



Artificial Intelligence and the Protection of Fundamental Rights: Challenges to Constitutional Rights in the Era of Automated Decision-Making

1. Asma Rehman

Advocate, Lahore High Court, Lahore; Visiting Faculty Member, Law Department, Punjab University, Gujranwala Campus, Pakistan.

Email: Asmarehman014@gmail.com

2. Gulshan Rehman

Advocate High Court & Tax Consultant, Irfan & Farhad Law Associates, Lahore, Pakistan.

Email: Grehman136@gmail.com

3. Dr. Hafiz Muhammad Abrar Awan

Assistant Professor, Department of Urdu Encyclopedia of Islam, University of the Punjab, Lahore, Pakistan.

Email: abrara.uei@pu.edu.pk

ORCID: <https://orcid.org/0000-0002-5345-8424>

Abstract

Artificial intelligence and automated decision-making systems are increasingly used in public administration, policing, employment, education, healthcare, welfare distribution, immigration, financial services, credit assessment, and other fields in which decisions may substantially affect individuals' rights and opportunities. While these technologies offer significant advantages in efficiency, consistency, prediction, and large-scale data processing, they also create complex constitutional challenges. The central concern is no longer merely whether an automated system functions accurately, but whether its design, data, deployment, and institutional use remain compatible with fundamental rights. This study examines the constitutional implications of artificial intelligence through a comparative legal and human-rights perspective, focusing particularly on equality and non-discrimination, privacy and informational autonomy, human dignity, procedural fairness, the right to reasons, access to an effective remedy, transparency, and meaningful human oversight. It analyses major international and regional regulatory instruments, including the General Data Protection Regulation, the European Union Artificial Intelligence Act, the Council of Europe Framework Convention on Artificial Intelligence, UNESCO's Recommendation on the Ethics of Artificial Intelligence, and the OECD AI Principles, together with selected judicial developments such as SCHUFA, SyRI, and Bridges. The study argues that existing fundamental rights remain fully applicable in the age of automated governance, but their effective protection requires institutional and procedural adaptation. High-impact automated systems must therefore be governed by legality, proportionality, non-discrimination, data protection, explainability, traceability, human supervision, rights-impact assessment, and effective mechanisms of review. The article concludes that technological efficiency cannot justify the displacement of constitutional responsibility and that ultimate legal accountability must remain attributable to human institutions.

Keywords: Artificial Intelligence, Fundamental Rights, Automated Decision-Making, Constitutional Law, Algorithmic Bias, Privacy, Due Process, Human Oversight, AI Governance

Introduction

Artificial intelligence has altered the structure of modern decision-making. Public bodies and private institutions increasingly use algorithmic systems to rank applicants, assess risk, detect suspected fraud, determine eligibility, identify individuals, evaluate creditworthiness, predict



behaviour, recommend medical interventions, allocate resources, and support administrative decisions. The constitutional significance of this transformation lies not simply in the sophistication of the technology but in the growing capacity of computational systems to influence decisions affecting liberty, equality, privacy, reputation, livelihood, access to public services, and other legally protected interests.

A decision once made entirely by an identifiable public official may now result from several interconnected processes: collection of personal data, automated classification, predictive modelling, risk scoring, algorithmic recommendation, and final human approval. In such circumstances, the formal identity of the final decision-maker may reveal little about the actual distribution of decisional power. An official may technically sign a decision while relying almost completely upon an automated output that the official neither created nor fully understands.

This creates a fundamental problem for constitutional law. Traditional constitutional safeguards generally assume that the exercise of public power is identifiable, reasoned, reviewable, and attributable to a responsible institution. Artificial intelligence can complicate each of these assumptions. If an individual does not know that an algorithm affected a decision, cannot understand why the system classified the person in a particular way, cannot inspect or correct the underlying data, and cannot identify who bears responsibility for error, the practical protection offered by constitutional rights may become weaker even though those rights formally remain unchanged.

The issue must nevertheless be approached with balance. Artificial intelligence can also advance legitimate public purposes. It may reduce administrative delay, identify fraud, increase consistency, assist medical diagnosis, improve access to services, and enable authorities to analyse volumes of information that would otherwise be impossible to process. The constitutional question is therefore not whether AI should simply be accepted or prohibited. The more precise question is how legal systems can benefit from automated technologies while preserving the supremacy of fundamental rights, human dignity, democratic accountability, and the rule of law.

