LEX GAZETTE
International Multidisciplinary Research Journal

ISSN: 3049-3536 (O)

VOLUME: II, ISSUE: II
APRIL (2026)

Lex Gazette International Multidisciplinary Research Journal

DOUBLE BLIND PEER REVIEWED

OPEN ACCESS

Justice Or Just Code? Envisioning Criminal Justice in AI-Driven Law Enforcement Through The Lens Of Ethics And Morality

Deeksha Dewangan
Student at Institute of Law and Legal Studies, SAGE University, Indore
Contact at: dewangandeeksha698@gmail.com

ABSTRACT

The deployment of Artificial Intelligence in criminal justice systems is arguably one of the most significant technological changes in the 21st century. AI is gradually transforming the entire spectrum of law enforcement in how it identifies, targets, and processes individuals. The technologies include predictive policing algorithms, facial recognition surveillance, and automated risk assessment tools that significantly affect bail and sentencing decisions. Proponents tout these technologies as capable of delivering efficiency, objectivity, and improved crime-fighting results on a scale never before imaginable. However, several fundamental questions about the technology remain in the eyes of the critics, especially concerning algorithmic bias, violations of the right to a fair trial, privacy infringements, and dehumanisation of the concept of justice. This paper takes an in-depth look at the use of Artificial Intelligence in criminal justice from the perspectives of both ethics and morality. It employs well-established philosophical concepts such as Rawlsian justice, deontological ethics, and utilitarian theory to assess whether algorithmic systems are able to bring about justice or only serve to perpetuate and worsen existing social inequalities in a coded form. Moreover, it discusses the revolutionary prospective role of AI in justice systems, considering its advantages against its ethical threats, and sets out a model for ethical regulation. The paper ends by arguing that justice, as a moral and social institution, cannot be equated with computational logic without intentional ethical considerations, participatory design, and strong democratic control.

KEYWORDS

Artificial Intelligence, Criminal Justice, Algorithmic Bias, Predictive Policing, Ethics, Morality, Facial Recognition

1. INTRODUCTION

Justice is a moral concept at its most basic level. Underlying it are principles of fairness, equality, accountability, and human dignity – ideals that democratic societies have long sought to achieve through law. Within this context of Artificial Intelligence, we have the administration of justice involving algorithms, machine learning models, and other data-driven systems, operating at a scale and pace not possible for any human institution. What this paper asks is not whether AI is useful in criminal justice – empirically, it often is – but rather whether a system capable of moral reasoning can, however partially, be replaced by a system based on statistical inference.

Law enforcement agencies around the world show an overwhelming eagerness to use AI tools. In the USA, Chicago and Los Angeles have tested predictive policing programs for crime-solving. COMPAS serves as a risk assessment algorithm that courts use for bail and sentencing. In the UK, the Metropolitan Police used facial recognition cameras in public spaces. In India, the automated facial recognition system (AFRS) and the crime and criminal tracking network and systems (CCTNS) indicate the significant investment of the state in AI-led policing. These systems are not experimental setups at the margins; they are parts of real-life operational systems.

However, the rapid deployment has occurred at a faster pace than the development of ethical, legal, and regulatory frameworks. There have been documented incidents of algorithms being racially biased, wrongful arrests due to faulty facial recognition, and a lack of transparency with respect to black-box decision-making tools.[1] These have all led to urgent calls from civil society, legal academics, and ethicists for scrutiny. The sense of urgency is not simply an academic one. The consequences are immediate and irreversible when an algorithm indicates a person may pose a high recidivism risk, and that output is used by a judge to deny bail.

This paper has eight sections. This article first places the contemporary use of artificial intelligence (AI) in the criminal justice system in context. It evaluates whether and how AI can deliver justice and whether it can do so in the future. It then goes on to a systematic evaluation of the ethical and moral challenges that may be posed by such technology. Through the lens of justice theory, the paper weighs benefits against risks. It offers a suite of policy and governance

[1] Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information (2015).

recommendations aimed at ensuring that the adoption of AI enhances – rather than undermines – the moral architecture of justice.

2. CURRENT USE OF AI IN CRIMINAL JUSTICE

The deployment of AI in criminal justice spans several interconnected domains, each representing both a technical innovation and a site of ethical contestation. Understanding what these tools do — and how they do it — is essential before any normative evaluation can proceed.

