that the judge might want to provide to explain and encapsulate their decision. One judge told us that she could,

using the system, do 6 cases in the time she had previously taken to do one.

9. After my visit, the Shenzhen judges provided me with a copy of their White Paper describing the system they are using. Let me be clear, the system is remarkable. First, it is a single system for all types of civil case, which is something I had always thought would be difficult to achieve. Secondly, it is based on entirely digital machine-readable files that include every aspect of each case and access to relevant legal materials. Thirdly, it provides the judge with the opportunity to refer quickly to the originals of all those materials. Fourthly, it provides, the judges say, for both data security and hallucination protections.

10. The White Paper that the Shenzhen Intermediate Peoples' Court issued in early October includes the following:

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.

11. It continues:

By taking over formatting, data entry, and proofreading, the System frees judges from mechanical work for the contested heart of the case. Formally, documents are produced strictly to standard, with fewer human errors and defects. Substantively, the reasoning is fuller and answers each contested point. As one lawyer reported: "My client read the reasoning in the judgment and said he felt taken seriously. As the losing party, his sense of grievance was noticeably eased." The System has now assisted in drafting 678,000 judicial documents.

How will judicial decision-making be affected?

12. It seems to me that some of the statements in the White Paper and the reality of the way the Shenzhen System works are in conflict with the stark fact that AI only saves the time of human lawyers and judges if those lawyers and judges actually accept what it says. If they read everything they would otherwise have read, it takes the judge the same time as it always took them. Likewise, if they write themselves everything they previously wrote.
13. The problem is, therefore, more complex in the real world than simply saying AI can never replace judges or judicial decision-making for five main reasons.
14. First, as court backlogs continue to grow in many domestic court legal systems, there will or may be considerable **public pressure** to use AI to reduce those backlogs and to decide cases more quickly. The use of AI in the judicial decision-making process, even without it being ultimately responsible for the decision itself, will resolve more disputes more quickly and more cheaply than was possible before.
15. Secondly, the effect of this pressure to reduce backlogs and to decide cases more quickly, will or may be that judges, using AI, will have less and less time to read the case papers and will become more and more reliant on machine-generated fact summaries, lists of issues, and indicative outcomes. Time and economic pressures will simply not allow judges to check the work of the machine or to be inquisitive by looking at all the core documents themselves, even though they are readily available within the systems.
16. Thirdly, parties to smaller legal disputes are already demonstrating by their actions that they are or may be willing to **agree** to their disputes being resolved by machines rather than by human judges. We already see 90% of the 60 million small eBay disputes every year being decided without human intervention and, more importantly, without human complaint. The American Arbitration Institution - ICDR - has introduced an arbitration process by machine, and nobody

seems to be suggesting that the parties are acting inappropriately if they agree to it.

17. Fourthly, machines will soon be able to offer reliable solutions to legal problems. At that stage, in small cases at least, the parties on both sides of a legal dispute will use those machines to tell them what the judicial outcome is likely to be. As they are both advised of these outcomes, the parties will feel economic pressure to settle at the level predicted by the machine rather than waiting for a decision made – whether or not that decision is guided by the same machines or programmes that the parties are using.
18. Fifthly, in my opinion, it is very difficult to articulate the boundary between mechanical or algorithmic tasks such as transcription and case management and summaries and chronologies, on the one hand and tasks that comprise part of the decision-making process on the other hand. That distinction is more easily said than actually defined. I am troubled by the idea of the machine drawing up lists of important facts, when the judge has not read all the facts. Surely, a judge using the Shenzhen System could be tempted not to check if any relevant facts are missing. I am troubled by the idea of a machine drawing up lists of relevant issues. Surely, a judge might be tempted, if time was short, not to check if any relevant issues are missing. These steps are all critical parts of the judicial decision-making process. If the judge can do twice or three times the number of cases that they could do before, surely they cannot decide so easily what they absolutely need to read.

Is there a slippery slope?

19. It seems to me, therefore, that we may be on the edge of a slippery slope. It is obvious that simple small cases can and will be decided by AI if the parties agree. It is obvious the parties will agree because it will be quicker and cheaper. But how do the courts draw the line between small cases and more important ones?

20. It may be a solution to say that first instance and/or simple and/or small cases may be resolved with the help of AI, but appeals may not. But will judges then have the experience or expertise to do the work for appeals that is not required at first instance. In any event, we surely cannot stop parties consenting to have their disputes resolved privately in the way they want if they give fully informed consent – at least in the civil environment.
21. I do not have a solution, but I think a serious international debate is needed so that the problem I have described does not become a very serious one. It would be very undesirable to see arguments between jurisdictions about whether decisions assisted by AI were enforceable in another jurisdiction on the basis that they were not properly to be regarded as decisions by independent and impartial tribunals.

SIFoCC's AI project

22. The Standing International Forum of Commercial Courts is attempting to distil good practice in the use of AI in commercial decision-making from across its 71 jurisdictions. It is attempting to identify uncontroversial use cases for AI, and to say where caution is required because of the risk of inadvertently reducing judicial involvement in decision-making.
23. SIFoCC's AI project will likely end up developing model rules for the use of AI. I think such Model Rules would be valuable to all jurisdictions, whether or not they are enthusiastic adopters of AI to assist judicial decision-making.
24. There is a wide spectrum of judicial approaches to AI, ranging from the Shenzhen System at one end, to judges in some jurisdictions that reject its use entirely. An example of that was a judge from the Philippines who said last month at a conference in Seoul that AI was “an irresponsible lawyer and a catastrophic judge”.

Conclusions

25. I hope I have given you some food for thought.
26. It has been a pleasure speaking to this important gathering as Master of the Rolls now for three consecutive years. This is a vibrant industry event that will ensure that the use and adoption of new technologies is informed and collaborative. Long may that last.

GV