The Challenges of Artificial Judicial Decision-Making for Liberal Democracy

Christoph K. Winter (ID)

Abstract The application of artificial intelligence (AI) to judicial decision-making has already begun in many jurisdictions around the world. While AI seems to promise greater fairness, access to justice, and legal certainty, issues of discrimination and transparency have emerged and put liberal democratic principles under pressure, most notably in the context of bail decisions. Despite this, there has been no systematic analysis of the risks to liberal democratic values from implementing AI into judicial decision-making. This article sets out to fill this void by identifying and engaging with challenges arising from artificial judicial decision-making, focusing on three pillars of liberal democracy, namely equal treatment of citizens, transparency, and judicial independence. Methodologically, the work takes a comparative perspective between human and artificial decision-making, using the former as a normative benchmark to evaluate the latter.

The chapter first argues that AI that would improve on equal treatment of citizens has already been developed, but not yet adopted. Second, while the lack of transparency in AI decision-making poses severe risks which ought to be addressed, AI can also increase the transparency of options and trade-offs that policy makers face when considering the consequences of artificial judicial decision-making. Such *transparency of options* offers tremendous benefits from a democratic perspective. Third, the overall shift of power from human intuition to advanced AI may threaten judicial independence, and with it the separation of powers. While improvements

C. K. Winter (✉)
Faculty of Law, Instituto Tecnológico Autónomo de México (ITAM), Ciudad de México, Mexico

Department of Psychology, Harvard University, Cambridge, MA, USA

Legal Priorities Project, Cambridge, MA, USA
e-mail: christoph_winter@fas.harvard.edu

regarding discrimination and transparency are available or on the horizon, it remains unclear how judicial independence can be protected, especially with the potential development of advanced artificial judicial intelligence (AAJI). Working out the political and legal infrastructure to reap the fruits of artificial judicial intelligence in a safe and stable manner should become a priority of future research in this area.

The application of artificial intelligence (AI) to judicial decision-making has already started. Significant progress has been made, not only in the United States, most prominently with regard to bail decisions,[1] but also in Russia[2] and Mexico.[3] China has placed more than 100 robots in courts offering legal advice to the public,[4] and Estonia is piloting a program in which small scale civil suits are decided by an algorithm.[5] Furthermore, a recent cross-cultural survey suggests that legal scholars believe that, on average, almost 30% of judicial decision-making will be carried out by AI in only 25 years' time, tripling their estimate that the current role of AI accounts for less than 10% of judicial decision-making.[6] Against this background, it seems plausible to assume that future advances in AI will revolutionize the judicial sector. To many, AI not only promises greater fairness, justice, and legal certainty, but it may finally satisfy the legal requirements of the internationally accepted concept of a *fair trial*, in particular with regards to access to justice, as recognized by, among others, Article 6 of the European Convention on Human Rights (ECHR), Article 10 of the Universal Declaration of Human Rights (UDHR), and Articles 7 and 25 of the African Charter on Human and Peoples' Rights (ACHPR).[7]

At the same time, issues of discrimination and transparency have emerged in the context of bail decisions, putting liberal democratic principles under pressure. Although intense debates about the criminal justice system and demands for what has been coined "algorithmic fairness" have been ushered, further challenges that arise from the shift in authority from human intuitions to artificial intelligence within the judicial system remain unclear and neglected in the discourse. As of now, there

---

[1] Angwin et al. (2016).

[2] Zavyalova (2018).

[3] In Mexico, the Expertius system is advising judges and clerks regarding the determination of whether the plaintiff is or is not eligible for pension. See Carneiro et al. (2015).

[4] World Government Summit (2018). The robot named Xiao Fa has a vaguely humanoid appearance and provides simple legal advice, such as how to bring a lawsuit or retrieve case histories, verdicts, and laws.

[5] Niiler (2019).

[6] Martínez and Winter (2021).

[7] For various rights associated with a fair trial, see also the Sixth Amendment to the United States Constitution.

has been no systematic analysis of the risks to liberal democracy from implementing AI into judicial decision-making. This article therefore sets out to fill this void by identifying and engaging with challenges arising from artificial judicial decision-making.

In this way, this analysis not only offers a first attempt to systematize these challenges, but it also outlines a number of crucial issues which demand further research. Throughout the analysis, I will apply a comparative perspective between human thought processes and artificial intelligence,[8] with an emphasis on ethical issues arising from (Sect. 1) large-scale discrimination and (Sect. 2) a potential lack of transparency within the judiciary. Indeed, I will argue that these issues pose some short- and medium-term threats. However, drawing most notably on recent research by Kleinberg and colleagues concerning bail decisions,[9] I am confident that related technical and ethical issues can be solved to yield outcomes better than those from human decision-making, from a wide range of philosophical perspectives on discrimination. While the lack of transparency in AI decision-making, narrowly construed, does pose risks which ought to be addressed, AI can also increase transparency in another—arguably more important—domain, which I will refer to as *transparency of options*. The far greater long-term threat for liberal democratic values lies in (Sect. 3) the overall shift of power from human intuition to advanced AI and, more precisely, in the possibly accompanying threats to judicial independence and the separation of powers. Accordingly, an AI-based judiciary might significantly contribute to the rise of digital authoritarianism.

Before engaging with these questions, it is necessary to clarify and limit the scope of this analysis. A common framework within the discourse is the distinction between whether AI will *supplement* or *replace* human beings.[10] Will AI help human judges to detect biases or make better decisions, or will it replace human judges altogether? Furthermore, the discussion in law and policy is usually focused on *narrow* applications of so-called AI that use machine learning or simple decision trees to supplement or partially replace human decision-making. I will concentrate on these partial replacements and corresponding issues of discrimination and transparency in the first two sections of this chapter. However, in the final section I will move on to more advanced forms of AI that could replace the vast majority of human judicial decision-making, thereby going beyond present applications, possibilities, and the main focus of literature. Even though this contemplates advanced forms of AI understood as significantly exceeding current capabilities, one should note that this would not require superintelligence that "greatly exceeds the cognitive performance of humans in virtually all domains of interest."[11] Instead, as I will explain

---

[8] Note that this approach, as intuitive as it may sound, is not that common. Indeed, most critics of the application of AI towards the judicial sector focus on the disadvantages of present AI while overlooking the tremendous shortcomings of human decision-making.

[9] Kleinberg et al. (2018, 2019).

[10] See, among others, Sourdin (2018).

[11] See generally, Bostrom (2014); in a similar vein, Bostrom (2006), p. 11.

later, it would require the development of an *advanced artificial judicial intelligence* (AAJI) with more limited domain knowledge and capabilities.

A prevailing argument among researchers and politicians alike seems to indicate that one should focus on present issues rather than trying to solve those which *may* occur in the future.[12] However, research about future developments, such as the implementation of more advanced forms of AI into the judiciary and its entailing risks is crucial because outlining future risks may provide valuable information on what sort of developments ought to be followed with particular care. If, for instance, great risks for liberal democracy are likely to occur in the future as a result of the application of AI, one may try to find ways in the present to mitigate such risks, if one considers liberal democracy to be a desirable political system. To put it simply: it can be beneficial to avoid an encounter with a predator in the first place, rather than trying to run away when it is already in front of you.

1   From Human Bias to Algorithmic Fairness

This analysis will first focus on equal treatment for all citizens which, arguably, is *the* distinctive part of *liberal* democracies.[13] Much of the academic[14] and popular[15] debate on AI in the judiciary has been dominated by issues of illegitimate discrimination based on sex or race. This is understandable, given that Google's face recognition algorithm labeled black people as gorillas,[16] algorithms employed in job selection decisions favored white males,[17] and white people benefitted more than black people from the COMPAS[18] recidivism algorithm,[19] among others. It is clear

---

[12] See, e.g., Reiling (2018): "Don't spend time thinking about things that don't exist yet." Cf. also Sourdin and Cornes (2018), p. 113, who seem to classify such considerations as "unhelpful". While I agree that, for the foreseeable future, "human intelligence may (rather) be supplemented by technological advances", it seems almost naïve to consider only the more plausible scenarios and reject considerations of how one should act in case less likely scenarios unfold. Instead, such concerns and related research seem well justified not only from an expected value perspective, but from any reasonable version of the precautionary principle. Cf. also the emerging argument on bridging near- and long-term concerns about AI: Baum (2018, 2020); Cave and Ó hÉigeartaigh (2019); Prunkl and Whittlestone (2020).

[13] Mukand and Rodrik (2020).

[14] See, among others, Chander (2017); Chen (2019); Hacker (2018); Kleinberg et al. (2019); Sourdin (2018), pp. 1128–1129.

[15] See, e.g., Levin (2016) and Smith (2016).

[16] Zhang (2015).

[17] See, e.g., the well-documented case regarding application for a medical school in the UK: Lowry and Macpherson (1988).

[18] COMPAS stands for "Correctional Offender Management Profiling for Alternative Sanctions" and was developed by Northpointe (now Equivant). The COMPAS recidivism algorithm is used by U.S. courts in many states to assess the likelihood of a defendant becoming a recidivist.

