
DEEPFAKE REGULATION AND DIGITAL CONSTITUTIONALISM: HOW MUCH STATE CONTROL IS CONSTITUTIONALLY JUSTIFIED?

Kajal Srivastava, Advocate
Amity University

ABSTRACT

Deepfakes have become one of the most challenging issues in recent years. These are AI-generated or heavily manipulated videos, images, and audio clips that can make a person appear to say or do something they never actually said or did. While the technology has creative and entertainment uses, it has also been widely misused to create non-consensual intimate images, spread false political messages, and commit fraud. As a result, governments across the world have started introducing laws to regulate deepfakes.

This raises an important constitutional question: how much control can the state exercise over this technology without crossing the limits set by fundamental rights? This article looks at the problem through the idea of digital constitutionalism. Digital constitutionalism basically means that the same constitutional values that limit state power in the physical world, such as free speech, privacy, dignity, and accountability, should also apply in the digital space.

The article first explains what deepfakes are and why they create new legal difficulties. It then examines the regulatory frameworks that have developed so far, especially in the United States and the European Union. These frameworks include criminal penalties, civil remedies, mandatory labelling requirements, and obligations placed on online platforms. The central argument is that some amount of state regulation is clearly justified, particularly when it comes to non-consensual intimate deepfakes and deliberate attempts to mislead voters during elections. At the same time, overly broad or vaguely worded laws can easily interfere with protected speech, satire, and artistic expression.

The article therefore argues that any regulatory framework dealing with deepfakes must remain carefully limited. It should be guided by principles of necessity, proportionality, and narrow tailoring so that the state can address real harms without unnecessarily restricting free expression. In short, the question is not whether regulation is needed, but how much regulation is constitutionally acceptable in a democratic society.

INTRODUCTION

Deepfakes, realistic images, videos, or audio created or altered by artificial intelligence, have moved from a mere technological curiosity to a widespread political and social concern. What started as experimental tools for entertainment purposes can now be used by almost anyone having access to a smartphone and a free software. The same technology that offers space for creative expression, and sometimes harmless parody can also produce non-consensual content, misleading political posts, or fraudulent messages that looks and sounds authentic.

The more this synthetic media has become easier and cheaper to create, governments around the world have now begun introducing laws to regulate them. Some frauds target election-related deepfakes, others focus on non-consensual sexual imagery, and still others require labelling or disclosure. At the same time, these rules raise a basic constitutional question: how far can the state go in controlling this new form of expression without undermining privacy, free speech, and other fundamental rights?

This article uncovers the tension created by digital constitutionalism – the ideology that core constitutional principles must continue to guide state power in the digital sphere. The central question is pretty straightforward: how much state control over deepfakes is constitutionally justified? The balance is created when we understand that while targeted regulation of the most harmful uses is legitimate and necessary, poorly tailored restrictions risk going too far and should be carefully confined.

The article first explains the nature of deepfakes and the main harms they cause, then reviews current regulatory approaches, and finally assesses the constitutional boundaries that should constrain state action.

CONCEPTUAL FOUNDATIONS

DEFINING DEEPFAKES AND SYNTHETIC MEDIA

Before engaging with the legal or constitutional issues surrounding deepfakes, it is important to first clarify what the term actually encompasses. In public discussions, “deepfake” is often used broadly to describe almost any manipulated image or video. However, technical and legal scholarship generally adopts a more specific understanding of the term. In its strict sense, a deepfake refers to media that has been created or modified through deep learning technologies,

including Generative Adversarial Networks (GANs) and, increasingly, diffusion models. “Growing sophistication of the GAN approach is sure to lead to the production of increasingly convincing deep fakes.”¹ These technologies can produce highly convincing representations of individuals appearing to say or do things that they never actually said or did. This makes deepfakes distinct from earlier forms of media manipulation, such as photoshop editing or dubbing, which depend primarily on manual alteration rather than algorithmic systems trained on large datasets containing information about a person’s face, voice, or mannerisms.

TECHNICAL CHARACTERISTICS

The legal difficulty posed by deepfakes arises not merely from their falsity, but from their growing ability to appear almost indistinguishable from genuine recordings, particularly to an untrained viewer. Three characteristics are particularly important in this regard.

