
DEEPPAKES AND THE LAW: HOW INDIA'S REGULATORY RESPONSE COMPARES, AND WHERE IT STILL FALLS

SHORT

Arun, MVN University, Aurangabad, Haryana

ABSTRACT

Between September 2023 and October 2025, five different Indian actors, a classical singer, and a stock exchange chief executive each had to approach a High Court individually to obtain an injunction against an AI-generated video or voice clone impersonating them, because no single statute gave any of them a direct cause of action against the specific harm a deepfake causes. This paper examines how Indian law arrived at that position, and what has changed since. It traces the gap between the Information Technology Act, 2000's cheating and obscenity provisions, none of which was drafted with synthetic media in mind, and the personality rights jurisprudence the Delhi and Bombay High Courts built case by case to fill that gap, beginning with the ex-parte omnibus injunction granted to Anil Kapoor in September 2023 and extending through the interim relief granted to Akshay Kumar two years later. It then turns to India's first deepfake-specific regulatory instrument, the amendments to the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021 notified in February 2026, which introduce a statutory definition of synthetically generated information and impose continuous labelling and metadata obligations on intermediaries. The paper places this framework alongside two different regulatory choices made elsewhere: the European Union's Article 50 of the Artificial Intelligence Act, which imposes a disclosure duty on the deployer of a deepfake rather than a labelling duty on the platform that hosts it, and the United Kingdom's decision to criminalize the creation, not merely the sharing, of non-consensual intimate deepfakes. It argues that India's labelling-first model addresses the harm regulators have found easiest to define, misleading synthetic content in general circulation, while leaving the harm that has generated the most litigation, non-consensual and often sexualized impersonation of identifiable individuals, to be handled through civil litigation that only a well-resourced plaintiff can realistically bring. The paper closes with a proportionate set of legislative proposals addressed specifically to that gap.

I. Introduction

In September 2023, the actor Anil Kapoor approached the Delhi High Court because AI tools were being used to generate his voice, face, and catchphrase for commercial products he had never endorsed, and no single Indian statute gave him a direct legal claim against that specific conduct. He got an ex-parte, near-universal injunction instead, framed around personality rights and passing off, doctrines the Indian judiciary had built up over two decades for a different set of problems entirely.¹ Over the following twenty-five months, the actor Jackie Shroff, the actor Akshay Kumar, the singer Arijit Singh, the entrepreneur Ankur Warikoo, and the managing director of the National Stock Exchange each found themselves in the same position, filing broadly similar suits before the Delhi or Bombay High Courts because the same statutory gap that Kapoor had encountered was still there.² This paper takes that pattern, six separate plaintiffs, six separate lawsuits, one recurring gap, as its starting point, because it says something more precise about the state of Indian law than the general observation that deepfakes are a growing problem. It says that as of the end of 2025, India's courts, not its legislature, were doing the work of defining what a deepfake violates, on a case-by-case basis, for whichever plaintiff had the resources and standing to bring a commercial civil suit.

That has now changed, at least in part. On February 10, 2026, the Ministry of Electronics and Information Technology notified amendments to the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021, introducing India's first statutory definition of synthetically generated information and a set of labelling and metadata obligations that apply across the board, not only to the kind of commercial impersonation the litigation above involved.³ This paper's central argument is that the new rules and the litigation that preceded them are addressed to two different problems that happen to share a technology, and that treating the 2026 amendment as a complete answer to what the personality rights cases were actually about would be a mistake. The amendment is well suited to the problem of a viral video that misleads an audience about whether an event occurred. It is considerably less well suited to the problem of a named individual being sexualized, defamed, or commercially

¹Anil Kapoor v. Simply Life India & Ors., CS(COMM) 652/2023 (Del. H.C., Sept. 20, 2023) (order of Prathiba M. Singh, J.).

²Jaikishan Kakubhai Saraf alias Jackie Shroff v. The Peppy Store & Ors., 2024 SCC Online Del 3664 (May 15, 2024); Akshay Kumar v. John Doe & Ors. (Bombay H.C., Oct. 15, 2025) (interim order of Arif S. Doctor, J.); Arijit Singh v. Codible Ventures LLP (Bombay H.C. 2024).

³Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, G.S.R. 120(E) (Feb. 10, 2026).

exploited by content generated in their likeness without their consent, which is the problem every one of the cases in the previous paragraph actually involved.

