International Journal of Creative and Open Research in Engineering and Management
ISSN: 3108-1754 (Online)
Volume 02 Issue 09 September-2026 | Impact Factor: 3.5

Deepfakes and the Indian Constitution: Rethinking Privacy, Reputation and Freedom of Expression in the Generative AI Era

Author: **VATTEM MURALI KRISHNA**
Affiliation: Aurora University Email: muralikrishna1968@yahoo.com

How to Cite this Article:

KRISHNA, V. M. (2026). Deepfakes and the Indian Constitution: Rethinking Privacy, Reputation and Freedom of Expression in the Generative AI Era. International Journal of Creative and Open Research in Engineering and Management, 02(9), 1-9. https://doi.org/10.55041/ijcope.v2i9.158

https://doi.org/10.55041/ijcope.v2i9.158

Abstract

Generative artificial intelligence has transformed the production and circulation of audio-visual content. Among its most controversial applications are deepfakes—synthetically generated or manipulated images, videos, voices and other media capable of making individuals appear to say or do things that never occurred. Although synthetic media can have legitimate applications in education, entertainment, satire, political communication and artistic expression, malicious deepfakes may seriously interfere with individual privacy, dignity, reputation, autonomy and democratic discourse.

India does not presently regulate deepfakes through a single comprehensive statute. Instead, the legal response is distributed across the Constitution of India, the Information Technology Act 2000, the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules 2021, the Bharatiya Nyaya Sanhita 2023 and the Digital Personal Data Protection Act 2023, together with judicially developed principles concerning privacy, dignity, reputation and freedom of expression. The regulatory position has also evolved significantly with the 2026 amendments to the Information Technology Rules concerning synthetically generated information.

This article critically examines whether India's existing constitutional and statutory framework is capable of addressing deepfake harms without creating disproportionate restrictions on legitimate expression. It argues that deepfakes require a rights-based regulatory approach rather than a technology-specific prohibition. The article proposes a constitutional model based on legality, legitimate aim, necessity, proportionality, procedural safeguards, transparency, rapid remedies and platform accountability. It further argues that the legal distinction should not be between 'real' and 'synthetic' content alone, but between lawful synthetic expression and synthetic content that causes identifiable and legally cognizable harm.

Keywords: Deepfakes, Generative AI, Article 19, Article 21, Privacy, Reputation, Freedom of Expression, Artificial Intelligence, Intermediary Liability, Digital Rights.

1. Introduction

Artificial intelligence has progressively changed the manner in which information is created, communicated and consumed. Generative AI systems can now produce realistic images, voices, videos and text with limited technical expertise. Deepfake technology represents one of the most significant consequences of this development because it can manipulate a person's identity and create apparently authentic representations of conduct that never occurred.

The legal significance of deepfakes lies not merely in their falsity. A fabricated image or video may affect several legally protected interests simultaneously. A sexually explicit deepfake may intrude upon privacy and dignity; a fabricated statement by a political candidate may influence electoral discourse; a manipulated video of a professional may damage reputation; and an impersonation may facilitate fraud or identity theft.

The constitutional difficulty is therefore two-sided. On one side stands the individual's right to privacy, dignity, autonomy and reputation. On the other stands the constitutionally protected freedom of speech and expression under Article 19(1)(a). A regulatory framework that removes all synthetic content merely because it is artificial could suppress satire, parody, artistic creation, political criticism and other legitimate forms of expression.

The Indian constitutional framework consequently requires a careful balancing exercise rather than an absolute preference for either privacy or expression.

The Supreme Court's decision in *Justice KS Puttaswamy (Retd) v Union of India* recognised privacy as a constitutionally protected right connected with dignity, liberty and autonomy.[1] The judgment significantly altered the constitutional understanding of informational and decisional privacy in India. The Court subsequently emphasised that an interference with privacy must satisfy legality, a legitimate State aim and proportionality.[2]

At the same time, freedom of speech and expression remains a foundational democratic right. Article 19(1)(a) protects expression, while Article 19(2) permits only constitutionally recognised reasonable restrictions. The challenge posed by deepfakes is therefore not simply whether harmful synthetic content should be removed, but whether the legal mechanism used to remove it is sufficiently precise, necessary and proportionate.

This article examines that constitutional problem.

2. Understanding Deepfakes in the Generative AI Era

The term 'deepfake' generally refers to digitally manipulated or synthetically generated media in which artificial intelligence techniques are used to imitate or replace a person's face, body, voice or other characteristics.