- First, the gap between synthetic and authentic media has narrowed considerably. Earlier deepfakes often contained obvious signs of manipulation, such as unnatural blinking, distorted facial features, or blurred edges, but advances in the underlying technology have made many of these flaws far less noticeable. The trajectory across 2017 to 2026 is clear: autoencoders dominated the early "FaceSwap"-style era, GANs took over for high-fidelity face synthesis through 2022, and diffusion models now lead on virtually every quality benchmark for image and video. Voice cloning is its own track, dominated by autoregressive transformers and flow-matching models.²
- Second, deepfake creation is no longer limited to specialized laboratories or users with sophisticated computing resources; consumer applications now allow such content to be produced within minutes.
- Third, and most importantly from a legal perspective, deepfakes can be created, reproduced, and circulated on a scale that traditional forms of forgery could not achieve. A single fabricated video may be copied and distributed across multiple platforms at a speed that can easily outpace existing mechanisms for detection and removal.

¹ Bobby Chesney & Danielle Citron, *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security*, 107 California Law Review 1753(2019).

² DuckDuckGoose, *How Are Deepfakes Made? A Technical Guide for 2026* (Apr. 28, 2026), <https://www.duckduckgoose.ai/blog/how-deepfakes-are-made>.

Taken together, these features have contributed to what some scholars describe as an “evidentiary crisis,” in which the growing possibility of manipulated media lowers confidence of people not only in fabricated recordings but also in authentic ones. This broader effect is commonly referred to as the “liar’s dividend.” “Put simply: a skeptical public will be primed to doubt the authenticity of real audio and video evidence. This skepticism can be invoked just as well against authentic as against adulterated content. Hence what we call the liar’s dividend: this dividend flows, perversely, in proportion to success in educating the public about the dangers of deep fakes”³

DIGITAL CONSTITUTIONALISM: MEANING AND CORE PRINCIPLES

If deepfakes represent the emerging challenge, digital constitutionalism provides a framework for understanding how that challenge might be addressed. Broadly understood, digital constitutionalism involves applying constitutional values such as rights, limitations on power, and accountability to the digital environment. This becomes particularly relevant as the internet is no longer shaped by states alone. Private technology companies, online platforms, algorithms, and transnational digital firms now exercise significant forms of power, often alongside or beyond traditional state institutions.

EXTENSION OF CONSTITUTIONAL VALUES INTO THE DIGITAL SPHERE

A central premise of digital constitutionalism is that the digital environment should not be regarded as a space beyond constitutional principles. In the same way that values such as freedom of speech, privacy, and due process place limits on the exercise of state power in the physical world, these principles must also have relevance within digital spaces. Digital constitutionalists therefore seek to embed such values into digital governance through binding legal measures, including legislation and judicial doctrines, and softer forms of regulation, such as platform charters, self-regulatory frameworks, and international initiatives, including movements advocating an Internet Bill of Rights. This does not necessarily mean that constitutional provisions designed for the physical world should simply be applied to the internet without alteration. Instead, the argument is that constitutional methods of reasoning, particularly proportionality, necessity, and protection against arbitrary exercises of power, should guide the use of authority in the digital sphere, irrespective of whether that authority is

³ Bobby Chesney & Danielle Citron, *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security*, 107 California Law Review 1753(2019).

exercised by a government institution or through an algorithm responsible for content moderation.

THE HARMS OF DEEPPAKES AND THE CASE FOR REGULATORY INTERVENTION

INDIVIDUAL HARMS

The argument for regulating deepfakes is not simply based on a general fear of technological change, it is grounded in an increasing body of evidence showing that these technologies can cause serious harm to individuals. Before considering their broader effects on society or democratic institutions, it is important to examine the injuries suffered at the individual level, as these are often that require first prompt legislative action. Unlike certain technology-related harms that remain uncertain or difficult to identify, the consequences of deepfakes for individuals are often immediate, tangible, and difficult to reverse once the material has gone viral.

NON-CONSENSUAL INTIMATE DEEPPAKES AND PERSONALITY RIGHTS

Among the most serious and extensively documented forms of deepfake abuse is the non-consensual production of sexually explicit material, with women being disproportionately targeted. Research examining deepfake material available online has consistently indicated that a substantial proportion of such content is pornographic and created without the knowledge or consent of the person depicted. In many cases, photographs taken from social media platforms are used to create the fabricated content. The resulting harm extends beyond the conventional notion of reputational damage. It represents a direct interference with personality rights, which encompass an individual's interests in their name, image, likeness, and voice.

