

International Journal of Law Research, Education and Social Sciences

Open Access Journal – Copyright © 2026 – ISSN 3048-7501
Editor-in-Chief – Prof. (Dr.) Vageshwari Deswal; Publisher – Sakshi Batham



This is an Open Access article distributed under the terms of the Creative Commons Attribution-Non-Commercial-Share Alike 4.0 International (CC-BY-NC-SA 4.0) License, which permits unrestricted non-commercial use, distribution, and reproduction in any medium provided the original work is properly cited.

The Right to Privacy in the Age of Artificial Intelligence: Constitutional Protection and Emerging Legal Challenges in India

Shreya Taneja^a

^aDepartment of Laws, Panjab University, Chandigarh, India

Received 05 August 2026; Accepted 02 September 2026; Published 05 September 2026

Artificial Intelligence is increasingly involved in activities that require the collection, analysis and use of information relating to individuals. Facial recognition, predictive policing, targeted advertising, healthcare, recruitment and financial assessment are only some examples. AI systems can process vast amounts of information, identify patterns and generate predictions about individuals in ways that traditional technologies could not. While these capabilities offer significant social and economic benefits, they also raise difficult questions about privacy, autonomy and accountability. The Supreme Court's decision in Justice K.S. Puttaswamy (Retd.) v Union of India¹ recognised privacy as a fundamental right and identified dignity, autonomy and informational privacy as important constitutional concerns. The Digital Personal Data Protection Act, 2023 (DPDP Act) and the Digital Personal Data Protection Rules, 2025 (DPDP Rules) have subsequently created an important statutory framework for the protection of digital personal data. However, AI presents challenges that extend beyond conventional data collection. Profiling, algorithmic discrimination, biometric surveillance, automated decision-making, deepfakes and AI-generated inferences raise questions that cannot always be answered through notice and consent alone. This article examines the constitutional foundations of privacy in India and their relevance to AI-driven data processing. It evaluates the DPDP framework, particularly its treatment of consent,

¹ Justice KS Puttaswamy (Retd) and Anr v Union of India and Ors (2017) 10 SCC 1

data processing and accountability, while considering the constitutional implications of State surveillance and algorithmic decision-making. It argues that privacy protection in the AI era must move beyond the question of who possesses personal information and address what can be inferred, predicted and decided about an individual through that information. India therefore requires a rights-based and risk-sensitive approach to AI governance that protects constitutional values without unnecessarily restricting technological innovation.

Keywords: *artificial intelligence, right to privacy, article 21, informational privacy, data protection.*

INTRODUCTION

Artificial Intelligence has moved from being an emerging technology to becoming part of everyday social and economic life. Search engines predict preferences, digital platforms recommend content, financial institutions use automated systems for fraud detection, healthcare providers increasingly use algorithmic tools, and governments are exploring technological solutions for identification and surveillance. There is little doubt that AI can produce substantial benefits. It can assist doctors in analysing medical information, improve public-service delivery, detect fraudulent transactions and process large datasets rapidly. The legal difficulty begins when these same capabilities are used to analyse, classify and predict individuals. AI systems depend heavily on data. Photographs, financial records, location histories, online behaviour, biometric identifiers and other information can be combined to produce conclusions that an individual may never have expressly provided.

This changes the nature of the privacy problem. The question is no longer limited to whether information was collected without permission. It is also whether a person should be subjected to predictions or classifications generated from information that they provided for an entirely different purpose.

For example, an individual may provide information to access a digital service. An AI system may subsequently analyse patterns in that information to predict preferences, behaviour or risk. The person may never have directly disclosed those characteristics. This creates a tension between technological capability and individual autonomy.

India's constitutional framework provides an important foundation for addressing this tension. In *Justice K.S. Puttaswamy (Retd.) v Union of India*, a nine-Judge Bench unanimously recognised privacy as a fundamental right. The Court's reasoning placed dignity, liberty and autonomy at the centre of privacy protection and recognised informational privacy as an important aspect of the right.