Deepfakes may include:

1. face-swapping videos;
2. AI-generated photographs;
3. voice cloning;
4. synthetic speech;
5. manipulated political speeches;
6. fabricated interviews;
7. sexually explicit synthetic images;
8. impersonation of public figures; and
9. composite audio-visual material combining genuine and synthetic content.

The technology itself is not inherently unlawful. A filmmaker using synthetic media with appropriate consent and disclosure does not necessarily create a legal wrong. Similarly, parody or satire may involve deliberate alteration of a person's image without being defamatory or otherwise unlawful.

The legal problem emerges when synthetic content is used to deceive, exploit, impersonate, sexually abuse, defame, threaten, harass or manipulate individuals or public institutions. This distinction is important because a technology-based prohibition may be constitutionally overbroad. The appropriate legal question should therefore be:

What harm does the synthetic content cause, to whom, by what means, and with what degree of intent or recklessness?

[1] . *Justice KS Puttaswamy (Retd) v Union of India*, (2017) 10 SCC 1.

[2] . *Justice KS Puttaswamy (Retd) v Union of India*, (2017) 10 SCC 1, especially the discussion concerning legality, legitimate State aim and proportionality.

3. Constitutional Foundations

3.1 Article 21: Privacy, Dignity and Autonomy

Article 21 provides that no person shall be deprived of life or personal liberty except according to procedure established by law.

The modern interpretation of Article 21 extends substantially beyond physical existence. In *Maneka Gandhi v Union of India*, the Supreme Court established that the procedure affecting life and personal liberty must satisfy standards of fairness, reasonableness and non-arbitrariness.[3]

The constitutional recognition of privacy reached a decisive stage in *Puttaswamy*. The nine-judge Bench recognised privacy as a fundamental right and connected it with dignity, autonomy and liberty.[4]

Deepfakes directly engage this constitutional framework because an individual's identity can itself become the raw material for synthetic manipulation.

A person's face, voice, gestures and personal characteristics are increasingly capable of being converted into digital data and reproduced through AI systems. When such characteristics are used without consent in circumstances causing substantial interference with personal autonomy or dignity, the constitutional concept of privacy becomes directly relevant.

Privacy in the deepfake context therefore has at least three dimensions:

(a) Informational privacy

This concerns control over personal information, including images, biometric characteristics and voice-related data.

(b) Decisional autonomy

A person should retain meaningful control over how their identity is represented and used.

(c) Privacy of personality

The misuse of one's image, voice or likeness may interfere with the individual's ability to control personal identity and dignity.

Deepfake regulation must therefore recognise that privacy is not limited to secrecy. It may also concern control over one's identity and personal representation.

3.2 Article 19(1)(a): Freedom of Speech and Expression

Article 19(1)(a) guarantees freedom of speech and expression.

Digital communication falls within the contemporary understanding of expression. In *Shreya Singhal v Union of India*, the Supreme Court struck down s 66A of the Information Technology Act 2000 because the provision imposed an unconstitutional restriction on speech.[5]

[3] . *Maneka Gandhi v Union of India*, (1978) 1 SCC 248.
[4] . *Justice KS Puttaswamy (Retd) v Union of India*, (2017) 10 SCC 1.
[5] . *Shreya Singhal v Union of India*, (2015) 5 SCC 1.

The decision demonstrates an important principle for AI regulation: the State cannot suppress expression merely because the expression is capable of causing discomfort, annoyance or disagreement. Restrictions must fall within the constitutional framework of Article 19(2).

Deepfake regulation therefore creates a difficult question.

A manipulated video may be:

- malicious impersonation;
- political satire;
- parody;
- artistic expression;
- journalistic illustration;
- fictional entertainment; or
- deliberate disinformation.

Treating all such material identically would ignore the constitutional value of context.

4. Reputation as a Constitutional Interest

Reputation occupies an important position within Indian constitutional jurisprudence.

In *Subramanian Swamy v Union of India*, the Supreme Court upheld the constitutional validity of criminal defamation and recognised reputation as an important aspect of an individual's dignity.[6]

The constitutional tension between reputation and free speech is especially visible in deepfake cases.

Consider a fabricated video showing a public official accepting a bribe. If the video is presented as authentic, it may seriously damage reputation and public trust. At the same time, the mere fact that the video is synthetic cannot automatically make it unlawful because satire and political criticism may also involve manipulated imagery.

