

**INTERNATIONAL JOURNAL OF LAW
MANAGEMENT & HUMANITIES**
[ISSN 2581-5369]

Volume 9 | Issue 5

2026

© 2026 *International Journal of Law Management & Humanities*

Follow this and additional works at: <https://www.ijlmh.com/>

Under the aegis of VidhiAagaz – Inking Your Brain (<https://www.vidhiaagaz.com/>)

This article is brought to you for free and open access by the International Journal of Law Management & Humanities at VidhiAagaz. It has been accepted for inclusion in the International Journal of Law Management & Humanities after due review.

In case of **any suggestions or complaints**, kindly contact support@vidhiaagaz.com.

To submit your Manuscript for Publication in the **International Journal of Law Management & Humanities**, kindly email your Manuscript to submission@ijlmh.com.

The Rise of AGI and ASI: Threats to Humanity and the Legal Overhaul India Needs

AMIT SINGH SANDHU* AND VELICHETY MANIKHANTA SESHA SAI ANIRUDHA**

ABSTRACT

Artificial intelligence has moved in a few short years from narrow tools to general-purpose systems that write, reason and act. The people building these systems now speak openly of Artificial General Intelligence (AGI) and, beyond it, Artificial Superintelligence (ASI). This Article traces that evolution, sets out the present and future threats it poses to human beings, and examines how five Indian statutes (the Information Technology Act, 2000, the Digital Personal Data Protection Act, 2023, the Bharatiya Nyaya Sanhita, 2023, the Bharatiya Sakshya Adhinyam, 2023 and the Bharatiya Nagarik Suraksha Sanhita, 2023) fall short. It proposes specific amendments to each, along with cross-cutting institutional reforms, so that as machines grow more capable, responsibility and control stay with people.

Keywords: Artificial Superintelligence, required suggestions in IT Act, DPDP Act, BNS, BSA

I. INTRODUCTION: A QUESTION WE CAN NO LONGER POSTPONE

A few years ago, most of us met artificial intelligence through small conveniences. It finished our sentences in email, recommended a song, or told us the fastest route through Hyderabad traffic. Nobody lost sleep over it. Today the same technology drafts legal notices, writes working software, passes professional exams, clones a grandmother's voice from a ten-second clip, and holds conversations that many people find hard to tell apart from talking to a person.

That jump did not take a generation. It took roughly four years. And the people building these systems openly say that the next step is Artificial General Intelligence, or AGI: a machine that can learn and reason across almost any task at least as well as a capable human. Some go further and speak of Artificial Superintelligence, or ASI, a system that would outthink the best human minds in every field at once.

For a country like India, this is not a distant, abstract debate reserved for Silicon Valley. We have more than a billion people online or coming online, a young workforce that depends heavily on IT and business-process services, a justice system already carrying a backlog of

* Author has pursued Law from School of Law, Aurora Deemed to be University, Telangana, India.

** Author has pursued Law from School of Law, Aurora Deemed to be University, Telangana, India.

crores of cases, and democratic institutions that run on public trust. Each of these is exposed to the choices made about AI in the next few years.

The law is where a society writes down what it will and will not tolerate. India has recently rebuilt its criminal law framework through the Bharatiya Nyaya Sanhita (BNS),¹ the Bharatiya Nagarik Suraksha Sanhita (BNSS)² and the Bharatiya Sakshya Adhinyam (BSA).³ It has passed the Digital Personal Data Protection Act, 2023 (DPDP Act)⁴ and continues to lean on the Information Technology Act, 2000 (IT Act).⁵ These are serious pieces of work. Yet almost none of them were drafted with a thinking machine in mind. They assume that behind every harmful act there is a human hand, that evidence is created by people, and that data is processed for purposes a person can name.

This Article asks two plain questions. First, how did we get here, and what could AGI and ASI realistically mean for human beings? Second, what exactly should India change in its five key statutes so that the law keeps pace with a technology that is moving faster than any Parliament can sit? The aim is not to spread fear, and not to cheer blindly either. It is to think carefully, in Indian terms, about a future that is arriving whether we are ready or not.

