DEEPFAKE CRIMES AGAINST WOMEN IN INDIA: TOWARDS A VICTIM-CENTRIC FRAMEWORK OF CRIMINAL LIABILITY, DIGITAL DIGNITY AND PLATFORM ACCOUNTABILITY

Dr. Cumaran Nadaradjan, B.A. LL.B., LL.M., Ph.D. (Law)
Advocate, Bar Council of Delhi

ABSTRACT

Generative artificial intelligence has transformed the architecture of technology-facilitated abuse against women. A perpetrator no longer requires possession of an authentic intimate photograph or recording to manufacture sexual humiliation, reputational injury or coercive control. Publicly available photographs, fragments of video and samples of speech can be converted into realistic synthetic images, audio and video that falsely depict an identifiable woman in intimate, sexual or otherwise compromising circumstances. The resulting injury is not adequately understood as misinformation alone. It implicates privacy, dignity, sexual autonomy, reputation and an individual's ability to exercise meaningful control over the digital representation of her identity.

Indian law now contains a significant but dispersed body of rules capable of addressing different stages of deepfake abuse. The Bharatiya Nyaya Sanhita, 2023, the Information Technology Act, 2000, the Bharatiya Sakshya Adhiniyam, 2023, constitutional privacy jurisprudence, data-protection law and intermediary regulation may each become relevant according to the conduct involved. The 2026 amendments to India's intermediary rules mark an especially important regulatory development by expressly addressing synthetically generated information and strengthening obligations concerning unlawful synthetic content, transparency, provenance and intermediary due diligence. Yet recognition of synthetic media as a regulatory category does not itself create a coherent remedial architecture for victims.

This Article argues that the central weakness of the existing framework is remedial fragmentation. The law separates criminal liability, platform responsibility, evidence, privacy and regulatory compliance, while the victim experiences creation, publication, amplification, repeated circulation, evidentiary loss and reputational injury as one continuing episode of abuse. To address that disconnect, the Article develops "synthetic dignity harm" as

an analytical concept describing serious injury caused by the non-consensual appropriation of an identifiable person's identity to manufacture a harmful synthetic representation. It further proposes a Victim-Centric Deepfake Response Framework (VCDRF) organised around five interconnected duties: prevention, expedited containment, evidentiary preservation, victim protection and accountable redress.

The Article additionally proposes recognition of Non-Consensual Synthetic Intimate Content (NCSIC) as a distinct regulatory category; develops differentiated responsibility for creators, commissioners, initial publishers, knowing redistributors, platforms and deliberately facilitating technological services; and identifies the authenticity paradox, persistence gap and removal-preservation paradox as structural problems that conventional digital-remedy mechanisms do not adequately resolve. It concludes that the next stage of Indian deepfake regulation should move beyond asking whether particular content is authentic or artificial. The deeper legal question is whether an individual possesses meaningful protection against another person's non-consensual technological manufacture and exploitation of her identity.

Keywords: Artificial Intelligence; Deepfakes; Women; Synthetic Media; Digital Dignity; Privacy; Bharatiya Nyaya Sanhita; Non-Consensual Synthetic Intimate Content; Platform Accountability; Synthetically Generated Information; Intermediary Liability.

I. INTRODUCTION: WHEN IDENTITY BECOMES GENERATIVE MATERIAL

Artificial intelligence has changed not only the production of information but the conditions under which identity itself can be represented. For most of the history of photography and recorded sound, fabrication demanded time, technical expertise, access to editing tools and, often, source material that already resembled the desired result. Generative systems alter that equation. A handful of ordinary photographs or a short voice sample may now provide sufficient raw material for a convincing synthetic representation. The legal significance of this change is not that deception is new. It is that the cost of manufacturing a believable false version of another human being has fallen dramatically.

For women, the consequences are especially acute where synthetic technology is used to manufacture nudity, sexual conduct, intimate speech, humiliating scenarios or other compromising representations. The victim need never have created an intimate image. She need never have participated in the depicted conduct. The perpetrator can manufacture the event and then attach its social consequences to a real person. This inversion distinguishes

synthetic sexual abuse from many traditional forms of non-consensual intimate-image distribution.

The conventional legal instinct is to ask which existing offence applies: obscenity, voyeurism, stalking, defamation, intimidation, cheating, forgery, privacy violation or unlawful electronic publication. Those questions remain necessary. But they do not fully describe the victim's problem. The same episode may involve acquisition of photographs, generation of synthetic material, anonymous upload, algorithmic amplification, repeated reposting, extortion, workplace consequences and evidentiary disappearance. Statutes divide these events into doctrinal categories; the victim experiences them as a single continuum.

This Article therefore advances a victim-centred method. It does not argue that India has no law capable of reaching deepfake abuse. That proposition is increasingly inaccurate, particularly after the 2026 amendments to the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules expressly addressed synthetically generated information. The more difficult question is whether the present combination of criminal law, intermediary regulation, evidence law, privacy principles and data governance can operate together with sufficient speed and coherence to protect the person whose identity has been appropriated.

The argument proceeds in four stages. First, it conceptualises the injury as synthetic dignity harm: a serious interference with representational autonomy caused by the non-consensual technological appropriation of an identifiable person's attributes. Second, it examines the Indian legal architecture, including the Bharatiya Nyaya Sanhita, the Information Technology Act, constitutional privacy jurisprudence, the Bharatiya Sakshya Adhiniyam, data protection and the 2026 synthetic-information rules. Third, it identifies structural failures that persist even when individual legal provisions are available. Finally, it proposes the Victim-Centric Deepfake Response Framework, or VCDRF, as an organising model for prevention, containment, evidence preservation, victim protection and accountable redress.

II. FROM DEEPFAKES TO GENDERED SYNTHETIC ABUSE

The term "deepfake" is commonly used for realistic media generated or materially altered through artificial intelligence or related computational techniques. As a legal category, however, the label can be imprecise. Synthetic media includes harmless entertainment,

dubbing, accessibility tools, educational reconstructions, parody and artistic experimentation. Regulation cannot sensibly begin from the premise that artificiality itself is wrongful. The legal inquiry must focus on the relationship among realism, identity, consent, purpose, knowledge and harm.