[19] See Larson et al. (2016).

that improvements are needed. Yet, it is a separate question whether these issues justify the abolishment or omission of incorporation of AI into the judicial sector. After all, the much longer history of human intelligence (HI) in the judiciary seems to indicate that human decision-making has not done justice to all groups either. While it may be possible to mitigate discrimination resulting from implicit human biases to some degree,[20] we need to identify how feasible it is for AI to overcome such biases compared to HI. If there is a sufficiently good chance that AI is even better positioned to do so in the long term, one should not reject entirely the idea of artificial judicial decision-making due to discrimination, but rather concentrate on how to mitigate discrimination by AI. For AI to play an important role in the judiciary of the future, it does not have to be perfect, but only *better* than HI. Indeed, we may even be obligated to adopt it in such cases.

Pertaining to discrimination, the law works remarkably similar across jurisdictions in distinguishing between *direct* and *indirect* discrimination, also understood as disparate *treatment* and disparate *impact* respectively.[21] Direct discrimination is typically understood as treating a person less favorably than another person in a comparable situation on the basis of any protected ground, such as sex, racial or ethnic origin, religion, disability, age, or sexual orientation.[22] Hence, the disadvantageous treatment is based on the possession of specific characteristics. Here, cases turn on the analysis of the causal link between the protected ground and the less favorable treatment,[23] as well as the comparability of persons in similar situations. In some jurisdictions, direct discrimination can be justified;[24] in others it cannot.[25]

---

[20] See, e.g., Carnes et al. (2015), Lai et al. (2014) and Devine et al. (2012).

[21] See, e.g., Article 2 (2) of the EU Racial Equality Directive; ECtHR, Biao v. Denmark [GC], No. 38590/10, 24 May 2016, para. 89–90. For the US approach, see definitions provided in McGinley (2011), p. 626.

[22] See, e.g., the very similar definitions offered under EU Law by Article 2 (2) of the EU Racial Equality Directive, by the case law of the ECtHR in ECtHR, Biao v. Denmark [GC], No. 38590/10, 24 May 2016, para. 89; and ECtHR, Carson and Others v. the United Kingdom [GC], No. 42184/05, 16 March 2010; and the US approach in Bolling v. Sharpe, 347 U.S. 497, 499, 74 S. Ct. 693, 694, 98 L. Ed. 884 (1954); Brown v. Board of Education, 347 U.S. 483 (1954); Washington v. Davis, 426 U.S. 229, 239, 96 S. Ct. 2040, 2047, 48 L. Ed. 2d 597 (1976). Some jurisdictions favor a non-exhaustive list of these grounds which means that further grounds can be added if deemed necessary. See, e.g., German Basic Law Article 3.

[23] To establish such a causal link, one must answer the following question: "Would the person have been treated less favorably had they been of a different sex, race, age, or in any converse position under any one of the other protected grounds?" If the answer is "yes", then the less favorable treatment is caused by the grounds in question.

[24] See, e.g., for the United Kingdom, Professor John Pitcher v. Chancellor, Masters and Scholars of the University of Oxford and Saint John the Baptist College in the University of Oxford [2019] UK Employment Tribunals 3323858/2016. Even for the most protected categories in the US like race and religion, direct discrimination is justifiable if it passes strict scrutiny. Note, however, that this is rarely the case. See Fisher v. Univ. of Texas at Austin, 570 U.S. 297, 310, 133 S. Ct. 2411, 2419, 186 L. Ed. 2d 474 (2013).

[25] For example, in the UK, direct discrimination can only be justified regarding age and disability. See United Kingdom, Equality Act 2010, Part 2 Chapter 2 Sections 13 and 19.

Indirect discrimination takes place when an apparently neutral criterion, provision, or practice would put persons protected by some general prohibition of discrimination (e.g., those having a protected characteristic) at a disadvantage compared to others.[26] If such discrimination occurs, it will have to be justified, that is, the provision in question will be upheld only if it has a legitimate aim, and the means of achieving that aim is necessary and appropriate. In contrast to direct discrimination, indirect discrimination is based on apparently neutral criteria which are not formally prohibited. Nevertheless, the consequence of both direct and indirect discrimination is essentially the same: an individual who belongs to a protected group is disadvantaged.

Despite the fact that the application of these definitions can be tricky depending on the exact facts of the case and the evidence available, they work rather well in the context of algorithms, as they explicitly state what criteria are directly taken into account and which ones are not. For example, in the context of bail decisions, algorithms which do not directly take race, sex, religion, or other prohibited factors into account for assessing the flight risk of an individual do not discriminate directly.[27] However, as mentioned earlier, the COMPAS recidivism algorithm had worse effects for people of color in comparison to white people. Thus, the algorithm may indirectly discriminate by taking into account apparently neutral criteria such as the zip code, level of education, or a poor credit rating when predicting the flight risk. Since the abovementioned factors have often been influenced by previous discrimination,[28] the common worry that AI may *perpetuate* discrimination, even if the algorithm does not explicitly take race or sex into account, is justified.

However, the COMPAS case, which has strongly shaped the public discourse on AI in the judiciary, has limited value for informing long-term policy, as the exact algorithm is undisclosed and kept as a trade secret.[29] While the case shows that it is certainly possible to build algorithms that indirectly discriminate against protected groups (and perhaps discrimination is even likely, absent safeguards in development and application), our primary concern is whether it is possible to build an algorithm or AI that is beneficial from different philosophical points of view in comparison to human decision-making, even with respect to issues of discrimination. In this vein, recent research by Kleinberg and colleagues shows that such improvements are not

---

[26]United Kingdom, Equality Act 2010. Cf. also the case law of the ECtHR: ECtHR, D.H. and Others v. the Czech Republic [GC] (No. 57325/00), 13 November 2007, para. 184; ECtHR, Opuz v. Turkey (No. 33401/02), 9 June 2009, para. 183; ECtHR, Zarb Adami v. Malta (No. 17209/02), 20 June 2006, para. 80.

[27]If not the sole determent for accessing whether pretrial release is to be allowed, the flight risk is in many jurisdictions a crucial factor. See, e.g., State of New Hampshire v. Christina A. Hill (2019) Supreme State Court NH 2018-0637. In other states in the US, the likelihood that the defendant will be arrested for or convicted of a crime also matters, cf. Dabney et al. (2017), p. 408; Karnow (2008), p. 1.

[28]See discussions of structural racism, among others, Wallace et al. (2017) and Hammer (2018).

[29]Cf. Chohlas-Wood (2020). I will discuss the apparent lack of transparency in more detail in Sect. 2.

merely a utopian vision, but feasible to build today.[30] Again focusing on bail decisions, the machine learning algorithm developed by Kleinberg and colleagues, which used gradient-boosted decision trees, was trained with a large dataset of 758,027 defendants who were arrested in New York City between 2008 and 2013. The dataset included the defendant's prior rap sheet, the current offense, and other factors available to judges for making the decision. When tested on over 100,000 different cases, the algorithm proved to be significantly better than human judges at predicting whether defendants would fail to appear or be re-arrested after release, which—depending on one's policy preferences—can come with diverging yet strong benefits. More precisely, simulations showed failure to appear and re-arrest ("crime") reductions of up to 24.7% and no less than 14.4% with no change in jailing rates, or jailing rate reductions up to 41.9% and no less than 18.5% without any increase in crime rates.[31] Crucially from the perspective of preventing discrimination, the algorithm was able to achieve crime reductions while simultaneously *reducing* racial disparities in all crime categories.[32] The adoption of the algorithm would thus put policymakers in the rather comfortable position of choosing between the options of releasing thousands of people pre-trial without adding to the crime rate *or* preventing thousands of crimes without jailing even one additional person— whilst reducing racial inequalities. Needless to say, these are not the only options, but they illustrate what kind of trade-offs policymakers will have to make with regards to the balancing of crime and detention rates, as well as the disproportionate imprisonment of minorities, particularly black and Hispanic males in the United States, if AI replaces HI. In fact, Sunstein argues that the algorithm developed by Kleinberg and colleagues "does much better than real-world judges (. . .) *along every dimension that matters*".[33]

Importantly, these results are not unique to New York City, as Kleinberg and colleagues were also able to obtain qualitatively similar findings in a national dataset.[34] However, one might be tempted to argue that it remains uncertain whether these results can be achieved in jurisdictions outside the US. Could it be that there is just something off with judicial decision-making in the US? Might there be a bug in the system which does not exist in other jurisdictions when it comes to bail decisions? After all, common and civil law systems vary,[35] and legal education in the United States is fundamentally different from most places, even other common law systems such as Australia, the United Kingdom, or India. Despite this, I fail to see why such differences would lead to significantly different outcomes because,

---

[30] Kleinberg et al. (2018).

[31] Kleinberg et al. (2018), p. 241.

[32] Kleinberg et al. (2018), pp. 237–238.

