

THE DUTIES AND JOYS OF DIGITAL JUDGING, ARTIFICIAL JUSTICE? THE INTEGRITY PROBLEM FACING JUDICIAL AI¹

Author: The Right Honourable Dame Siobhan Keegan, Lady Chief Justice of Northern Ireland

Introduction

The issue of AI represents one of the most exciting – but also potentially existential – developments in the legal sector. Exciting because there are undoubtedly ways to tap the potential of artificial intelligence (AI) to streamline and overall improve the administration of justice which we should embrace. Existential (without being too dramatic about it) because that same potential could cause serious problems within the justice system and disadvantage those who rightly seek to rely on the courts as a standard-bearer for justice and respect for the rule of law.

I come to this debate with what I hope can be described as some healthy scepticism. That is for several reasons. First, AI is not a one size fits all. Whilst AI may assist with discovery and non-contentious business, it is hard to see it being successfully utilised in areas of discretionary decision-making. Second, there is the issue of what is within the AI space in terms of information and who or what private interests controls that. Third, there is the issue of regulation, and more importantly for me as a Chief Justice, guidance for judges. Fourth, there is a need for transparency. Fifth, there is a worry of de-skilling and changing the essence of decision-making even subtly where, of course, two judges can validly come to different conclusions on the same issue, for instance in relation to interpretation of a rule.

We have already heard fascinating insights from my colleagues on the measures already being taken towards digital justice. However, I ask myself, what are the rule of law implications on justice by AI? Admittedly, this is a vast topic; and one that certainly could not be covered in a brief 30-minute presentation. As such, I intend to give only a brief overview of the salient points of discussion. To start, perhaps predictably given the subject matter, I thought I might ask one very popular large language model chatbot to engage with this issue. After all, if it provides a satisfactory answer, I can spare you the trouble of this lecture. So, I asked my judicial assistant to engage ChatGPT with the question ‘do you think the pros of using AI in the justice system outweigh the cons?’ It provided the following answer:

The use of AI in the justice system offers several substantial advantages that outweigh the associated drawbacks. Firstly, AI significantly enhances efficiency by quickly analysing legal documents, precedents, and evidence, potentially reducing case backlogs and ensuring more timely justice.

AI’s role in predictive policing is another critical benefit. It assists law enforcement in deploying resources more effectively by identifying crime patterns, ultimately bolstering public safety. Moreover, AI can help reduce bias within the justice system. When properly designed and implemented, it operates based on predefined rules and objective data, diminishing the influence of human biases in legal proceedings.

¹ Based on a paper originally delivered at the National Judicial Conference, Dublin, 17 November 2023.

Cost savings are also noteworthy, as the automation of routine legal tasks and research can lead to considerable financial benefits for the justice system, which can then be reinvested where they are most needed. Lastly, AI enhances security by identifying potential threats and vulnerabilities, contributing to the overall safety of the justice system.

While concerns regarding AI bias and potential job displacement are valid, these challenges can be addressed through stringent oversight, regulation, and continuous refinement of AI algorithms. When AI is thoughtfully implemented and closely monitored, it has the potential significantly to enhance the justice system's efficiency, delivering more timely and equitable justice and ensuring a safer and more secure society.²

Perhaps unsurprisingly, the chatbot came out in favour of its inclusion in the justice system. And while it provides a fairly robust response, we must delve deeper into the question at hand to ascertain the full picture. I will undertake this task by recourse to the following structure. First, I will discuss the rule of law as an ideal, so that we might properly apply the doctrine's principles to the AI-sphere. Then, I will approach the question of AI as a congruent judge, congruency being a necessary feature of any society governed by the rule of law. From there, I will outline some applications of AI in the judicial process and finally, I will outline what seem to me to be the hurdles that a truly AI-driven justice model must overcome before briefly commenting on the impending EU draft Act on Artificial Intelligence.