Gendered synthetic abuse illustrates why this distinction matters. A clearly fictional image that no reasonable viewer would associate with an actual person presents a different legal problem from a realistic sexual image constructed around an identifiable woman's face. The second representation derives its harmful force from recognisability. It asks the audience to connect fabricated sexual conduct to a real individual. The identity of the victim becomes the mechanism through which the content acquires value to the abuser.

Synthetic abuse also disrupts assumptions embedded in older intimate-image frameworks. Traditional image-based abuse often begins with an authentic photograph or recording that was obtained or distributed without consent. Deepfake abuse can begin with an ordinary portrait, a professional headshot, a social-media photograph or publicly accessible video. Consent to the existence of the source photograph is not consent to its transformation into sexual material. Treating those forms of consent as interchangeable would convert ordinary participation in digital life into a waiver of sexual and representational autonomy.

The harm is intensified by digital persistence. Once a file is posted, copies may be downloaded, mirrored, compressed, cropped, re-encoded or distributed through private channels. The first successful takedown may therefore have limited practical effect. Search engines may continue to surface derivative pages. Anonymous accounts may re-upload the material. The victim may spend months proving the same falsity to different services. The architecture of the internet converts a single act of fabrication into a potentially recurring injury.

These features justify treating non-consensual synthetic intimate material as more than a subset of generic misinformation. The relevant interests include sexual autonomy, dignity, reputation, safety, privacy, equality and control over the social meaning attached to one's identity. A legal response that concentrates only on whether viewers were deceived may miss the point. Even where every viewer knows that an image is fabricated, the use of a woman's identity to create sexualised content can remain degrading, exploitative and coercive.

III. SYNTHETIC DIGNITY HARM AND REPRESENTATIONAL AUTONOMY

Privacy is indispensable to the analysis, but privacy alone does not capture the full injury. A deepfake may be made entirely from photographs already visible to the public. The perpetrator may reveal no true secret about the victim. Yet the act can still be profoundly invasive because it appropriates identity to construct a false intimate representation. The wrong lies not only in disclosure but in manufacture.

This Article uses the term "synthetic dignity harm" to describe the non-consensual use or algorithmic reconstruction of an identifiable person's attributes to create a realistic false representation that subjects that person to sexualisation, humiliation, coercion, impersonation or comparable serious injury. The concept is offered as an analytical framework, not as a claim that Indian courts have already recognised a freestanding cause of action under that name.

Four elements are central. First, the representation must concern an identifiable person. Second, the representation must appropriate attributes that anchor the synthetic content to that person, such as face, voice, likeness or other identifying features. Third, it must materially falsify conduct, speech, nudity, intimacy or circumstances. Fourth, the use must create or foreseeably expose the person to serious dignitary, sexual, coercive or reputational harm. These elements distinguish synthetic dignity harm from trivial editing and ordinary creative transformation.

Representational autonomy provides the normative foundation. Individuals cannot possess absolute ownership over every depiction of themselves; such a principle would collide with journalism, biography, criticism, satire, art and public discourse. But the opposite position is equally untenable: technological capacity cannot itself create an unrestricted entitlement to manufacture realistic sexual or humiliating versions of another person. The law requires a boundary between legitimate representation and serious non-consensual identity appropriation.

The concept is particularly useful in sexual deepfake cases because the victim's injury does not depend upon the truth of the depicted act. A woman may be harmed precisely because the event is false and has been made to appear true. This produces what the Article calls the authenticity paradox: traditional categories may become harder to apply as the underlying sexual event becomes less authentic, even though the victim's representational injury remains substantial. Modern protection should therefore focus on the non-consensual association of an identifiable person with fabricated intimate conduct, not on whether an authentic intimate act ever occurred.

IV. CONSTITUTIONAL FOUNDATIONS: PRIVACY, DIGNITY AND EQUALITY

The Supreme Court's privacy jurisprudence supplies an important constitutional foundation. Justice K.S. Puttaswamy (Retd.) v. Union of India recognised privacy as a constitutionally protected interest connected with dignity, liberty and autonomy. The significance of that reasoning for synthetic media extends beyond secrecy. Privacy includes dimensions of personhood and control that become increasingly important when technology can reproduce identity independently of the individual.1

A face or voice is not merely a data point. It is also a social identifier through which others recognise, judge and interact with a person. When generative technology permits third parties to attach fabricated conduct to those identifiers, informational autonomy acquires a representational dimension. The constitutional value of dignity supports an interpretation of statutory protections that takes seriously the harm of forced synthetic sexualisation.

Article 14 is also relevant at the level of policy design. A facially neutral technology may produce gendered patterns of harm. The constitutional commitment to equality does not require the law to pretend that all manifestations of a technology create identical social consequences. A proportionate framework can recognise the specific severity of non-consensual sexual deepfakes while remaining available to victims of any gender.2

At the same time, Article 19(1)(a) requires precision. Synthetic media can be expressive. Political satire, cinema, parody and artistic works may employ imitation and transformation. A victim-centred framework should therefore be tied to clearly articulated harms, knowledge, intent, non-consent where relevant, and realistic identification. The object is not to suppress artificial expression; it is to address serious identity-based abuse.3

Constitutional values are most useful here as interpretive principles. They should guide statutory interpretation, procedural design and legislative reform without collapsing every private act of deepfake abuse into a direct constitutional claim. The practical task is to construct remedies that respect expression while ensuring that dignity is not rendered meaningless by the

1Justice K.S. Puttaswamy (Retd.) v. Union of India, (2017) 10 S.C.C. 1 (recognising privacy as a constitutionally protected right associated with dignity, liberty, and autonomy).
2INDIA CONST. arts. 14, 21.
3INDIA CONST. art. 19(1)(a).

ease with which identity can be digitally appropriated.

