

Should legal disputes be decided by artificial, rather than human means?

Artificial intelligence (AI) has the potential to resolve some legal disputes more accurately, and with greater consistency, than humans are able to. Practical advantages, such as speed and cost of resolution, have already led to its limited adoption in some countries. This is likely to increase as AI continues to evolve rapidly, both through development of the algorithms with which it is programmed and through machine learning. But it is unlikely at any point that AI will be able to update fast enough to reflect ever-changing and infinitely complex real world circumstances, or to master the (also ever-changing) linguistic subtlety upon which much legal argument is based. Many disputes involve unique circumstances which cannot be fully encapsulated by pre-set algorithms: human oversight will be necessary. Furthermore, AI's inherent lack of transparency means that it will fail to meet the principle of open justice, except in limited circumstances: both users and subjects of AI may not understand the decisions it makes. This would risk society losing confidence in its legal system. While some disputes should be determined by AI, even these must be subject to human checks. In many cases, however, AI's best use will be as an expert advisor (or "co-bot", perhaps) for human arbitrators. Society is likely to consider it morally unacceptable for life-changing decisions (where the death penalty could be imposed, for example) to be made by a "robot judge". Ultimately, law evolves to reflect changes in society and life itself; an inanimate entity such as AI will be incapable of comprehending the infinite complexity of life in the conceivable future. It would thus be mistaken to devolve all power to AI for deciding legal disputes.

Practicing lawyers have made increasing use of AI over the last decade. A survey of the websites of the UK's 12 largest solicitors¹ indicates that this will continue, albeit with AI typically being used for back office functions and organising and categorising "Big Data" sources. It is often used for the retrieval of precedents for use in upcoming cases, for example, or for reviewing contracts. AI can perform many such tasks far more quickly, accurately and comprehensively than human solicitors would find possible. While programming and other costs are not insignificant, these solicitors all note that AI is allowing them to deploy more staff in areas where greater specialist skills are required, while automating more routine tasks. While AI has been contributing to dispute resolution for some years, it is now starting to come into use in the decision making process itself, having performed well in a series of robot v human contests. In 2018, for example, a LawGeex AI system achieving 94% accuracy in reviewing a set of non-disclosure agreements, compared with 85% by a group of legal experts in that field. Even more strikingly, AI took 26 seconds to review each agreement, while the humans required 51 minutes on average². Cambridge-based Case Cruncher Alpha also had notable success in such competitions. In China, there are now over 100 robots in use in courts³ while they have also been introduced in countries as diverse as Colombia⁴ and Estonia⁵. However, their use remains limited to that of an advisor in high-volume, low-value cases, with a human judge in attendance making any final decision. Nevertheless, AI's use here has clear practical benefits and may help improve the consistency of penalties or sentences being given. The UK too has recognised these advantages, with the Civil Justice Council recommending as long ago as 2015 the introduction of Her Majesty's Online Court to resolve some disputes where the value is under

¹ <https://www.linklaters.com/en/insights/online-services/artificial-intelligence>

² <https://www.cicerosoftware.com/en/blog-robot-lawyers-lawyers-compete/>

³ <https://www.worldgovernmentsummit.org/observer/articles/could-an-ai-ever-replace-a-judge-in-court>

⁴ https://www.garrigues.com/en_GB/new/artificial-intelligence-international-arbitration-legal-prediction-awards-issued-robots

⁵ <https://www.lexisnexis.ca/en-ca/ihc/2019-06/from-estonian-ai-judges-to-robot-mediators-in-canada-uk.page>

£25,000⁶. However, the implementation of this has been delayed until at least 2023⁷, partly due to the technology available remaining inadequate. This highlights the difficulties that the sheer complexity and holistic nature of many legal disputes poses AI.

The current usage of AI is limited to relatively simple and clearly defined situations, where organising known facts, which may not in themselves be disputed (in some divorces, for example) is what is required. There are many examples of “bots” being used outside the court system too - the donotpay.com site being one of the most well-known. Originally designed to help the public to obtain refunds for parking tickets wrongfully issued, donotpay.com has now extended its scope to help receive compensation where it is due from airlines, or to force companies to cancel subscription services when they seem reluctant to do so.

AI has some clear practical benefits which are leading to its increased adoption. While its capabilities are insufficient to deal with most decision-making tasks at present, the pace of technological change (estimated in Intel co-founder Gordon Moore’s “Moore’s Law” as a doubling of processing power every 2 years) should not be underestimated. It is clear that these capabilities will advance to a position where it could be used far more widely in future - at least theoretically.

However, there are overwhelming arguments against allowing AI alone to decide many legal disputes. These arguments may be classified broadly as practical, legal and moral.

Practically, an important difficulty is that the world is constantly evolving in unexpected ways and AI will itself need time to absorb and process new laws which may themselves be passed rapidly. In response to the coronavirus pandemic, a new UK Coronavirus Act was swiftly drawn up in early March and given royal assent on 25th March 2020. This will potentially lead to many legal disputes for which there are no precedents, and to as yet unforeseen consequences. Both programmers updating AI algorithms and machine learning will struggle to incorporate this new information in the short term. Similarly, programmers will more generally find it difficult to ensure that changes made to an AI system are not having unintended consequences on the algorithms which are already in place – particularly as the majority of programmers will not be lawyers. Neither will they be perfect individuals, so there is significant potential for their conscious or unconscious biases to be incorporated in an AI system (just as there are in any human system). Large and complex AI systems could also well be open to cyberattacks which have already led to major financial losses for many organisations; any such malicious attack on the legal infrastructure (possibly not identified for many years afterwards) could have extremely damaging consequences on the legal system.