The ideal of the rule of law

Turning first to the rule of law... as an ideal, few other concepts have been so fiercely defended by lawyers than the rule of law. From academics to practitioners, from positivists to critical theorists, the rule of law has been steadfastly upheld. As Sir Thomas Bingham once said, it has become the closest thing to a universal legal religion.³

Despite this reverence, or maybe more accurately because of this reverence, the rule of law as an ideal has been incredibly difficult to define. Conceptions run from thick, all-encompassing definitions to thinner, more formalistic rationalisations. However, amongst these conceptions, there may be said to be a common thread. That is as follows. The rule of law's core of certainty should be understood, as all the conditions that the law must meet in order to be legitimate are morally binding on its subjects.⁴ To my mind this formulation captures both the procedural and the substantive aspects of the rule of law.

From my own perspective, an influential account of what the rule of law is and how it orders judicial process comes from the late Lord Bingham. His book, simply titled 'The Rule of Law', is essential reading for all students of law – from university level right through to the upper echelons of the judiciary. It is a conception that, to my mind, has not been improved upon since its publication. In the book, Lord Bingham takes us through his case for the rule of law, expounding eight guiding principles. They are invariably:

² We asked ChatGPT to explicitly take a side in this argument. We further limited its response to 200 words.

³ Tom Bingham, *The Rule of Law* (Penguin Books 2010).

⁴ John Tasioulas, 'The Rule of Law' in Tasioulas, *The Cambridge Companion to the Philosophy of Law* (Cambridge University Press 2020)

- 1) The law must be accessible and so far as possible intelligible, clear and predictable.
- 2) Questions of legal right and liability should ordinarily be resolved by application of the law and not the exercise of discretion.
- 3) The laws of the land should apply equally to all, save to the extent that objective differences justify differentiation.
- 4) Ministers and public officers at all levels must exercise the powers conferred on them in good faith, fairly, for the purpose for which the powers were conferred, without exceeding the limits of such powers and not unreasonably.
- 5) The law must afford adequate protection of fundamental human rights.
- 6) Means must be provided for resolving, without prohibitive cost or inordinate delay, bona fide civil disputes which the parties themselves are unable to resolve.
- 7) The adjudicative procedures provided by the state should be fair. [And that;]
- 8) The rule of law requires compliance by the state with its obligations in international law as in national law.⁵

These principles, Lord Bingham argued, are necessary in order to ensure ‘that all persons and authorities within the state, whether public or private, should be bound by and entitled to the benefit of laws publicly and prospectively promulgated and publicly administered in the courts.’⁶

Lord Bingham’s tenets admittedly make the case for a thick conception of the rule of law. This standard has not been universally attained. We should be alive to the reality that only 46% of the world’s population lives under the protection of the law.⁷ Further, Richard Susskind has made the claim that in countries such as India and Brazil there is a backlog of cases of 30 million and 100 million respectively.⁸ Therefore, one can reasonably understand the political allure of AI and the argument that, in streamlining the judicial process, technology actually enhances the rule of law. Central to this argument, proponents believe, is that AI will be able to produce congruent judgments; thus, satisfying the requirements for the rule of law. I, myself, am not entirely convinced.

AI as a congruent judge

The need for congruency is well accepted within the rule of law; law cannot be said to be properly applied unless legal standards are followed. Lon Fuller described it as ‘the most complex of all the desiderata that make up the internal morality of the law’ and that it ‘may be destroyed or impaired in a great many of ways’ from mistaken interpretation to bribery,

⁵ Bingham (n.3) see also Bingham Centre for The Rule of Law <<https://binghamcentre.biicl.org/our-vision>> accessed 9th November 2023.

⁶ Bingham (n 3) 8.

⁷ John Tasioulas, ‘The Rule of Algorithm and the Rule of Law’ (2023) 3 *Vienna Lectures on Legal Philosophy*; OECD ‘Government at a Glance 2021’ accessed 9 November 2023.