The development of personality rights has traditionally been associated with cases involving the unauthorized commercial use of an individual's identity. Applying these principles to deepfakes, however, presents a more difficult but increasingly important legal question. In such cases, the victim's identity is appropriated not necessarily for commercial benefit, but often for humiliation, harassment, or degradation. Courts and legislatures, in a number of jurisdictions, have therefore begun to recognize the need of the hour, which is to treat this conduct as a distinct legal wrong, rather than relying solely on broader concepts of privacy or obscenity.

Existing legal categories frequently fail to capture the particular nature of the violation: the victim's likeness has been artificially placed into content in which they never participated and to which they never consented.

REPUTATIONAL, PSYCHOLOGICAL AND ECONOMIC INJURY

The effects of deepfakes extend beyond non-consensual intimate imagery and can produce several forms of harm that reinforce one another. Reputational harm may arise when a fabricated recording portrays an individual making statements or engaging in conduct that negatively affects their professional or social standing. The problem becomes particularly serious because highly realistic deepfakes can make it difficult for the person concerned to demonstrate that the recording is false. Consequently, the damage to reputation may continue even after the content has been exposed as fabricated or a correction has been issued.

Psychological harm is also becoming increasingly visible through victim accounts and emerging clinical research. Individuals subjected to deepfake abuse may experience anxiety, heightened vigilance, and a continuing sense of violation similar to the effects associated with other forms of harassment and image-based abuse. The uncertainty surrounding the circulation of the material can intensify this distress, particularly when victims do not know how widely it has been shared or who may have accessed it.

Economic consequences can arise alongside these reputational and psychological effects, especially where an individual's employment or livelihood depends upon professional credibility, public confidence, or personal branding. A fabricated recording may lead to the loss of employment, professional relationships, business opportunities, or client trust before the falsity of the material can be established. The interaction between these forms of harm makes the legal problem particularly complex. Psychological distress may be aggravated by continuing reputational or financial consequences, while the victim's inability to control the circulation or removal of the fabricated material can itself become a significant part of the injury.

DEFAMATION, PRIVACY AND, RIGHT OF PUBLICITY

Existing civil law provides some avenues through which victims of deepfakes may seek redress, although these remedies are often incomplete. Defamation law may apply where a

deepfake communicates a false factual claim that harms an individual's reputation. However, its application becomes complicated because traditional defamation principles generally focus on false statements of fact. It is less settled whether a fabricated video should be understood as making a factual assertion in the same manner as a written or spoken statement, or whether it should instead be treated as a visual depiction that requires a different form of legal analysis.

Privacy law provides another possible avenue, particularly through concepts such as intrusion upon seclusion and the disclosure of private facts. Yet these doctrines were largely developed to address the misuse or exposure of genuine private information, rather than situations in which entirely fictional events are created and falsely associated with an individual. The right of publicity, which historically developed to protect individuals from unauthorized commercial exploitation of their identity, may offer a somewhat more adaptable framework because deepfakes inherently involve the unauthorized use of a person's identity or likeness. Some legislative proposals and enacted laws have consequently sought to extend publicity protections to synthetic representations, including circumstances where the use is not strictly commercial.

Despite these developments, none of these legal doctrines was originally designed with generative artificial intelligence in mind. As a result, individuals seeking legal remedies may have to rely on established principles and extend their application to forms of conduct that could not have been anticipated when those doctrines were first developed.

CRIMINAL FRAUD AND RELATED OFFENCES

When deepfake technology is used primarily as a means of obtaining financial benefit rather than for harassment or humiliation, existing criminal fraud provisions provide a more direct legal framework. Examples include voice-cloning scams in which offenders impersonate relatives to deceive elderly individuals, synthetic-video schemes in which executives are imitated to authorize financial transfers, and fabricated “proof of life” communications used in extortion or kidnapping-related fraud. Such incidents are increasingly being reported across different jurisdictions and demonstrate that deepfake technology can facilitate conventional forms of deception on a much larger scale.