This study adopts a comparative constitutional and human-rights approach. It does not assume that all jurisdictions define constitutional rights identically. Instead, it examines a group of rights and principles that recur across constitutional and international legal systems: equality, non-discrimination, dignity, privacy, procedural fairness, transparency, effective remedy, proportionality, and accountable exercise of public power.

Section One: Artificial Intelligence, Automated Decision-Making and the Constitutional Framework of Fundamental Rights

The constitutional analysis of artificial intelligence should begin with conceptual clarity. AI, automated decision-making, algorithmic decision support, and profiling are related but not identical concepts. Their legal effects also vary significantly. A system that recommends entertainment content does not ordinarily raise the same constitutional concerns as one used to determine welfare eligibility, assess criminal risk, identify a person through facial recognition, or rank candidates for public employment. The seriousness of constitutional scrutiny should therefore correspond to the nature, context, and consequences of the automated system.

1. Concept, Forms and Scope of Artificial Intelligence and Automated Decision-Making

Artificial intelligence is a broad technological category rather than a single method. It includes systems using machine learning, statistical modelling, knowledge-based approaches, natural-language processing, computer vision, predictive analytics, and other computational techniques. For constitutional analysis, however, the crucial issue is not merely the technical



architecture of a system but the function it performs and the degree to which its output can influence human decisions.

The OECD definition emphasizes the machine-based and inferential nature of AI:

“An AI system is a machine-based system that ... infers, from the input it receives, how to generate outputs.”¹

The complete definition includes predictions, content, recommendations, and decisions capable of influencing physical or virtual environments. It also recognizes that AI systems may possess different degrees of autonomy and may exhibit varying levels of adaptability after deployment.

This definition is particularly useful for legal analysis because it directs attention toward outputs capable of influencing decisions. The constitutional concern begins to intensify when AI moves from merely processing information to generating classifications, recommendations, predictions, scores, or decisions that materially affect an identifiable person.

A second important distinction concerns fully automated decision-making. Article 22 of the General Data Protection Regulation provides:

“The data subject shall have the right not to be subject to a decision based solely on automated processing.”²

That protection applies, subject to specified qualifications and exceptions, where the decision produces legal effects or similarly significant effects upon the individual.

The distinction between fully automated and AI-assisted decisions is important but should not be treated mechanically. If a human officer simply confirms an algorithmic output without genuine scrutiny, the presence of a human signature may provide little meaningful protection. Consequently, constitutional analysis should investigate actual decisional influence rather than merely formal institutional arrangements.

An AI system may therefore operate at several levels. It may provide information to a human decision-maker; rank available choices; produce a risk score; recommend a decision; establish a presumptive outcome that a human may override; or independently determine an outcome. The closer an automated output comes to determining the final result, the greater the justification for heightened procedural safeguards.

2. Fundamental Rights and the Constitutional Responsibility of the State

Fundamental rights are intended to protect individuals against arbitrary, discriminatory, disproportionate, or otherwise unlawful exercises of power. Their precise wording varies between constitutional systems, yet certain principles are widely recognized in domestic constitutional law and international human-rights law. These include equality, personal liberty, dignity, privacy, procedural fairness, access to justice, freedom of expression, and protection against arbitrary state interference.

The emergence of artificial intelligence does not render these traditional rights obsolete. On the contrary, it creates new circumstances in which established rights must be applied.

¹ Organisation for Economic Co-operation and Development (OECD), *Recommendation of the Council on Artificial Intelligence*, OECD/LEGAL/0449, adopted May 22, 2019, amended May 3, 2024; OECD, *Explanatory Memorandum on the Updated OECD Definition of an AI System*, OECD Artificial Intelligence Papers, no. 8 (Paris: OECD Publishing, 2024), <https://doi.org/10.1787/623da898-en>.

² European Parliament and Council, *Regulation (EU) 2016/679 of 27 April 2016*, General Data Protection Regulation, OJ L 119 (May 4, 2016), art. 22(1).



UNESCO places rights and dignity at the centre of AI governance:

“The protection of human rights and dignity is the cornerstone of the Recommendation.”³

This statement identifies human dignity not as an optional ethical consideration but as a foundational principle governing AI development and deployment. It means that efficiency, prediction, and institutional convenience must remain subordinate to rights protection.