⁸ Richard Susskind, *Online Courts and the Future of Justice* (Oxford University Press 2019) ch 28.

prejudice, indifference, or stupidity.⁹ The argument in favour of artificial intelligence is that AI-generated decisions will be at least as congruent as human-made decisions, if not more so. Eugene Volokh styles the argument in the following way: ‘The question is not whether an AI judge actually follows rules at some deep level; the question is whether an AI judge’s opinions persuade observers who expect opinions to be consistent with the legal rules. Rule-following is as rule-following does.’¹⁰

Essentially, Volokh asks us to trust AI to the extent that it passes the Turing Test. Put simply, if AI can persuade a panel of expert lawyers to at least the same degree as a human, AI is safe to use. What matters, Volokh argues, is the result, not the process.¹¹ Another iteration of the AI as a congruent judge argument has been put forward by Cass Sunstein, who argues that AI-generated decisions would be less-susceptible to what he calls the ‘noise’ of human life.¹² Noise, according to Sunstein, consists of ‘unwanted variability in judgments’¹³ which is categorised as:

- ‘occasion’ noise, where the same judge is influenced by irrelevant considerations in particular cases;
- ‘level’ noise where judgments in the system, as a whole, exhibit unwanted variability; or
- ‘pattern’ noise where the variability across judges does not take a general form, but where instead there are various patterns of variability among judges.¹⁴

Sunstein, along with Kahneman and Sibony, argue that AI will be immune to noise; it will never be impacted by the hungry judge effect (if it indeed exists),¹⁵ nor will it act according to a human judge’s own bias. AI, the argument goes, avoids the reality that ‘in some cases, noise can be counted as a rights violation, and in general, legal systems all over the world should be making much greater efforts to control noise.’¹⁶

I see this as a dangerous argument. Hyper-focus on congruency and the need for silence misunderstands the role of the judge. Attempts to capture the essence of legal decision-making into a step-by-step algorithm fail to take account that complex questions of law depend on complex instances of human fact. As Lord Sales has elsewhere observed, this results in potential immunity ‘to any appeal to have regard to extenuating circumstances, or to any plea for mercy.’¹⁷ To illustrate this point, Professor Tasioulas aptly draws from the work of the celebrated author, Kazuo Ishiguro. Tasioulas observes that toward the end of ‘Klara and the Sun’ – itself a meditation on what it means to experience humanity – the eponymous narrator, an ‘Artificial Friend,’ speculates that human dignity has its source not in something within us, but in the love of others for us.¹⁸

The central plotline in the novel revolves around Klara and her relationship with ‘her human’, Josie, a chronically unwell child. Klara is employed by Josie’s mother to ‘become’ Josie, so

⁹ Lon Fuller, *The Morality of Law* (Yale University Press 1964) 81.

¹⁰ Eugene Volokh, ‘Chief Justice Robots’ (2019) 68 *Duke Law Journal* 1135, 1161.

¹¹ *ibid.*, 1189.

¹² Cass Sunstein, ‘Governing by Algorithm? No Noise and (Potentially) Less Bias’ (2022) 71 *Duke Law Journal* 1175.

¹³ *ibid.* 1178.

¹⁴ Kahneman, Sibony and Sunstein, *Noise: A Flaw in Human Judgment* (Harper Collins Publishers 2021) 365-367.

¹⁵ Danziger, Levav and Avnaim-Pesso, ‘Extraneous factors in judicial decision’ (2011) 108 *PNAS* 17 6889.

¹⁶ Kahneman, Sibony and Sunstein (n 14) 359-360.

¹⁷ Lord Sales, ‘Algorithms, Artificial Intelligence and the Law’ (The Sir Henry Brooke Lecture, 12 November 2019).

¹⁸ Kazuo Ishiguro, *Klara and the Sun* (2021); John Tasioulas, ‘Artificial Intelligence, Humanistic Ethics’ (2022) 151 *Daedalus* 2 232-240.

that when Josie passes, the mother might still have her daughter. This is a task the Artificial Friend feels unable to complete. At the end of the novel, Klara (having instead cured Josie), observes:

Mr Capaldi believed there was nothing special inside Josie that couldn't be continued. ... But I believe now he was searching in the wrong place. There was something very special, but it wasn't inside Josie. It was inside those who loved her. That's why, I think now, Mr Capaldi was wrong and I wouldn't have succeeded.¹⁹

Thus, for Klara, AI-droids could never fully become human no matter how hard they tried, because at the end of the day, they could never replicate the depth of experience and feeling that humans share for each other. Just as Klara was unable to understand what it truly meant to be human, we should tread carefully, I think, before supplanting human decision-makers with artificial intelligence in the judicial process.