These cases generally fit within existing fraud and impersonation offences because they involve familiar elements of criminal conduct, particularly deception, impersonation, and

financial loss. The more difficult issue arises at the evidentiary stage. Successful prosecution may depend on establishing that the audio, video, or other digital material was synthetically generated. This requires reliable forensic expertise and technical capabilities that many law enforcement agencies are still developing.

Taken together, these forms of individual harm, affecting dignity, psychological well-being, reputation, livelihood, and financial security, provide an important empirical and legal basis for regulatory intervention. They also demonstrate a central limitation of the existing legal framework: no single established category of law is capable of fully addressing the range and cumulative nature of harm that deepfakes can produce.

MAPPING EXISTING REGULATORY APPROACHES

Governments across the world have responded to deepfakes in various different ways. In the United States, most activity has happened at the state level, with a growing number of laws targeting both non-consensual intimate imagery and election-related deepfakes. The European Union has adopted a more uniform approach centered around labelling and transparency. Across these frameworks, a few common tools keep showing up: civil liability, criminal penalties, watermarking, mandatory labelling, and duties imposed on online platforms.

UNITED STATES

“Deepfake legislation in the U.S. is advancing swiftly to combat the rising risks associated with synthetic media, addressing critical areas such as cybersecurity, privacy, election integrity, and intellectual property. Federal and state lawmakers are enacting and refining laws to curb the misuse of deepfake technology, focusing on issues like fraud, defamation, election manipulation, and non-sexual explicit content.”⁴ With the evolving technology, it is essential to formulate proper legislations to make sure that these powerful tools are being used responsibly. Federal laws like *TAKE IT DOWN Act (2025)*, *Federal Trade Commission Act (FTC Act)*, and *US Criminal Code (18 U.S.C.)* govern and regulate deepfakes and artificial intelligence in the United States.

The U.S. Congress passed the TAKE IT DOWN Act in May 2025. It is the first major federal statute pinpointing non-consensual intimate imagery (NCII) including AI-developed

⁴ Kalpan, C. (2025, January 28). *What legislation protects against deepfakes and synthetic media?* HALOCK.

“deepfakes”. This Act mentions criminal penalties for the intentional disclosure of non-consensual intimate visual depictions. Violations involving adults are punishable by up to two years’ imprisonment, while offences involving minors may attract imprisonment of up to three years. The Act also criminalizes certain threats involving such material.⁵

The proposed DEEPFAKES Accountability Act adopts a multi-layered penalty framework. It provides for criminal punishment of up to five years’ imprisonment, or fine, or both for specified aggravated violations, while failure to make the prescribed disclosure or intentional alteration of disclosure may attract civil penalties of up to \$150,000 per record or alteration. The bill further creates a private right of action allowing victims to seek damages for harmful deepfake representation.⁶

EUROPEAN UNION

The EU Artificial Intelligence Act addresses deepfakes primarily through mandatory transparency rather than outright prohibition. Article 50 requires that providers of AI systems manipulating and generating images, videos, or audios content indicate that content as artificially generated in a machine-readable format.⁷ Deployers of systems that produce deepfakes, particularly those depicting real persons, must clearly disclose to viewers that the material is not authentic. The purpose is to inform the public. Limited exceptions apply to purely satirical, fictional, artistic, or analogous works, provided the synthetic nature is still made clear where appropriate and the exception does not overshadow the protective aim. These transparency obligations largely became applicable from 2 August 2026. By prioritizing labelling over content bans, the AI Act aims to preserve expressive space while enabling citizens and platforms with the ways to discount and detect manipulated media.

Under the Digital Services Act, large online platforms and search engines bear broader due-diligence duties to identify systemic risks, including the dissemination of harmful synthetic media that may undermine public discourse, or individual dignity. DSA also requires platforms to reduce the risks of exposing citizens to harmful or illegal content.⁸ These obligations

⁵ TAKE IT DOWN Act, Pub. L. No. 119-12, Section 2(a), 139 Stat. 55, 58-59 (2025) (codified at 47 U.S.C. Sections 223(h)(4), (6)).

⁶ DEEPFAKES Accountability Act of 2023, H.R. 5586, 118th Cong. Sections 6(f)(1) – (2), (g)(1) – (2) (2023).