The OECD similarly requires AI actors to:

“respect the rule of law, human rights and democratic values.”⁴

This formulation extends across the AI lifecycle and includes dignity, autonomy, privacy, equality, fairness, diversity, and social justice. It demonstrates that rights protection should begin during system design rather than only after harm has occurred.

The Council of Europe Framework Convention likewise aims to ensure that AI lifecycle activities remain:

“fully consistent with human rights, democracy and the rule of law.”⁵

The Convention was opened for signature on 5 September 2024 and represents the first legally binding international treaty specifically addressing artificial intelligence, human rights, democracy, and the rule of law.

These developments reflect a fundamental constitutional principle: technology does not create a zone in which legal responsibility disappears. If a government delegates a function to software, contracts with a private AI supplier, or relies upon an algorithmic recommendation, the underlying constitutional duties of the public authority do not ordinarily vanish merely because the method of exercising power has changed.

State responsibility therefore possesses both a negative and a positive dimension. The state must itself avoid using AI in ways that unjustifiably interfere with fundamental rights, and it may also be required, depending upon the applicable constitutional system, to establish legal safeguards governing private uses of AI that substantially affect protected rights.

3. Automated Decisions and the Transformation of Constitutional Accountability

The introduction of automated systems into public decision-making changes the practical structure of accountability. Traditional administrative decisions can ordinarily be attributed to an identified office, officer, or institution. AI-assisted decisions may involve software vendors, data providers, model developers, public authorities, human reviewers, and interconnected technical systems. Responsibility can therefore become institutionally fragmented.

The Court of Justice of the European Union emphasized in the SCHUFA litigation that an automatically generated score may fall within the legal framework governing automated

³ UNESCO, “Recommendation on the Ethics of Artificial Intelligence,” official overview of the 2021 Recommendation, under “Human rights and human dignity,” <https://www.unesco.org/en/articles/recommendation-ethics-artificial-intelligence>; UNESCO, *Recommendation on the Ethics of Artificial Intelligence* (Paris: UNESCO, 2021).

⁴ OECD, *Recommendation of the Council on Artificial Intelligence*, Principle 1.2, “Respect for the rule of law, human rights and democratic values, including fairness and privacy.”

⁵ Council of Europe, *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law*, CETS No. 225, opened for signature September 5, 2024.



decision-making where that score plays a decisive role in a later decision concerning an individual.⁶

The importance of this approach lies in its rejection of purely formal analysis. If a bank, employer, or public body technically makes the final decision but routinely relies upon an automated score, constitutional and legal scrutiny should focus upon the real influence of that score.

The Council of Europe's AI framework similarly examines activities throughout the AI lifecycle rather than isolating only the final stage of decision-making. This broader perspective recognizes that risks may originate in data collection, design, training, testing, deployment, or institutional reliance.

The essential constitutional question is therefore not merely, "Who signed the decision?" More searching questions are required: Who designed the system? What data shaped its output? What criteria were embedded in the model? Did the human reviewer possess authority and competence to reject the output? Could the affected individual challenge the input data? Who is legally responsible if the output was discriminatory, inaccurate, or arbitrary?

Automated governance should consequently be analysed as a chain of legally relevant actions rather than as a single technical event.

Section Two: Major Constitutional Challenges Created by Automated Decision-Making

Artificial intelligence can threaten fundamental rights in ways that are different in form, though not necessarily in principle, from conventional administrative conduct. Discrimination can become statistical rather than explicit; surveillance can operate continuously and at scale; decisions may become difficult to explain; and institutional responsibility may be distributed across public and private actors. These characteristics create particular challenges for constitutional guarantees.

1. Equality, Non-Discrimination and Algorithmic Bias

Equality is among the rights most directly implicated by automated decision-making. AI systems learn or infer patterns from data, but data are not socially neutral. Historical datasets may reflect existing inequalities, institutional practices, underrepresentation, policing patterns, hiring preferences, or unequal access to services. If these patterns are reproduced algorithmically, historical discrimination can acquire the appearance of technical objectivity.

The OECD Principles expressly identify:

"non-discrimination and equality"⁷

among the values that AI actors should respect.

UNESCO similarly requires that:

"AI actors should promote social justice, fairness, and non-discrimination."⁸

These principles are particularly significant because algorithmic bias may emerge even without an expressly discriminatory instruction.