Intermezzo: Will AI be all that bad?

But perhaps I am rushing to my judgment. After all, as we have heard from Chief Justice O'Donnell, AI is already being implemented to varying degrees across judicial systems in the Council of Europe. Indeed, it is fair to say that there has been considerable excitement engendered to the development of this technology. For example, one judicial colleague made quite the stir recently when he described ChatGPT as 'jolly useful'.²⁰ In a speech to the Law Society of England and Wales, Lord Justice Birss spoke of the 'real potential' for AI's future use in the legal sector, explaining that:

I thought I would try it. I asked ChatGPT can you give me a summary of this area of law, and it gave me a paragraph. I know what the answer is because I was about to write a paragraph that said that, but it did it for me and I put it in my judgment. It's there and it's jolly useful.²¹

The Court of Appeal judge assured his audience that he took full responsibility for his use of AI, and that he only included the passage in his judgment because he recognised the answer as being acceptable. While I have every faith that Lord Justice Birss accurately applied the law (indeed, on his account the system passed a personal Turing Test), not everybody will so rigorously complete their digital due diligence. We need only look to two lawyers in America who chanced their arm by asking ChatGPT to draft legal submissions for them.²² In its answer, the chatbot cited non-existent cases that were spotted not by the attorneys, but by the presiding judge. The moral of the story is that the pair received a \$5,000 fine.

From these two examples, I think we can understand the peaks and troughs of the AI argument. Already, the technology has developed to a stage where it can accurately identify core legal principles when asked. But the system is not infallible. On the contrary, AI is not

¹⁹ Kazuo Ishiguro, 'Klara and the Sun' (Faber & Faber 2021).

²⁰ Hibaq Farah, 'Court of Appeal Judge Praises 'Jolly Useful' ChatGPT After Asking it for Legal Summary' *The Guardian* (15 September 2023) <<https://www.theguardian.com/technology/2023/sep/15/court-of-appeal-judge-praises-jolly-useful-chatgpt-after-asking-it-for-legal-summary>> accessed 22 March 2024.

²¹ John Hyde and Bianca Castro, 'Solicitor condemns judges for staying silent on 'woeful' reforms' *Law Gazette*, (14 September 2023) <https://www.lawgazette.co.uk/news/solicitor-condemns-judges-for-staying-silent-on-woeful-reforms/5117228.article> accessed 22 March 2024.

²² Molly Bohannon, 'Judge Fines Two Lawyers for Using Fake Cases from ChatGPT' *Forbes* (22 June 2023) <<https://www.forbes.com/sites/mollybohannon/2023/06/22/judge-fines-two-lawyers-for-using-fake-cases-from-chatgpt/>> accessed 22 March 2024.

yet smart enough to know when it is relying on legal fact or legal fiction. Lawyers rejoice: our jobs are safe for now! There is also the issue of the information base. AI can only draw from what it is fed and that needs to be accurate. But, given the remarkable development of AI technology in the last few years, it is plausible that there will soon come a stage when AI will accurately rely on caselaw to provide solutions to complex human questions. And given that chatbots can already pass a contract law exam,²³ it may be said that this day may come sooner rather than later. How, then, might we implement AI in the justice system?

The first area where AI may prove effective is the digitalisation of dispute resolution. This is an area where mediation and settlement are commonplace and one that has already experienced some move toward online adjudication. Indeed, I am looking forward to Geoffrey Vos MR's talk later in the day, which will discuss the digitalisation of court processes in England and Wales. We are now open to modernising by way of paper light courts (at least in Northern Ireland) if not entirely paperless courts, which is an illustration of quite substantial change. Beyond dispute settlement claims, how else may AI assist in the judicial process? There are some obvious advantages to using AI in summarising disclosure documents. It is not uncommon, as a Court of Appeal judge, to receive thousands of pages in affidavit evidence that parties rely on to advance their arguments. Given the far greater capacity of AI to analyse and synthesise large swathes of information, may we one day reliably expect AI to provide concise summaries and key paragraphs of bundled documents that may be realistically digestible for the human judge?