The legal analysis should therefore examine:

1. whether the representation is presented as fact;
2. whether the subject is identifiable;
3. whether the content is substantially false;
4. whether the publisher knew or reasonably should have known of the falsity;
5. whether actual or legally cognisable reputational harm resulted;
6. whether the content falls within a recognised exception or legitimate expressive purpose; and
7. whether the restriction imposed is proportionate.

The central constitutional problem is consequently not 'deepfake versus truth' but **harmful deception versus protected expression**.

5. Existing Indian Statutory Framework

India currently addresses deepfake-related harms through several overlapping legal instruments rather than one comprehensive AI statute.

[6] . *Subramanian Swamy v Union of India*, (2016) 7 SCC 221.

5.1 Information Technology Act 2000

The Information Technology Act contains several provisions potentially applicable to deepfake misuse.

Section 66C addresses identity theft. Section 66D addresses cheating by personation using a computer resource or communication device. Section 66E addresses violation of privacy. Sections 67 and 67A regulate obscene and sexually explicit material in electronic form. Section 69A provides a framework for blocking public access to specified information, while section 79 establishes the statutory framework concerning intermediary liability and safe harbour subject to specified conditions.[7]

These provisions can be useful where deepfakes involve impersonation, privacy violations, sexual content or other unlawful conduct.

However, they were not designed specifically for generative AI. The absence of a dedicated deepfake offence creates difficulties in cases where a synthetic representation causes serious reputational or psychological harm but does not neatly fall within existing offences.

5.2 Bharatiya Nyaya Sanhita 2023

The Bharatiya Nyaya Sanhita 2023 has become relevant to synthetic media offences.

Section 356 deals with defamation and applies to imputations made through words, signs or visible representations.[8] This language is significant because a manipulated image or video may constitute a visible representation capable of affecting reputation.

The BNS also contains provisions dealing with cheating by personation and forgery-related conduct that may become relevant depending upon the manner in which synthetic material is created and used.

Section 353 may also become relevant where false or misleading statements, rumours or reports fall within its statutory conditions.

Nevertheless, these provisions generally address the consequences or surrounding conduct rather than the technological phenomenon of generative AI itself.

6. Information Technology Rules and the 2026 Regulatory Development

The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules 2021 provide an important regulatory layer.

The Rules impose due-diligence obligations upon intermediaries and require them to address categories of unlawful or prohibited information. The regulatory framework underwent an important development in February 2026 with amendments relating to **synthetically generated information**.

The updated framework introduces stronger obligations concerning identification and traceability of permissible synthetic content and strengthens intermediary responsibilities concerning unlawful AI-generated material.[9] This development represents an important movement from a largely reactive approach towards a more structured platform-governance model.

[7] . Information Technology Act 2000, ss 66C, 66D, 66E, 67, 67A, 69A and 79.
[8] . Bharatiya Nyaya Sanhita 2023, s 356.
[9] . Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules 2021, as amended on 10 February 2026.

The Government has expressly identified deepfakes and AI-generated audio, video and text as emerging harms affecting privacy, dignity and reputation.[10] However, increased intermediary obligations also raise constitutional concerns. If platforms are encouraged to remove content rapidly to avoid liability, there is a risk of excessive removal of lawful speech. Automated systems may not reliably distinguish between:

- malicious deepfakes;
- parody;
- satire;
- legitimate political commentary;
- fictional content;
- educational demonstrations; and
- journalistic investigation.

Consequently, speed of removal must be accompanied by procedural safeguards.

7. Deepfakes and the Right to Privacy

The relationship between deepfakes and privacy is particularly strong where an individual's face, voice or personal information is used without consent. A deepfake may convert an individual's identity into a form of digital raw material. This raises a question that traditional privacy law did not have to confront in the same manner:

Does an individual have a constitutional interest in controlling the synthetic reproduction of their identity?

The answer should generally be yes, particularly where the use is intrusive, deceptive or harmful.

The constitutional foundation can be developed from the combined principles of privacy, dignity and autonomy recognised in *Puttaswamy*. However, a right to control identity cannot become an absolute right to prevent all representations of a person.

Public figures, politicians, actors and other individuals regularly appear in news reporting, criticism, parody and artistic works. Constitutional protection must therefore distinguish between legitimate reference to a person's identity and exploitative synthetic impersonation. A proportionality-based approach provides a workable method.

8. Deepfakes and Reputation

Deepfakes create an unusual reputational problem because the evidence may appear visually convincing even when it is entirely false.