V. CRIMINAL LIABILITY UNDER THE BHARATIYA NYAYA SANHITA, 2023

The Bharatiya Nyaya Sanhita, 2023, in force from 1 July 2024, supplies the principal general criminal-law framework. Deepfake conduct should not be analysed by searching for a single offence called "deepfake." Liability depends on the acts performed, the mental element, the nature of the representation, the method of communication and the consequences intended or caused.4

Depending on the facts, synthetic abuse may intersect with provisions concerning stalking, voyeuristic conduct, criminal intimidation, cheating by personation, forgery or false electronic records, defamation and other offences. A sexual deepfake sent directly to a victim with a threat presents a different offence matrix from a fabricated video uploaded for reputational sabotage. A cloned voice used to induce payment presents yet another. Technology describes the instrument; criminal law must still identify the prohibited conduct.

The distinction among actors is equally important. The person who generates the material may not be the person who commissions it, uploads it or causes it to go viral. A legal system that focuses exclusively on the technical creator risks ignoring deliberate procurement and knowing redistribution. Conversely, imposing criminal responsibility on every accidental recipient would disregard ordinary requirements of culpability. The proper approach is differentiated responsibility.

The creator who intentionally manufactures non-consensual synthetic intimate content occupies the most direct position. A commissioner who procures that creation for revenge or harassment should not escape scrutiny merely because another person operated the software. An initial publisher who knowingly introduces the material into a network contributes to the foreseeable expansion of harm. A knowing redistributor who circulates it for humiliation or coercion may be materially different from a person who receives it without knowledge of its character.

The BNS also matters in cross-border cases. Synthetic content can be generated abroad, hosted elsewhere and directed at a person in India. Jurisdictional analysis will depend upon the

4Bharatiya Nyaya Sanhita, No. 45 of 2023, INDIA CODE (2023) (brought into force July 1, 2024).

statutory nexus and facts, but modern criminal law cannot assume territorial neatness in a networked environment. Effective enforcement will often require cooperation with platforms and foreign service providers, preservation of technical records and legally valid requests for identifying information.5

The remaining gap is conceptual and procedural rather than absolute. Existing offences can reach many manifestations of deepfake abuse, but they were not all designed around the distinctive structure of fabricated intimacy. Where liability turns on concepts historically associated with authentic observation or recording, courts and legislators must avoid an authenticity trap in which the absence of a real sexual event weakens protection against a fabricated one.

VI. THE INFORMATION TECHNOLOGY ACT AND SYNTHETIC ABUSE

The Information Technology Act, 2000 remains central because deepfake abuse is created, transmitted and stored through computer resources. Provisions dealing with identity misuse, personation, privacy, obscene material, sexually explicit electronic content and intermediary liability may be relevant according to the facts. The Act should be read alongside the BNS rather than treated as a self-contained deepfake code.6

Section 79's conditional safe-harbour architecture is especially important. Platforms do not become automatically liable merely because users upload unlawful material. At the same time, statutory protection is connected to compliance with due-diligence obligations and the legal framework governing knowledge and response. This balance reflects a broader constitutional concern: intermediaries cannot realistically pre-adjudicate every user communication, but victims also cannot be left without meaningful mechanisms when serious unlawful content is identified.7

The Supreme Court's decision in Shreya Singhal v. Union of India remains foundational to the understanding of intermediary liability and lawful restrictions on online speech. Deepfake regulation must preserve that concern with procedural legitimacy. A system that encourages platforms to remove anything controversial upon vague accusation can chill lawful expression.

5Bharatiya Nyaya Sanhita, No. 45 of 2023, § 1(5), INDIA CODE (2023) (providing for specified extraterritorial application, including offences targeting a computer resource located in India).
6Information Technology Act, No. 21 of 2000, §§ 66C, 66D, 66E, 67, 67A, INDIA CODE (2000).
7Information Technology Act, No. 21 of 2000, § 79, INDIA CODE (2000).

A system that requires victims of obvious synthetic sexual abuse to obtain prolonged adjudication before any containment can render remedies practically useless. The challenge is calibrated procedure.8

The IT Act's legacy categories continue to matter even after synthetic-media-specific rules. Deepfake conduct may involve more than one wrong at once: identity appropriation, sexually explicit publication, intimidation and fraudulent personation can coexist. The 2026 regulatory turn does not displace these provisions; it provides an additional layer of platform-focused governance.

VII. THE 2026 SYNTHETICALLY GENERATED INFORMATION REGIME

India's regulatory position changed materially in February 2026 when amendments to the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules expressly addressed synthetically generated information. The updated framework and accompanying government materials recognise the need for rules tailored to realistic artificially or algorithmically created or altered audio-visual information. This development makes it increasingly inaccurate to describe Indian law as wholly silent on deepfakes.9

The regulatory importance of the amendments lies in their shift from generic content governance toward synthetic-media-specific due diligence. The framework addresses unlawful synthetic information, transparency and provenance, and enhanced responsibilities for significant social media intermediaries. Government explanatory materials also identify sexually explicit deepfakes and AI-generated non-consensual intimate imagery as serious examples of harmful synthetic use.10

Transparency is valuable but limited. Labelling a permissible synthetic video can reduce deception. Provenance metadata can assist users, platforms and investigators in understanding how content was produced. But a label cannot transform non-consensual sexual exploitation into lawful expression. A watermark stating that an image is AI-generated does not restore the victim's autonomy or erase humiliation. Transparency mitigates deception; it does not cure

8Shreya Singhal v. Union of India, (2015) 5 S.C.C. 1.
9Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, Gazette of India, G.S.R. 120(E) (Feb. 10, 2026), corrected Feb. 26, 2026.
10Ministry of Electronics & Information Technology, Government of India, FAQs on Amendments to the IT Rules, 2021 Relating to Synthetically Generated Information (Feb. 10, 2026).

abuse.

The 2026 rules should therefore be understood as a regulatory floor rather than the completion of the project. Their effectiveness will depend on implementation, platform design, grievance procedures, evidentiary preservation and the relationship between synthetic-content obligations and existing criminal law. The central policy question is no longer whether synthetic information should be recognised, but how recognition can be translated into effective victim outcomes.11

VIII. PLATFORM ACCOUNTABILITY AND THE PERSISTENCE GAP

Platforms are not merely passive locations in the social life of a deepfake. Their architecture determines discoverability, recommendation, virality, searchability and recurrence. This does not justify automatic liability for user conduct, but it explains why platform procedure is central to victim protection.

