Artificial Intelligence, the Courts, the Legal Profession and the Legal Process

Lord Sales

UK Supreme Court

Keynote speech

Conference on AI and the Legal Profession, DAC Beachcroft, 18 November 2026

I am grateful for the opportunity to speak to you today on the impact that Artificial Intelligence (AI) is having on the legal profession and the legal system and about the opportunities, challenges and the risks associated with it. The impact of AI on the legal profession, as on so many areas of commercial life, has moved to the top of the agenda for legal professionals in line with the amazing speed with which the power of AI tools has developed. I do not think it is an exaggeration to say that we find ourselves living through a knowledge revolution which is swiftly developing into a legal revolution.

I will address three topics: AI and the courts; AI and the legal profession: and AI and the legal process. Each of them interacts with the others. But it is fair to say that it is the impact of AI on the legal profession which is at the forefront in this jurisdiction and which is producing effects in the other areas. Accordingly, I am very conscious that you all, both legal professionals and the insurance professionals working in the legal services area, are likely to have a highly developed appreciation of some of the things I will speak about. Nonetheless, I hope it will be of value if I try to set out a view of the developing picture as seen from the cloistered, but increasingly concerned, environment of the Supreme Court.

It is difficult to think of a subject about which there is simultaneously so much excitement, so much anxiety and so much confident prediction. I am not a computer scientist. What I want to do is to stand back from the technology and consider its implications for the legal system in general terms. We have moved with incredible speed through three phases, from relatively basic digital tools, to increasingly powerful tools based on large language models (LLMs), to agentic AI, where AI agents can be set off to perform tasks according to their own "intelligence" without the need for continuing human supervision or prompting. My particular concern is not so much what AI can do today, but what the next generation of AI tools may be able to do tomorrow and what their arrival may mean for the institutions and professions which sustain the rule of law.

There is a thread running through the three topics of AI and the courts, the legal profession and the legal process. AI promises something of immense value: greater access to the courts and to legal advice, achieved more quickly and at far less cost and, in some respects, perhaps achieving better results overall.

But the growth and development of AI also pose a rather uncomfortable question. What happens to a legal system when the technology upon which lawyers and judges increasingly depend begins to

perform the very activities through which lawyers and judges traditionally acquire their expertise? There is a risk that they may become de-skilled. That risk seems to me to be one of the central questions which should be being thought about now and which will have to be addressed by both the courts and the legal profession.

The courts

Let me start with the courts. At the Supreme Court we approach AI with considerable interest, but also with considerable caution. There are three obvious reasons for that caution. First, there is the problem of the **accuracy** of results which research using AI and LLMs can achieve, and where it falls short. This is also, of course, something which legal professions have to be very careful about.

Generative AI can produce propositions, quotations and authorities which look entirely convincing and yet are entirely false. Lawyers have already discovered, sometimes to their considerable embarrassment, that an invented authority delivered by an AI system can look remarkably like a real one.

The danger is particularly acute in law because so much of what we do depends upon authority. A judicial decision has to be accurate. A statutory provision attributed to Parliament must really be in the statute. A quotation attributed to a judge must really have been said by that judge. A case which is cited must actually exist. Plausibility is not enough.

The second problem is **confidentiality**. Judges work with draft judgments, internal communications, parties' confidential material and, most importantly, engage in private judicial deliberations. Putting such material into an open AI system may expose information which must remain confidential. That has important consequences for the way courts adopt the technology. It makes secure systems, effective information barriers and careful governance into requirements rather than optional extras. Secure bespoke AI environments can make an important difference, provided information-sharing functions are appropriately controlled. But this has costs implications for financially stretched courts.

The third problem is more profound. It is **legitimacy**. Suppose an appellate judge asks an AI system to summarise the parties' written cases. Many people, I suspect, would regard that as a sensible use of technology, at least for preliminary preparatory work. Suppose the judge asks it to identify whether the draft judgment has failed to address one of the parties' arguments. Again, there may be relatively little objection.

Now suppose the judge asks AI to research a discrete legal question. Perhaps one becomes slightly more cautious. Then suppose the judge asks it to prepare the first draft of the judgment. And finally suppose that the machine is asked: “Here are the papers. Who wins?”

At some point as we travel along that spectrum, something important has changed. But precisely where? That is not merely a technological question. It is a constitutional and social question.