2.1 Predicting Crime Contagion

Predictive policing is an algorithmic tool that forecasts where a crime is likely to occur or who is likely to commit a crime. Tools like PredPol (now branded as Geolitica) use historical crime data along with geographic and other patterns to allocate police resources to certain "high-risk" places or people. The whole premise that past behaviour predicts future behaviour, a claim that criminologists say does not hold water, either with evidence or on normative grounds. Critics say these systems are trained on data from biased policing that has been historically done. That data is used to predict where those crimes would occur again using prediction technology.[2]

2.2 Risk Evaluation Algorithms

Risk assessment tools such as COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) and the Public Safety Assessment (PSA) are employed at various stages of the criminal justice process, including pretrial detention, sentencing, and parole decisions. These tools generate scores estimating the likelihood that a person will reoffend. COMPAS was the subject of an analysis published by the investigative journalists at ProPublica in Broward County, Florida, in 2016.[3] The analysis found that Black defendants were almost twice as likely to be falsely flagged as high-risk compared to White defendants. It also found that White defendants were more likely to be falsely flagged as low-risk. The research sparked a worldwide debate concerning fairness in algorithmics and the evidentiary standards that ought to govern AI evaluations in court cases.

[2]Ales Završnik, Algorithmic Justice: Algorithms and Big Data in Criminal Justice Settings, 18 Eur. J. Criminology 623 (2021).

[3]Julia Angwin et al., Machine Bias, ProPublica (May 23, 2016), https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

2.3 Facial Recognition Technology

Facial recognition technology (FRT) is used to identify or verify a person from a digital image or video frame. Law enforcement uses facial recognition technology to identify suspects from footage, track individuals in public places, and verify identities at the border. The accuracy of the technology is not consistent. Research by the National Institute of Standards and Technology (NIST) shows that the best facial recognition algorithms have heightened misidentification rates for dark-skinned faces, especially those of dark-skinned women, compared to light-skinned faces.[4] This difference in accuracy has resulted in several wrongful arrests. This has been documented in the case of Robert Williams, Michael Oliver, and Nijeer Parks, who is a Black man wrongly arrested in the US as a result of FRT.

Automated Surveillance And Social Media Monitoring.

Outside of these headline-grabbing uses, AI is also extensively deployed for social media monitoring on a grand scale, license plate recognition, drone monitoring, and gang database management. The law enforcement tools reach into both digital and public spaces in ways that raise unsettling questions about privacy, association, and the presumption of innocence. Police Forces in India Have Deployed AI-Surveillance Tools at Protests: As per civil liberties groups, the deployment of artificial intelligence-based systems at public gatherings will deter free speech and political dissent.

3. EFFICIENCY OF AI IN DELIVERING JUSTICE

Proponents of AI in criminal justice make a compelling case rooted in efficiency, consistency, and the demonstrable fallibility of human judgment. It is important to engage with this case seriously before subjecting it to ethical critique.

3.1 Velocity and Magnitude

Artificial intelligence systems can utilize vast amounts of data at a speed and scale that no human being can match, from criminal records to social media activity and beyond. This computational power creates a true operational edge in complex fraud investigations, money laundering cases, or counter-terrorism activities. Work that takes humans months to sort through can happen in just a few hours with the help of AI, saving lives.

[4]Patrick Grother, Mei Ngan & Kayee Hanaoka, U.S. Dep't of Commerce, Nat'l Inst. of Standards & Tech., NISTIR 8280, Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects (2019).

3.2 Reducing Human Biases.

One of the main arguments made for the deployment of AI in the criminal justice system is that it can reduce inherent biases in humans. Scientists who study human behaviour, in particular psychologists and criminologists, have long recognised that judges are not immune to influences which have nothing to do with the merits of a case. These influences include things like the defendant's race and/or appearance, the time of day, and how hungry the judge may be at that point in time. The hungry judge phenomenon, detailed by Danziger, Levav, and Avnaim-Pesso, had the finding that Israeli parole board judges became less favourable as the day wore on.[5] The argument goes that if AI tools can provide a more consistent baseline, they can lessen such arbitrary variance.

3.3 Optimization of Resources

Law enforcement suffers from a constant lack of funds. Using AI systems for crime mapping, optimising patrol functions, and prioritising cases allows for directing limited resources to high-impact work. AI case triage and document review can aid jurisdictions with large backlogs to speed up the pace of justice and lessen the time defendants spend in pre-trial detention while improving throughput.