Conventional notice-and-takedown mechanisms often operate at the level of a particular URL, post or account. Synthetic intimate abuse exposes the weakness of that model. Once the content has been downloaded and redistributed, the victim may repeatedly report materially identical copies. The Article calls the resulting mismatch the persistence gap: the distance between removal of an identified instance and meaningful reduction of continuing circulation.

A victim-centred system should therefore distinguish source removal, distribution interruption, recurrence mitigation, search visibility and perpetrator accountability. These are not interchangeable remedies. Removing the first post may stop one pathway but leave copies intact. Suspending an account may not affect mirrored content. Delisting may reduce discoverability without eliminating the underlying file. Effective response requires a combination proportionate to the circumstances.

For confirmed unlawful non-consensual synthetic intimate material, platforms should consider technically feasible recurrence-mitigation tools, including privacy-preserving fingerprinting or matching mechanisms, subject to safeguards. Such systems must not become indiscriminate filters for controversial speech. Review procedures, context-sensitive exceptions and mechanisms for correcting false matches are essential.

11Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021, as amended Feb. 10, 2026 and corrected Feb. 26, 2026.

Recommendation systems raise a further issue. Where a platform has reliable notice that particular content constitutes prohibited synthetic intimate abuse, continued algorithmic amplification can deepen the injury. Accountability should therefore consider not only hosting but the platform's response after legally relevant knowledge, while preserving the statutory conditions governing intermediary protection.

IX. DATA PROTECTION AND THE LIMITS OF A DATA-CENTRIC MODEL

The Digital Personal Data Protection framework is relevant because deepfake systems may process facial images, voice samples and other digital personal data associated with identifiable individuals. Questions of lawful processing, consent, purpose and responsibility can therefore arise in the broader ecosystem of synthetic-media generation.12

Yet deepfake injury cannot be reduced to data processing. A photograph may have been lawfully available to the public, but the synthetic sexual representation constructed from it can still be profoundly abusive. The legal interest changes when data is transformed into a fabricated social claim about the person. Data protection asks, among other things, when information about a person may be processed; synthetic dignity asks when technology may be used to manufacture a harmful false version of that person.

The distinction is especially important where source material is publicly accessible. Public visibility should not be treated as universal consent to downstream synthetic sexualisation. A person who posts a professional portrait does not thereby authorise its use in a nudification service. Legal analysis must separate access to source data from permission to create the resulting representation.

Data protection is therefore one component of a layered response. It can regulate parts of the input and processing chain, while criminal law, platform regulation and civil or injunctive remedies address the harmful output and its circulation.

X. EVIDENCE IN THE AGE OF SYNTHETIC REALITY

The Bharatiya Sakshya Adhiniyam, 2023, also in force from 1 July 2024, modernises India's general evidence framework and recognises electronic and digital records. Deepfake disputes will increasingly test the difference between the existence of a digital file and the reliability of

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

what that file purports to represent.13

Two opposite dangers arise. The first is false acceptance: a fabricated recording may be treated as authentic because it appears visually persuasive. The second is false rejection: genuine evidence may be dismissed merely because modern technology makes fabrication possible. The second problem is sometimes described as the liar's dividend. As synthetic media becomes common knowledge, a wrongdoer can attempt to manufacture doubt around authentic evidence.14

Courts should therefore avoid treating any single automated detector as dispositive. Forensic assessment may consider provenance, metadata, source history, file integrity, editing traces, device records, platform logs, expert analysis, chain of custody and corroborating evidence. Detection systems may assist, but they can produce false positives and false negatives and may degrade as generation techniques evolve.

The evidentiary problem intersects directly with takedown. Victims understandably seek immediate removal, yet investigators may need account data, timestamps, original files and distribution records. This is the removal-preservation paradox: the faster the content disappears, the greater the risk that evidence disappears with it. A sound framework should allow public access to be restricted while relevant evidence is preserved under lawful safeguards.

Evidence preservation should also minimise secondary victimisation. Authorities should not require repeated unnecessary copying or viewing of synthetic intimate material. Secure handling, controlled access and confidentiality are not merely administrative preferences; they are part of a victim-centred justice process.

XI. THE VICTIM'S JOURNEY AND REMEDIAL FRAGMENTATION

A doctrinal map can make the legal system appear more coherent than it feels to a victim. The lived sequence begins differently. A woman learns that an image is circulating. She may not know who created it, where it originated, whether it has been copied, or which legal category applies. She must first prove to a platform that she is the person depicted and that the content

13Bharatiya Sakshya Adhiniyam, No. 47 of 2023, INDIA CODE (2023) (brought into force July 1, 2024).
14See Anvar P.V. v. P.K. Basheer, (2014) 10 S.C.C. 473; Arjun Panditrao Khotkar v. Kailash Kushanrao Gorantyal, (2020) 7 S.C.C. 1 (addressing electronic evidence under the predecessor evidentiary regime).

is synthetic or non-consensual. She may then approach police, seek legal advice, contact search engines, notify employers or family, and repeat the same explanation across institutions.

Each institution sees a fragment. The platform sees a content-policy complaint. Police see a possible offence. A court sees evidence and legal elements. A data-protection authority may see processing. A search engine sees indexing. The victim sees one attack. This institutional disaggregation is what the Article calls remedial fragmentation.

Fragmentation creates delay, duplication and evidentiary risk. A platform may remove material before police obtain records. Police may ask the victim to identify an anonymous creator whom only the platform can trace. A search engine may delist one page while mirrors remain. A victim may obtain a finding that the content is false but still face continuing association between her name and the fabricated material.

The solution is not a single super-agency. It is procedural coordination around predictable victim needs. Those needs are prevention where feasible, rapid containment, preservation of evidence, confidentiality, identification of culpable actors and remedies capable of reducing recurrence. The VCDRF is designed around that sequence.