The authority of a court depends in part upon public acceptance that the person who exercises the power of the state by deciding the dispute has actually exercised judgment. At the constitutional level, that consideration is especially powerful. Constitutional and final appellate adjudication requires choices between arguments and principles, assessments of institutional competence, and

decisions which sometimes contribute to the development of the law itself. The judge is accountable because the judgment is theirs, and their reasons are open to public examination and criticism.

An AI system cannot bear constitutional responsibility. That does not mean that it cannot assist somebody who can. That distinction between assistance and responsibility seems to me extremely important.

The lower courts and the access-to-justice paradox

One should be cautious about reasoning from the Supreme Court to the whole judicial system. It is likely that the most transformative effect of AI will initially occur elsewhere.

For decades one of the great problems of civil justice has been cost. People may possess perfectly good legal rights yet have no economically rational means of enforcing them. AI has the potential to alter that equation dramatically. Someone who cannot afford a solicitor may now be able to describe a grievance to an AI system and receive something which resembles legal advice. The system may help identify a cause of action. It may explain procedure. It may turn the grievance into something resembling a plausible legal claim. That is, potentially, an extraordinary advance in access to justice.

But there is a paradox. If AI substantially reduces the cost of producing claims, it does not automatically reduce the cost to the justice system of deciding them. A claimant may take seconds to generate 30 pages of apparently plausible legal argument. Unfortunately, it does not take an unaided judge seconds to discover that substantial parts of it are nonsense.

So improved access to Ai-sourced legal advice and AI production of claims is already leading to a tsunami of claims. Courts already struggling with limited resources are likely to face an enormous increase in plausible-looking legal material.

Something similar applies to lawyers. Generative AI abolishes one of the natural constraints on litigation: the expense of producing words. Historically, asking a solicitor or barrister to formulate six alternative arguments cost money. Producing a 100-page submission required considerable human effort. An AI tool has no such inhibition. The marginal cost of the 101st page is effectively zero.

For insurers of lawyers this may be a point worth watching. The professional risk created by AI is not necessarily that lawyers will do too little. It may sometimes be that it enables them to do far too much, far too rapidly, without sufficient human discrimination.

AI deciding cases

How should the courts respond to the coming tsunami of claims?

AI may itself have to be enlisted by the courts to become part of the solution. Initially that can occur at relatively uncontroversial levels: administration, document management, summarisation, legal research, routine case management and the organisation of court files.

But if caseload pressures continue to rise, the pressure to go further will be considerable. AI may have to be used to triage cases. It may identify apparently straightforward claims. It may prepare material for a judge. It may produce draft orders. And eventually it may draft judgments.

At that point we confront a question which is uncomfortable but which I think lawyers and judges have to discuss openly: must adjudication always be human?

In an article I have published about the possible use of AI and its compatibility with article 6 of the European Convention on Human Rights[1] I have explored the proposition that the answer may not necessarily be yes in all cases.

Article 6 speaks of an independent and impartial tribunal. A tribunal is characterised by its judicial function and by the procedural and substantive guarantees surrounding that function. It is at least arguable that one should not simply insert into the Convention an additional requirement that a tribunal must, in every circumstance, consist of a human being.

That proposition may initially sound startling. But consider a low-value commercial dispute involving no disputed oral evidence and a straightforward issue of contractual interpretation. Imagine an AI system which can determine that dispute in one hour for £50. Imagine that the human alternative costs the parties £20,000 and provides a judgment in 18 months. Society will come to ask, which is better access to justice? The answer is not self-evident.

Article 6 itself includes the right to a determination within a reasonable time. Rights have to be practical and effective rather than theoretical. Speed and cost are not vulgar considerations external to justice. In some circumstances they are important aspects of justice. The existing jurisprudence on digital justice and access to court provides conceptual room for technological procedures where they improve the effective administration of justice.

There may be contexts in which experimentation is easier than others. Parties might voluntarily agree that a particular class of dispute should be determined using AI. Properly informed and voluntary waiver of rights already has a recognised place within the article 6 jurisprudence.

AI might also operate as a specialist first-instance decision-maker subject to review by a human court. Existing Strasbourg jurisprudence recognises that specialist administrative decision-making can be compatible with article 6 where appropriate judicial supervision exists. That suggests an interesting possible model for fields such as regulatory, planning or welfare determinations. Low-value civil disputes present another obvious possible field of experimentation. By contrast, more serious criminal proceedings are plainly a very different matter.