Traditional defamation law assumes that a defamatory statement can be identified and assessed. AI-generated media complicates this assumption because the false representation may consist of a synthetic event rather than a conventional verbal statement.

For example, a deepfake may depict:

- a doctor accepting a bribe;
- a lawyer making an offensive statement;
- a businessperson engaging in criminal conduct;
- a politician making inflammatory remarks; or

[10] . Ministry of Electronics and Information Technology, Government of India, 'Government Strengthens Regulatory Framework to Address AI-Generated Deepfakes' (6 August 2026).

- a private individual appearing in sexually explicit material.

In each situation, the synthetic representation may lower the person in the estimation of others.

Section 356 of the BNS is therefore potentially significant. Its definition of defamation expressly includes visible representations.[11] Nevertheless, the law should distinguish between factual deception and protected parody. A satirical video clearly presented as satire should not automatically be treated in the same manner as a deliberately fabricated video presented as authentic evidence.

9. Deepfakes and Freedom of Expression

The strongest constitutional objection to aggressive deepfake regulation is the possibility of censorship.
Generative AI can be used for:

- political satire;
- artistic expression;
- film-making;
- historical reconstruction;
- education;
- research;
- commentary;
- parody; and
- creative experimentation.

A law that prohibits synthetic content merely because it is false could have serious implications for Article 19(1)(a).

Indian constitutional law already recognises that restrictions on speech must satisfy Article 19(2). In *Shreya Singhal*, the Supreme Court distinguished protected discussion and advocacy from speech crossing legally permissible limits.[12]

The same conceptual distinction should apply to AI-generated media.

Protected synthetic expression

A fictional or satirical AI video clearly identified as such may ordinarily receive stronger protection.

Harmful synthetic deception

A fabricated video deliberately presented as authentic to deceive voters, destroy a person's reputation, facilitate fraud or sexually exploit an individual should receive substantially weaker protection. This distinction is central to constitutionally sustainable regulation.

10. Intermediary Liability and Platform Responsibility

Deepfakes are often distributed through social-media platforms rather than directly through traditional publishers.

Platforms therefore occupy a critical position in the regulatory ecosystem.

The challenge is to impose meaningful responsibility without transforming platforms into private censorship authorities.

[11] . Bharatiya Nyaya Sanhita 2023, s 356(1).
[12] . *Shreya Singhal v Union of India*, (2015) 5 SCC 1.

A balanced framework should require platforms to provide:

1. accessible reporting mechanisms;
2. rapid handling of high-risk deepfake complaints;
3. clear labelling of permissible synthetic media;
4. mechanisms for preserving evidence;
5. transparent content-removal policies;
6. appeal mechanisms;
7. protection against repeated re-uploading of unlawful material; and
8. appropriate cooperation with lawful investigations.

The 2026 amendments to the IT Rules move Indian regulation in this direction by strengthening requirements around synthetic information, labelling and traceability.[13]

However, platform responsibility should not eliminate judicial oversight.

11. The Role of the Digital Personal Data Protection Act 2023

The Digital Personal Data Protection Act 2023 creates another layer of relevance because AI systems may process personal data in generating synthetic content.

The Act is designed to regulate processing of digital personal data and recognises the individual's interest in protecting personal data.[14]

Where a deepfake system uses identifiable personal information, images, recordings or other personal data, questions concerning lawful processing, consent and data governance may arise.

Nevertheless, the DPDP Act should not be treated as a complete deepfake statute.

Its principal concern is the processing of digital personal data, whereas deepfake harms may also involve reputation, fraud, sexual exploitation, political manipulation and public misinformation.

The statutes therefore operate as complementary rather than interchangeable instruments.

12. The Constitutional Proportionality Test

The most appropriate constitutional mechanism for deepfake regulation is proportionality.

The *Puttaswamy* jurisprudence requires restrictions on privacy to satisfy legality, legitimate State aim and proportionality.[15]

Applied to deepfakes, a regulatory measure should satisfy at least four questions.

First: Is there a valid legal basis?

Restrictions should be grounded in law rather than uncertain executive directions or vague platform policies.

[13] . Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules 2021, as amended on 10 February 2026.
[14] . Digital Personal Data Protection Act 2023, ss 4–8 and related provisions.
[15] . *Justice KS Puttaswamy (Retd) v Union of India*, (2017) 10 SCC 1.

Second: Is the objective legitimate?

Protecting privacy, dignity, reputation, public order, electoral integrity and individuals from fraud are legitimate objectives.