The central challenge today is therefore not whether privacy applies to AI. The more difficult question is how constitutional privacy principles should operate when technology can continuously analyse, predict and influence individuals.

CONSTITUTIONAL FOUNDATION OF THE RIGHT TO PRIVACY

The Constitution of India does not expressly contain a separate provision titled the "right to privacy". Nevertheless, the Supreme Court has developed privacy protection through the interpretation of fundamental rights, particularly Articles 14, 19 and 21.

Article 14 protects equality and guards against arbitrary State action.² Article 19 protects important freedoms necessary for individual expression and participation in society, while Article 21 protects life and personal liberty.³ Privacy interacts with all three.

The constitutional position was decisively clarified by the nine-Judge Bench in *Puttaswamy*. The Court rejected the view that privacy was merely a statutory or common-law interest and recognised it as a constitutionally protected fundamental right. Privacy was linked with dignity, individual choice and the ability of a person to control fundamental aspects of life.

For AI governance, informational privacy is particularly significant. Modern technologies allow information from different sources to be aggregated. Information that appears insignificant when viewed separately may become highly revealing when combined with other data.

Informational privacy should therefore not be understood merely as secrecy. It also concerns the individual's ability to exercise meaningful control over how information relating to them is processed and used.

² Constitution of India 1950, art 14; *E P Royappa v State of Tamil Nadu and Anr* (1974) 4 SCC 3

³ *Ibid* arts 19 and 21

Privacy, however, is not absolute. State interference may be constitutionally permissible where it satisfies the requirements developed in privacy jurisprudence, including legality, legitimate State purpose and proportionality. This becomes particularly important when AI is used for surveillance or law-enforcement purposes.

INFORMATIONAL PRIVACY AND AI-GENERATED INFERENCES

AI makes informational privacy more complicated because data processing does not end with collection. AI can identify patterns in apparently ordinary information and generate predictions about an individual.⁴ A system may infer preferences, behaviour or risk from information that does not reveal those characteristics.

This raises an important question: should an individual have a privacy interest in an inference generated about them even where the underlying information was lawfully obtained?

A narrow approach to data protection might focus primarily on the original information. A constitutional understanding of privacy, however, is concerned with dignity and autonomy. If an inference significantly affects a person's life, treating that inference as legally irrelevant simply because it was generated rather than directly collected could create a significant gap in protection.

This does not mean that every AI prediction should be prohibited. Prediction is one of the principal benefits of AI. The stronger argument is that the legal scrutiny should increase with the consequences of the inference. A recommendation about entertainment is fundamentally different from a prediction that affects employment, credit, insurance or access to public services. Privacy law must therefore increasingly consider not only data possession, but also data consequences.

⁴ Sandra Wachter and Brent Mittelstadt, 'A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI' (2019) 2019(2) Columbia Business Law Review 494
<https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3248829> accessed 10 July 2026

THE DIGITAL PERSONAL DATA PROTECTION ACT 2023

The DPDP Act is India's principal legislation governing the processing of digital personal data.⁵ It establishes concepts including the 'Data Principal' and 'Data Fiduciary' and regulates processing through specified statutory requirements.

The Act provides for grounds for processing under Section 4, notice under Section 5 and consent under Section 6.⁶ Section 7 deals with certain legitimate uses. Section 8 establishes general obligations of Data Fiduciaries, while Section 10 imposes additional obligations on Significant Data Fiduciaries.⁷ Sections 11 and 12 provide rights relating to access, correction and erasure, and Section 13 provides for grievance redressal.⁸ These provisions are relevant to AI because AI systems frequently rely on personal data for training, testing, deployment or operation.⁹ However, the DPDP Act should not be treated as a complete AI-governance statute.

Data protection and AI governance overlap, but they address different problems. A data-protection framework primarily regulates the processing of personal data.¹⁰ AI governance may additionally involve questions of discrimination, accuracy, explainability, safety, reliability and the consequences of automated decision-making.¹¹ For example, an algorithm may process information lawfully but still produce systematically discriminatory outcomes.¹² The legality of the original data processing would not, by itself, answer the constitutional or legal question created by the discriminatory result.