⁷ (Regulation (EU) 2024/1689, 2024, Art. 50(1))

⁸ European Parliament & Council of the European Union. (2022). *Regulation (EU) 2022/2065 (Digital Services Act)*.

complement the AI Act's labelling rules by requiring platforms to assess and address the spread of deepfakes as part of their risk-management systems.

Beyond the United States and the European Union, a growing number of jurisdictions, including the United Kingdom, Australia, China, and several Latin American states, have begun to regulate deepfakes through a mix of criminal prohibitions (especially for non-consensual intimate imagery) and transparency or labelling requirements. Approaches vary in strictness and emphasis, but the common pattern is a hybrid of punitive measures for the most harmful uses and disclosure rules for broader synthetic content.

CONSTITUTIONAL CONSTRAINTS ON STATE CONTROL

The state has a valid interest in tackling the harms caused by deepfakes and synthetic media. But that interest does not give it unlimited power. In India, any restriction on speech or online activity must respect the Constitution, particularly the right to free speech under Article 19(1)(a) and the right to privacy and dignity under Article 21.⁹

Free speech is not absolute and it is ensured by reasonable restrictions in the interests of public order, decency, morality, and the sovereignty and integrity of India, provided under Article 19(2). Although, the Supreme Court has made it clear that these restrictions cannot be too broad or vague. In *Shreya Singhal v. Union of India*, the Court struck down Section 66A of the IT Act precisely because vague language created a chilling effect; people became afraid to speak freely online for fear of arbitrary action.¹⁰ This same principle applies to deepfake regulation as well. Laws that use open-ended terms like “inauthentic,” “misleading”, or “harmful” content are at the risk of being used to silence satire, curb political criticism, or parody. False or exaggerated speech about public figures, even when created with AI, often remains protected political expression.

At the same time, privacy and dignity set important limits in the opposite direction. The Supreme Court in *Justice K.S. Puttaswamy (Retd.) v. Union of India* recognized privacy as a fundamental right under Article 21.¹¹ Non-consensual intimate deepfakes, which use a person's face or voice without permission to create sexual or humiliating content, clearly violate this

⁹ The Constitution of India, 1950, Art. 19(1)(a)

¹⁰ *Shreya Singhal v. Union of India*, (2015) 5 SCC 1 (India)

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

right. Here the state has stronger grounds to act. The harm is concrete and personal, not merely the expression of an unpopular opinion.

India's current regulatory approach reflects this tension. The Digital Personal Data Protection Act, 2023 treats facial images and voice recordings as personal data. This means that carving someone's photos or voice to train a deepfake model can attract obligations under the Act.¹² Recent amendments to the IT Rules have also introduced specific duties on platforms regarding "synthetically generated information," including labelling and faster takedown requirements.¹³ These measures are attempts to respond to real harms. Yet they also raise constitutional questions. Broad exemptions for the state under the DPDP Act, the absence of clear journalistic protections, and the power to compel rapid removal of content can easily tip into overreach. When rules are drafted too widely, platforms tend to remove more content than necessary, and ordinary users begin to self-censor.¹⁴

The constitutional boundary is therefore relatively clear. The state may regulate to prevent concrete, non-speech harms, especially non-consensual intimate imagery and clear cases of fraud or impersonation. It may also require transparency, such as clear labelling that a video is AI-generated. What it cannot do is create a general power to decide which political or satirical deepfakes are "acceptable." Any restriction must be narrowly tailored, based on clear criteria, and open to judicial scrutiny. Ambiguous and overbroad rules that chill legitimate speech or expand executive discretion without adequate safeguards will struggle to survive constitutional challenge.

In short, digital technologies create new problems, but they do not rewrite the Constitution. State control remains justified only to the extent it is necessary, precise, and respectful of the rights to speak and to be left alone.

TOWARDS A CONSTITUTIONALLY JUSTIFIED REGULATORY FRAMEWORK

One does not need a completely new set of rules to deal with deepfakes. One can work with the laws already in existence, but they must be used carefully.

¹² Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023).

¹³ Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026

¹⁴ Indumugi C. and Naman Kumar, "Supreme Court Issues Notice on Constitutional Challenge to the Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025," Internet Freedom Foundation, February 16, 2026.

First, the different types of harm should be segregated. Deepfakes that show people in sexual situations without their permission, or that are used for cheating and fraud, cause real damage. These can be dealt with under existing criminal laws and the right to privacy. The law should make it easier for ordinary people, not just famous persons, to get quick resolution and compensation.