The paper proceeds in six parts. Part II sets out briefly what a deepfake is and why it resists straightforward legal categorization. Part III traces the Indian statutory and regulatory response, from the gaps in the Information Technology Act, 2000 through the 2026 amendment to the Intermediary Guidelines. Part IV examines the personality rights litigation that filled the gap before the 2026 amendment existed and asks what it can and cannot do going forward. Part V compares India's labelling-based model to the European Union's disclosure-based model under Article 50 of the Artificial Intelligence Act and the United Kingdom's criminalization-based model under the Online Safety Act, 2023 and its 2026 extension. Part VI draws out why the current Indian framework, taken as a whole, still falls short of the specific harm the litigation in Part IV reveals, and Part VII proposes a proportionate set of changes before the paper concludes.

II. What a Deepfake Actually Is, and Why the Law Struggles with It

A deepfake, in the sense this paper uses the term and consistently with the definition the Indian and European regulatory instruments discussed below have both now adopted, is audio, image, or video content that has been generated or manipulated using a computer system in a way that makes it appear authentic when it is not.⁴ Two features of that definition explain why lawmakers everywhere have struggled to regulate the underlying technology as cleanly as they might like. First, the definition is entirely agnostic about purpose. The same underlying technique that produces a fraudulent video of a bank executive authorizing a wire transfer also produces a film studio's de-aged actor, a language-dubbing tool that keeps an actor's lip movements in sync, and a harmless entertainment filter that swaps two friends' faces in a video for a joke. A rule broad enough to capture the fraud will, unless carefully drawn, also capture the film studio and the joke. Second, and more consequentially for the litigation this paper examines in Part IV, the harm a deepfake causes frequently has nothing to do with whether the audience was deceived at all. A viewer who understands perfectly well that a sexually explicit video of a named actress is synthetic can still cause that actress serious, tangible harm by

⁴Compare Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules 2021, r. 2(1) (wa), as inserted by the 2026 Amendment (n 3) (defining "synthetically generated information"), with Regulation (EU) 2024/1689 (Artificial Intelligence Act), art. 3(60) (defining "deep fake").

creating, sharing, or merely possessing it, because the harm lies in the non-consensual use of her identity and likeness, not in anyone's mistaken belief that the video is authentic.

This second point deserves emphasis because it is precisely where a labelling-based regulatory response, however well designed, runs out of usefulness. A label that says content is synthetically generated solves the deception problem. It does nothing at all for the person whose face, voice, or body has been used without consent, because that person's objection is not that viewers might be fooled; it is that the content was made and circulated using their identity in the first place. Part III traces how India's regulatory response has, so far, been built almost entirely around the first problem, and Part IV shows why the second problem has had to be litigated separately, case by case, through doctrines never designed for it.

III. India's Statutory and Regulatory Response

A. The Pre-2025 Position: Gaps in the IT Act and Criminal Law

Before 2025, no Indian statute used the word deepfake or any equivalent term, and a lawyer advising a deepfake victim had to work with tools built for other purposes. Section 66D of the Information Technology Act, 2000 penalizes cheating by personation using a computer resource, but it requires the kind of fraudulent inducement associated with online scams, not the reputational or dignitary harm caused by a synthetic video that deceives no one about its subject's identity even as it misrepresents what that person supposedly said or did.⁵ Section 66E, which penalizes the capture, publication, or transmission of an image of a person's private area without consent, requires an actual image of the actual person, a requirement a synthetic composite does not straightforwardly satisfy. Section 67 and Section 67A, covering obscene and sexually explicit material respectively, apply readily enough to sexualized deepfakes once they are published, but neither section was drafted with the specific, identity-based harm of non-consensual synthetic pornography in mind, and neither offers a mechanism for rapid takedown independent of the ordinary, and often slow, criminal process.⁶ The Bharatiya Nyaya Sanhita, 2023, which replaced the Indian Penal Code, similarly offers defamation, cheating, and forgery provisions that can be stretched to cover particular deepfake scenarios but that were, again, not drafted with the technology in view.

⁵Information Technology Act 2000, No. 21 of 2000, INDIA CODE, sec. 66D.

⁶Information Technology Act 2000 (n 5) ss. 67, 67A.