THE DIGITAL PERSONAL DATA PROTECTION RULES, 2025

The regulatory framework was supplemented by the Digital Personal Data Protection Rules, 2025.¹³ The Rules were notified by the Ministry of Electronics and Information Technology in

⁵ Digital Personal Data Protection Act 2023

⁶ *Ibid* ss 4–6

⁷ *Ibid* ss 7–10

⁸ *Ibid* ss 11–13

⁹ OECD, 'Recommendation of the Council on Artificial Intelligence' (22 May 2019) OECD/LEGAL/0449

¹⁰ EU General Data Protection Regulation 2016, art 1

¹¹ EU Artificial Intelligence Act 2024, arts 9, 10, 13 and 14

¹² Solon Barocas and Andrew D Selbst, 'Big Data's Disparate Impact' (2016) 104 California Law Review 671 <<https://www.cs.yale.edu/homes/jf/BarocasSelbst.pdf>> accessed 10 July 2026

¹³ Digital Personal Data Protection Rules 2025

November 2025 with the Gazette notification dated 13 November 2025. Their provisions have a phased commencement rather than becoming operative entirely on one date.

The Rules provide greater operational detail to the statutory framework. Among other matters, they address requirements concerning notices, security safeguards and other implementation mechanisms.

The phased commencement is legally important. The DPDP Act itself also provides for different commencement dates for different provisions, and the November 2025 notification specifies the relevant phases.

The Rules are relevant to AI where AI systems process digital personal data. However, compliance with data-protection requirements does not automatically establish that an AI system is fair, accurate or constitutionally permissible.

Suppose an AI system lawfully processes personal data but incorrectly classifies a person as a high-risk customer. The resulting problem may involve not only privacy but also accuracy, fairness, transparency and access to an effective remedy. AI governance must therefore be viewed as broader than data protection alone.

THE LIMITATIONS OF CONSENT

Consent is an important element of privacy protection, but AI demonstrates its limitations. In theory, an individual can decide whether to provide personal information. In practice, digital services often involve lengthy and technical privacy notices that users cannot realistically understand in full. This creates a distinction between formal consent and meaningful consent.¹⁴ A user may agree to processing without appreciating that information could later be combined with other datasets or analysed to generate behavioural predictions.

The problem becomes more serious where the service is practically necessary. If refusing data processing means losing access to an important financial, educational or digital service, the individual's choice may be more theoretical than genuine.

¹⁴ Daniel J Solove, 'Privacy Self-Management and the Consent Dilemma' (2013) 126 Harvard Law Review 1880 <https://scholarship.law.gwu.edu/cgi/viewcontent.cgi?article=2093&context=faculty_publications&httpsredir=1&referer> accessed 10 July 2026

AI also creates a temporal difficulty. At the time data is collected, organisations may not know every future use to which machine-learning systems could put it. Models may evolve, datasets may be combined, and new patterns may be discovered.

Consent should therefore not become a substitute for responsible data governance. Principles such as purpose limitation, security, proportionality and accountability remain important. Where AI processing has significant consequences, individuals should also have meaningful mechanisms to understand and challenge its effects.

AI PROFILING AND INDIVIDUAL AUTONOMY

AI-powered profiling is one of the clearest examples of the changing privacy landscape. Profiling involves analysing information to evaluate or predict aspects of an individual's behaviour, preferences or characteristics. Some profiling is relatively harmless, such as recommending films or music. The situation changes when profiling influences significant decisions. An algorithm may classify a person as a credit risk, potential fraudster or unsuitable candidate. If such classifications affect employment, financial services or access to public benefits, the consequences can be serious. The person may not even know that the classification exists. This raises a deeper concern about autonomy. Privacy is not simply about preventing information from reaching another person. It also protects the space in which individuals can make choices without being constantly observed, predicted and classified. Continuous profiling may therefore affect the individual's ability to control how they are represented in digital environments.