3.4 The Efficiency Argument: What Limits?

But efficiency, however attractive as a method, cannot be the measure of justice. A system that provides transactional processing at high speed but systematically disadvantages certain groups is not delivering justice — it is simply automating injustice at scale. The efficiency argument also tends to muddle the difference between optimising a system and improving one. Speeding injustice does not make it more just; indeed, it can make challenging it more exasperating. Justice. Efficiency is a value, but it is not the only or highest value in the administration of justice.

4. FUTURISTIC ROLE OF AI IN JUSTICE SYSTEMS

Looking beyond current applications, AI promises to reshape criminal justice in ways that are both transformative and deeply uncertain. Envisioning this future requires intellectual honesty about both its potential benefits and its inherent risks.

[5]Shai Danziger, Jonathan Levav & Liora Avnaim-Pesso, Extraneous Factors in Judicial Decisions, 108 Proc. Nat'l Acad. Sci. 6889 (2011).

4.1 Smart-proof Analysis.

When it comes to more sophisticated forensic analysis such as DNA sequencing, other digital forensics, audio and video authentication, as well as behavioural pattern analysis, future AI systems will easily outperform anything available right now. Improving conviction reliability would significantly reduce wrongful convictions and could also solve cold cases that have stumped humanity for ages. Tools for legal research that use AI will also make high-quality assistance more accessible. This may bridge the justice gap between those who can afford proper lawyers and those who can't.

4.2 Predicting and Preventing Wrongdoings

Some legal theorists and technologists envisage a future in which AI systems will identify at-risk people for crime before any crime is committed, so social help can be provided. The idea often referred to as 'pre-crime' prevention after Philip K. Dick's fictional conceit raises deep philosophical questions about whether justice is capable of being prospective, whether one can be culpable for statistical probabilities, and whether this is compatible with the basic presumption of innocence. A justice system that would punish potential rather than action, while technically possible, would represent a categorically different animal from any existing system of law.

4.3 Virtual courts assisted by AI-based adjudication.

The pandemic brought virtual proceedings in courts, and now several jurisdictions discuss AI-assisted adjudication for minor offences, traffic violations, and small claim cases. Estonia has tested an artificial intelligence judge for small claims. While these applications might promote access to justice and reduce costs, they also risk depersonalising adjudication in ways that may undermine the rights of defendants to be heard, to confront evidence, and to receive individualised consideration of circumstances.

4.4 AI to rehabilitate

AI can assist in rehabilitation processes, which are aggressive and of the least disputable or compatible. The potential of AI tools is significant when it comes to personalizing educational and vocational programs for incarcerated individuals, identifying mental health needs and modeling effective pathways for reintegration. Using this application involves the use of restorative justice principles, as it concentrates on the restoration rather than prediction or punishment.

5. ETHICAL AND MORAL CHALLENGES OF AI JUSTICE

The deployment of AI in criminal justice raises ethical and moral challenges of the first order - challenges that cut to the heart of what justice means and who it is for.

5.1. Algorithmic Bias and Structural Discrimination

The extensively documented challenge is algorithmic bias, the situation where AI systems produce discriminatory outputs that disadvantage the racially, economically, and socially marginalised. The flaw is systematic and not accidental or deviant. AI systems learn from historical data, and historical criminal justice data reflects centuries of discriminatory policing, prosecution, and incarceration. A recidivism prediction tool trained on data from a system that has over-policed Black communities will reproduce those patterns in its outputs. In this context, the algorithm accurately demonstrates the bias of historical injustice.

John Rawls writes in A Theory of Justice that just institutions are those that rational actors would choose behind a 'veil of ignorance' — including the people who don't know their social position, race or wealth.[6] No rational person behind such a veil would pick a justice system that systematically disadvantaged them because of immutable characteristics of identity. AI systems with embedded racial and class bias fail the Rawlsian test outright.