XII. THE VICTIM-CENTRIC DEEPFAKE RESPONSE FRAMEWORK (VCDRF)

The Victim-Centric Deepfake Response Framework is proposed as an organising model rather than a replacement statute. It can guide legislative reform, platform procedure, law-enforcement protocols and judicial remedies. Its five duties correspond to the stages at which the present system most often fails victims.

First is the duty of prevention. Services capable of generating realistic representations of identifiable people should take reasonable and proportionate measures against predictable attempts to create clearly unlawful non-consensual intimate material. Prevention must be risk-sensitive and should not require universal monitoring of lawful creative activity. The stronger the service's knowledge that its design is being used for sexual identity abuse, the stronger the case for targeted safeguards.

Second is expedited containment. Victims should have an accessible reporting route specifically capable of handling synthetic impersonation and intimate abuse. A complaint should permit rapid identification of the content, the person represented and the basis of non-

consent. Where the material falls within a clearly prohibited category, review should be prioritised. Containment may include restriction, removal, disabling access or other proportionate steps under applicable law.

Third is evidentiary preservation. Removal should not erase the pathway to accountability. Platforms should preserve relevant records when serious synthetic abuse is credibly reported, subject to lawful retention limits, privacy safeguards and valid investigative process. The objective is to preserve what may identify the creator, uploader or dissemination chain without leaving the harmful content publicly accessible.

Fourth is victim protection. Complaint systems, police procedures and courts should minimise unnecessary reproduction of intimate synthetic material. Victim identity should receive appropriate confidentiality protection. Processes should recognise that the harm can extend to employment, education, family relations and physical safety even after falsity is established.

Fifth is accountable redress. Depending on the law and facts, redress may include criminal investigation, injunctive relief, account action, recurrence mitigation, correction of search visibility, compensation where available, and appeal when a platform fails to act. The objective is not merely deletion of one file. It is restoration of meaningful control to the person whose identity has been appropriated.

XIII. NON-CONSENSUAL SYNTHETIC INTIMATE CONTENT (NCSIC)

The Article proposes recognition of Non-Consensual Synthetic Intimate Content as a distinct regulatory category. NCSIC should encompass realistic synthetically generated or materially manipulated intimate or sexually explicit audio-visual material that represents an identifiable natural person and is created or disseminated without that person's consent to the intimate synthetic representation.

The definition should not require proof that the depicted body, nudity or sexual act was authentic. That requirement would recreate the authenticity paradox. The law should focus on whether an identifiable person has been non-consensually represented in intimate synthetic content, not whether the represented event existed in physical reality.

A carefully drawn NCSIC category can also protect expression. It should exclude wholly fictional persons who are not reasonably identifiable as real individuals, and it should be subject

to appropriate public-interest, law-enforcement, judicial, research, reporting and other lawful-use safeguards. Context matters, particularly where synthetic material is used to expose or discuss abuse rather than to perpetrate it.

The category would improve procedural clarity. Platforms could create dedicated reporting pathways; police could use standard preservation protocols; courts could assess urgency without forcing victims to fit fabricated intimacy into concepts designed for authentic recordings. The purpose is not to create a technology-specific moral offence but to recognise a recurring structure of identity-based harm.

XIV. A DIFFERENTIATED LIABILITY SPECTRUM

Deepfake regulation should resist the temptation to identify a single responsible actor. Synthetic abuse often involves a chain. The creator generates the file. A commissioner may procure it. An uploader introduces it to a network. Redistributors expand its reach. A commercial service may facilitate generation. A platform may receive notice and decide how to respond.

The malicious creator ordinarily bears the clearest responsibility where intent and identity appropriation are established. A commissioner who deliberately procures synthetic sexual content for revenge, harassment or extortion should be assessed according to participation and applicable criminal doctrines. The initial publisher materially increases exposure. A knowing redistributor may contribute to continuing harm where dissemination is deliberate and culpable.

Platform responsibility should remain connected to statutory duties, knowledge and control rather than strict liability for every user upload. Safe harbour and victim protection need not be opposites. A platform can retain protection for good-faith hosting while being expected to comply with clearly defined due-diligence and response obligations.

AI service providers require similarly careful treatment. General-purpose model development should not automatically create liability for every downstream misuse. But a service specifically designed, marketed or knowingly operated to create non-consensual sexual representations of identifiable persons presents a different case. Technological neutrality should not become a formal label that shields deliberate facilitation.

A differentiated spectrum therefore aligns responsibility with conduct, knowledge, intent,

capability and control. It avoids both over-criminalisation and under-accountability.

XV. COMPARATIVE LESSONS WITHOUT LEGAL TRANSPLANTATION

Comparative law can illuminate regulatory techniques, but India should not mechanically import foreign rules. Different jurisdictions allocate responsibility among criminal law, privacy, platform regulation, intellectual property and civil remedies in different ways. The useful question is functional: which techniques solve problems that Indian victims currently encounter?

European approaches to platform governance and AI transparency illustrate the value of provenance, systemic risk management and procedural accountability. United Kingdom developments concerning intimate-image abuse show the importance of defining prohibited conduct without requiring victims to prove every traditional privacy assumption. United States federal and state initiatives demonstrate growing attention to digital replicas and non-consensual intimate imagery, though constitutional and federalism structures differ significantly from India. South Korean responses to sexually explicit deepfakes illustrate the use of targeted criminal measures in a context of acute gendered digital abuse.

Three comparative lessons are especially transferable. First, transparency rules should coexist with prohibition of clearly abusive synthetic content. Second, victim remedies must operate quickly enough to matter in a viral environment. Third, liability should distinguish malicious creators and knowing disseminators from neutral infrastructure and lawful expression.

India's constitutional structure, statutory framework and scale of digital-platform use justify a locally tailored model. The 2026 SGI amendments already provide an Indian regulatory foundation. Reform should build upon that foundation rather than imitate foreign statutes word for word.