Historical bias arises where data encode earlier patterns of unequal treatment. Representation bias occurs where particular communities are insufficiently represented. Proxy discrimination may arise where seemingly neutral variables correlate strongly with protected characteristics.

⁶ Court of Justice of the European Union, *SCHUFA Holding (Scoring)*, Case C-634/21, Judgment of December 7, 2023, ECLI:EU:C:2023:957.

⁷ OECD, *Recommendation of the Council on Artificial Intelligence*, Principle 1.2.

⁸ UNESCO, *Recommendation on the Ethics of Artificial Intelligence*, paras. 28–29; see also UNESCO's official overview under "Fairness and non-discrimination."



Design bias may arise from the selection of targets, thresholds, variables, or institutional definitions of risk.

The constitutional difficulty is therefore wider than intentional discrimination. The absence of discriminatory motive does not necessarily eliminate unequal or unjustified effects.

A rights-based analysis should examine whether differential outcomes are justified, whether error rates vary across groups, whether proxy variables are operating, whether data are sufficiently representative, and whether less discriminatory alternatives are available.

Transparency alone cannot resolve these concerns. A fully transparent system may remain discriminatory. Equality therefore requires substantive evaluation of effects as well as procedural visibility.

2. Privacy, Informational Autonomy and Human Dignity

Artificial intelligence depends heavily upon data. Contemporary systems can combine information collected for different purposes, infer characteristics individuals never directly disclosed, and generate profiles concerning health, behaviour, financial reliability, relationships, preferences, or perceived risk.

UNESCO states:

“Privacy must be protected and promoted throughout the AI lifecycle.”⁹

This principle recognizes that privacy risks arise not only at the moment of data collection but also during analysis, inference, sharing, reuse, and decision-making.

The SyRI litigation in the Netherlands provides an important illustration. The Hague District Court concluded that the legislative framework governing the system failed to satisfy the requirements of Article 8 of the European Convention on Human Rights. The court identified concerns arising from a system that was:

“insufficiently transparent and verifiable.”¹⁰

The case demonstrates that a legitimate objective such as fraud prevention does not automatically justify extensive automated data processing. The method must still satisfy requirements of legality, proportionality, necessity, transparency, and protection of private life.

Biometric surveillance presents an additional concern. In *Bridges v Chief Constable of South Wales Police*, the Court of Appeal examined the:

“lawfulness of the use of live automated facial recognition technology.”¹¹

Facial recognition may permit identification or tracking of persons at scale, raising questions involving privacy, data protection, freedom of movement, proportionality, and potentially freedom of assembly.

Privacy in the AI era therefore extends beyond secrecy. It includes the ability to resist unjustified profiling, disproportionate surveillance, uncontrolled reuse of data, and the construction of predictive identities that individuals cannot meaningfully understand or contest.

⁹ UNESCO, *Recommendation on the Ethics of Artificial Intelligence*, paras. 32–34; see also UNESCO’s official overview under “Right to privacy and data protection.”

¹⁰ District Court of The Hague, *NJCM c.s. v State of the Netherlands (SyRI)*, Judgment of February 5, 2020, ECLI:NL:RBDHA:2020:865; Rechtspraak, “SyRI legislation in breach of European Convention on Human Rights,” February 13, 2020.

¹¹ Court of Appeal of England and Wales, *R (Bridges) v Chief Constable of South Wales Police* [2020] EWCA Civ 1058, Judgment of August 11, 2020.



Human dignity reinforces these concerns. Individuals should not be reduced entirely to statistical categories such as “high risk,” “unreliable,” or “undesirable” where such classifications materially affect their legal rights or social opportunities.

3. Procedural Fairness, the Right to Reasons and Effective Legal Remedy

Procedural fairness may represent one of the most serious constitutional challenges created by automated decision-making. A person cannot effectively challenge an adverse decision if the operative reasons remain inaccessible.

Where an automated score, classification, or prediction affects welfare, employment, credit, immigration, education, or another significant interest, the individual may need to know whether automated processing was used and what role it played.

The GDPR provides, in relevant circumstances:

“the right to obtain human intervention.”¹²

It also protects the opportunity to express a point of view and contest the decision.

The EU Artificial Intelligence Act requires relevant high-risk systems to be:

“sufficiently transparent to enable deployers to interpret a system’s output.”¹³

This requirement is significant because legal accountability depends upon some capacity to understand how an automated output is being used.