[33] Sunstein (2019), p. 2.

[34] Kleinberg et al. (2018), p. 241.

[35] Even though I maintain that these differences are still being very much exaggerated; for an insightful analysis on this topic, see Pejovic (2001).

first of all, the law works remarkably similarly with regards to bail decisions.[36] Second, Sunstein points out that the analysis of the data suggests that cognitive biases can explain why AI so clearly outperforms HI.[37] This is a relevant and important insight given that biases occur worldwide irrespective of the jurisdiction.[38] To be precise, Sunstein argues that judges make two fundamental mistakes that significantly influence their overall performance.[39] They treat high-risk defendants as if they were low-risk when their current charge is relatively minor, and they treat low-risk people as if they were high-risk when their current charge is especially serious.[40] Thus, in each of these cases, judges seem to assign too much value to the current offense in comparison to other relevant factors including the defendant's prior criminal record, age, and employment history.[41] Sunstein refers to this phenomenon as the *current offense bias*.[42]

He then links the current offense bias to the well-known *availability bias*, i.e., the tendency to overestimate the likelihood of an event occurring in the future if examples can easily be brought to mind.[43] Even though one may be a bit skeptical as to the degree of relatedness between the current offense bias and the availability bias,[44] the overall point stands: judges in different jurisdictions are likely to make the same mistakes because there is no reason to assume that Asian, African, or European judges will not suffer from the current offense biases. This is especially so because the related availability bias seems to be a general trait of the mind trying to access the probability of an event occurring based on associative distance.[45]

We have now seen that crime and detention rates can be significantly decreased without perpetuating or even increasing racial discrimination. This itself could be viewed as a welcome development, as the absolute numbers of minorities in jail would decrease. But does that mean that, if we adopt Kleinberg and colleagues' algorithm, the percentage of African Americans and other minorities in prison will remain the same? If we assume that human decision-making will eventually decrease discrimination, and thereby reduce not only the absolute but also the relative number of African Americans in prison, the question arises whether algorithms would lag behind HI after some time has passed. It seems plausible to assume that both absolute

---

[36]Baughman (2017), p. 15.

[37]Sunstein (2019), p. 501.

[38]Of course, not all identified biases occur globally in the same way. However, as we shall see, chances are that the bias in question will occur not only in the US. See also Dhami and Ayton (2001) showing that human adjudicators in the UK follow simple heuristics.

[39]Sunstein (2019), p. 502.

[40]*Ib.*

[41]*Ib.*

[42]*Ib.*

[43]Kahneman and Tversky (1982).

[44]Sunstein (2019), p. 502 views them as "close cousins".

[45]Tversky and Kahneman (1973). This having been said, the cross-jurisdictional application of the current offense bias is ultimately an empirical issue that requires empirical confirmation.

and relative detention rates for minorities will decrease in the future, if a human-centered judiciary follows humanity's track record on moral progress.[46] However, it will take time to develop the necessary and crucial cognitive features to do so more fully, such as through enhancement of the abilities to override implicit biases.[47] At the same time, it is possible to instruct an algorithm in a way that it produces morally or socially desirable outcomes, whatever these may be. For instance, in one scenario, Kleinberg and colleagues instructed the algorithm to maintain the same detention rate while equalizing the release rate for all races. Given this, the algorithm was still able to reduce the failure to appear and re-arrest ("crime") rate by 23%.[48] In another simulation, the algorithm was instructed to produce the same crime rate that judges currently achieve. Under these conditions, a staggering 40.8% fewer African Americans and 44.6% fewer Hispanics would have been jailed.[49]

We might wonder whether these results really indicate that human decision-making *ought* to be *supplemented* or even (partially) *replaced* by AI from the perspective of discrimination. If it is all about cognitive biases—whether they are related to implicit racial biases or about current offense bias—maybe we can try to improve human decision-making before we get rid of it entirely. Indeed, it is plausible to mitigate some biases to some degree. For instance, the strength of the availability bias can even be reduced by simply being aware of it.[50] Additionally, I have previously suggested different institutional and procedural changes as well as mandatory training in behavioral economics and cognitive biases for members of the judiciary in order to mitigate biases.[51] However, it does not seem feasible to eliminate them altogether this way.[52] To put it simply, in order to completely get rid of cognitive biases at this day and age, one must get rid of human decision-making.[53]

---

[46] Pinker (2011, 2018).

[47] I leave aside other notable problems these approaches might bring about, such as concerns related to privacy and freedom of thought. I also leave aside the question to what degree discriminatory outcomes within the US criminal justice system are currently influenced by implicit racial bias, or whether other factors, such as education, poverty, and higher police presence in African American communities are the main drivers (while acknowledging that factors such as these are in turn influenced by structural discrimination). For a critical note on the role of implicit bias, see, for instance, Oswald et al. (2013).

[48] Kleinberg et al. (2018).

[49] Kleinberg et al. (2018).

[50] Gigerenzer (1991). For different alleviation strategies regarding the hindsight bias in negligence ascriptions, *see* Kneer and Skoczeń (forthcoming). For more on countering implicit biases, see Teal et al. (2012), Boscardin (2015), Ingriselli (2015) and Shaked-Schroer (2008).

[51] Winter (2020).

[52] *Ib.*

[53] *Ib.* Note, however, that this does not imply that an AI-based judiciary does not face *any* of such risks. Of course, the design of the AI itself as well as the selection of training data depend once again on human decision-making. This said, it may be easier to avoid biases within said design and selection process than to debias a human-centered judiciary, as indicated by the results of Kleinberg and colleagues' study.

To summarize, Kleinberg and colleagues' findings enormously strengthen the case for AI in the judiciary from the perspective of decreasing discrimination. This is especially so because it is possible to reduce crime and detention rates at the same time. Although there is much more to say about human decision-making, cognitive biases, and the role of empathy therein,[54] contrary to the present popular belief, AI *can* be a "force for (. . .) equity."[55] While it is important to caution that Kleinberg and colleagues' algorithm was merely concerned with bail decisions, it seems reasonable to assume that it would generally be easier to prove whether direct discrimination occurred under algorithmic than human judicial decision-making. In fact, if the decision procedure is transparent, it would be so obvious that it is unlikely to occur in the first place. The preferred trade-offs among crime, detention, and discrimination would still have to be made, but regardless of how one decides, the result would be an improvement from HI. Having said this, it is important to note—as experienced in the COMPAS case—that the exact trade-offs made may not always be visible. So, should one reject the implementation of AI, even if it can reduce discrimination, for lacking transparency?

2   From Transparency of Procedures to Transparency of Options

The relationship between democracy and transparency is highly contested. Contrary to widespread intuition, it may even be the case that authoritarianism leads to more transparency than democracy. This may be so because greater vulnerability to public disapproval could make democratically elected officials more inclined to promote opacity or withhold information compared to their autocratic counterparts, who have to worry less about public perception.[56] The fact that democratic governments have incentives to obfuscate their policies can be illustrated by the recent remarks of Germany's Minister for the Interior Horst Seehofer, who recently stated that "you have to make laws complicated".[57] Despite the fact that the statement itself was arguably not well thought-through from a political standpoint, and Seehofer consequently shortly thereafter tried to argue that he was being "slightly ironic", the statement does get at the heart of the tension between democracy and transparency. If passing a specific law is *better* for citizens than any of the alternatives, from the

---

[54] Empathy is often raised as an argument against AI. But note that it may not always be the compassionate, nice concept we often take it to be and may even be responsible for much of discrimination. See generally Bloom (2016). For the role of empathy in judicial decision-making, see Bandes (2009), Chin (2012), Colby (2012), Weinberg and Nielsen (2012), Lee (2014), Glynn and Sen (2015) and Negowetti (2015).

[55] Kleinberg et al. (2018), p. 241.

[56] Hollyer et al. (2011).

[57] For more information regarding the context of his remarks, see Das Gupta and Fried (2019).

perspective of elected officials, but the law in question is very unpopular, should it still be adopted? And if so, is it permissible to communicate it in a way that makes re-election more probable, despite the fact that citizens disagree with the measure being taken, even if they ultimately benefit from the new law?

The argument that democracy may *in fact* not always be the ideal engine for transparency does not entail that transparency is not a crucial factor thereof from a *normative* perspective. Of course, transparency broadly understood helps to fight corruption, promote trust in public institutions, and contribute to the public discourse.[58] Yet, in order to answer more specific questions regarding the state of transparency when it comes to artificial judicial decision-making from a democratic perspective, one first needs to consider the function of transparency within a democracy. Interestingly, despite the visibility of transparency as a concept in the public discourse, surprisingly little attention has been paid to its underlying purpose.[59] Gupta even calls it an "overused but under-analyzed concept".[60] The basic argument for transparency being an essential part of *any* democratic system runs along the following lines:[61]

1. Democracy requires citizens to take an active part in politics.
2. For citizens to be able to participate in politics, they need to be able to make political judgments based on relevant information.
3. In order to make political judgments based on relevant information, citizens need to have access to that relevant information.
4. In order for citizens to have access to relevant information, such information must be transparent.