B. The 2021 Intermediary Guidelines and the 2023 Amendment

The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021 supplied the first indirect regulatory hook, requiring intermediaries, as part of their due diligence obligations, not to host content that is patently false, misleading, or impersonates another person, and a 2023 amendment sharpened that obligation by adding an explicit due diligence requirement that intermediaries make reasonable efforts to identify synthetically or artificially morphed images of an individual and take such content down once notified.⁷ That 2023 amendment was a real, if modest, improvement, but it operated only as a takedown obligation triggered by notice, with no accompanying definition of what counted as synthetic content, no labelling requirement for content that had not yet been flagged, and no mechanism addressed to the deepfake's creator rather than the platform hosting it after the fact. It was, in substance, an extension of the ordinary notice-and-takedown model that already governed other categories of unlawful content, applied to a category of content the Rules did not yet define with any precision.

C. The 2026 Amendment: Synthetically Generated Information

The amendment MeitY notified on February 10, 2026, following a draft released in October 2025 and an extended public comment period, is a considerably more developed instrument. It inserts a formal definition of synthetically generated information into Rule 2(1)(wa), covering any information artificially or algorithmically created, generated, modified, or altered using a computer resource in a manner that makes it reasonably appear authentic or true.⁸ Around that definition, the amendment builds three distinct obligations. First, a labelling obligation: rather than the fixed-size watermark the October 2025 draft had originally proposed, covering at least ten per cent of a visual frame or the first ten per cent of an audio clip's duration, the final rules require synthetically generated information to be labelled clearly and prominently, with visual content carrying a visual label and audio content carrying an audio disclosure, continuously for the duration of the content rather than only at the outset.⁹ Second, a traceability obligation: intermediaries must, where feasible, embed permanent metadata or a unique identifier capable

⁷Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules 2023 (amending r. 3(1)(b) of the 2021 Rules).

⁸Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules 2021, r. 2(1)(wa), as inserted by the 2026 Amendment (n 3).

⁹Freshfields Bruckhaus Deringer, India Targets Deepfakes and AI-Generated Content: Key Changes Under MeitY's 2026 Amendments to the IT Rules (Feb. 20, 2026).

of tracing the computer resource used to generate or alter the content, and must not permit that label or metadata to be removed or modified. Third, a due diligence obligation on significant social media intermediaries specifically, requiring deployment of technology-based measures, including automated tools, to identify synthetically generated information at the point of upload rather than relying solely on user notice after the fact.

Two design choices in the 2026 amendment are worth flagging because they shape everything the rest of this paper argues about the framework's limits. The first is that the amendment is addressed to intermediaries, not to the individual who creates or uploads a deepfake in the first place; its obligations are due diligence obligations for platforms, backed by the threat of losing the safe harbour Section 79 of the Information Technology Act otherwise confers, rather than direct offences the creator of harmful synthetic content personally commits. The second is that the amendment is content-neutral in the sense described in Part II: a label applies equally to a misleading political video and to a non-consensual sexualized deepfake of a private individual, and the obligation is the same, clear and continuous labelling, regardless of which kind of harm is actually in play. Both choices make sense given the amendment's own stated purpose, addressing intermediary accountability for misleading synthetic content at scale, but neither choice speaks directly to the identity-based harm at the centre of the litigation examined next.

D. The Road to the 2026 Amendment: A Court-Prompted Rulemaking Process

The 2026 amendment did not emerge from the ordinary legislative process alone. A public interest litigation before the Delhi High Court, filed amid growing concern over deepfake-driven misinformation and financial fraud, prompted the central government to constitute a dedicated committee to examine the regulatory response to synthetic media, and the government's own status report filed before that Court recorded that the committee held multiple meetings addressing AI detection tools, forensic analysis, watermarking standards, and the accuracy and limitations of existing detection methods, before MeitY released its October 2025 draft amendment.¹⁰ The sequence is worth noting for what it reveals about how this particular rule came to exist: a judicial proceeding, rather than an independent legislative initiative, was the proximate trigger for the executive rulemaking process examined in Part III.C, a pattern this paper has already observed once, in the personality rights litigation of Part

¹⁰See Explanatory Note to the Draft Amendment to the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021 (MeitY, Oct. 22, 2025)

IV, and will observe again in Part VI's account of the framework's remaining gaps. Indian deepfake regulation to date has been substantially reactive, built in response to litigation and public controversy rather than anticipatory of a harm regulators identified independently, and that reactive posture is itself a partial explanation for why the resulting framework, examined in the round in Part VI, addresses the harms that had already become visible through litigation more completely than it addresses harms that have not yet produced a prominent lawsuit.