XVI. PROPOSED INDIAN LEGISLATIVE AND PROCEDURAL REFORMS

The first reform is definitional clarity. Indian law should expressly recognise non-consensual synthetic intimate content involving identifiable persons without requiring an authentic intimate event. The definition should be technologically neutral enough to survive changes in generation methods.

Second, complaint architecture should be unified at the victim-facing level. Platforms should provide a visible route for synthetic intimate abuse, and government guidance should standardise the minimum information necessary for a credible complaint. Victims should not need technical expertise to invoke protection.

Third, takedown and preservation should be paired. A lawful mechanism should permit immediate restriction of clearly prohibited content while preserving relevant technical records for investigation. Retention should be limited, secure and subject to legal process.

Fourth, recurrence mitigation should be available for confirmed unlawful content. Exact or near-identical reuploads should not force the victim to restart the entire complaint process. Technical tools should include safeguards against overblocking and allow legitimate evidentiary, journalistic or judicial uses.

Fifth, the law should differentiate roles. Creation, procurement, initial publication, knowing redistribution, commercial facilitation and post-notice platform conduct should not be treated as legally identical. Liability should correspond to culpability and control.

Sixth, confidentiality should be built into procedure. The victim's name and the synthetic intimate material should not become more widely distributed through the very process intended to provide relief. Police, prosecutors, courts and platforms should adopt minimisation principles for viewing, copying and disclosure.

Seventh, forensic capacity must improve. Investigators and judges require continuing education concerning provenance, metadata, model limitations, detection uncertainty and the liar's dividend. A confident-looking automated score should never replace reasoned evidentiary assessment.

Eighth, remedies should address reputation and persistence. Depending on legal authority, courts should be able to craft proportionate orders concerning removal, recurrence, search visibility, preservation and identification of culpable actors. The goal should be practical restoration, not symbolic victory after the content has become permanently replicated.

Ninth, lawful expression requires explicit safeguards. Satire, parody, artistic transformation, accessibility uses, research, reporting and evidence should not be swept into an overbroad prohibition merely because synthetic techniques are involved. Harm, identifiability, consent, context, knowledge and purpose should guide the boundary.

Finally, regulatory review should be continuous. Generative technology evolves quickly. Rules should be evaluated against measurable victim outcomes: response time, recurrence, evidence preservation, false removal, accessibility of complaint systems and successful coordination with lawful investigations.

XVII. DIGITAL DIGNITY IN THE GENERATIVE ERA

The deepest legal problem presented by deepfakes concerns the boundary of personhood in a world where identity is reproducible. A face can be synthesised, a voice cloned, gestures simulated and speech manufactured. These capabilities make it possible to separate the social appearance of a person from the person's own conduct.

The law cannot respond by granting absolute ownership over identity. Public figures must remain open to criticism and satire; journalism must be able to report; artists must be able to transform; courts must be able to use demonstrative material. But meaningful liberty also requires protection against serious non-consensual appropriation. A society in which anyone can manufacture convincing sexual or criminal representations of another person without effective remedy would impoverish, rather than expand, expressive freedom.

Representational autonomy offers a balanced principle. It does not prohibit representation. It protects against serious synthetic appropriation where identity is used as the raw material for sexual exploitation, coercion, fraud or comparable harm. The principle is compatible with the constitutional values of dignity, autonomy and expression because it focuses on the character of the interference rather than the mere use of artificial intelligence.

This perspective also clarifies why women's deepfake abuse is not a niche technology issue. It concerns participation in digital society. If ordinary photographs can be weaponised into fabricated sexual content, the burden may fall on women to withdraw from public visibility, professional networking and online expression. Effective protection therefore supports, rather than restricts, equal digital participation.

XVIII. CONSENT, SOURCE MATERIAL AND THE PROBLEM OF TRANSFORMATIVE MISUSE

Consent is one of the most easily misunderstood ideas in synthetic-media disputes. A woman may consent to being photographed, to uploading a photograph to a public profile, or to appearing in a video. None of those decisions necessarily includes consent to transform her

likeness into an intimate or sexually explicit synthetic representation. The legal analysis should therefore distinguish consent to the source material from consent to the transformation, and both from consent to publication of the transformed output. Treating these as a single act of consent would permit a narrow permission to become an unlimited licence over identity.

This distinction becomes particularly important when a perpetrator uses material taken from an open social-media account. Public accessibility changes the practical expectation that strangers may view or copy a photograph, but it should not erase the legal significance of purpose. The fact that an image can be downloaded does not answer whether it may be algorithmically reconstructed into fabricated nudity. A purpose-sensitive analysis is more consistent with autonomy: it recognises that permission to participate in ordinary digital life does not amount to permission to be synthetically sexualised.

The same reasoning applies to voice cloning. A speech, interview, lecture or public video may supply sufficient audio to reproduce a recognisable voice. The source may be entirely lawful. The harmful act arises when the voice is made to say words the speaker never uttered, particularly where the fabrication is used for sexual humiliation, fraud, intimidation or reputational sabotage. In such cases, the law should ask not only whether the source was public but whether the downstream representation materially altered the meaning attributable to the identifiable person.

A consent architecture for NCSIC should therefore be representation-specific. Consent should relate to the intimate synthetic representation itself and, where relevant, its dissemination. It should be freely given and capable of being evaluated in context. A general platform term buried in a standard-form agreement should not be treated as meaningful permission for third parties to manufacture sexualised representations of identifiable users. The more serious the interference with dignity and sexual autonomy, the stronger the justification for requiring clear and specific consent.

This approach also prevents a harmful burden shift. Without it, victims may be asked why they placed photographs online, as though digital visibility were the legal cause of abuse. The relevant inquiry should instead focus on the conduct of the person who transformed or disseminated the material. Digital participation is now integral to education, employment, professional networking and social life. Law should not force women to choose between public participation and protection from synthetic sexual exploitation.

XIX. PROCEDURAL DESIGN: FROM COMPLAINT TO EFFECTIVE REMEDY

Substantive offences cannot protect victims if procedure moves more slowly than the technology. A useful deepfake response must therefore be designed around time. The first hours after discovery may determine whether a synthetic image remains confined to a small group or is replicated across platforms. The legal system should distinguish urgent containment from final adjudication. Temporary restriction of apparently unlawful intimate synthetic material can be subject to subsequent review without requiring the victim to prove every element of a criminal offence before the first protective step is taken.