The central question in this framework is what exactly counts as *relevant* information. Do citizens need to know *who* made the decision in question and *how*, or do they also need to be aware of alternative routes which could have been taken? What information should be available to citizens when it comes to decisions by the judiciary, rather than the executive or legislature?

If, as mentioned above, democracy requires citizens to take an active part in politics, then relevant information could be any information helpful for making

---

[58] Note, however, that some research points in the direction that high levels of transparency may also have negative consequences. See, e.g., Fox (2007), Licht (2011, 2014) and Moore (2018). The scope of this paper forces me to solely focus on the liberal democratic perspective of transparency, and I will leave aside other theories, perspectives, and approaches to the matter which may or may not shift one's opinion, particularly on the desired degree of transparency and noteworthy exceptions.

[59] Moore (2018).

[60] Gupta (2008).

[61] See, among others, Dahl (1971), who argues that any conception of democracy requires the free flow of information in order to make informed choices, compared to a minimalist approach to democracy as taken by Schumpeter (1942) or, more recently, Przeworski (2000), who define democracy as a regime in which the executive and the legislature are both filled by "contested elections".

political choices such as who or what to vote for, what to protest, and when and how to engage in public discourse more generally. As Bellver and Kaufmann argue, the information provided needs to account for the performance of public institutions because transparency is a tool to facilitate the evaluation of public institutions.[62] More importantly, not only the decisions made by directly or indirectly elected individuals, but state actions more generally, including decisions and judgments made by the judiciary,[63] need to be transparent in order for citizens to be sufficiently informed.

In order to evaluate public institutions carefully, citizens ideally need to know (or at least be able to know) not only the outcome of the decision itself, the people involved, and their respective roles in the process but—importantly—also alternative options which could have been taken. If, for instance, a decision is not very popular, but the alternatives are significantly worse, then ideally this would be communicated, just as it should be transparent if a decision was made which intuitively sounds good, but which arguably had much better alternatives from the perspective of citizens. Accordingly, in order to guarantee that citizens receive relevant information to evaluate public institutions, they should be able to learn what kind of decision was made, the procedures leading to it, and what other options were available.[64] To put it simply, *transparency about options* may be just as important as transparency about procedures and actors from the liberal democratic point of view.

Since it would be an improvement from a liberal democratic point of view if options were more frequently transparent, some AI—as illustrated by the clear trade-offs Kleinberg and colleagues' recidivism algorithm produces—has strong advantages over HI in this regard. While nowadays citizens and policymakers alike have to trust their (often unreliable) intuitions to assess the situation, the clarity of the trade-offs, for example between public safety, discrimination, and detention rates, as in the case of bail decisions, will become clearer with the implementation of AI. For instance, this means that citizens will not only be able to see the racial composition of the population denied bail, but will also know the consequences of lowering or increasing this rate. In sum, citizens will have better access to highly relevant information in order to participate in the political process.[65] Indeed, for better or worse, it will be much more difficult for politicians like Seehofer to use complicated laws to obscure the political decision and evaluation process. This is the case even if programs do not use explicitly programmed decision-making (e.g., a simple decision tree or a more complex algorithm), but instead use machine learning, including deep

---

[62] Kaufmann and Bellver (2005), p. 5; cf. also Florini (1999), p. 5.

[63] As noted by Liptak (2008), the United States is the only country that elects a *significant* portion of its judges directly.

[64] Cf. also Licht and Licht (2020), who distinguish among (a) transparency that informs about final decisions or policies, (b) transparency with regards to the process resulting in the decisions, and (c) transparency about the reasons on which the decision is based.

[65] Sunstein (2019), p. 7 even argues that the clarity of tradeoffs "may be the most important point" with regards to the recidivism algorithm.

learning, due to the fact that the relevant trade-offs can still be transparent by evaluating actual and simulated impact, regardless of the complexity of machine learning.

Yet most commentators in the debate on AI and transparency take a different stance and focus. Instead of concentrating on the new levels of transparency regarding the tradeoff being made, they rightly point out that the algorithm itself is often unknown.[66] Case in point, the COMPAS risk assessment algorithm, which is now used for bail and sentencing decisions in more than 20 jurisdictions within the United States, is a protected trade secret and thus remains a black box.[67] This is already problematic because without knowing the algorithm or machine learning training set, one cannot understand and challenge the decision, in case the design is flawed. From a liberal democratic point of view, this is even more troublesome. If citizens are not able to obtain vital information to evaluate the operations of the judiciary and related procedural norms, they cannot participate in the democratic process in a meaningful way, aside from pointing out that the algorithm should be made public. It may not be unethical *per se* to outsource typical tasks of the judiciary to private companies, yet it becomes highly questionable if this process leads to the lack of transparency witnessed in the COMPAS case.

Hence, while AI has strong potential benefits with regards to the clarity of tradeoffs, the current lack of transparency (of procedures) in many cases and the potential lack thereof in the future has led scholars to become increasingly skeptical.[68] However, in order for it to be sufficient reason to reject the implementation of AI, one would have to show that, first, HI is in fact better suited to tackle these issues and, second, the lack of this kind of transparency is more harmful from the liberal democratic perspective than the lack of transparency regarding the above-mentioned trade-offs. I will consider these in turn.

It may be argued that human judicial decision-making is much more transparent than artificial judicial decision-making because citizens can identify human judges. They can read their arguments and decide for themselves whether the judgment was reasonable or not. In short: human judges *explain* why they have come to the decision they came to. Although one can find plenty of ongoing research with regards to the explainability of AI, present applications certainly lag behind human capabilities to explain.[69] However, there are notable problems with this line of reasoning.

---

[66] In this vein, among others, Floridi et al. (2018), O'Neil (2016) and Wachter et al. (2017).

[67] Piovesan and Ntiri (2018).

[68] For instance, the number one recommendation of AI Now Institute's annual report states that "core public agencies, such as those responsible for criminal justice, healthcare, welfare, and education (i.e., "high stakes" domains) should no longer use 'black box' AI and algorithmic systems." AI Now (2017), p. 1. The report emphasizes that 'black box' systems are especially vulnerable to skewed training data and the replication of human biases due to a lack of transparency (p. 15). Cf. also Završnik (2020).

[69] Cf. Waltl and Vogl (2018) and Barredo Arrieta et al. (2020). The technical challenges of explainable AI (XAI) are also referred to as the *interpretability problem*.

First, giving human explanations does not equal full transparency of procedures. Explaining judgments is important and can provide a basis for appeal, but an explanation does not necessarily reveal the true underlying cognitive processes. The processes of human (judicial) decision-making are just as unknown as the proceedings of some applications of AI. Following a computational theory of cognition,[70] one might even argue that we can at least know the algorithm or data set with regards to artificial decision-making, even if we may not understand its operations, whereas the algorithm running human decision-making remains unknown. It is crucial to note that while a computational perspective of cognition certainly strengthens this argument, it does not rest upon it. Whatever philosophy of mind one might prefer, human decision-making remains a black box as well.

Second, human explanations are not only prone to error and cognitive biases, but also at risk of rationalizing underlying motivations and preferences.[71] Intuitively, one might expect that this general tendency might be mitigated by expert knowledge and education, yet as studies have shown over and over again, experts rarely perform better than lay-people.[72] Accordingly, the worry is that even though one might know and understand the official reasons given, they might be misleading and the risk remains that the ultimately decisive reasons are still unknown. This is not only important for democratic evaluations but might also decrease the chances of a potential successful appeal.

Third, depending on the complexity of the individual judgment, one might question whether laypeople can fully understand judgments anyway. Surely, one might counter-argue that lawyers will be able to understand the judgment and advise their client accordingly. However, from a liberal democratic perspective, there is not much benefit in lawyers navigating cases if the vast majority of citizens are unable to comprehend and—based on this—evaluate judicial norms and behavior correctly. Arguably, overconfidence may even lead laypeople to think they understand the judgment and give them a false sense of understanding.

These theoretical and practical concerns question the assumption that HI really does provide more liberal-democratically relevant transparency than AI. However, let us assume that this analysis is either mistaken or that there are other overriding arguments which speak in favor of HI-transparency in this regard.[73] Consequently, one would still have to investigate whether the lack of human explanations for

---

[70] Piccinini (2016) and Chalmers (2011).

[71] Cushman (2020) with further references.

[72] See the comprehensive list covering research on military leaders, engineers, accountants, doctors, real estate appraisers, option traders, psychologists, and lawyers presented in Guthrie et al. (2001); see also Meadow and Sunstein (2001). See generally Kahneman and Tversky (1983). Research specifically focused on judicial decision-making can be found in Englich et al. (2006), Wistrich et al. (2015), Rachlinski and Wistrich (2017), Wistrich and Rachlinski (2018), Struchiner et al. (2020) and Winter (2020).