A separate but related strand of executive guidance addressed deepfakes specifically in the electoral context, an application of the technology this paper has not otherwise discussed but that generated some of the earliest official concern. Ahead of the 2024 general election, MeitY issued advisories in November and December 2023 directing significant social media intermediaries to identify and label AI-generated and deepfake content and to act on user reports within a defined period, advisories the Ministry framed explicitly around preserving electoral integrity rather than around the personality rights concerns that were, at almost exactly the same time, working their way through the Delhi High Court in Anil Kapoor's suit.¹¹ Those advisories were non-binding and preceded any statutory definition of synthetic content, which meant compliance depended entirely on individual platforms' own willingness to act, but they mark the point at which the Indian executive first treated deepfakes as a distinct regulatory category worth addressing in its own right, roughly two years before the 2026 amendment gave that category a definition with legal force.

IV. Filling the Gap Through Litigation: The Personality Rights Cases

A. Anil Kapoor and the Omnibus Injunction

Anil Kapoor's suit against Simply Life India and a further fifteen named defendants alleged a cluster of related harms: unauthorized use of his image and voice to falsely endorse a motivational speaking business, the sale of ringtones and GIFs built from his likeness, cybersquatting on domain names using his name, and the circulation of AI-morphed pornographic content bearing his face.¹² Justice Prathiba M. Singh granted an ex-parte, omnibus injunction restraining not only the named defendants but, in substance, the world at large from using Kapoor's name, voice, image, or likeness through any technological tool,

¹¹Ministry of Electronics and Information Technology, Press Release (Nov. 7, 2023); Ministry of Electronics and Information Technology, Press Release (Dec. 26, 2023).

¹²Anil Kapoor (n 1).

including artificial intelligence, GIFs, or face-morphing, for any unauthorized commercial purpose. The order's reasoning is worth quoting for how directly it names the problem this paper's Part II describes: the Court observed that the technological tools now freely available make it possible for any unauthorized user to imitate a celebrity's persona using artificial intelligence, and that a celebrity is entitled to object to his image, voice, or likeness being portrayed in a dark or degrading manner regardless of whether any viewer is actually deceived by the result.¹³

The doctrinal foundation for this relief was personality rights and the tort of passing off, categories the Delhi High Court had already developed over roughly two decades in cases involving Amitabh Bachchan, the cricketer Sourav Ganguly, and others whose names or images had been used in unauthorized advertising.¹⁴ What made Anil Kapoor's suit different from that earlier line of authority was not the doctrine applied but the technology the doctrine now had to reach: earlier personality rights cases involved a defendant using a plaintiff's actual name or actual photograph without permission, while the AI-generated content in Kapoor's case did not use any authentic recording of him at all. It synthesized a new one. The Court's willingness to extend personality rights doctrine to that scenario was significant precisely because nothing in the doctrine's prior development had specifically contemplated it.

B. Jackie Shroff, Akshay Kumar, and the Pattern That Followed

The cases that followed Anil Kapoor's suit were largely confirmed rather than complicated the pattern it set. In May 2024, Justice Sanjeev Narula granted Jackie Shroff an ad-interim injunction against, among other defendants, an unlicensed AI chatbot built on his persona and a set of accounts distributing what the order described as distorting videos of the actor, and directed the Department of Telecommunications and the Ministry of Electronics and Information Technology to have the infringing links blocked.¹⁵ In October 2025, Justice Arif S. Doctor of the Bombay High Court granted Akshay Kumar ex-parte interim relief against a cluster of websites and John Doe defendants operating AI voice-cloning tools, including one site the actor's counsel demonstrated could generate speech in his voice from any typed text, after a video circulated in which a synthetic version of the actor appeared to make inflammatory

¹³Anil Kapoor (n 1) (order, para. 25).

¹⁴See Titan Industries Ltd. v. Ramkumar Jewellers, 2012 SCC OnLine Del 2382; Sourav Ganguly v. Tata Tea Ltd., 2019 SCC OnLine Cal 1409.