The philosopher John Searle’s Chinese Room Argument⁸ drew a distinction between the syntax and the semantics of language and argued that AI could never fully replicate human understanding. While AI may recognise words in themselves, it will always struggle to understand, or even recognise, the multiple linguistic subtleties upon which many legal disputes turn. In essence, many disputes are too complex to be simplified in digital format. If AI cannot understand arguments in their human form (instead merely simulating this), it is unlikely to be able to administer justice fairly to the satisfaction of the humans who need it. An analogy with the recent introduction of the Video Assistant Referee (VAR) in football is interesting as, even when rules were relatively simple and well understood, the attempt to implement them, using artificial means, has been widely deemed unsatisfactory and inconsistent by players and fans alike. VAR, applied strictly, lacks any capacity for discretion.

⁶ <https://www.judiciary.uk/wp-content/uploads/2015/02/Online-Dispute-Resolution-Final-Web-Version1.pdf>

⁷ <https://www.lexology.com/library/detail.aspx?g=10651420-b002-4626-8a01-4c88411be9ff>

⁸ In “Minds, Brains, and Programs” (1980)

This complexity, and AI's inherent lack of transparency, are at the heart of major counterarguments to its use in the legal system. Centrally, the principle of open justice – whereby the law should be clearly comprehensible to all, with decisions clearly communicated – is unlikely to be satisfied by AI, except in limited circumstances. Re and Solow-Niederman have identified “a range of concerns relating to [AI's] tendency to make the legal system more incomprehensible, data-based, alienating, and disillusioning”⁹. Were this to be the case, using AI alone would in many cases very likely lead to society losing confidence in legal systems which it has used for centuries to provide an agreed framework for it to operate within. Another academic, Sourdin¹⁰, notes “questions such as who makes the decision, and who possesses the legal authority to make such a decision. Is it the computer programmer, the policymaker, the human decision-maker or the computer or automated system itself” – these cannot easily be answered and raise further crucial problems, such as how and to whom appeals might be made.

Re and Solow-Niederman ask “can a prearranged decision procedure really incorporate an idea like mercy or develop fact-sensitive balancing of mitigation factors in a criminal case?” The principle of equitable justice might mean a recently unemployed husband stealing medicine for his sick wife could be treated more leniently than a drug dealer stealing drugs to sell on the street, when it came to a human judge passing sentence. Whereas AI, certainly for the foreseeable future, would not be capable of taking every extenuating circumstance into account. Particularly as every case will have unique surrounding circumstances; and the more complex the case the more complex those circumstances are likely to be.

But, even if developments in AI made it possible in the distant future to overcome all of the practical and legal issues outlined, there are critical moral problems which will remain impossible to solve.

The UK's legal system (and that of most other countries') has had human interaction woven into it ever since it was created, or naturally evolved from primitive, informal justice methods. This not only enables a judge to exercise discretion, or mercy, but it is also considered to be a fundamental right for litigants to be able to make their case personally – and to feel that they are being heard and understood. Similarly, a jury decision making process places much emphasis on human interaction and argument. If AI displaced all such human interaction, litigants would be likely over time to become disaffected, eventually losing faith in the system of justice itself.

Where very serious cases were involved, with very serious potential consequences (in some countries, the death penalty for example) it will be considered morally unacceptable for life-changing decisions to be made by a “robot judge”. It is inconceivable that the degree of technological development that would be necessary to overcome these reservations is within human reach.

Finally, privacy issues, both for individuals and for countries when national security matters might be involved, further complicate the landscape for AI. The extent to which these are concerning depends partly on political ideology: in China today, for example, all citizens are subject to “social scoring” based on their internet presence – with low scores affecting individuals' rights to have a passport or to get a loan. This would be considered too intrusive in many other countries. In any event, AI may find itself denied information which it had in previous instances needed to make a fair decision. The

⁹ “Developing Artificially Intelligent Justice”, Stanford Technology Law Review 242 (2019)

¹⁰ “Judge v Robot?” AI and Judicial Decision-Making”, UNSW Law Journal vol 41

withholding of information for reasons of privacy will continue to be seen as a necessary human “judgement call” – albeit clearly capable itself of abuse. This is an area where politics and law meet.

The extent to which AI will improve in the future should not be underestimated. This will increase its use, which is currently limited to relatively routine, standardised tasks. Its ability to assimilate and organise data is already vastly superior to any human’s. However, it will always struggle to understand the linguistic complexity which is central to much legal argument, or the complexity of life itself, which is also often central to such argument. Without humans being part of legal dispute resolution, it is fundamentally impossible for the principles of open justice or of discretion to be met. Lower cost and greater speed will tempt rule makers, but AI’s most appropriate use is as an expert advisor to human judges.