[73] For instance, one might argue that AI reasoning is sufficiently alien to human cognition that humans will always (think that they) understand the thought processes of another HI better than the operations of an AI.

individual judgments is more harmful than the lack of transparency regarding the above-mentioned trade-offs. Should humans have the *possibility* to understand judgments in individual cases, or should they rather be able to clearly see the necessary trade-offs behind rules, which not only shape society but also the very individual judgments they care to understand? Should they understand the purpose of the rules, its underlying trade-offs and the explanations given by the ones they voted for (legislature) or the application of them made by often not democratically accountable judges in selected scenarios? Even though the advantages of having access to human judicial explanations ought not to be downplayed, I am afraid that the transparency of the trade-offs in question might simply be even more important.[74]

3   From Separation of Powers to Judicial Dependence

While improvements with regards to discrimination (Sect. 1) and transparency (Sect. 2) seem possible, and related issues have been receiving the appropriate academic attention in the past few years, medium- and long-term threats arising from the application of advanced AI within the judiciary for liberal democracy have been almost completely neglected.[75] Hence, this part of the analysis aims to raise awareness of these risks. More precisely, I will make the point that advanced artificial judicial intelligence may threaten judicial independence and the separation of powers more generally. Before I proceed with the analysis by outlining why the development of artificial judicial intelligence ought to be followed closely from the point of view of liberal democracy, it is necessary to provide a brief explanation for this seemingly science-fiction scenario.

3.1   The Possibility of an Advanced Artificial Judicial Intelligence

Advanced artificial judicial intelligence (AAJI) can be defined as an artificially intelligent system that matches or surpasses human decision-making in all domains relevant to judicial decision-making. Crucially, this does not require the development of an artificial general intelligence (AGI) while still avoiding the need for hybrid human-AI judicial systems. Such a system goes significantly beyond the current state of the art; with it, humans would have the ability to outsource

---

[74]Assuming that human judicial explanations are more important than transparent trade-offs, one would still have to show that this also justifies higher incarceration rates—both generally and especially regarding minorities.

[75]One exception which will be discussed below is the work by Michaels (2019).

decision-making within the judicial sector entirely if they wished to do so. Although such a development is often considered by legal researchers to be extremely unlikely in the medium- or even long-term future,[76] machine learning researchers think that even the development of an AGI, i.e., an artificially intelligent system that matches or surpasses human decision-making in *all* relevant domains,[77] is much closer than the common sense among jurists seems to indicate.[78] For instance, experts believe there is a 50% chance of AI outperforming humans in all tasks in 45 years.[79] With regards to specific activities, they predict that AI will outperform humans fairly soon. This includes translating languages (by 2024), writing high-school essays (by 2026), driving a truck (by 2027), working in retail (by 2031), writing a bestselling book (by 2049), and working as a surgeon (by 2053).[80]

Surely, one can debate whether it is even possible for an AI to engage in legal reasoning. Crootof, for instance, points out that "the judgment we value in a common law process is a distinctively human skill."[81] In this regard, she further argues that "[g]iven their sensitivity to changing social norms, human judges are uniquely able to oversee legal evolution and ensure that our judicial system 'keep [s] pace with the times.'" However, if one is primarily concerned about detecting (changing) social norms, AI may arguably be better equipped to do so than HI. While AI can be trained with vast data sets from diverse sources and be more easily and broadly adapted with changing norms, HI relies on personal interpretations of comparatively little and highly selected information from media and interactions with again a very selected group of people. Hence, AI not only has the power to process a greater amount of relevant data, but it is also less likely to fall for the ubiquitous confirmation bias and other misleading selection mechanisms that lead one to paint an inaccurate picture of existing social norms.[82]

This having been said, following a Dworkinian interpretation of law, one could certainly question whether even an advanced AI would be up for the task of "moral reasoning" as part of the judicial decision-making process.[83] At the same time,

---

[76] See Martínez and Winter (2021).

[77] On the notion of AGI, see Wang and Goertzel (2007). Other terminologies being used to describe the same or a very similar technologies are "strong AI", "human-level AI", "true synthetic intelligence", and "general intelligent system".

[78] For expert surveys on AGI timelines, see Baum et al. (2011), Müller and Bostrom (2016), Grace et al. (2018) and Gruetzemacher et al. (2019).

[79] Grace et al. (2018).

[80] Grace et al. (2018).

[81] Crootof (2019), p. 238, citing Kaminski (2019).

[82] Again, assuming that it is easier to avoid biases within the AI design and data selection process compared to debiasing a human-centered judiciary. See also supra note 53 and accompanying text.

[83] Note, however, that even this is far from clear. For instance, one might argue that Dworkin's (2011) description of moral reasoning as "the interpretation of moral concepts" (p. 102) would not necessarily exclude advanced AI, which would presumably be as or even more capable than humans at the "integration of background values and concrete interpretive insights" (p. 135). Curiously, Martínez and Winter (2021) in a global survey of law professors found no significant

chances that the judicial system can be outsourced would arguably increase if one would follow Oliver Wendell Holmes' prediction theory of law.[84] Holmes famously stated that "[t]he prophecies of what the courts will do in fact, and nothing more pretentious, are what I mean by the law."[85] The purpose of this section is not to convince the reader of the merits of a specific legal theory, but rather to emphasize the high degree of jurisprudential uncertainty.[86] If we can agree that AI would be able to take over the judiciary from *some* legal theoretical perspectives—even if one personally does not share those views or considers them unlikely to be correct—we have strong reason to start thinking about the accompanying consequences and potentially relevant safeguards. In other words, the mere possibility of an AAJI should lead us to take its implications seriously.

3.2   Advanced AI & Judicial Independence

Even though one might argue that a liberal democratic system may in theory be achieved and sustained without the separation of powers, this separation of powers has come to be a cornerstone of any liberal democratic system around the world. The separation of the judicial, legislative, and executive branches of government serves a vital purpose in minimizing the concentration of power and maintaining checks and balances across the government. While checks and balances differ in their configuration across jurisdictions, any significant influence on a given power structure (e.g., branch of government) will likely destabilize the system. Relatedly, this destabilization can take very different forms and outcomes in different systems. For instance, parliamentary democracies in Western Europe may respond very differently to the implementation of an AAJI in comparison to the presidential systems in the Americas. Additionally, those jurisdictions and cultures which consider their judicial branch to have some legislative-like lawmaking powers, such as the European

---

effect of law professors' views regarding the role of moral reasoning in judicial decision-making on their views on how much of judicial decision-making should be carried out by AI. Furthermore, they found that legal academics who endorse positivism favored only marginally higher ratings for the percentage of judicial decision-making that should be carried out by AI than those who endorsed natural law.

[84] *See* Holmes (1897), pp. 457, 461.

[85] Holmes (1897), p. 461.

[86] Discussions on "jurisprudential uncertainty" understood as "normative uncertainty with respect to legal theory" (Winter et al. 2021, p. 97) have only just started. See Winter et al. (2021) and Winter (2022). The related debate in ethics has been getting more attention recently. See, among others, Gustafsson and Torpman (2014), Lockhart (2000), MacAskill (2014), MacAskill et al. (2020), Tarsney (2018); see also Barry and Tomlin (2019) who apply moral uncertainty to different issues in criminal law theory, including sentencing and criminalization theories.

Union[87] and the United States,[88] may evaluate the threats and opportunities imposed by AAJI very differently than those who consider their judiciary as primarily politically neutral (even if this may in fact not be the case, as legal realism tells us), such as the United Kingdom[89] or Germany.[90] This fact can be illustrated by the intensity of discussions taking place during the selection procedure of Supreme Court Justices. While appointments in the United States can occupy media attention for weeks,[91] there is comparatively little debate in Germany or the UK, which is unsurprising if the judiciary is considered to apply laws neutrally. Finally, the crucial question is not *whether* but *how* the interaction and power dynamics between branches of government may be threatened by an AAJI. Will the judiciary become more or less powerful? How would this affect the interaction with and powers of the legislature and executive? Such questions are important from any liberal democratic perspective because, first of all, authoritarianism tends to favor a weak judiciary. Second, an independent judiciary is an important ally in the fight for minority rights and preventing the "tyranny of the majority".

Michaels argues that the current human-based judiciary is necessary to maintain the separation of powers, for only a human-centered judiciary raises ample attention to the law from the involvement of judges, lawyers, law professors, and so on.[92] For Michaels, it seems likely that the legal community would diminish if artificial judicial intelligence were to replace human judges,[93] and without a legal community, humans would pay little attention to the law.[94] He further argues that such lack of sufficient attention from the legal community might have devastating effects. First, it would be hard to imagine any public response to abuses of authority or concentration of power across the three branches.[95] Second, there would generally be "little

---

[87] Barnard (2019), p. 377.

[88] Cf. Michaels (2019), p. 1098 argument in this regard: "Courts exercise an important lawmaking and policymaking function when they interpret the law so as to resolve legal questions, and it is beneficial for such interpretation to take place in the context of concrete factual disputes."