¹⁵Jackie Shroff (n 2).

remarks about a religious figure. Justice Doctor's order described the manipulation as so sophisticated that distinguishing the fake from the authentic had become virtually impossible, and grounded relief in Article 21 of the Constitution alongside the Copyright Act, 1957, in addition to the personality rights doctrine Anil Kapoor had already established.¹⁶ A parallel Bombay High Court order the previous year had gone further still in one respect, directing social media intermediaries to remove deepfake videos of the National Stock Exchange's managing director that were being used to promote fraudulent stock recommendations, and to disclose the identity of the accounts responsible within three weeks, treating the deepfake there as much as a trademark and securities-fraud problem as a personality rights one.¹⁷

C. What These Cases Cannot Do

Six cases across two High Courts, all decided in the plaintiff's favour, might look like a settled and adequate body of law. This paper argues it is neither, for three reasons that follow directly from how each of these suits had to be brought. First, every one of them was a commercial civil suit filed by a plaintiff with the resources to retain senior counsel, identify infringing content across dozens of platforms, and litigate on an urgent, ex-parte basis. Personality rights doctrine, as currently applied, protects the commercial and dignitary interests of people whose name and likeness already carry recognised commercial value; it offers considerably less obvious a route for a private individual, with no public persona to protect and no pre-existing commercial value in their likeness to point to, who is the target of a non-consensual sexualized deepfake.

Second, every one of these orders was directed at specific, identified or identifiable defendants and specific infringing content, and while several included John Doe formulations to reach unknown parties, the practical burden of identifying new infringing content as it appears, and returning to court to have it blocked, still falls on the plaintiff. A wealthy film star can maintain a standing legal team to do this. Most people cannot, and the 2026 Intermediary Guidelines amendment, addressed as it is to platform-level labelling rather than to a private right of action, does not obviously change that calculus for a private individual weighing whether litigation is a realistic option at all.

¹⁶Akshay Kumar (n 2).

¹⁷High Court Directs Social Media Intermediaries to Take Prompt Action on Deepfake Content Impersonating NSE Chief, All India Radio News (July 21, 2024) (order of R.I. Chagla, J.).

Third, and most fundamentally, personality rights doctrine is built around commercial exploitation and dignitary harm to a recognizable public figure. It says relatively little about the deepfake scenario that has generated the most public alarm internationally and that the United Kingdom's legislative response, examined in Part V, treats as its central concern: the creation of non-consensual, sexually explicit synthetic imagery of a private individual with no public profile at all, where the harm is not economic exploitation of a persona but a direct violation of sexual privacy and dignity. Indian personality rights litigation to date has been driven almost entirely by public figures for exactly the reason one would expect, that they are the plaintiffs equipped to bring it, which means the body of precedent Part IV.A and IV.B describe has developed around the harm public figures experience rather than the harm private individuals, and particularly women and girls, are most often reported to experience.

V. Comparative Models: Disclosure, Transparency, and Criminalisation

A. The European Union: Article 50 and the Deployer's Disclosure Duty

Article 50 of the European Union's Artificial Intelligence Act, due to apply from August 2, 2026, takes a structurally different approach from India's 2026 amendment, even though both instruments use the word label. Article 50(2) requires the provider of a generative AI system, meaning the entity that builds and offers the underlying technology, to ensure its outputs are marked in a machine-readable format detectable as artificially generated. Article 50(4) separately requires the deployer of that system, meaning whoever actually uses it to create a deepfake, to disclose to the audience that the content has been artificially generated or manipulated.¹⁸ The deployer-facing disclosure obligation is the more significant point of contrast with the Indian framework: it attaches personal legal responsibility to the specific person who created and published a given deepfake, in addition to whatever obligations fall on the intermediary hosting it, and European Commission guidance issued alongside the Regulation confirms that a platform's machine-readable marking under Article 50(2) does not, on its own, discharge a deployer's separate disclosure duty under Article 50(4).¹⁹

India's 2026 amendment, by contrast, places its principal obligations on the intermediary rather than the individual deployer who actually created the content, an approach that is

¹⁸Regulation (EU) 2024/1689 (Artificial Intelligence Act), arts. 50(2), 50(4).

¹⁹European Commission, Guidelines on Transparency Obligations for Providers and Deployers of Certain AI Systems (2026).

administratively simpler to enforce, since there are far fewer large platforms than there are individual users, but that leaves the individual who actually produced a non-consensual deepfake facing no direct statutory labelling obligation at all. Whether that gap matters in practice depends on how large a share of harmful deepfakes are produced and consumed within the reach of a platform sophisticated enough to run the automated detection the amendment requires, as against content produced and circulated through smaller channels, encrypted messaging, or direct device-to-device sharing that a platform-level rule cannot see in the first place.