5.2 Opacity, Due Process, and the Right to a Fair Hearing

Proprietary and opaque AI tools, including some of the most impactful ones in the criminal justice system, present a significant problem. The developers of these tools keep their algorithms as trade secrets, thus preventing third-party review. Such a lack of transparency not only conflicts with the basic idea of due process rights but also violates them. In Mathis v. United States as well as in other constitutional rulings, it has been recognized that defendants can "challenge the evidence against them". However, practically, how can a defendant challenge a risk score computed by an unseen algorithm, employing inputs that are not fully disclosed and verified using a completely inaccessible methodology? The famous Wisconsin case State v. Loomis (2016) brought out the issue quite clearly.[7] The defendant was criticizing his sentence, partly relying on a COMPAS risk assessment, and claimed that he was not in a position to effectively challenge the tool's method. The Wisconsin Supreme Court, although

[6]John Rawls, A Theory of Justice (1971).

[7]State v. Loomis, 881 N.W.2d 749 (Wis. 2016).

supporting the sentence, recognized the discomfort caused by the use of secret algorithms in the criminal process. If anything, the court indicated that while the due process framework may have to be altered to suit the AI era, the change should be aimed at enhancing transparency rather than reducing it.

5.3 Privacy, Surveillance, and the Presumption of Innocence

Extensive AI surveillance systems can degrade the presumption of innocence by regarding whole populations as potential suspects. People's faces are recognised in a mass facial recognition surveillance system, social media is monitored, and predictive geofencing is a method through which individuals who have committed no offence become the subject of state scrutiny at all times. This is not just a privacy issue in the limited sense; it is a deep disruption of the moral framework of liberal justice, which assumes at the start that citizens are free unless it is proven that they are not. Immanuel Kant, the great philosopher who argues through deontological ethics that people should be treated as ends in themselves and not as means to achieving some external purpose,[8] would consider such surveillance as something that is really not allowable on a moral level. To keep an eye on a person without a reasonable cause is to reduce their body and conduct to mere data input for the use of the state - that is, to turn them into instruments of social control rather than recognizing them as individuals endowed with dignity and autonomy.

5.4 Dehumanisation and the Moral Dimension of Judgment

One of the biggest moral issues AI in the criminal justice system could face is the dehumanization of judicial decisions. Besides following rules, justice requires a judge, a jury, or even a police officer to use their moral decision-making abilities, their empathy and understanding of the context, and most importantly, their recognition of the full humanity of the person who is before the law. For example, a judge who decides on a sentence will weigh the defendant's background, the situation, the level of regret demonstrated, family situation, and even the chances of the person's rehabilitation. An algorithm just considers the 'variables'. These are two entirely different things, and mixing them up is essentially reducing justice to a mathematical operation. Martha Nussbaum's view, based on the concept of capabilities, sees justice from the perspective of individuals' being complex human beings with various human

[8]Immanuel Kant, Groundwork of the Metaphysics of Morals (Mary Gregor trans., Cambridge Univ. Press 1998) (1785).

needs that cannot be broken down any further; in other words, not as mere elements of data.[9] If risk scores take over from the narrative, probability from judgment, and efficiency from dignity, then something fundamental to justice is lost.

6. BENEFITS VS. RISKS: THE JUSTICE DILEMMA

Any honest assessment of AI in criminal justice must resist the temptation of easy conclusions in either direction. The picture is genuinely complex, and the ethical challenge is to navigate that complexity with rigour.

6.1 The Case for Cautious Optimism

There are real advantages to AI-assisted justice that deserve consideration. In fact, human decision-makers in criminal justice are flawed beings prone to bias, tiredness, feelings, and even corruption at times. If AI tools can help highlight sentencing inconsistencies, identify potential wrongful convictions via re-examining evidence, or point to crime-related communities for policing, these will actually contribute to justice. Besides, the other choice that we have is not an ideal, error-free human judgement but the existing one with all its documented issues. An AI tool that is well-designed and regularly tested for performance and is capable of lessening sentencing disparity or allowing more rapid handling of cases for pre-trial detainees may, if you look at it as a whole, be a more just solution than the present situation, even if only slightly so.

6.2 The Gravity of the Risks

Human freedom is one of the deepest interests of a person. Any AI system that produces a false positive - that is, wrongly deciding that someone is a criminal risk, mistakenly identifying them in a crime through facial recognition, or wrongly labeling them as a flight risk - can lead that person to be arrested, detained, prosecuted, or imprisoned. These are not statistical disruptions but serious violations of human dignity which have a further impact on the victim's employment, family, mental health, and civic life. Also, systemic discrimination in AI, as already observed, is not the only factor that disproportionately affects certain groups of people according to race, poverty, and historical disadvantage. It is not only unfair to mistreat the same people again each time - it is a moral disaster that perpetuates the existing injustice while being hidden under the guise of technical neutrality. The falsehood of algorithm objectivity

[9]Martha C. Nussbaum, Creating Capabilities: The Human Development Approach (2011).

contributes to the fact that human prejudice can be disguised by human-like science, making both fetching and challenging discrimination more difficult.