The OECD additionally emphasizes the availability of information enabling persons adversely affected by AI to:

“challenge its output.”¹⁴

These principles reveal three related dimensions of procedural fairness.

First, notification is important. Individuals cannot challenge automated influence if they do not know it occurred.

Second, explanation must be meaningful. Merely releasing source code may not enable an ordinary individual to understand a decision. Relevant information may include decisive factors, important data, the role of automation, known limitations, and available avenues of review.

Third, review must be effective. A human reviewer who simply confirms the original output without independent consideration does not provide genuine procedural protection.

Trade secrecy and intellectual property may legitimately protect certain commercial information, but such interests should not automatically defeat judicial or regulatory scrutiny where an automated system materially affects constitutional rights.

Section Three: A Legal and Constitutional Framework for Protecting Rights in the Age of AI

The constitutional response to artificial intelligence must extend beyond general ethical declarations. Rights must be translated into institutional duties, procedural safeguards, independent oversight, and enforceable remedies. The seriousness of regulation should correspond to the seriousness of possible harm.

1. Human Oversight, Transparency and Algorithmic Accountability

Human oversight is frequently proposed as a safeguard, but the mere presence of a human somewhere in the decision-making process is insufficient.

¹² Regulation (EU) 2016/679, art. 22(3).

¹³ European Parliament and Council, *Regulation (EU) 2024/1689 of 13 June 2024* laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L, 2024/1689 (July 12, 2024), art. 13(1).

¹⁴ OECD, *Recommendation of the Council on Artificial Intelligence*, Principle 1.3, “Transparency and explainability.”



The EU AI Act requires appropriate human supervisors of relevant high-risk systems to be capable, in specified circumstances, of:

“disregard[ing], overrid[ing] or revers[ing] the output.”¹⁵

This captures the difference between symbolic and meaningful oversight. A reviewer who lacks power to reject the automated recommendation cannot provide an effective constitutional safeguard.

UNESCO makes the underlying principle even clearer:

“AI systems do not displace ultimate human responsibility and accountability.”¹⁶

This principle should guide the allocation of responsibility across the AI lifecycle.

Meaningful oversight therefore requires authority, competence, information, and time. Human decision-makers must understand relevant limitations and should not treat algorithmic outputs as presumptively infallible.

Transparency also operates at several levels. Institutional transparency concerns whether AI is being used. Individual transparency concerns the role of AI in a particular decision. Technical transparency concerns system performance, relevant data, limitations, and error rates. Governance transparency concerns procurement, auditing, institutional responsibility, and complaint procedures.

The constitutional purpose of transparency is not disclosure for its own sake. Its purpose is to support accountability, contestability, correction, and effective legal review.

2. Risk-Based Regulation and Fundamental-Rights Impact Assessment

Not all AI systems create equal constitutional risks. A recommendation system used for entertainment does not warrant the same degree of regulation as an AI system influencing welfare eligibility, employment, criminal justice, healthcare, or immigration decisions.

The EU AI Act expressly establishes an:

“assessment of the impact on fundamental rights”¹⁷

for certain deployments of high-risk AI systems.

A meaningful fundamental-rights impact assessment should identify the system’s purpose, legal basis, affected groups, relevant rights, possible discriminatory effects, privacy consequences, accuracy limitations, human oversight mechanisms, alternative measures, and available remedies.

The Council of Europe’s AI governance framework similarly adopts a risk-oriented approach by requiring attention to the seriousness and likelihood of adverse effects on human rights, democracy, and the rule of law.

Such assessments should occur before deployment rather than only after widespread harm has occurred. They therefore convert rights protection from a purely reactive model into a preventive one.

A significant recent legal development must also be recognized. Regulation (EU) 2026/1744, adopted on 8 July 2026, amended the EU AI Act as part of the Digital Omnibus on AI and revised aspects of its implementation framework. The consolidated AI Act now reflects that amendment.¹⁸

¹⁵ Regulation (EU) 2024/1689, art. 14(4)(d).

¹⁶ UNESCO, *Recommendation on the Ethics of Artificial Intelligence*, para. 36; see also UNESCO’s official overview under “Human oversight and determination.”

¹⁷ Regulation (EU) 2024/1689, art. 27(1).