I imagine Lord Justice Birss would be in favour, at the very least.

Going a step further, there have, perhaps unsurprisingly, been attempts to model judicial decision-making. Aletras *et al* have already built a model that, they claim, will correctly predict decisions of the European Court of Human Rights with an accuracy rate of 79%.²⁴ Such a development represents a step towards justice by AI. Indeed, the findings made in the report make for exciting reading – and there are undoubtedly ways to nurture a technology that can predict outcomes based on facts and procedure in order to develop a system that reliably outputs 'good' law. Given that this technology is still in its infancy, we must watch such developments with interest. However, the above example only tells half the story. John Morison and Adam Harkens at Queen's University Belfast have compellingly argued in their 2019 article titled, 'Re-engineering justice? Robot Judges, Computerized Courts and (semi) automated legal decision-making', that studies such as the Aletras one just mentioned are useful to the extent that the algorithm will deal with issues on appeal.²⁵ This means that the tools were reliant on a number of set givens, already defined by a combination of procedural norms and human input and interaction.

As such, the authors argue that the true reality of the situation is that:

the automation of such judging activities through the implementation of algorithmic tools in precisely this manner, as tools or assistants, rather than newly functioning

²³ Emily Hinkley, 'Westminster Uni Academic uses ChatGPT to Pass Contract Law Exam' *Legal Cheek* (28 March 2023) <<https://www.legalcheek.com/2023/03/westminster-uni-academic-uses-chatgpt-to-pass-contract-law-exa>> accessed 22 March 2024.

²⁴ Nikolaos Aletras, Dimitrios Tsarapatsanis, Dabuek Preotiuc-Pietro and Vasileios Lamos, 'Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing Perspective' (2016) *Peer J Comput. Sci.*

²⁵ Morison and Harkens, 'Re-engineering justice? Robot Judges, Computerized Courts and (semi) automated legal decision-making' (2019) 39 *Legal Studies* 618.

independent systems, recognises that at least for the present moment, fully automated justice as a practical reality is unlikely.²⁶

That being so, it remains that there is clearly potential within the system. Indeed, such programmes, so long as the element of human oversight remains, may in fact be of assistance as a method of triaging cases, in an effort to deal with rising caseloads. However, what remains clear, is that judges will continue to have a role in safeguarding the rule of law. These examples demonstrate to my mind, the possibility that AI might be—in certain instances—a force for good. But in such cases, AI is truly being used to assist the judge; thus, representing what I will call semi-automated justice. Dare we dream further, and suggest that AI might one day write judgments in our stead?

Towards a theory of justice for artificial intelligence

I have already set out in brief the existing attempts at justifying the use of AI in the judicial process. At the moment, the arguments put forward by Volokh, Susskind and others broadly reflect the traditional Austin-style ‘command backed by sanction’ theory of justice.²⁷ This might be a rather crude approximation, but the point stands that AI cannot yet deal with the complex social process inherent in our legal system.²⁸ But suppose it could. Suppose AI could, as Professor Zuckerman has suggested: ‘be better at operating the principles of adjudication advocated by Professor Ronald Dworkin, which involve the integration into legal thinking of widely held moral and political principles.’²⁹

What steps need we take to ensure that such a process may happen? How might we preserve the rule of law in the era of artificial intelligence? First, I would point to the minimum standard of the rule of law. In this regard, I accept that Lord Bingham’s conception of the rule may be too arduous. After all, ‘the more we front-load the rule of law with additional—especially substantive—ideals, the higher the bar it sets, and the bleaker the prospects are of AI tools serving as effective means of realising the rule of law.’³⁰ Where others³¹ have relied on thinner conceptions of the rule of law—such as those expounded by Fuller or Raz³²—to query whether we could truly envisage justice by AI, I would seek to strike a slightly different course. To my mind, we should not shy away from a ‘low bar’ conception of the rule of law when querying the suitability of AI in the judicial process. Rather, I would suggest that if we are to take the possibility of justice by AI seriously, we might begin from the original position.³³ As such, adopting a framework broadly Rawlsian in outlook, I would make the following observations.³⁴

Transparency and publicity

Rawls’ theory of justice aimed to identify principles for the governance of major institutions that can be justified to people despite variation in their beliefs about what a good or perfect

²⁶ *ibid* 21.