So I doubt that the future choice will be between "AI judges" and "human judges". The more realistic question is how we divide functions between systems and humans at different stages and for different sorts of dispute.

The human in the loop

That leads to the fashionable idea of keeping a "human in the loop". It sounds reassuring. The AI tool produces an answer. The judge checks it. Human responsibility is therefore preserved. But I think we should be rather careful about seeking reassurance in this idea.

Imagine a judge with 50 cases to deal with this week. The AI tool generates a beautifully written draft judgment for each of them. It identifies the facts, cites the relevant authorities and produces a conclusion which appears entirely sensible. The judge can spend two hours reconstructing the

[1] Lord Sales, 'Human Rights, Artificial Intelligence and Judging' [2096] Public Law 569.

analysis independently, or press “approve”. The human is certainly in the loop; but is the human performing any useful function there?

The more efficient and apparently reliable AI becomes, the greater the danger of automation bias: the tendency to accept the machine's answer precisely because the machine usually gets things right. The “human in the loop” can gradually become the human rubber stamp. If a junior judge under intense caseload pressure receives an apparently perfect draft at the press of a button, the institutional incentives push towards acceptance rather than genuine independent scrutiny.

This is why AI competence cannot simply mean knowing how to operate the software. True AI competence requires the human being to possess enough independent expertise to know when the machine is wrong.

Which brings me to the legal profession.

The legal profession

Law firms are already moving considerably faster than courts. That is not surprising. The incentives are powerful.

AI can conduct research, summarise large numbers of documents, analyse contracts, produce first drafts, compare provisions, construct chronologies and perform many other tasks which previously absorbed substantial amounts of lawyers' time. That offers clients faster service. It offers firms higher productivity. And it can reduce cost dramatically.

Those pressures have competitive effects. Once one firm can provide a good product in two hours using AI, another firm will find it difficult to charge for 20 hours of junior lawyer time to produce an equivalent result. So the question whether lawyers want to use AI will become increasingly irrelevant. In many fields market competition will effectively require them to do so.

That has obvious implications for professional liability and professional indemnity insurance. AI changes the mechanism through which mistakes occur, but it does not remove professional responsibility. If an AI system invents an authority and a barrister puts it into a skeleton argument, “the computer made it up” is unlikely to be a defence which greatly impresses the judge. If the AI generated advice to a client proves to be wrong, the client will not regard that as acceptable. The professional obligation is still the lawyer's.

That means that firms need systems which distinguish at least three things: tasks that may safely be delegated to AI; tasks for which AI may be used but its output must be independently verified; and tasks in which its use may be inappropriate altogether.

Insurers have an obvious interest in how that governance develops. Questions which were once matters of IT procurement are becoming questions of risk management. What system did the firm authorise? What confidential information could the system receive? What audit trail exists? What verification was required? Who had responsibility for checking the product? Were lawyers trained in the particular risks of the system? Those questions look remarkably like the questions which insurers have always asked about professional systems and supervision. The technology is new. The underlying problem is not.

De-skilling

There is another problem which I think is more profound and more difficult. It is the problem of the de-skilling effect of AI.

AI is already good at the work currently done by trainees and junior lawyers, and is improving. From a client's perspective, that sounds excellent. Why pay a junior lawyer for ten hours to review documents or research authorities if an AI system can do it in ten minutes? But now consider the question from the perspective of the profession 20 years later. Where does the experienced senior lawyer come from? Who will be the lawyer with the critical faculties required to examine the product which AI generates?

The apparently mundane tasks undertaken by junior lawyers are not merely a less efficient means of producing legal services. They are also part of the process by which lawyers learn. You read 50 cases and gradually develop a sense of which cases matter. You draft a note and somebody more experienced pulls it apart, and you learn from the experience. You conduct disclosure and learn how facts emerge untidily from documents. You make bad points and discover why they are bad. Eventually, if all goes well, something called professional judgment develops.

The less efficient human process has been producing two different things at once: today's legal work and tomorrow's lawyer.

AI allows us to separate those products. That creates a serious problem. A law firm which eliminates expensive junior work may improve its productivity this year while quietly failing to generate and develop its own human capital for ten years to come. The judiciary faces the same problem. If judges cease to research, analyse and write because machines increasingly perform those functions, critical faculties will atrophy. If younger lawyers never perform those tasks themselves, the future judiciary may contain people who have never acquired those faculties in the first place.