A standardised complaint should permit a victim or authorised representative to identify the relevant content, provide evidence of identity, explain the absence of consent and indicate whether there is a threat, extortion demand or safety concern. Platforms should avoid demanding the very intimate source material that the victim says never existed. Where the allegation concerns a fully fabricated sexual image, insisting upon an authentic comparison image may deepen the violation and misunderstand the nature of the complaint.

The complaint process should also generate a preservation event. Once a sufficiently credible report of serious synthetic intimate abuse is received, relevant technical information should not disappear merely because public access is restricted. Preservation can include the reported file, upload time, account identifiers, available IP or device information, hashes or matching signatures, and records of enforcement action, subject to applicable law and proportionality. Access to preserved information should remain governed by lawful process rather than being made available indiscriminately.

Inter-agency coordination is equally important. Victims should not have to translate their experience separately into the language of every institution. Model protocols could explain when police should seek preservation, when platforms should retain records, how courts may handle confidential exhibits, and how search or indexing remedies can complement source removal. A coordinated protocol would not change the elements of criminal offences; it would reduce the procedural gaps through which accountability is presently lost.

Appeal and correction mechanisms are necessary on both sides. A person whose lawful satire or journalism is removed should have access to review. A victim whose credible complaint is rejected should receive a reasoned pathway for escalation. Procedural fairness strengthens

rather than weakens victim protection because it makes rapid measures more defensible and reduces incentives for platforms to adopt either indiscriminate removal or institutional inaction.

XX. COMMERCIAL FACILITATION AND AI-SERVICE RESPONSIBILITY

The emerging market for generative services complicates the traditional distinction between speaker and intermediary. Some services are general-purpose systems capable of many lawful uses. Others may be configured, promoted or monetised around the production of sexualised representations of real people. Legal responsibility should take this difference seriously. A provider should not be liable merely because a versatile tool can be misused, but deliberate commercial facilitation of predictable identity-based sexual abuse presents a stronger case for intervention.

Relevant factors may include the provider's stated purpose, design choices, knowledge of repeated misuse, availability of safeguards, response to complaints, monetisation model and degree of control over outputs. A service that advertises the ability to "undress" identifiable persons or encourages users to upload photographs of real women is not situated like a general image editor used incidentally by a wrongdoer. Functional analysis is preferable to labels such as "platform," "developer" or "AI company," because the same label can conceal very different forms of control.

Reasonable safeguards need not require perfect prevention. No technical measure can eliminate all misuse. The legal standard should instead consider whether the provider took proportionate steps in light of foreseeable risk. Those steps might include restrictions on certain prompts or workflows, controls on the processing of identifiable real persons in sexual contexts, abuse reporting, preservation after credible complaints, and action against accounts repeatedly engaged in prohibited conduct.

Care is nevertheless required to avoid imposing a general obligation to inspect all private creative activity. Such a regime could be technologically unrealistic and constitutionally problematic. Risk-based obligations are more defensible when directed at services or features whose operation creates a demonstrable and specific danger. Regulation should target facilitation of unlawful conduct rather than artificial intelligence as a category.

Commercial responsibility also matters because synthetic abuse can be industrialised. When

generation is offered at scale, a single service can lower the cost of producing thousands of abusive representations. The law's response must therefore consider not only individual perpetrators but the business models that knowingly convert another person's identity into a commodity for sexual exploitation.

XXI. SEARCH, RECOMMENDATION AND THE AFTERLIFE OF A DEEPFAKE

Removal from the original host does not necessarily end a deepfake's social life. Search results, cached descriptions, screenshots, reposts, discussion threads and recommendation systems can preserve the association between the victim's name and the fabricated content. The law should therefore recognise an afterlife of synthetic abuse: reputational injury can persist after the original file is no longer readily available.

Search visibility is especially important because many people encounter information through queries about a person's name rather than through the original upload. A proportionate remedy may therefore need to address indexing of confirmed unlawful material or obvious mirrors, while preserving legitimate reporting about the existence of the incident. The objective is not to erase public discussion of deepfake abuse. It is to prevent the abusive representation itself from remaining permanently discoverable merely because copies have migrated.

Recommendation systems create a related but distinct concern. A platform may not create the content, yet automated distribution can expose it to audiences far beyond the uploader's followers. Once a platform has reliable notice that particular material falls within a prohibited category, continued recommendation may aggravate harm. Due diligence should therefore include reasonable steps to prevent confirmed unlawful synthetic intimate content from being algorithmically promoted.

These measures must remain content-specific. The existence of a deepfake controversy can itself be newsworthy. Journalists, researchers and advocates may need to describe or analyse the event. A victim-centred approach should distinguish circulation of the abusive file from discussion of the abuse. That distinction protects both dignity and public discourse.

The persistence gap ultimately demonstrates why deletion cannot be the sole metric of success. A meaningful remedy asks whether the victim's exposure is decreasing, whether identical copies are recurring, whether evidence is preserved, whether perpetrators can be identified and

whether the association is being amplified by platform systems. These outcome-oriented questions should guide future regulatory evaluation.

XXII. INSTITUTIONAL CAPACITY, FORENSICS AND JUDICIAL CAUTION

Deepfake regulation will succeed or fail partly through institutional capacity. Police officers, prosecutors, judges and platform reviewers require sufficient technical literacy to understand both the capabilities and limitations of synthetic-media tools. Overconfidence is dangerous in both directions. A convincing video should not be assumed authentic merely because it looks natural, and a genuine recording should not be rejected merely because artificial generation is possible.

Forensic expertise should be treated as one part of an evidentiary mosaic. Detection models can assist investigators, but their performance depends on the type of media, compression, generation technique and availability of reference data. Courts should be cautious about categorical conclusions based on opaque scores. Where authenticity is disputed, reasoning should consider provenance, metadata, corroboration, acquisition history, device information and expert methodology.