²⁷ John Austin, *The Province of Jurisprudence Determined* (London, John Murray 1832).

²⁸ Morison and Harkens (n 25); Hamish Ross, *Law as a Social Institution* (Hart Publishing 2001).

²⁹ Adrian Zuckerman, ‘Artificial Intelligence in the Administration of Justice’ in Higgins (ed) *The Civil Procedure Rules—20th Anniversary Conference, 2019* (Oxford Academic 2020).

³⁰ Tasioulas (n 7) 4.

³¹ In general, see Tasioulas (n 7).

³² Fuller (n 9); Joseph Raz, *The Authority of Law* (Oxford University Press 1979).

³³ John Rawls, *A Theory of Justice* (revised ed. 1999).

³⁴ For a more sustained, thoughtful discussion, see Iason Gabriel, ‘Toward a Theory of Justice for Artificial Intelligence’ (2022) 151(2) *Daedalus* 218.

society would look like. Just as citizens should be able to recognise the rationality of their governance, so too should they be able to recognise the rationality of justice by AI.

As such, I worry about the development of AI-generated judicial decisions. If, as Volokh has contended, what matters is the result and not the process,³⁵ I fear we could lose faith in how judicial decisions are reached. The reality of justice by AI is that, before we can implement justice by algorithm, we need to build an algorithm that can consider the wide variety of scenarios presented in any single case. But we lawyers cannot build that code. Thus, we would need to defer to specialists, who themselves do not have the required legal background to ensure the adequate protection of the principles governing justice. This judicial-technical catch-22 would lead to what has been termed ‘black box justice’—justice arrived at by way of reasons not accessible to the vast majority of individuals.³⁶

Any instance of black box justice clearly offends Bingham’s first principle. Indeed, it contravenes any conception of the rule of law, no matter how thick or thin you rationalise it. I am reminded of the dicta of Lord Diplock in *Attorney General v Leveller Magazine*; that: ‘the way that courts behave cannot be hidden from public ear and eye this provides a safeguard against judicial arbitrariness or idiosyncrasy and maintains the public confidence in the administration of justice.’³⁷ The same must be true in an AI-driven sphere. Let me illustrate this point by way of example. In the case of *State of Wisconsin v Loomis*,³⁸ Mr Loomis was arrested while driving a car involved in a shooting. He was sentenced to six years’ imprisonment. The circuit court referred to a risk assessment tool using a specific algorithm (known as COMPAS) developed by the private company Northpointe, which labelled Mr Loomis as having a high likelihood of recidivism. This was identified as one of several factors relevant to the determination of the length of his sentence. Mr Loomis challenged this decision on the basis that it violated his constitutional right to due process.

While not an example of justice by AI, one can see how similar challenges might be brought with the advent of fully automated, or even partially automated, AI decision-making. In that case, the Wisconsin Supreme Court concluded that the impugned risk assessment tool can be used in sentencing but placed certain limitations and cautions upon its use. In particular, to ameliorate concerns about the lack of transparency, the court introduced a ‘warning label’ to be included in all future pre-sentencing reports containing the COMPAS risk assessments advising the assigned judge of certain deficiencies inherent in the risk assessment.

The Loomis example further shows the influence that the private sector will have in a justice system driven by AI. It was a private organisation that decided Mr Loomis was at a risk of recidivism, not an independent judge. In the normative exercise that is judicial decision-making, why are we to trust that private bodies will accurately reflect the views of the judiciary? As Greenstein has argued:

the transfer of governance to technology, such as witnessed in the Loomis case, brings with it a monopoly in terms of access to the technology. It is essentially only those who govern that have the resources to produce or purchase the technology

³⁵ Volokh (n 10).

³⁶ Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press 2015).

³⁷ [1979] AC 440, at 450.

³⁸ *State v Loomis* 881 NW2d 749 (Wis 2016) 754 (US).

that is used to make decisions about citizens. This continually increasing imbalance is disempowering the governed in favour of those who govern.³⁹

This is a significant hurdle that must be overcome before one can even consider the possibility of AI having any impact on judicial decision-making.