[89] See, for instance, the still relevant observation by Denning (1963): "According to the British conception it is vital that the judges should be outside the realm of political controversy. They must interpret the law and mould it to meet the needs of the times, but they cannot bring about any major alterations in policy. It is only thus that judges can keep outside the sphere of politics" (p. 300).

[90] Kischel (2013) analyzes the election procedure of German constitutional court justices from a comparative perspective and finds that it is "necessary for its proper [politically] neutral functioning."

[91] For an overview of the Kavanaugh hearings in 2018, see Grynbaum (2018) and Bowden (2018).

[92] Michaels (2019), p. 1096.

[93] Michaels (2019), p. 1096; cf. also Volokh (2019) and Richard and Solow-Niederman (2019).

[94] Michaels (2019), p. 1096. It should be noted that Michaels does not specify how advanced the AI would have to be for such a scenario to occur. However, it is clear that he has a less capable AI in mind than what has been defined as an AAJI in this article. Presumably, his argument would apply all the more in the case of AAJI.

[95] Michaels (2019), p. 1096.

incentive to construct high-quality legal arguments if there was no possibility that doing so could shape the result."[96]

However, even assuming that the legal community would in fact diminish or be significantly reduced, it is not clear that this *itself* would necessarily lead to great power imbalances due to decreased human attention. On the contrary, there may be some significant accompanying advantages resulting from the shift in focus. The attention nowadays spent on specific cases may shift to law-making processes and the analysis of the overall consequences of the laws in question. While society is often occupied with extraordinary individual cases of little systemic value for the long-term, with AAJI, humans would be able to focus on the evaluation of abstract principles and rules rather than on our intuitive judgments about specific, often highly politicized, cases. Given that such perceptual intuitions are especially prone to cognitive biases,[97] offer little normative guidance, and have limited value for the whole of society, this may even nudge the public discourse onto a more beneficial path. Instead of focusing on extraordinary individual cases, attention might shift to the much more important trade-offs between detention, crime rates, and discrimination discussed in the first part of this article—potentially leading to lower incarceration, higher security, and less discrimination. Furthermore, the shift in focus from the judiciary onto the legislature (and executive) may carry the advantage of greater public accountability of the legislature, which would again profit from the transparency of the aforementioned trade-offs.

Although one can debate whether the adoption of AAJI would in fact lead to decreased attention on the judiciary, and whether such development would necessarily go hand in hand with a loss of power, the judiciary may nonetheless lose a significant share of its power and, eventually, its independence. This is ultimately so because taking the human out of the equation faces risks of democratic legitimacy. This itself may lead to a reduction of judicial power. Additionally, with reduced democratic legitimacy, it is difficult to imagine that judicial discretion would remain a decisive factor in many cases. On the contrary, I assume that legislatures would use a democratically weak judiciary to justify more and more specific rulings, leaving little discretion for the AAJI.[98] With power shifting from the judiciary to the legislature, judicial independence and separation of powers more generally would be under threat. The graphic below outlines the threat to judicial independence resulting from the implementation of AAJI (Fig. 1).

If democratic legitimacy were compromised in this way, the people might even welcome more powerful legislatures (and executives). Such a development could be accelerated by the improved accountability of the legislature resulting from greater

---

[96]Michaels (2019), p. 1097, see also p. 1084. Although this point is not crucial for the argument that follows, it may well be possible to have adversarial AIs arguing on behalf of each party, with an AAJI adjudicating.

[97]The terminology "perceptual intuition" goes back to Sidgwick (1907).

[98]See also Stern et al. (2021), p. 550 for a similar point on judicial authority and independence in the Chinese context. For discussion of how judicial independence in China may have increased, in part due to the use of judicial AI, see Deng (2019), pp. 242–243, 248–249.

Fig. 1 AAJI, legitimacy, and judicial independence

transparency of the aforementioned trade-offs. At last, the legislature may increasingly be perceived as responsible for the courts' rulings. Case in point, if the AAJI does not decide as the public wishes, people may blame the legislature for—in their eyes—a poor design of the laws in question, or poor decision to use a biased AAJI, rather than blaming the AAJI for poor implementation or interpretation of such laws. From this perspective, even governing parties who strongly support judicial independence might give in to public pressure to limit AAJI's discretion over time.

This is not a definite outcome, and it is highly uncertain where this shift in power would ultimately lead. However, the adoption of an AAJI may pose a great risk for judicial independence which arguably ought not to be taken before necessary safety mechanisms will have been generated. While there are tremendous efforts invested into the development of more advanced AI, including AGI and AAJI, little attention has been paid to creating the necessary political infrastructure that would allow society to harness the potentially enormous benefits of AI without risking the collapse of judicial independence. Although I disagree with Michaels regarding the reason why one might expect a shift in power (attention vs. legitimacy), this analysis strengthens the case for AAJI affecting the separation of powers. More precisely, I argue that it poses a great threat to judicial independence, which ought not be taken lightly.

4   Conclusion

As part of the legal community, we might be inclined to improve the judicial system as much as possible and defend it from potential threats, even if those threats accompany extraordinary improvements in other domains. As academics, we want to completely understand AI reasoning before we recommend its implementation, even if we do not fully understand human reasoning either. And as humans, we tend to prefer the status quo, even when change would be net-positive.[99] All these motivations might explain why so many of us have a strong aversion to the

---

[99]Cf. Samuelson and Zeckhauser's (1988) analysis of the status quo bias. Eidelman and Crandall (2012) further note that the status quo bias creates "barriers to cognitive and social change" (p. 270).

implementation of AI into the judicial sector, but, unfortunately, motivations are not ideal truth-tracking processes.[100] At least when it comes to issues of discrimination and transparency, this analysis makes clear that current applications could come with great advantages. Whereas this proposition can be defended from a wide range of normative perspectives on discrimination, the improvements with regards to transparency depend on how or whether one favors transparency of options over transparency of procedures and agents. Having said this, one should bear in mind that the existence of the possibility of such improvements does not mean that there is no risk of large-scale discrimination or a lack of transparency. As seen, the fact that such AI is *possible* does not mean that AI used *in practice* will be more transparent or less discriminatory.

While the analysis hitherto focused on existing technological capabilities, the final chapter introduced the concept of AAJI, defined as an artificially intelligent system that matches or surpasses human decision-making in all domains relevant to judicial decision-making. Concentrating on the potential power shift from the judiciary to the legislature, it has been argued that the adoption of an AAJI threatens judicial independence—and with that one of the central foundations of liberal democracy.

Needless to say, this attempt to outline the challenges of artificial judicial decision-making for liberal democracy is far from complete, and further research is needed in many regards. First, it is necessary to identify what other challenges the adoption of an AAJI may bring about, such as the risk that algorithms may grant laws a new kind of permanency and thereby create an additional barrier to legal evolution.[101] Second, it will be crucial to investigate how one can uphold liberal democratic values more generally and, in particular, the separation of powers with an AI-based judiciary. The short-term benefits AI could offer over HI with regards to access to justice, transparency, and fairness are so enormous that one might overlook the long-term threats imposed by AAJI. Working out the political and legal infrastructure to reap the fruits of artificial judicial intelligence in a safe and stable manner should become a priority of future research.

References

AI Now (2017) AI Now 2017 Report. New York University, New York. ainowinstitute.org/AI_Now_2017_Report.pdf

Angwin J et al (2016) Machine bias. ProPublica. www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

Bandes SA (2009) Empathetic judging and the rule of law. Cardozo Law Rev de novo 2009:133–148

---

[100]On the contrary, given the intuitive appeal of these motivations, one might consider debunking arguments.

[101]Crootof (2019) refers to this phenomenon as "technological-legal lock-in".

Barnard C (2019) The substantive law of the EU, 6th edn. Oxford University Press, Oxford, New York

Barredo Arrieta A et al (2020) Explainable Artificial Intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. Inf Fusion 58:82–115. https://doi.org/10.1016/j.inffus.2019.12.012

Barry C, Tomlin P (2019) Moral uncertainty and the criminal law. In: Alexander L, Ferzan KK (eds) Palgrave handbook of applied ethics and the criminal law. Palgrave Macmillan/Springer Nature, Cham, Switzerland

Baughman SB (2017) The bail book: a comprehensive look at bail in America's criminal justice system. Cambridge University Press, Cambridge

Baum S (2018) Reconciliation between factions focused on near-term and long-term artificial intelligence. AI Soc 33(4):565–572. https://doi.org/10.1007/s00146-017-0734-3

Baum S (2020) Medium-term artificial intelligence and society. Info 11(6):290–305. https://doi.org/10.3390/info11060290

Baum SD, Goertzel B, Goertzel TG (2011) How long until human-level AI? Results from an expert assessment. Tech Forecast Soc Change 78:185–195. https://doi.org/10.1016/j.techfore.2010.09.006

Bloom P (2016) Against empathy: the case for rational compassion. Ecco Press, New York

Boscardin C (2015) Reducing implicit bias through curricular interventions. J Gen Intern Med 30(12):1726–1728. https://doi.org/10.1007/s11606-015-3496-y