Third: Is the measure necessary?

The State should demonstrate why the particular restriction is required and whether less restrictive alternatives are available.

Fourth: Is the measure proportionate?

The benefit achieved must justify the burden imposed upon freedom of expression.

This framework prevents the State from adopting an unnecessarily broad prohibition.

13. The Need for a Harm-Based Deepfake Framework

India would benefit from moving towards a **harm-based rather than technology-based regulatory model**.

The proposed framework could classify deepfakes into four broad categories.

Category I – Low-risk synthetic content

Examples:

- entertainment;
- artistic experimentation;
- clearly fictional material;
- educational demonstrations.

Regulatory approach: disclosure and voluntary labelling.

Category II – Context-sensitive content

Examples:

- political satire;
- parody;
- journalism;
- commentary.

Regulatory approach: contextual safeguards and human review.

Category III – Harmful deceptive content

Examples:

- deliberate impersonation;
- fraudulent representations;
- fabricated evidence;
- targeted reputational attacks.

Regulatory approach: rapid removal, investigation and civil/criminal remedies where statutory elements are satisfied.

Category IV – High-risk abusive content

Examples:

- non-consensual sexual deepfakes;
- extortion;
- targeted harassment;
- threats;
- fraud;
- organised disinformation.

Regulatory approach: immediate protective measures, preservation of evidence, investigation and strong sanctions.

Such classification would allow the law to distinguish between creativity and abuse.

14. Procedural Safeguards

Deepfake regulation must not become purely automated.

An individual whose content is removed should, subject to urgent-risk exceptions, have access to:

- notice of removal;
- reasons for the decision;
- an opportunity to contest the decision;
- human review;
- an appeal mechanism; and
- restoration where the material is lawful.

Similarly, a person whose likeness has been misused should have an effective mechanism for rapid complaint and removal. This is particularly important because false positives can be constitutionally harmful.

A journalist reporting on a deepfake may use the manipulated content as evidence. A researcher may reproduce it to analyse misinformation. A satirist may intentionally manipulate a public figure's image. An automated system may not understand these distinctions.

15. Evidence and Authentication Problems

Deepfakes also create significant evidentiary challenges.

Courts traditionally rely upon digital evidence through statutory rules concerning electronic records. AI-generated material increases the importance of establishing:

1. provenance;
2. metadata;
3. source device;
4. chain of custody;
5. authenticity;
6. manipulation history; and
7. expert analysis.

The increasing availability of sophisticated synthetic media means that courts should not automatically assume that visual or audio material is genuine merely because it appears convincing.

At the same time, a claim that content is a deepfake should not automatically defeat otherwise admissible evidence. The legal system therefore requires stronger forensic capabilities and expert standards for authentication.

16. Deepfakes and Democracy

The constitutional problem extends beyond individual rights. Deepfakes can affect democratic decision-making by creating false statements attributed to candidates, public officials or political organisations. A fabricated political speech released shortly before an election may spread faster than a correction.

This creates an asymmetry:

creation may take minutes, but correction may take days.

The traditional legal system is often slower than digital virality.

Consequently, deepfake regulation requires mechanisms capable of responding rapidly while retaining procedural fairness. However, electoral regulation must be particularly careful. Government authorities should not receive unrestricted power to determine what political content is 'false' because such power could itself threaten democratic freedom. Independent review and judicial safeguards are therefore essential.

17. Comparative Constitutional Lessons

International approaches provide useful but limited guidance. The European Union has increasingly focused on transparency and risk-based regulation of artificial intelligence. The United States, by contrast, presents a stronger constitutional speech-protection tradition, making broad content-based restrictions particularly difficult. India occupies a distinctive constitutional position.

Its framework simultaneously protects:

- dignity;
- privacy;
- personal liberty;
- reputation;
- free expression; and
- reasonable restrictions in the public interest.

India therefore does not need to replicate another jurisdiction's approach mechanically. A uniquely Indian framework can be constructed around constitutional proportionality and rights-based platform governance.

18. Proposed Legal Reforms

This article proposes the following reforms.

18.1 A dedicated statutory definition of deepfakes

Indian law should clearly define deepfakes and distinguish them from ordinary digital editing.

18.2 Harm-based classification

The law should classify synthetic media according to risk and harm rather than banning AI-generated content generally.

18.3 Rapid remedy for victims

Victims should have access to a specialised and time-bound complaint mechanism.