Fundamental rights and bias

To my mind, compliance with the rule of law means compliance with fundamental human rights. Key to this consideration is the requirement that justice must be effectively and impartially administered. Such requirements are recognised by Lord Bingham and are a necessary component to protect against arbitrariness. Thus, where Sunstein might argue that AI reduces the possibility of bias, practice might suggest otherwise. After all, one need only look to how AI has been used in criminal profiling to realise the errors of our past may yet haunt us in our future. In 2016, the investigative journalist organisation ProPublica revealed how a computer programme had scored Brisha Borden, an 18-year-old black woman charged with petty theft, as at a higher risk of re-offending than Vernon Prater, a 46-year-old white man who had previously been convicted of armed robbery.⁴⁰

Sadly, the phenomenon is all too easy to explain. For the criminal justice system in America has historically punished black offenders more harshly than white offenders. Where human adjudicators are capable of critical reflection and are alive to the historical prejudices commonplace in older decisions, AI technologies only see information in binaries. Based on its analysis of decades of human-made decisions, it can crudely discern the formula: if white, then less likely to re-offend; if non-white, then more likely to re-offend. In the case of Borden and Prater, the opposite was true. Two years later, Borden had not been charged with any crimes. Prater had been sentenced to an eight-year prison term for breaking into a warehouse and stealing thousands of dollars' worth of electronics.

As such, we may conclude that if AI is to decide decisions based on centuries of human-made decisions, the binary nature of AI has the capacity to reflect our own failures on a massive scale. On a more abstract level, we must accept that the requirements of justice will not be met by the adoption of processes that do not discriminate against people. Instead, 'a just society will aim to eliminate the impact of a wide range of unchosen features on [citizens'] life prospects.'⁴¹ As such, if AI is to truly apply the rule of law, it must really do so within the limits of Rawls' difference principle. In other words, to ensure justice, AI-generated decisions must, I suggest, work to the greatest advantage of the least-advantaged members of society.

The role of the courts more generally

Finally, I raise the candid observation made originally by Morison and Harkens who argue that it is important to see courts within their context as part of a constitutional landscape.⁴² They point to the seminal UK Supreme Court decision in *Unison*, which demonstrates that actual and effective access to justice, and the procedure that the courts and tribunals provide, is not merely a public service like any other, but a key part of the rule of law and the fabric

³⁹ Stanley Greenstein, 'Preserving the rule of law in the era of artificial intelligence (AI)' (2022) 30 *Artificial Intelligence Law* 291-313.

⁴⁰ Julia Angwin et al, 'Machine Bias' *Pro Publica* (23 May 2016) <<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>> accessed 22 March 2024.

⁴¹ Gabriel (n 34) 225.

⁴² Morison and Harkens (n 25).

of rights.⁴³ If we are to expect AI to play a role in this process, it must do so while respecting the complex social and constitutional role that the court currently plays.

Enacting regulation: a word on the EU's Artificial Intelligence act

From the preceding, the need for robust regulation is patently clear. Síoifra O'Leary, President of the European Court of Human Rights, in the 2022 MacDermott Lecture spoke about the difficulties with regulation of internet and social media. I think that sentiment can read across to the AI debate. Judge O'Leary said this:

The court has sought to grapple with the 'conflicting realities' (a term used in *Delfi v Estonia*) to which the internet and new technologies give rise. It has recognised, on the one hand, that user-generated expressive activity on the internet provides an unprecedented platform for the exercise of freedom of expression. On the other hand, the internet can act as a forum for the speedy dissemination of unlawful forms of speech which may remain persistently online.⁴⁴

We should welcome, therefore, the news that the European Parliament is now considering regulation. The Parliament's draft Artificial Intelligence Act is the first attempt to enact a horizontal regulation for AI. As the Parliament has recognised, 'AI technologies are expected to bring a wide array of economic and societal benefits to a wide range of sectors' but there is a real concern about 'freedom of expression, human dignity, personal data protection and privacy.'⁴⁵

The draft Act has the following proposed objectives. It would: (i) ensure that AI systems placed on the EU market are safe and respect existing EU law; (ii) ensure legal certainty to facilitate investment and innovation in AI; (iii) enhance governance and effective enforcement of EU law on fundamental rights and safety requirements applicable to AI systems, and (iv) facilitate the development of a single market for lawful, safe, and trustworthy AI applications and prevent market fragmentation.