Bostrom N (2006) How long before superintelligence. Linguist Phil Investig 5:11–30

Bostrom N (2014) Superintelligence: paths, dangers, strategies. Oxford University Press, Oxford

Bowden J (2018) Timeline: Brett Kavanaugh's Nomination to the Supreme Court. The Hill. https://thehill.com/homenews/senate/410217-timeline-brett-kavanaughs-nomination-to-the-supreme-court

Carneiro D et al (2015) Online dispute resolution: an artificial intelligence perspective. AI Rev 41:211–240. https://doi.org/10.1007/s10462-011-9305-z

Carnes M et al (2015) Effect of an intervention to break the gender bias habit for faculty at one institution: a cluster randomized, controlled trial. Acad Med 90(2):221–230. https://doi.org/10.1097/ACM.0000000000000552

Cave S, Ó hÉigeartaigh S (2019) Bridging near- and long-term concerns about AI. Nat Mach Intel 1:5–6. https://doi.org/10.1038/s42256-018-0003-2

Chalmers D (2011) A computational foundation for the study of cognition. J Cogn Sci 12(4):325–359. https://doi.org/10.17791/jcs.2011.12.4.325

Chander A (2017) The racist algorithm? Mich Law Rev 115(6):1023–1045

Chen D (2019) Machine learning and the rule of law. In: Livermore M, Rockmore D (eds) Law as data. Santa Fe Institute Press, Santa Fe, pp 433–441

Chin D (2012) Sentencing: a role for empathy. Univ Penn Law Rev 160(6):1561–1584

Chohlas-Wood A (2020) Understanding risk assessment instruments in criminal justice. Brookings. www.brookings.edu/research/understanding-risk-assessment-instruments-in-criminal-justice

Colby TB (2012) In defense of judicial empathy. Minn Law Rev 96:1944–2015

Crootof R (2019) "Cyborg justice" and the risk of technological-legal lock-in. Columbia Law Rev Forum 119:233–251

Cushman F (2020) Rationalization is rational. Behav Brain Sci 43:E28. https://doi.org/10.1017/S0140525X19001730

Dabney D et al (2017) American bail and the tinting of criminal justice. Harv J Crime Justice 56(4):397–418. https://doi.org/10.1111/hojo.12212

Dahl R (1971) Polyarchy: participation and opposition. Yale University Press, New Haven

Das Gupta O, Fried N (2019) Seehofer redet über Gesetzestrick—hinterher spricht er von Ironie. Süddeutsche Zeitung www.sueddeutsche.de/politik/seehofer-datenaustauschgesetz-1.4479069

Deng J (2019) Should the common law system welcome artificial intelligence? A case study of China's same-type case reference system. Georgetown Law Tech Rev 3(2):223–280

Denning L (1963) The function of the judiciary in a modern democracy. Pak Horizon 16(4):
299–305

Devine P et al (2012) Long-term reduction in implicit race bias: a prejudice habit-breaking
intervention. J Exp Soc Psychol 48(6):1267–1278. https://doi.org/10.1016/j.jesp.2012.06.003

Dhami MK, Ayton P (2001) Bailing and jailing the fast and frugal way. J Behav Decis Mak 14:141–
168. https://doi.org/10.1002/bdm.371

Dworkin R (2011) Justice for hedgehogs. Harvard University Press, Cambridge

Eidelman S, Crandall C (2012) Bias in favor of the status quo. Soc Personal Psychol Compass 6(3):
270–281. https://doi.org/10.1111/j.1751-9004.2012.00427.x

Englich B et al (2006) Playing dice with criminal sentences: the influence of irrelevant anchors on
experts' judicial decision making. Personal Soc Psychol Bull 32(2):188–200. https://doi.org/10.
1177/0146167205282152

Floridi L, Cowls J, Beltrametti M et al (2018) AI4People—an ethical framework for a good AI
society: opportunities, risks, principles, and recommendations. Minds Mach 28:689–707.
https://doi.org/10.1007/s11023-018-9482-5

Florini A (1999) Does the invisible hand need a transparent glove? The politics of transparency.
Paper presented at Annual World Bank Conference on Development Economics, Washington,
D.C.

Fox J (2007) Government transparency and policymaking. Pub Choice 131(1/2):23–44

Gigerenzer G (1991) How to make cognitive illusions disappear: beyond "heuristics and biases."
Eur Rev Soc Psych 2(1):83–115. doi:https://doi.org/10.1080/14792779143000033

Glynn AN, Sen M (2015) Identifying judicial empathy: does having daughters cause judges to rule
for women's issues? Am J Pol Sci 59(1):37–54. https://doi.org/10.1111/ajps.12118

Grace K, Salvatier J, Dafoe A et al (2018) Viewpoint: when will AI exceed human performance?
Evidence from AI experts. J AI Res 62(1):29–754. https://doi.org/10.1613/jair.1.11222

Gruetzemacher R, Paradice D, Lee KB (2019) Forecasting transformative AI: an expert survey.
arXiv:190108579 [cs]

Grynbaum MM (2018) Kavanaugh hearings on TV offer riveting drama to a captive nation.
New York Times. https://www.nytimes.com/2018/09/27/business/media/kavanaugh-blasey-
ford-hearing-tv.html

Gupta A (2008) Transparency under scrutiny: information disclosure in global environmental
governance. Glob Envt Pol 8(2):1–7. https://doi.org/10.1162/glep.2008.8.2.1

Gustafsson JE, Torpman O (2014) In defence of my favourite theory. Pacific Phil Q 95:159–174.
https://doi.org/10.1111/papq.12022

Guthrie C et al (2001) Inside the judicial mind. Cornell Law Rev 86(4):777–830

Hacker P (2018) Teaching fairness to artificial intelligence; existing and novel strategies against
algorithmic discrimination under EU law. Common Market Law Rev 55(4):1143–1185

Hammer P (2018) Detroit 1967 and today: spatial racism and ongoing cycles of oppression. J Law
Soc 18(2):227–235

Hollyer et al (2011) Democracy and transparency. J Pol 73(4):1191–1205. https://doi.org/10.1017/
s0022381611000880

Holmes O (1897) The path of the law. Harv Law Rev 10:457–478. https://doi.org/10.2307/1322028

Ingriselli E (2015) Mitigating jurors' racial biases: the effects of content and timing of jury
instruction. Yale Law J 124(5):1690–1745

Kahneman D, Tversky A (1982) Intuitive prediction: biases and corrective procedures. In:
Kahneman D, Slovic P, Tversky A (eds) Judgment under uncertainty: heuristics and biases.
Cambridge University Press, Cambridge, pp 414–421

Kahneman D, Tversky A (1983) Extensional versus intuitive reasoning: the conjunction fallacy in
probability judgment. Psych Rev 90(4):293–315. https://doi.org/10.1037/0033-295X.90.4.293

Kaminski M (2019) Binary governance: lessons from the GDPR's approach to algorithmic
accountability. South Cal Law Rev 92:1529–1616

Karnow C (2008) Setting bail for public safety. Berkeley J Crim Law 13(1):1–30

Kaufmann D, Bellver A (2005) Transparenting transparency: initial empirics and policy applications. MPRA Paper 8188. University Library of Munich, Germany

Kischel U (2013) Party, pope, and politics? The election of German constitutional court justices in comparative perspective. Int J Const Law 11:962–980. https://doi.org/10.1093/icon/mot040

Kleinberg J et al (2018) Human decisions and machine predictions. Q J Econ 133(1):273–293. https://doi.org/10.1093/qje/qjx032

Kleinberg J et al (2019) Discrimination in the age of algorithms. J Leg Anal 10:113–174. https://doi.org/10.1093/jla/laz001

Kneer M, Skoczeń I (forthcoming) Outcome effects, moral luck and the hindsight bias. Cognition. https://doi.org/10.2139/ssrn.3810220

Lai C et al (2014) Reducing implicit racial preferences: I. A comparative investigation of 17 interventions. J Exp Psych Gen 143(4):1765–1785. https://doi.org/10.1037/a0036260

Larson et al (2016) How we analyzed the COMPAS recidivism algorithm. ProPublica. www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm

Lee RK (2014) Judging judges: empathy as the litmus test for impartiality. Univ Cincinnati Law Rev 82(1):145–206

Levin S (2016) A beauty contest was judged by AI and the robots didn't like dark skin. Guardian. www.theguardian.com/technology/2016/sep/08/artificial-intelligence-beauty-contest-doesnt-like-black-people

Licht JD (2011) Do we really want to know? The potentially negative effect of transparency in decision making on perceived legitimacy. Scand Pol Stud 34:183–201. https://doi.org/10.1111/j.1467-9477.2011.00268.x

Licht JD (2014) Policy area as a potential moderator of transparency effects: an experiment. Pub Admin Rev 74(3):361–371. https://doi.org/10.1111/puar.12194