On the other hand, videos made for jokes, satire, art, or political criticism should still be protected under free speech. If the law treats every AI-made video as dangerous, it can create fear and stop people from speaking freely. This is the same problem the Supreme Court saw when it struck down Section 66A in the *Shreya Singhal*¹⁵ case. Clear definitions and exceptions for honest uses are necessary.

The changes made to the IT Rules in 2026 are a step in the right direction. They define synthetic content, ask platforms to put clear labels on AI-generated videos, and require faster removal of clearly illegal material. Labelling is better than banning content outright because it lets people know what they are watching without stopping the speech.

A few more simple steps can help. Give ordinary people an easy way to stop the misuse of their face or voice. Treat face and voice data as sensitive under the Digital Personal Data Protection Act. Put the main rules in a proper law passed by Parliament instead of only in government notifications. And keep platforms safe from liability if they remove illegal content quickly when they are told about it.

The basic idea is simple: any rule must be necessary, clear, and not go further than needed. Regulation should guide technology, not block free speech or privacy.

CONCLUSION

Deepfakes are a real burning issue in India. They can be used to harm a person's privacy, damage reputations, spread false information during elections, and even commit fraud. At the same time, free speech and the right to privacy are basic rights given by the Constitution. Any attempt to control deepfakes must respect both these rights.

India has already started taking steps. The 2026 changes to the IT Rules require platforms to

¹⁵ *Shreya Singhal v. Union of India*, (2015) 5 SCC 1 (India)

label AI-generated content and remove illegal material more quickly. The *Digital Personal Data Protection Act* gives some protection to personal data like face and voice. Courts have also started protecting personality rights in cases involving misuse of a person's image. These are useful measures.

However, the work is not finished. If the rules are written in vague language, they can create fear among ordinary users and content creators. People may stop sharing satirical videos or political comments because they worry about getting into trouble. This kind of chilling effect is dangerous for democracy. The Supreme Court has already warned against such vague restrictions in the past.

A better approach is to focus only on clear and serious harms. Non-consensual intimate deepfakes and deepfakes used for fraud should be dealt with strictly. At the same time, satire, parody, art, and political criticism should remain free. Labelling AI-generated content is a good middle path because it informs viewers without blocking speech. Rules should be clear, limited, and based on proper law passed by Parliament rather than only on government notifications.

In the end, technology will keep changing. New tools will keep appearing. The Constitution, however, remains the same. It asks the state to protect citizens from harm, but it also asks the state not to go beyond what is necessary. If India can keep this balance, protecting people from real damage while protecting free speech and privacy, then regulation of deepfakes can work in a fair and constitutional way. The goal should not be to control every piece of online content. The goal should be to stop genuine harm without weakening the rights that make open discussion possible in a democracy.

BIBLIOGRAPHY**A. Books and Journal Articles**

- Chesney, B., & Citron, D. (2019). Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review*, 107, 1753–1819.
- Kalpan, C. (2025, January 28). What legislation protects against deepfakes and synthetic media? HALOCK.

B. Online Articles and Web Resources

- C., I., & Kumar, N. (2026, February 16). Supreme Court issues notice on constitutional challenge to the Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025. Internet Freedom Foundation.
- DuckDuckGoose. (2026, April 28). How are deepfakes made? A technical guide for 2026. <https://www.duckduckgoose.ai/blog/how-deepfakes-are-made>

C. Legislation and Regulations

- Constitution of India. (1950). Article 19(1)(a).
- Digital Personal Data Protection Act, 2023, Act No. 22 of 2023.
- Deepfakes Accountability Act of 2023, H.R. 5586, 118th Cong. §§ 6(f)(1)–(2), 6(g)(1)–(2) (2023).
- Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026.
- Regulation (EU) 2022/2065 of the European Parliament and of the Council (Digital Services Act).
- Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act), Art. 50(1).
- TAKE IT DOWN Act, Pub. L. No. 119-12, § 2(a), 139 Stat. 55, 58–59 (2025) (codified at 47 U.S.C. §§ 223(h)(4), (6)).

D. Judicial Decisions

- Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 SCC 1.
- Shreya Singhal v. Union of India, (2015) 5 SCC 1.