The law should consequently pay greater attention to high-impact profiling, particularly where algorithmic classifications have significant effects on rights or opportunities.¹⁵

ALGORITHMIC BIAS AND ARTICLE 14

The assumption that algorithms are automatically objective is incorrect. AI systems learn from datasets. Where datasets reflect historical inequalities or discriminatory patterns, an AI system may reproduce or amplify those patterns. This creates a potential constitutional concern under Article 14, especially where AI is used by public authorities.¹⁶

¹⁵ EU Artificial Intelligence Act 2024, arts 6 and 9

¹⁶ Constitution of India 1950, art 14; *Shayara Bano v Union of India and Ors* (2017) 9 SCC 1

A government decision cannot escape constitutional scrutiny simply because it was generated with the assistance of software rather than made entirely by a human officer.¹⁷ If an automated system produces arbitrary or discriminatory outcomes, responsibility cannot disappear into the technology.

AI systems may involve several actors: developers create models, vendors provide software, data providers supply information and government departments or private organisations deploy the systems. The individual affected by the final decision, however, experiences only one outcome.

Technological complexity should not therefore create an accountability vacuum. Where AI is used for high-impact public decision-making, authorities should be capable of demonstrating that appropriate safeguards exist against serious error, discrimination and arbitrary outcomes.

FACIAL RECOGNITION AND BIOMETRIC PRIVACY

Facial recognition is among the most sensitive applications of AI. It can transform a person's facial characteristics into data capable of identifying or comparing that person against a database. Unlike passwords, biometric characteristics cannot simply be replaced once compromised.

Large-scale facial recognition can therefore create risks of continuous identification and surveillance. The State may have legitimate reasons to use facial recognition, including crime investigation or identification of missing persons. But a legitimate objective does not eliminate constitutional limits. In cases involving State surveillance, the interference must be legally authorised and must satisfy constitutional requirements of reasonableness and proportionality.

Facial-recognition systems should therefore be accompanied by safeguards relating to legal authority, purpose, retention, access controls, accuracy and mechanisms for challenging erroneous identification. False positives are particularly important. An innocent person incorrectly identified as a suspect may face consequences for liberty, reputation and personal security.

¹⁷ *State of West Bengal v Anwar Ali Sarkar* AIR 1952 SC 7

AI-ENABLED SURVEILLANCE AND ARTICLE 21

AI has the potential to transform conventional surveillance. CCTV footage, biometric information, location data and other digital records can be analysed together to identify patterns that would previously have required extensive human effort. This raises serious concerns about mass surveillance. In *People's Union for Civil Liberties v Union of India*,¹⁸ the Supreme Court treated telephone tapping as a serious invasion of privacy and emphasised the need for safeguards against abuse. The Court recognised that surveillance powers could not be exercised without satisfying the statutory conditions governing interception.

The principle becomes even more significant when AI dramatically increases the State's capacity to analyse information. A surveillance system capable of identifying individuals across different places and times may affect not only privacy but also freedom of expression and association. People may alter lawful behaviour if they believe that their movements, relationships or activities are constantly capable of being monitored. The constitutional issue is therefore not limited to data collection. It concerns how surveillance technology affects the practical enjoyment of fundamental freedoms.

AUTOMATED DECISION-MAKING AND FAIR PROCESS

AI may increasingly influence decisions concerning employment, financial services, insurance, education and public benefits. The central problem is accountability.

Suppose an individual is denied a service because an automated system classifies them as high-risk. If the person receives no meaningful explanation, it becomes difficult to determine whether the decision resulted from accurate information or an erroneous algorithmic assumption. This creates a procedural fairness concern. Meaningful transparency does not necessarily require disclosure of proprietary source code. It may instead require notification that AI materially contributed to a consequential decision, information about relevant factors, an opportunity to correct inaccurate personal information and human review where appropriate.

The underlying principle is straightforward: automation should change the method of decision-making, not eliminate responsibility for the decision.

¹⁸ *People's Union for Civil Liberties v Union of India and Ors* (1997) 1 SCC 301

This is particularly important when AI is used by public authorities. Administrative power cannot become immune from legal scrutiny merely because its exercise is technologically assisted.