Licht KD, Licht JD (2020) Artificial intelligence, transparency, and public decision-making: why explanations are key when trying to produce perceived legitimacy. AI Soc 35(4):917–926. https://doi.org/10.1007/s00146-020-00960-w

Liptak A (2008) U.S. voting for judges perplexes other nations. New York Times. www.nytimes.com/2008/05/25/world/americas/25iht-judge.4.13194819.html

Lockhart T (2000) Moral uncertainty and its consequences. Oxford University Press, Oxford

Lowry S, Macpherson G (1988) A blot on the profession. Br Med J 296(6623):657–658. https://doi.org/10.1136/bmj.296.6623.657

Macaskill W (2014) Normative uncertainty. Dissertation, Oxford University

MacAskill W, Bykvist K, Ord T (2020) Moral uncertainty. Oxford University Press, Oxford

Martínez E, Winter CK (2021) Artificial intelligence in the judiciary: a global survey of legal academics. [Manuscript in preparation]

McGinley A (2011) Ricci v. DeStefano: diluting disparate impact and redefining disparate treatment. Nevada Law J 12(3):626–639

Meadow W, Sunstein C (2001) Statistics, not experts. Duke Law J 51:629–646

Michaels AC (2019) Artificial intelligence, legal change, and separation of powers. Univ Cincinnati Law Rev 88:1083–1103

Moore S (2018) Towards a sociology of institutional transparency: openness, deception and the problem of public trust. Sociology 52(2):416–430. https://doi.org/10.1177/0038038516686530

Mukand S, Rodrik D (2020) The political economy of liberal democracy. Econ J 130(627):765–792. https://doi.org/10.1093/ej/ueaa004

Müller VC, Bostrom N (2016) Future progress in artificial intelligence: a survey of expert opinion. In: Müller VC (ed) Fundamental issues of artificial intelligence. Springer International Publishing, Cham, pp 555–572

Negowetti NE (2015) Judicial decisionmaking, empathy, and the limits of perception. Akron Law Rev 47(3):693–751

Niiler E (2019) Can AI be a fair judge in court? Estonia thinks so. Wired. www.wired.com/story/can-ai-be-fair-judge-court-estonia-thinks-so

O'Neil C (2016) Weapons of math destruction: how big data increases inequality and threatens democracy. Broadway Books, New York

Oswald F et al (2013) Predicting ethnic and racial discrimination: a meta-analysis of IAT criterion studies. J Pers Soc Psychol 105(2):171–192. https://doi.org/10.1037/a0032734

Pejovic C (2001) Civil law and common law: two different paths leading to the same goal. Victoria Univ Wellington Law Rev 32(3):817–842

Piccinini G (2016) The computational theory of cognition. In: Müller V (ed) Fundamental issues of artificial intelligence. Synthese Library, vol 376. Springer, Cham, pp 203–221

Pinker S (2011) The better angels of our nature: why violence has declined. Viking Press, New York

Pinker S (2018) Enlightenment now: the case for reason, science, humanism, and progress. Viking Press, New York

Piovesan C, Ntiri V (2018) Adjudication by algorithm: the risks and benefits of artificial intelligence in judicial decision-making. Advocates' J 44:42–45

Prunkl C, Whittlestone J (2020) Beyond near- and long-term: towards a clearer account of research priorities in AI ethics & society. arXiv:2001.04335v2 [cs.CY]

Przeworski A (2000) Democracy and development: political institutions and well-being in the world, 1950-1990. Cambridge University Press, Cambridge

Rachlinski J, Wistrich A (2017) Judging the judiciary by the numbers: empirical research on judges. Ann Rev Law Soc Sci 13:203–229. https://doi.org/10.1146/annurev-lawsocsci-110615-085032

Reiling D (2018) What role for AI in a judge's decision-making process? Lecture presented at the European Commission for the Efficiency of Justice's Conference on "Artificial Intelligence at the Service of the Judiciary." Council of Europe. www.coe.int/en/web/cepej/justice-of-the-future-predictive-justice-and-artificial-intelligence

Richard M, Solow-Niederman (2019) Developing artificially intelligent justice. Stanford Tech Law Rev 22(2):242–289

Samuelson W, Zeckhauser R (1988) Status quo bias in decision making. J Risk Uncertainty 1:7–59. https://doi.org/10.1007/BF00055564

Schumpeter (1942) Capitalism, socialism and democracy. Harper & Brothers, New York

Shaked-Schroer N (2008) Reducing racial bias in the penalty phase of capital trials. Behav Sci Law 26(5):603–617. https://doi.org/10.1002/bsl.829

Sidgwick H (1907) The methods of ethics, 7th edn. Macmillan, London

Smith M (2016) In Wisconsin, a backlash against using data to foretell defendants' futures. New York Times. www.nytimes.com/2016/06/23/us/backlash-in-wisconsin-against-using-data-to-foretell-defendants-futures.html

Sourdin T (2018) Judge v Robot? Artificial intelligence and judicial decision-making. Univ New South Wales Law J 41(4):1114–1133. https://doi.org/10.53637/ZGUX2213

Sourdin T, Cornes R (2018) Do judges need to be human? The implications of technology for responsive judging. In: Sourdin T, Zariski A (eds) The responsive judge: international perspectives. Springer, Singapore, pp 87–119

Stern RE et al (2021) Automating fairness? Artificial intelligence in the Chinese Court. Columbia J Transnatl Law 59:515–553

Struchiner N, Almeida G, Hannikainen I (2020) Legal decision-making and the abstract/concrete paradox. Cognition 205:104421. https://doi.org/10.1016/j.cognition.2020.104421

Sunstein (2019) Algorithms, correcting biases. Soc Res: Int Q 86(2):499–511

Tarsney C (2018) Moral uncertainty for deontologists. Ethic Theory Moral Pract 21:505–520. https://doi.org/10.1007/s10677-018-9924-4

Teal C et al (2012) Helping medical learners recognise and manage unconscious bias toward certain patient groups. Med Educ 46(1):80–88. https://doi.org/10.1111/j.1365-2923.2011.04101.x

Tversky A, Kahneman D (1973) Availability: a heuristic for judging frequency and probability. Cogn Psychol 5:207–232. https://doi.org/10.1016/0010-0285(73)90033-9

Volokh E (2019) Chief justice robots. Duke Law J 68:1135–1192

Wachter S, Mittelstadt B, Floridi L (2017) Why a right to explanation of automated decision-making does not exist in the general data protection regulation. Int Data Privacy Law 7(2): 76–99. https://doi.org/10.1093/idpl/ipx005

Wallace M et al (2017) Separate and unequal: structural racism and infant mortality in the US. Health Place 45(3):140–144. https://doi.org/10.1016/j.healthplace.2017.03.012

Waltl B, Vogl R (2018) Increasing transparency in algorithmic- decision-making with explainable AI. Datenschutz und Datensicherheit 42:613–617. https://doi.org/10.1007/s11623-018-1011-4

Wang P, Goertzel B (2007) Introduction: aspects of artificial general intelligence. In: Wang P, Goertzel B (eds) Advances in artificial general intelligence: concepts, architectures and algorithms, Proceedings of the AGI Workshop 2006, pp 1–16. https://dl.acm.org/doi/10.5555/1565455.1565457

Weinberg JD, Nielsen LB (2012) Examining empathy: discrimination, experience, and judicial decisionmaking. South Cal Law Rev 85(2):313–351

Winter CK (2020) The value of behavioral economics for EU judicial decision-making. German Law J 21(2):240–264. https://doi.org/10.1017/glj.2020.3

Winter CK (2022) Metamoralisches Strafrecht. [Manuscript in preparation]

Winter CK, Schuett J, Martínez E, Van Arsdale S, Araújo R, Hollman N, Sebo J, Stawasz S, O'Keefe C, Rotola G (2021) Legal priorities research: a research agenda. Legal Prior Proj. https://doi.org/10.2139/ssrn.3931256

Wistrich A, Rachlinski J (2018) Implicit bias in judicial decision making, how it affects judgment and what judges can do about it. In: Redfield S (ed) Enhancing justice, reducing bias. ABA Book Publishing, Chicago, pp 87–130

Wistrich A et al (2015) Heart versus head: do judges follow the law or follow their feelings? Texas Law Rev 93:855–923

World Government Summit (2018) Could an AI ever replace a judge in court? www.worldgovernmentsummit.org/observer/articles/could-an-ai-ever-replace-a-judge-in-court

Završnik A (2020) Criminal justice, artificial intelligence systems, and human rights. ERA Forum 20:567–583. https://doi.org/10.1007/s12027-020-00602-0

Zavyalova V (2018) Save money on legal advice: AI is replacing lawyers in Russia. Russia Beyond. www.rbth.com/science-and-tech/327585-free-legal-advice-robolawyer

Zhang M (2015) Google photos tags two African-Americans as gorillas through facial recognition software. Forbes. www.forbes.com/sites/mzhang/2015/07/01/google-photos-tags-two-african-americans-as-gorillas-through-facial-recognition-software/#61fdae0c713d