6.3 A Framework for Weighing the Dilemma

One of the most reputable frameworks for weighing one's dilemma is based on the principle of proportionality. As the stakes are raised, the more severe the possible deprivation of liberty, the stricter the ethical standards that should regulate the role of AI. For such liberty-regulating decisions as pretrial detention, sentencing, and parole, the standards should be set very high. On the other hand, for lower-stakes administrative uses such as case management, resource allocation, and document review, the risk-benefit ratio is more favorable. It is a big mistake to consider all AI criminal justice applications as being equally effective since it will hide the morally relevant differences.

7. ETHICAL GOVERNANCE AND POLICY RECOMMENDATIONS

Addressing the ethical challenges of AI in criminal justice requires coordinated action across the domains of law, regulation, technology design, and democratic accountability. The following recommendations are grounded in the analysis developed above and informed by emerging international best practices.

7.1 Mandatory Transparency and Explainability

Any AI system used in criminal justice proceedings must be transparent and explainable. This requirement means that defendants, attorneys, and higher courts should be provided with enough information about the system's architecture, data usage, validation techniques, and a list of expected system limitations to be able to challenge the system effectively. Such provision should not just be recommended but made a statutory requirement. A defendant's constitutional due process rights must not be overshadowed by a technology company's claim of trade secrets. The EU AI Act's categorization of law enforcement AI as 'high-risk ', demanding conformity checks, transparency responsibilities, and human supervision, offers a framework that other countries should emulate and improve.[10]

[10]Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), 2024 O.J. (L 1689).

7.2 Mandatory Algorithmic Bias Auditing

Every AI tool employed in criminal justice must be periodically exposed to unbiased and thorough audits for identifying bias and any other discriminatory effects along racial, gender, socioeconomic, and even geographical factors. There should be an assurance of these audits' independence from both the technology vendor and the agency that deploys the tool, and the outcomes of these assessments should be open to public scrutiny. In cases where an audit unearthed significant disparities, the use of the respective tool should be halted until remedial measures for the identified bias are undertaken. It is incorrect to view this merely as a technical requirement since it is a moral one that derives from the promise of equal treatment.

7.3 Human-in-the-Loop Requirements

AI results should be considered only as one of many factors in making high-stakes decisions such as pretrial detention, sentencing, parole, and criminal charging, and should never be the sole deciding factor. The human decision-makers must have the final say and be fully responsible for the decisions they make. This 'human-in-the-loop' stipulation should be a meaningful change in role, not just a superficial one; courts and agencies should be able to evidence real consideration of the algorithmic recommendations rather than simply approving them. Along with the knowledge of AI tools, training programs for judges, prosecutors, and law enforcement officers must also cover their limitations and the biases that have been identified in them.

7.4 Prohibition of High-Risk Applications

Some AI applications in the criminal justice sector are so problematic that they should be banned outright. The total and unmitigated use of facial recognition to surveil people in public locations, without a specific and judicially authorized target, should be forbidden since it is at odds with the presumption of innocence and the right to privacy. Predictive policing software that identifies individuals (rather than being used for the deployment of resources at aggregate levels) should be put on hold until there is clear evidence that they do not lead to racially discriminatory results. Risk-based detention before the crime, which is mainly based on prediction through algorithms, should not be allowed.

7.5 The Indian Legal and Regulatory Context

India is one of the most unique places in the world to implement these recommendations. To start with, our criminal justice system faces several challenges, such as overcrowded jails, long