B. The United Kingdom: Criminalizing Creation and Sharing

The United Kingdom took a third and considerably more direct route, treating the specific harm this paper's Part IV.C identifies, non-consensual sexualized deepfakes, as a criminal law problem rather than a labelling or disclosure problem from the outset. The Online Safety Act 2023 added a new offence to the Sexual Offences Act 2003, criminalizing the sharing, or threatening to share, an intimate photograph or film, defined to expressly include a manufactured or manipulated image such as a deepfake, without a reasonable belief that the subject would consent.²⁰ That sharing offence was made a priority offence under the Online Safety Act's Schedule 7, which triggers proactive platform obligations under Ofcom's separate regulatory regime rather than the reactive, notice-based obligations the Indian and European frameworks both rely on. In January 2026, the Government went a step further, announcing that the mere creation of a non-consensual intimate deepfake, independent of whether it is ever shared, would itself become a criminal offence, closing a gap the earlier 2023 offence had left open by covering distribution but not production.²¹ Parliamentary evidence cited in support of that extension recorded that roughly two hundred seventy-five thousand intimate deepfake videos had appeared on major platforms in 2023 alone, generating over four billion views between them, a scale that the Government's own framing treated as justifying direct criminalization rather than a platform-facing labelling regime.²²

C. What India Has Chosen, and What It Has Not

Set against these two comparators, India's 2026 amendment reads as a considered but partial

²⁰Sexual Offences Act 2003, c. 42, s. 66B, as inserted by the Online Safety Act 2023, c. 50; Online Safety Act 2023 (Priority Offences) (Amendment) Regulations 2024, amending Sch. 7 of the Online Safety Act 2023.

²¹UK Government cracks down on deepfake image creation, Ministry of Justice press notice (Jan. 2026).

²²UK Tightens Laws on AI Generated Sexual Deepfakes, DPP Law (2026) (citing Parliament evidence).

response. It shares the European Union's preference for a regulatory, disclosure-based mechanism over the United Kingdom's criminal law approach, which is a defensible choice on its own terms, since criminalization carries its own well-documented risks of overreach and chilling effect on legitimate synthetic media use in film, dubbing, and satire. But India has not, even within the disclosure-based model it has chosen, adopted the European Union's deployer-facing obligation, and it has not adopted anything resembling the United Kingdom's specific, targeted criminal offence for non-consensual intimate deepfakes, the exact category of harm Part IV's case survey shows Indian litigants have most urgently needed a remedy for. India's personality rights litigation has, in effect, been doing informally and case by case what the United Kingdom's Sexual Offences Act amendment now does by direct statute and doing so only for the subset of victims able to litigate a commercial civil suit.

A note of caution is due before treating either comparator as a model of settled, working regulation, because both are, in practice, considerably younger and less tested than a first reading suggests. The European Union's own transparency obligations under Article 50 apply only from August 2026, and the Council has separately agreed a further, later deadline of December 2026 for the specific marking and detection obligations under Article 50(2), meaning the provision this paper treats as the more developed comparator model has barely begun to operate as binding law at the time of writing.²³ Commentary on India's own draft rules has separately observed that the European experience with AI-generated content during the 2024 European Parliament elections, where synthetic material circulated despite the Artificial Intelligence Act's disclosure requirements already being publicly announced, suggests that a labelling or disclosure obligation, however well drafted, can be circumvented by a determined bad actor with no interest in complying with it in the first place.²⁴ None of this is a reason to prefer inaction. It is a reason to treat labelling and disclosure obligations, whichever jurisdiction adopts them, as one layer of a response that also needs the criminal law backstop the United Kingdom has built and the access-to-justice mechanism Part VII proposes, rather than as a sufficient response in their own right.

VI. Why the Current Indian Framework Still Falls Short

Three gaps follow from the comparison in Part V, and each is specific enough to be addressed

²³GT Alert, Deepfakes, Chatbots, AI-Generated Text: European Commission Details Transparency Obligations Under the AI Act, Greenberg Traurig (June 2026).

²⁴See Regulating Deepfakes in India: A Tentative First Step, Ronin Legal Consulting (Oct. 2025).

without wholesale redesign of the framework India has already built.

The first is the absence of any direct, deepfake-specific criminal offence covering the non-consensual creation of sexualized or intimate synthetic content, as distinct from the general obscenity provisions in Sections 67 and 67A of the Information Technology Act. Those general provisions require proving that the content is obscene or sexually explicit in the ordinary sense, a threshold that a synthetic but non-explicit deepfake, such as a video falsely showing someone in a state of undress that stops short of the Act's threshold, or one built to embarrass rather than titillate, may not clear at all, even though the identity-based harm to the victim is substantially the same either way.