We might then arrive at the extraordinary situation in which we insist that an AI-generated judgment must be checked by a human, while progressively eliminating the training through which human beings acquire the capacity to check it. That is not a sustainable model.

Training for a different profession

So legal education and professional training may have to change. Paradoxically, as AI becomes better, some lawyers may need to be trained without it. There may need to be protected parts of legal education and professional development in which people are required to research, analyse, draft and reason for themselves. Law firms may need to identify a smaller cadre of lawyers in whom they invest much more heavily, because those lawyers will eventually provide the sophisticated human judgment required to supervise increasingly capable AI tools. The same may apply to judicial training.

The result may be economically uncomfortable. The number of lawyers required to produce the same quantity of legal work may fall sharply. But the level of investment required in the development of some of those lawyers may have to increase. Fewer practitioners may be needed overall, while an expert cadre has to receive more intensive training so that it retains the critical legal skills necessary to interrogate and control AI output.

That, I think, is a subject in which professional indemnity insurers have a legitimate interest too, because adequate supervision ultimately depends on adequate expertise.

AI in the legal process

I turn finally to the legal process itself.

We have spent years talking about the information explosion. AI changes its legal significance. Modern commercial activity produces astonishing quantities of data. Emails, Teams messages, contracts, transactional records, electronic accounts, recordings and other digital material can all become relevant to litigation. In a major commercial dispute, it may simply no longer be realistic to understand the evidential picture without machine assistance.

Lawyers are increasingly having to use AI for disclosure, document classification, chronology construction and the identification of patterns across enormous datasets. Once the lawyers do that, the courts have to understand what they are doing. A judge needs to know enough about the tools being deployed to determine whether the process is fair, whether relevant material has been missed and whether the methodology has distorted the material presented to the court.

AI therefore becomes not simply a tool used in litigation but potentially an object of litigation and subject to dispute. We will have arguments about models, datasets, prompts, validation, bias, disclosure and audit trails.

Further, where AI is used for decision-making – say by government and regulators - we should not assume that the machine is free from bias simply because it is a machine. AI systems may inherit bias from training data, model design and the social material from which they learn. At the same time, one should resist the opposite mistake of imagining that human decision-making provides some perfect, bias-free benchmark. Human beings have their own conscious and unconscious biases. One interesting possibility is that algorithmic systems may sometimes allow forms of systematic bias to be identified and tested more rigorously than human intuition permits.

The correct question is therefore comparative:

For this particular task, subject to these particular safeguards, which decision-making process performs better?

That is more useful than asking whether AI is perfect. AI is not perfect and nor are humans.

Speed, cost and legitimacy

This brings me to what I regard as the deepest issue. We tend to think that maintaining public confidence in the legal system requires preserving human decision-making. At present that is probably largely true, particularly where decisions affect liberty, family life, the home, financial survival or important constitutional rights. Society may be content for a machine to determine a small consumer dispute while expecting something very different when the decision threatens someone's liberty, home, children or solvency.

But public expectations are not static. Imagine a world in which almost every ordinary service gives an immediate answer. A bank transaction takes seconds. Translation is instantaneous. A sophisticated document can be produced in moments. Information is available immediately. You get

used to this speed of living and decision-making. And then you have a legal dispute and a court says: “Thank you for your claim. Your case will be heard in 22 months.”

At some stage, delay itself becomes a problem of legitimacy. A judicial system can lose public confidence because it uses machines excessively. But it can also lose public confidence because it refuses to use them.

It is crucial that we acknowledge that tension. As our expectations of speed change, the social value attached to human determination may fall while the value attached to timely and inexpensive determination rises. We have to find ways to preserve human judgment without turning “human judgment” into an excuse for an inaccessible, unaffordable or intolerably slow justice system.

Where should the line be drawn?

I do not think there is a single line. There will be a series of lines. They will depend upon the type of case, the interests at stake, the quality of the technology, the availability of review and, importantly, public expectations.

For routine administrative tasks, extensive AI use may quickly become normal. For information management, research and summarisation, the opportunities are considerable, subject to verification and confidentiality. For low-value civil disputes, we may eventually accept forms of automated determination, particularly where parties agree. For specialist administrative or regulatory questions, an AI first decision followed by access to human judicial review may become conceivable. At the other end of the spectrum, decisions involving fundamental rights, serious criminal liability or the development of constitutional principle have much stronger claims to substantive human judgment.