delays in pretrial detention, etc. First, it needs to be emphasized that the deployment of AI in criminal justice in India is not new. Currently, the police use such tools to identify and cabin potential criminals and to find a crime scene. For instance, a Facial Recognition System (AFRS) has been installed in 29 metropolitan cities, and CCTV footage is used alongside these tools for identification. Secondly, using such technologies without any check or balance or without any specific laws in AI governance is a matter of concern. Thirdly, the Indian constitution envisages rights that must be protected even in the use of AI, such as the right to equality before the law under Article 14, the right to freedom of expression under Article 19, and the right to life and personal liberty under Article 21. Fourth, the highest court of the country declared that privacy is a fundamental right under the Constitution in the landmark judgment of K.S. Puttaswamy v. Union of India.[11] This case can be relied upon to argue that regulation of AI is necessitated due to its capacity to violate fundamental rights. At last, there is a pressing need for a made-in-India AI governance framework that not only caters to the needs of the general AI governance but also addresses the specific issues of law enforcement that include fraud identification and investigation, to name some and balances civil rights and liberties. Such a law must also lay down the duties of transparency, the power of supervisory inspections, and standards of strict accountability and ways of enforcement.

7.6 Inclusive and Participatory Design

It is a must that AI systems in criminal justice are created through the lively involvement of the communities that are most impacted by them. The involvement of only technical experts and legal professionals is not enough. Rather, it should include persons who have experienced incarceration, community organizations, civil liberties advocates, and representatives of historically over-policed communities. Inclusive design is not only a procedural concern or a box that needs to be checked, but it is also a moral requirement since those who are the most affected by a system should have a say in the making of that system. Responsible AI international standards, including the UNESCO Recommendation on the Ethics of AI, recognize the principle of inclusive governance as the groundwork.[12]

[11]Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 SCC 1 (India).

[12]UNESCO, Recommendation on the Ethics of Artificial Intelligence (2021).

8. CONCLUSION

The question this paper puts forward - Justice or Just Code? - is a question that does not allow for a straightforward answer. AI technologies in criminal justice are not simply tools of evil or good. They are instruments that get their ethical character not so much from how advanced they are technically, but from what values are built into their design, the circumstances of their deployment, and the ways of ensuring that their use is governed. This study has shown that the path that AI deployment in criminal justice currently follows is not ethically sufficient. Tools that are developed on biased data, are not transparent, are deployed without oversight that is meaningful, and have a disproportionate impact on marginalized groups cannot be referred to as tools of justice even if they are technically efficient. Their description in the language of objectivity and science does not make them neutral; it only conceals their injustice and makes it more dangerous to challenge. However, this study also shows that the answer that AI should not be used at all in criminal justice is not what it ends up with. The aim is not to keep the status quo - a human-administered system that is itself deeply biased - but to require that any technological change of justice must comply with the highest ethical standards.

If artificial intelligence (AI) is to help achieve justice rather than simply be a tool of the government, it should be made transparent, regularly checked for fairness, released with caution, and controlled through democratic accountability. In the first place, justice is a moral matter. It involves values that a computer program, no matter how advanced, can never fully understand: compassion, respect for human dignity, appreciation of context, good judgment, and the awareness of our shared humanity. These values are not hindrances to good government - they are the very reason why government exists. A community that delegates its moral decisions to software is not making justice more efficient; it is making it less just. Lawmakers, technologists, ethicists, lawyers, and the public in general need to find another way, one that uses the great potential of artificial intelligence, properly guided by moral principles, to help bring about justice. It is a big and important task, it is a difficult one, and it will be the main legal and moral issue that defines us.

REFERENCES

[1] Ales Završnik, Algorithmic Justice: Algorithms and Big Data in Criminal Justice Settings, 18 Eur. J. Criminology 623 (2021).

[2] Julia Angwin et al., Machine Bias, ProPublica (May 23, 2016), https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

[3] Patrick Grother, Mei Ngan & Kayee Hanaoka, U.S. Dep't of Commerce, Nat'l Inst. of Standards & Tech., NISTIR 8280, Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects (2019).

[4] Shai Danziger, Jonathan Levav & Liora Avnaim-Pesso, Extraneous Factors in Judicial Decisions, 108 Proc. Nat'l Acad. Sci. 6889 (2011).

[5] John Rawls, A Theory of Justice (1971).

[7] Immanuel Kant, Groundwork of the Metaphysics of Morals (Mary Gregor trans., Cambridge Univ. Press 1998) (1785).

[9] State v. Loomis, 881 N.W.2d 749 (Wis. 2016).

[11] Martha C. Nussbaum, Creating Capabilities: The Human Development Approach (2011).

[14] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act), 2024 O.J. (L 1689).

[16] Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 SCC 1 (India).

[17] UNESCO, Recommendation on the Ethics of Artificial Intelligence (2021).