¹⁸ European Parliament and Council, *Regulation (EU) 2026/1744 of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the*



This demonstrates that AI regulation itself remains dynamic. Constitutional protection therefore requires continuing legislative, regulatory, technical, and institutional adaptation rather than reliance upon a single static legal instrument.

3. The Role of Legislatures, Courts and Administrative Authorities

Artificial intelligence governance cannot be delegated entirely to developers or private technology vendors. Where automated systems influence constitutionally significant decisions, democratic authorization, administrative accountability, judicial review, and independent supervision become necessary.

The Council of Europe describes its Framework Convention as the:

“first-ever international legally binding treaty in this field.”¹⁹

This characterization is important because it recognizes that AI governance is not simply a matter of voluntary technological ethics. It is increasingly becoming a matter of legal obligation.

Legislatures should establish clear legal authority for high-impact automated decision-making. Legislation should define permissible purposes, prohibited uses, safeguards, oversight requirements, remedies, and institutional responsibility.

Administrative bodies should maintain inventories of consequential AI systems, document their purposes, assess risks, preserve relevant records, monitor performance, and establish accessible complaint procedures.

Courts should examine not only the formal legality of a public institution’s conduct but also the actual influence exercised by automated systems. Judicial review may need to consider whether data were accurate, whether discrimination occurred, whether privacy interference was proportionate, whether human oversight was genuine, and whether effective challenge was possible.

The SyRI and Bridges cases demonstrate that courts can subject sophisticated technological systems to established principles of legality, proportionality, privacy, and administrative accountability.

Independent regulators may provide an additional layer of protection. Data protection authorities, equality bodies, ombuds institutions, auditors, and specialized AI regulators can investigate patterns of harm that may be difficult for individual litigants to identify.

The legal system should therefore distribute duties without permitting responsibility to disappear. Developers, vendors, deployers, officials, and institutions may hold different obligations, but the individual affected by an automated decision should always be able to identify an accountable legal actor.

Findings of the Study

The study establishes that artificial intelligence does not suspend existing constitutional rights. Equality, privacy, dignity, procedural fairness, legality, proportionality, accountability, and judicial review remain applicable when decisions are technologically mediated.

Automated decision-making can generate constitutional injury without malicious intention. Historical data, representation problems, proxy variables, inappropriate optimisation criteria, and excessive institutional reliance upon algorithmic outputs may produce discriminatory or arbitrary effects.

simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), OJ L, 2026/1744 (July 24, 2026).

¹⁹ Council of Europe, *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law*, CETS No. 225; Council of Europe, “The Framework Convention on Artificial Intelligence,” official overview.



The distinction between human and automated decision-making should be substantive rather than merely formal. A nominal human reviewer who mechanically confirms an algorithmic recommendation may not constitute meaningful human oversight.

Transparency is necessary but insufficient. An entirely transparent system may still be unlawful, discriminatory, inaccurate, or disproportionate. Transparency must therefore operate alongside substantive rights protection.

Privacy in the AI context includes not only confidentiality but also informational autonomy, profiling, inference, aggregation, surveillance, and control over consequential uses of personal data.

Procedural fairness becomes particularly important where automated systems are complex or opaque. Individuals require notification, understandable reasons, opportunities to correct inaccurate information, genuine human review, and effective mechanisms for challenging decisions.

Risk-based regulation provides a rational framework because constitutional safeguards should increase with the probability, scale, severity, and irreversibility of potential harm.

Contemporary international and regional instruments reveal significant convergence around human rights, dignity, fairness, privacy, transparency, human oversight, accountability, and effective remedies.

The central constitutional conclusion is that automation may assist lawful institutional decision-making, but it cannot eliminate or obscure human legal responsibility.