But even there, AI may assist with supporting tasks. In other words, the sensible unit of analysis may not be the case. It may be the function. We should ask what is being delegated, why it is being delegated, what could go wrong, whether the output is reviewable and who remains responsible.

That is also why the debate cannot be left entirely to judges. Deciding how far human judges should ultimately be replaced by machines is fundamentally a democratic question.

The role of insurers

Let me say something directly to the insurers here today. You occupy a more significant position in this transition than may immediately be apparent. Professional indemnity insurance does not merely compensate for professional error after it occurs. Insurance also influences professional behaviour by identifying, pricing and helping to manage risk. That makes the insurance profession a potentially important participant in the development of responsible AI practice.

You will increasingly need to distinguish between firms which simply say, “We use AI”, and firms which can answer the harder questions. Which systems? For what purposes? On whose data? Subject to what confidentiality protections? With what verification? Under whose supervision? With what training? What happens when the machine's answer conflicts with the lawyer's judgment? And is there an auditable record enabling the firm subsequently to establish what occurred?

One can envisage AI use eventually being treated rather like other aspects of professional risk management. Mature organisations will have policies, authorised systems, training, supervision, escalation procedures and incident reporting.

The paradox is that the widespread use of AI may make professional judgment more, rather than less, important from the insurer's perspective. If machines perform routine work, human error will increasingly occur at the point at which someone decides whether to trust a machine.

Conclusion

So where does this leave us? Certainly not, I think, with a choice between embracing AI enthusiasts and retreating to a world of fountain pens, leather-bound law reports and dictation. AI is going to become embedded in legal practice. It offers enormous potential benefits: lower costs, improved productivity, better management of information and, most importantly, greatly improved access to legal knowledge and potentially to justice itself. Courts should not reject those benefits. Lawyers cannot reject them. And clients will not allow lawyers to reject them.

But technological capability does not answer the normative question: what do we want a legal system to be? At the highest level, adjudication is not merely a service for generating legally accurate outputs. Law is one of the mechanisms through which society exercises power over human beings. That is especially visible in constitutional and administrative adjudication. A constitutional and administrative court has to take responsibility for its decisions, expose its reasoning to criticism and exercise judgment about principles and institutions which cannot simply be reduced to prediction from existing patterns. Technology can assist this process, but responsibility for constitutional judgment should remain with identifiable independent human judges.

In the wider justice system the picture may be more complicated. There may be substantial areas in which machine decision-making becomes appropriate, perhaps because it is faster, cheaper and sufficiently reliable. If all we care about is commercial acceptability, AI decision-making may provide commercially attractive ways forward and become acceptable. There may be hybrid models.

There will certainly be extensive machine assistance. And the boundaries will move as technology and society change. Our task is therefore not to defend every activity which lawyers and judges happen currently to perform. Some of those activities can perfectly well be performed by machines. The more important task is to identify the things for which human responsibility really matters. And then to ensure that we continue to produce human beings capable of exercising it.

That is why I regard de-skilling as one of the great hidden risks of this technological revolution. The danger is not simply that the machine gets the answer wrong. The greater long-term danger is that human beings gradually lose the ability to know whether it has got the answer wrong.

For courts, that means preserving independent judicial judgment and training judges who can exercise it. For law firms, it means maintaining professional expertise even when the immediate economic incentive is to automate its acquisition away. For insurers, it means recognising that AI governance, human competence and professional risk are becoming inseparable. And for the legal system as a whole, it means finding a new accommodation between three things which are all aspects of the rule of law: accessibility, efficiency and human responsibility.

For the decisions which most profoundly affect individuals and society, somebody must remain answerable for the exercise of legal power. For the foreseeable future, that somebody should be a human being. The challenge for lawyers and judges is therefore not to compete with artificial intelligence at the things artificial intelligence can do better. It is to use it intelligently while preserving, developing and demonstrating the value of the things for which human judgment remains indispensable.

If we can do that, AI should not be understood primarily as a threat to the legal profession or to the courts. It can become an extraordinarily powerful servant of the rule of law.

But we should make sure that it remains a good servant, and not our master. We should also recognize that that may be easier said than done.

Thank you.

[AI was used in the production of this lecture at one stage of the writing process]