THE PRIVACY OF AI-GENERATED INFERENCES

One of the most difficult emerging issues concerns information that an individual never directly supplied. AI can generate inferences from ordinary information. Browsing behaviour, purchasing patterns and location data may collectively produce predictions about preferences or behaviour. The legal question is whether the individual should have protection when such an inference materially affects them. A constitutional approach suggests that the consequences of the inference matter. If an AI-generated prediction influences employment, credit, insurance or access to a public service, the individual should have stronger grounds to demand accuracy, transparency and accountability. This does not require treating every prediction as legally equivalent to sensitive personal information. Instead, India could develop a risk-sensitive approach in which greater legal safeguards apply to more consequential uses of inferred information. Such an approach would preserve the benefits of predictive technologies while recognising that predictions can themselves affect autonomy and opportunity.

DEEPFAKES, IDENTITY AND PERSONALITY

Generative AI has introduced another category of privacy-related harm through deepfakes and synthetic media. Deepfakes can create realistic images, audio or videos depicting individuals in events that never occurred. The consequences may involve privacy, dignity, reputation, identity and personality interests. The speed at which synthetic content can be produced and redistributed makes conventional remedies difficult. Harmful content may spread widely before an affected individual can obtain effective relief. At the same time, regulation must respect legitimate expression.¹⁹ Satire, parody, artistic works and political commentary may involve synthetic media without necessarily creating unlawful harm. The legal challenge is therefore to distinguish legitimate creative or expressive uses from deceptive and harmful uses of another person's identity.

¹⁹ *Shreya Singhal v Union of India* (2015) 5 SCC 1

PRIVATE AI AND STATE AI

The constitutional implications of AI differ depending on who deploys the technology. Private entities operate primarily within statutory, contractual and regulatory frameworks. State authorities, by contrast, are directly constrained by fundamental rights. The distinction is particularly significant in areas such as surveillance, policing and welfare administration.

Where a private platform uses AI to recommend content, concerns may primarily involve data protection and consumer protection. Where a government authority uses AI to identify suspects or determine eligibility for a public benefit, Articles 14, 19 and 21 may become directly relevant. The State should not be able to avoid constitutional responsibility by outsourcing technological development to a private company.

If a public authority deploys an algorithm that produces arbitrary or discriminatory outcomes, the fact that the software was developed by a private contractor should not by itself prevent constitutional scrutiny. Technological outsourcing cannot become constitutional outsourcing.

THE BLACK-BOX PROBLEM

Some advanced AI systems are difficult to explain even to their developers. This creates a particular legal problem when their outputs affect important interests. A technologically sophisticated system cannot be treated as beyond accountability simply because its internal operation is complex.

High-impact AI systems should therefore be subject to stronger safeguards, potentially including algorithmic impact assessments, documentation of system objectives, testing of training data, bias and accuracy assessments, human oversight, periodic auditing and clear allocation of responsibility.

The regulatory burden should, however, depend upon the level of risk. An AI system recommending music should not be regulated in the same manner as a system used in criminal investigation, employment or access to essential public services. A risk-based approach can therefore protect fundamental rights without unnecessarily restricting beneficial innovation.

INDIA'S REGULATORY CHALLENGE

India now has important components of a privacy framework: constitutional protection, the DPDP Act and the DPDP Rules. Nevertheless, AI reveals several areas requiring further development.

1. Data protection does not automatically constitute comprehensive AI governance.
2. Questions concerning consequential automated decision-making remain broader than the basic processing of personal data.
3. Algorithmic discrimination requires particular attention where public power is involved.
4. Biometric surveillance raises risks that may require safeguards beyond ordinary data-processing obligations.
5. AI-generated inferences create questions about how the law should treat conclusions derived from personal information.
6. Effective enforcement requires regulators and institutions with sufficient technical expertise to understand and audit complex AI systems.

The appropriate response is not to discard the existing framework but to make it stronger.