The draft Act begins with a definition, that is largely based on the OECD's own definition of AI.⁴⁶ Article 3(1) of the draft Act therefore proposes that AI should be considered as: '...software that is developed with [specific] techniques and approaches [listed in Annex 1] and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with.'

Annex 1 lays out a list of techniques and approaches that are used today to develop AI. Accordingly, the notion of an 'AI system' would refer to a range of software-based technologies that encompasses 'machine learning', 'logic and knowledge-based' systems, and 'statistical' approaches.⁴⁷ To maximise safe use, the regulation would enact a risk-based approach that lays down different requirements and obligations for the development, placing on the market and use of AI systems in the EU. The approach is encapsulated within a

⁴³ [2017] UKHL 51.

⁴⁴ Síoifra O'Leary, 'Democracy, expression and the law in our digital age', 50th Annual MacDermott Lecture at Queen's University Belfast 2022 (2022) 73 *Northern Ireland Legal Quarterly* 162-183, 165.

⁴⁵ European Parliament, Briefing on Artificial Intelligence Act (June 2023), PE698.792.

⁴⁶ OECD, Recommendation of the Council on Artificial Intelligence OECD/LEGAL/0449 Adopted 22.05.2019; Amended 08.11.2023.

⁴⁷ (n 44).

‘pyramid of risks’, which categorises AI activities from ‘low and minimal risk’ to ‘unacceptable risk’.

Within these categories, the use of AI in law enforcement is deemed as a ‘high risk’ activity. As such, the EU has proposed ex-ante assessment before employing any AI model in law enforcement, requiring all systems to be registered in an EU-wide database. Moreover, articles 8 to 15 of the draft Act sets out that any system would have to comply with further risk management, testing, technical robustness, data training and data governance, transparency, human oversight, and cybersecurity requirements.

As a first step, the EU legislation is certainly encouraging. Of course, it is still in the draft stage, and I note that there are significant steps to be made before any such regulation comes into force. Overall, though, when viewed in line with further calls on the need for regulation—such as those made by the UK recently—the sentiments of Dragos Tudorache, the co-rapporteur of the parliamentary committee steering through the legislation, that now is the time to be more optimistic than pessimistic about AI regulation gain momentum.⁴⁸ Self-evidently, it is only through regulation and establishing effective checks and balances that we might safeguard the rule of law.

Conclusions: where does this take us?

The points I have made form only a drop in the ocean in the AI-judiciary debate. There are many concerns that I, and a great many others, share. However, some healthy scepticism should not overshadow the potential benefits of advancing technologies. What we have spoken about today represents possibly the most profound challenge in the recent history of the legal sphere. We do not know precisely where AI governance will take us, and how our role as judges will change. The discourse on AI also seems to be focused in large part on a cost-benefit analysis – with buzz words of ‘streamlining’, ‘efficiency’ and ‘cost-effectiveness.’ These are important considerations, and commitment to the rule of law does promote efficiency and expediency in justice.

However, we cannot lose sight of the true benefit of human agency in decision-making. For while law is, to some extent, an algorithmic process, it is a very human one too. To my mind, judges know when to be rigid with the law, but they also know how to display compassion, restraint, and reserve. In a debate that is encompassed by preference-based utilitarianism, we would do well to remember that, at the end of the day, judges are asked to deal with ordinary people at their most vulnerable.

In my experience, a little bit of humanity goes a long way.

⁴⁸ Lisa O’Carroll, ‘EU ‘in touching distance’ of world’s first laws regulating artificial intelligence’ *The Guardian* (24 October 2023) <<https://www.theguardian.com/technology/2023/oct/24/eu-touching-distance-world-first-law-regulating-artificial-intelligence-dragos-tudorache>> accessed 26 March 2024.