The second is the absence of any deployer-facing obligation comparable to Article 50(4) of the European Union's Artificial Intelligence Act. India's 2026 amendment regulates the intermediary that hosts synthetic content but says nothing directly to the individual who generated it, leaving that individual's conduct to be addressed, if at all, through the general criminal and tort law provisions Part III.A already showed were not built for this problem.

The third is a practical access-to-justice gap rather than a doctrinal one: personality rights litigation of the kind Part IV describes remains, as a matter of practice rather than formal legal rule, available mainly to plaintiffs who already have the commercial profile and legal resources the doctrine was originally built to protect. A private individual targeted by a non-consensual deepfake has, under the current framework, the 2026 amendment's takedown mechanism to fall back on, which requires proving synthetic origin and waiting on an intermediary's due diligence process, and comparatively little else that does not require replicating, at private expense, the kind of urgent commercial litigation Anil Kapoor and Akshay Kumar were able to bring.

VII. Towards a More Coherent Indian Response

Three proportionate changes, each addressed to one of the gaps identified in Part VI, would close most of the distance between India's current framework and the more complete models examined in Part V, without requiring India to adopt either comparator wholesale.

First, Parliament should consider a narrow, targeted amendment to the Information Technology Act, or a provision within the Bharatiya Nyaya Sanhita, creating a specific offence for the creation or non-consensual distribution of a sexualized or intimate synthetic depiction of an

identifiable person, defined broadly enough to avoid the explicit-content threshold problem identified in Part VI, along the lines the United Kingdom's 2026 extension of its Online Safety Act offences has already tested. This need not displace the existing personality rights doctrine; it would instead supply the criminal law backstop that doctrine, built for commercial and dignitary harm rather than sexual privacy, was never designed to provide.

Second, the 2026 Intermediary Guidelines amendment could be extended, by further amendment rather than fresh legislation, to impose a disclosure obligation on the deployer of a generative AI tool used to create synthetic content depicting an identifiable person, modelled on Article 50(4) of the European Union's Artificial Intelligence Act, so that legal responsibility for a harmful deepfake does not rest solely on the intermediary that happened to host it after the fact.

Third, and most directly responsive to the access-to-justice gap Part VI.C describes, the takedown mechanism the 2026 amendment already gives intermediaries should be paired with a low-cost, non-litigious complaint route for individuals who are not public figures and cannot realistically fund the kind of commercial suit Part IV's case survey describes, modelled on the kind of dedicated grievance mechanism India has already built for other categories of online harm. A framework that requires a private citizen to retain the same calibre of counsel Anil Kapoor retained, in order to obtain the same relief Anil Kapoor obtained, has not actually closed the gap between what public figures and private individuals can secure; it has simply formalized the gap Part IV's litigation already revealed.

VIII. Conclusion

This paper opened with six lawsuits filed over roughly two years, each brought by a different public figure against a different set of deepfake defendants, because Indian law offered no single statutory route any of them could rely on directly. The 2026 amendment to the Intermediary Guidelines is a genuine and overdue step toward closing that gap, and this paper has not argued otherwise. What it has argued is that the amendment closes a different, narrower gap than the one the litigation actually exposed: a labelling and traceability regime addressed to misleading synthetic content in general circulation, built around the intermediary rather than the individual deployer, does real work against the deception harm a deepfake can cause, and comparatively little against the identity-based, often sexualized harm that drove Anil Kapoor, Jackie Shroff, and Akshay Kumar to court in the first place.

The European Union's deployer-facing disclosure duty and the United Kingdom's direct criminalization of non-consensual intimate deepfakes both address that harm more squarely than India's current framework does, though this paper has also flagged that both comparators are themselves too new to be treated as proven successes rather than reasonable attempts. Neither model needs to be imported wholesale. But the specific gaps they fill, and India's framework does not yet fill, are narrow enough to close by amendment rather than by waiting for the next public figure to litigate the next case, and precise enough that failing to close them can no longer be attributed to the novelty of the underlying technology. The six plaintiffs who had to litigate their way to relief between 2023 and 2025 succeeded because each could afford to. The measure of whether India's regulatory framework has actually caught up with the technology it now defines by statute will be whether the seventh person in this position, whoever they turn out to be, and however little public profile or legal budget they have, can expect the same outcome without having to become a plaintiff first.