COMPARATIVE PERSPECTIVE: THE EUROPEAN UNION

International developments demonstrate a growing movement towards risk-based AI regulation. The European Union's Artificial Intelligence Act adopts a risk-based framework and imposes stronger obligations on high-risk AI systems. Article 6 identifies categories of high-risk systems, while Article 9 requires a risk-management system covering the lifecycle of high-risk AI.²⁰ India need not reproduce the European model. India's constitutional structure, institutional capacity, digital economy and regulatory priorities are different.

Nevertheless, certain principles have broader relevance: risk assessment, transparency, human oversight, accountability and protection against discriminatory outcomes. The comparative experience therefore provides useful guidance without requiring India to adopt a foreign regulatory model wholesale. AI regulation should ultimately be understood not only as a question of technological innovation but also as a question of rights and institutional power.

²⁰ EU Artificial Intelligence Act 2024, arts 6 and 9

RECOMMENDATIONS AND THE WAY FORWARD

India's approach should seek to protect constitutional rights while allowing responsible technological development.

Adopt a Risk-Based Approach: AI systems should be regulated according to their potential impact. High-impact systems should face greater scrutiny than low-risk applications.

Introduce Algorithmic Impact Assessments: Deployers of high-impact AI should assess possible effects on privacy, equality and other fundamental rights before deployment.

Strengthen Transparency: Individuals should be informed when AI materially contributes to consequential decisions concerning them.

Ensure Meaningful Human Oversight: Important decisions affecting liberty, employment, financial access or public benefits should remain subject to meaningful review.

Strengthen Biometric Safeguards: Facial recognition and other biometric technologies should operate within clearly defined legal boundaries, with appropriate oversight and retention limits.

Address Algorithmic Discrimination: High-impact systems should be tested for discriminatory patterns and unreasonable outcomes.

Recognise Consequential Inferences: The law should pay greater attention to significant conclusions generated from personal information where those conclusions affect rights or opportunities.

Improve Institutional Expertise: Regulatory bodies should develop the technical capacity necessary to evaluate and audit AI systems.

Provide Effective Remedies: Individuals should have practical mechanisms to challenge inaccurate information, unlawful processing and harmful automated decisions.

Preserve Innovation: Regulation should focus on disproportionate risks rather than treating AI technology itself as inherently harmful.

The objective should be responsible innovation within constitutional boundaries.

CONCLUSION

Artificial Intelligence has changed both the scale and the nature of the privacy problem. Traditional privacy concerns often centred on secrecy, physical intrusion and control over personal information. AI adds another dimension: information can now be continuously aggregated, analysed and transformed into predictions about individuals.

India already possesses a strong constitutional foundation. Justice K.S. Puttaswamy (Retd.) v Union of India²¹ established privacy as a fundamental right connected with dignity, liberty and autonomy. Its recognition of informational privacy provides an important basis for addressing AI-driven processing of personal information.

The DPDP Act, 2023 and the DPDP Rules, 2025 are significant developments, but they should not be treated as a complete answer to every AI-related legal problem. Data protection addresses an important part of the problem, while AI also raises questions concerning discrimination, transparency, accuracy, surveillance and automated decision-making.

The next generation of privacy disputes may therefore concern not merely what information is collected, but what AI systems infer, how they classify individuals and what consequences follow from those classifications.

Privacy in the AI era cannot be reduced to whether an individual clicked "accept". The constitutional objective must be broader: individuals should retain meaningful protection over their identity, information and autonomy even when decisions concerning them are increasingly assisted by machines.

AI can improve healthcare, public administration, financial services and countless other areas of life. Technological efficiency, however, cannot become a justification for sacrificing constitutional rights. The central principle should remain simple: machines may process information about people, but the law must continue to protect the people behind that information.

²¹ *Justice KS Puttaswamy (Retd) and Anr v Union of India and Ors* (2017) 10 SCC 1

The future of privacy in India will not depend upon stopping technological progress. It will depend upon ensuring that technological progress remains consistent with dignity, liberty, equality and constitutional accountability.