Training should also address the human consequences of evidence handling. Synthetic intimate images can be psychologically harmful even though the depicted event never occurred. Repeatedly requiring a victim to display or explain the material can produce secondary victimisation. Standard operating procedures should minimise unnecessary exposure and limit access to personnel who require the material for legitimate investigative or adjudicative purposes.

Judicial remedies may need technical specificity. An order simply directing “removal from the internet” may be impossible to implement and too broad to supervise. More precise orders can identify accounts, URLs, files, hashes, search results or categories of materially identical reuploads, while specifying preservation and lawful exceptions. Precision improves enforceability and reduces the risk of suppressing unrelated lawful speech.

Institutional capacity also requires periodic review. Generative systems change rapidly, and the indicators useful for detecting one generation of synthetic media may become obsolete. Continuing education should therefore be built into cybercrime and judicial training rather than

treated as a one-time response to a temporary technological phenomenon.

XXIII. A MODEL VICTIM-CENTRIC LEGAL TEST

The proposals developed in this Article can be translated into a practical legal test for serious synthetic identity abuse. The first question is identifiability: would a reasonable observer connect the representation to a particular natural person? The second is synthetic materiality: has artificial generation or material manipulation attributed conduct, speech, nudity or circumstances that are not an authentic representation of that person? The third is consent: did the represented person consent to the relevant transformation and, where applicable, dissemination?

The fourth inquiry concerns harm and context. Sexualisation, coercion, fraud, threats, reputational sabotage and targeted humiliation strongly support intervention, while satire, obvious fiction, public-interest reporting and legitimate research may point in a different direction. The fifth concerns culpability: what did the creator, commissioner, publisher or redistributor know and intend? The sixth concerns platform responsibility: after legally relevant notice or under applicable due-diligence rules, what reasonable control did the service possess and how did it respond?

This test does not replace statutory elements. Criminal liability must always be established under the offence charged. Its value is organisational. It prevents the analysis from beginning and ending with the technical question whether a file is "AI-generated." Artificiality is only one fact. The legal wrong emerges from the interaction of identity, falsification, consent, harm, knowledge and dissemination.

The same test can assist platforms and policymakers because it is technologically neutral. Whether a representation is generated by diffusion, face swapping, voice synthesis or a future technique is less important than what the representation does to an identifiable person. Technology-specific definitions become obsolete quickly; harm-centred concepts can remain stable while technical methods evolve.

A victim-centric legal test also reinforces proportionality. The strongest interventions are reserved for the clearest combinations of identifiability, non-consent, serious harm and culpability. Borderline expressive cases receive greater contextual scrutiny. This structure is

preferable to treating every synthetic representation as presumptively unlawful.

The framework developed across these sections is intentionally cumulative. Synthetic dignity harm identifies the protected interest; NCSIC supplies a precise category for fabricated intimate abuse; the authenticity paradox explains why older image-based concepts can misfire; the persistence gap shifts attention from single-post deletion to continuing circulation; the removal-preservation paradox connects urgent victim protection with evidentiary integrity; and the VCDRF brings these concerns together in a procedural architecture. None of these labels is presented as an already established doctrine of Indian law. Their purpose is to provide a coherent vocabulary through which existing law and future reform can respond to a form of abuse that crosses conventional doctrinal boundaries.

XXIV. CONCLUSION

Deepfake technology does not create the human impulse to deceive, humiliate or exploit. What it changes is the power available to the wrongdoer. Generative artificial intelligence can give a lie another person's face, voice and apparent body. It can manufacture an event that never happened and attach its consequences to someone who never participated in it.

For women targeted through synthetic sexual abuse, the resulting injury cannot adequately be described as false information. The perpetrator appropriates identity itself. The victim may be forced to defend herself against an event that never existed, prove the falsity of a body she never exposed, and pursue copies she never created. The technology transforms identity into an involuntary instrument of abuse.

Indian law has moved meaningfully toward recognition of this problem. The Bharatiya Nyaya Sanhita supplies criminal-law tools. The Information Technology Act continues to regulate important electronic offences and intermediary responsibility. The Bharatiya Sakshya Adhiniyam provides the modern evidentiary architecture. Constitutional privacy jurisprudence supplies a foundation in dignity and autonomy. Data-protection law addresses aspects of personal-data processing. Most significantly, the 2026 intermediary-rule amendments expressly recognise synthetically generated information and establish tailored regulatory duties.15

15See Bharatiya Nyaya Sanhita, No. 45 of 2023, INDIA CODE (2023); Information Technology Act, No. 21 of

But legal coverage is not the same as effective protection. The victim does not encounter deepfake abuse statute by statute. Her image is obtained; her identity is reconstructed; content is generated; it is uploaded, recommended, copied and forwarded; she discovers it; she reports it; one copy disappears; another appears; evidence may vanish; the creator may remain anonymous; and the association may survive long after the content has been proved false.

The Victim-Centric Deepfake Response Framework responds to that reality by beginning with the person rather than the file. Its five duties—prevention, expedited containment, evidentiary preservation, victim protection and accountable redress—are intended to connect the fragmented legal mechanisms that already exist and identify the areas where further reform is required.

The proposed category of Non-Consensual Synthetic Intimate Content addresses the authenticity paradox by making clear that protection should not depend upon whether the represented sexual event actually occurred. The concepts of the persistence gap and removal-preservation paradox expose why ordinary notice-and-takedown processes may be insufficient. Differentiated liability aligns responsibility with culpability, knowledge and control.

The future of deepfake regulation in India should therefore move beyond the narrow question whether a digital object is real or artificial. The deeper question is whether law can preserve meaningful human autonomy in an environment where identity can be reproduced without participation by the person to whom it belongs.

The answer should not be a prohibition on synthetic creativity. It should be a principled boundary around synthetic abuse. The ability to simulate a human being should never be mistaken for an unrestricted legal entitlement to appropriate that human being's identity.

2000, INDIA CODE (2000); Bharatiya Sakshya Adhiniyam, No. 47 of 2023, INDIA CODE (2023).