18.4 Stronger protection against sexual deepfakes

Non-consensual sexual synthetic media should receive particularly strong legal protection because of its severe impact upon dignity and privacy.

18.5 Traceability without universal surveillance

Traceability mechanisms should be designed to identify the origin of unlawful content without creating unrestricted surveillance of lawful users.

18.6 Human review

Automated content moderation should not be the final decision-maker in contested cases.

18.7 Judicial oversight

Blocking and removal powers affecting significant political or journalistic expression should remain subject to meaningful legal and judicial safeguards.

18.8 AI literacy

Government and educational institutions should promote public awareness regarding synthetic media and verification.

18.9 Evidence standards

Courts should develop clearer procedures for authentication and forensic examination of suspected deepfakes.

18.10 Platform transparency

Large platforms should publish periodic transparency reports concerning synthetic-content complaints, removals, appeals and restoration.

19. Rethinking the Balance Between Privacy and Expression

The central constitutional mistake would be to treat privacy and freedom of expression as mutually exclusive. Both rights support democratic society.

Privacy protects the individual from unwanted intrusion and manipulation. Expression enables individuals to criticise, create, communicate and participate in democratic life. Deepfakes expose the point at which these interests collide.

A person should not lose control over their identity merely because technology makes replication technically possible. Conversely, an individual should not be able to suppress every fictional, satirical or critical representation simply by claiming a privacy interest. The appropriate principle is therefore: **The law should regulate harmful synthetic conduct, not synthetic creativity itself.** This principle provides greater constitutional stability than a general prohibition.

20. Conclusion

Deepfakes represent a fundamental challenge to the assumptions underlying traditional privacy, reputation and speech law. They blur the distinction between genuine and fabricated reality while enabling individuals and organisations to manipulate identity at unprecedented speed and scale.

India's constitutional framework provides significant foundations for responding to this challenge. Article 21, as interpreted through *Puttaswamy*, protects privacy, dignity and autonomy. Article 19(1)(a) protects expression, while Article 19(2) permits constitutionally defined restrictions. Indian defamation law protects reputation, and the IT Act, BNS, DPDP Act and IT Rules provide multiple statutory mechanisms relevant to deepfake-related conduct.

The development of the 2026 IT Rules concerning synthetically generated information represents an important step towards a more structured regulatory response. Nevertheless, regulatory expansion must remain constitutionally disciplined.

The future of deepfake regulation should therefore be built around four principles:

legality, proportionality, accountability and effective remedy.

The objective should not be to create a legal system in which every artificial representation is treated as unlawful. Rather, the law should distinguish between legitimate synthetic expression and synthetic content that causes identifiable and substantial harm.

Ultimately, the constitutional question is not whether artificial intelligence should be permitted to create representations of reality. It is whether the law can preserve human dignity and democratic discourse while allowing technological innovation and freedom of expression to continue.

A rights-based, harm-sensitive and procedurally fair regulatory framework offers the most constitutionally sustainable path for India in the generative AI era.

Bibliography

A. Cases

Justice KS Puttaswamy (Retd) v Union of India (2017) 10 SCC 1.

Maneka Gandhi v Union of India (1978) 1 SCC 248.

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

Subramanian Swamy v Union of India (2016) 7 SCC 221.

Anuradha Bhasin v Union of India (2020) 3 SCC 637.

Justice KS Puttaswamy (Retd) v Union of India (Aadhaar) (2019) 1 SCC 1.

B. Legislation

Constitution of India 1950.

Information Technology Act 2000.

Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules 2021.

Bharatiya Nyaya Sanhita 2023.

Bharatiya Nagarik Suraksha Sanhita 2023.

Bharatiya Sakshya Adhiniyam 2023.

Digital Personal Data Protection Act 2023.

C. Government Materials

Ministry of Electronics and Information Technology, Government of India, *Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021*, updated 10 February 2026.

Ministry of Electronics and Information Technology, Government of India, *Government Strengthens Regulatory Framework to Address AI-Generated Deepfakes* (6 August 2026).

Ministry of Electronics and Information Technology, Government of India, *Government Strengthens Framework to Counter AI-Generated Deepfakes through Legal Safeguards, Platform Accountability and Citizen Protection* (30 July 2026).

D. Suggested Secondary Sources

Daniel J Solove, *Understanding Privacy* (Harvard University Press 2008).

Neil M Richards, *Why Privacy Matters* (Oxford University Press 2021).

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

Danielle Keats Citron, *Hate Crimes in Cyberspace* (Harvard University Press 2014).