Recommendations

1. High-impact automated decision-making by public authorities should rest upon a clear legal basis.
2. Fundamental-rights impact assessments should precede deployment of systems affecting liberty, equality, employment, welfare, healthcare, education, policing, immigration, or other significant interests.
3. Individuals should be informed when automated processing materially influences a consequential decision.
4. Human oversight should be genuine rather than ceremonial.
5. Human reviewers should possess adequate authority, knowledge, time, and access to information necessary to override automated recommendations where appropriate.
6. Public authorities should preserve traceable records of data, outputs, institutional decisions, and human interventions.
7. Automated systems should undergo regular testing for accuracy, discriminatory effects, differential error rates, security vulnerabilities, and performance deterioration.
8. Individuals adversely affected by automated decisions should receive sufficiently understandable reasons to enable effective challenge.
9. Procurement contracts should not permit commercial confidentiality to prevent lawful judicial or regulatory scrutiny.
10. Independent supervisory institutions should possess authority to investigate systemic algorithmic harms.
11. Courts, lawyers, public administrators, and regulators should receive specialized training concerning AI-assisted decision-making.
12. Legal frameworks should provide correction, review, compensation, or other appropriate remedies where unlawful automated decisions cause harm.
13. Automated systems affecting vulnerable populations should be subject to enhanced safeguards and periodic review.



14. Legislatures should distinguish between low-impact and high-risk applications and impose obligations proportionate to their constitutional consequences.

Conclusion

Artificial intelligence is changing not merely the speed of institutional decision-making but the architecture through which legal and administrative power is exercised. Its constitutional significance is particularly acute when automated systems influence employment, welfare, policing, healthcare, education, credit, immigration, or other matters closely related to dignity, liberty, equality, and personal autonomy.

Constitutional law should neither idealize human decision-making nor presume algorithmic systems to be inherently objective. Human officials may be inconsistent and biased, while well-designed technology may improve accuracy and efficiency. At the same time, automated systems may reproduce historic inequalities, amplify errors at scale, create new forms of surveillance, obscure reasons, and fragment institutional responsibility.

The appropriate constitutional response is therefore neither categorical prohibition nor uncritical technological adoption. It is rights-centred governance.

Equality requires examination of actual effects rather than programmer intentions alone. Privacy requires protection against disproportionate collection, inference, profiling, and surveillance. Human dignity requires caution where people are reduced to statistical categories. Procedural fairness requires notice, meaningful reasons, independent reconsideration, and effective remedies. Accountability requires legal responsibility to remain identifiable even when technological systems are technically complex.

International and regional developments increasingly reflect these principles. The GDPR regulates significant solely automated decisions; the European Union Artificial Intelligence Act establishes a risk-based regulatory structure; Regulation (EU) 2026/1744 has subsequently amended aspects of the Act's implementation framework; UNESCO and the OECD emphasize human-centred governance; and the Council of Europe has established a treaty framework directly connecting artificial intelligence with human rights, democracy, and the rule of law.

The decisive constitutional question of the automated age is therefore not simply whether machines may participate in decision-making. The deeper question is whether human institutions will ensure that every consequential exercise of automated power remains lawful, proportionate, reviewable, non-discriminatory, and subordinate to fundamental rights.

Bibliography

- Council of Europe. *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law*. CETS No. 225. Opened for signature September 5, 2024.
- Court of Appeal of England and Wales. *R (Bridges) v Chief Constable of South Wales Police*. [2020] EWCA Civ 1058. Judgment of August 11, 2020.
- Court of Justice of the European Union. *SCHUFA Holding (Scoring)*, Case C-634/21. Judgment of December 7, 2023. ECLI:EU:C:2023:957.
- District Court of The Hague. *NJCM c.s. v State of the Netherlands (SyRI)*. Judgment of February 5, 2020. ECLI:NL:RBDHA:2020:865.
- European Parliament and Council. *Regulation (EU) 2016/679 of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)*. OJ L 119, May 4, 2016.



- European Parliament and Council. *Regulation (EU) 2024/1689 of 13 June 2024* laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). OJ L, 2024/1689, July 12, 2024.
- European Parliament and Council. *Regulation (EU) 2026/1744 of 8 July 2026* amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI). OJ L, 2026/1744, July 24, 2026.
- Organisation for Economic Co-operation and Development (OECD). *Recommendation of the Council on Artificial Intelligence*. OECD/LEGAL/0449. Adopted May 22, 2019; amended May 3, 2024.
- Organisation for Economic Co-operation and Development (OECD). *Explanatory Memorandum on the Updated OECD Definition of an AI System*. OECD Artificial Intelligence Papers, no. 8. Paris: OECD Publishing, 2024. <https://doi.org/10.1787/623da898-en>.
- UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO, 2021.
- UNESCO. "Recommendation on the Ethics of Artificial Intelligence." Official overview of the 2021 Recommendation. <https://www.unesco.org/en/articles/recommendation-ethics-artificial-intelligence>.