II. HOW WE GOT HERE: FROM NARROW TOOLS TO GENERAL MINDS

A. The long road of narrow AI

The idea of a thinking machine is older than the computer itself. When Alan Turing asked in 1950 whether machines could think,⁶ he was really asking how we would ever know. For the next half-century, the answer was mostly “not yet.” Early AI was built by hand. Engineers wrote rules: if the patient has fever and a rash, consider measles. These “expert systems” worked in tight corners but collapsed the moment reality stepped outside the rulebook.

The turn came when researchers stopped telling machines what to think and started letting them learn from examples. Machine learning, and later deep learning with layered neural networks, allowed computers to find patterns in data that no human had written down. By the 2010s, these

¹ The Bharatiya Nyaya Sanhita, 2023, No. 45, Acts of Parliament, 2023 (India) [hereinafter BNS].

² The Bharatiya Nagarik Suraksha Sanhita, 2023, No. 46, Acts of Parliament, 2023 (India) [hereinafter BNSS].

³ The Bharatiya Sakshya Adhinyam, 2023, No. 47, Acts of Parliament, 2023 (India) [hereinafter BSA]. All three codes came into force on July 1, 2024.

⁴ The Digital Personal Data Protection Act, 2023, No. 22, Acts of Parliament, 2023 (India) [hereinafter DPDP Act].

⁵ The Information Technology Act, 2000, No. 21, Acts of Parliament, 2000 (India) [hereinafter IT Act].

⁶ A.M. Turing, *Computing Machinery and Intelligence*, 59 MIND 433 (1950), <https://doi.org/10.1093/mind/LIX.236.433>.

systems could recognise faces, translate languages and beat a world champion at Go.⁷ Still, each system was narrow. The program that could win at Go could not read a railway timetable.

B. The generative leap

The real surprise arrived with large language models. Trained on enormous amounts of text, and later on images, audio, video and code, these models learned something that looked a lot like general competence. The same model could summarise a High Court judgment, write a poem in Punjabi, debug Python and explain a tax rule. When public chatbots launched in late 2022, one of them, ChatGPT, reached an estimated one hundred million monthly users within two months.⁸ India quickly became one of the largest user bases in the world.⁹

Since then the pace has only increased. Models now “reason” through problems step by step, use tools, browse the web, write and run their own code, and operate as agents that can take actions on a computer with limited supervision. A system that once answered questions can now book tickets, file forms, negotiate with other software and manage long projects. That shift, from a tool that speaks to a system that acts, is the most important change of the last two years, and it is the one our laws are least prepared for.

C. What AGI would actually mean

AGI has no single agreed definition, which is part of the problem.¹⁰ Broadly, it means an AI that can perform most economically and intellectually valuable tasks as well as a skilled human, and that can learn new tasks without being rebuilt for each one. It would not need to be conscious or have feelings. It would simply need to be competent across the board.

Opinions on timing differ sharply. Some leaders of major AI labs have publicly suggested that AGI-level systems could appear within this decade.¹¹ Many academic researchers are more cautious and argue that current models still fail in brittle, surprising ways and lack real

⁷ See David Silver et al., *Mastering the Game of Go Without Human Knowledge*, 550 NATURE 354, 354 (2017), <https://doi.org/10.1038/nature24270> (noting that AlphaGo, a program built on deep neural networks, “became the first program to defeat a world champion in the game of Go”).

⁸ Krystal Hu, *ChatGPT Sets Record for Fastest-Growing User Base, Analyst Note Says*, GLOBE & MAIL (Feb. 1, 2023) (Reuters), <https://www.theglobeandmail.com/business/article-chatgpt-sets-record-for-fastest-growing-user-base-analyst-note-says/> (reporting a UBS estimate that ChatGPT reached 100 million monthly active users in January 2023, two months after its launch).

⁹ Jagmeet Singh, *India Has 100M Weekly Active ChatGPT Users, Sam Altman Says*, TECHCRUNCH (Feb. 15, 2026), <https://techcrunch.com/2026/02/15/india-has-100m-weekly-active-chatgpt-users-sam-altman-says/> (reporting that India is ChatGPT’s second-largest user base after the United States).

¹⁰ See Meredith Ringel Morris et al., *Levels of AGI for Operationalizing Progress on the Path to AGI* (arXiv:2311.02462, 2024), <https://arxiv.org/abs/2311.02462> (analysing the competing definitions of AGI in use).

¹¹ See, e.g., Dario Amodei, *Machines of Loving Grace* (Oct. 2024), <https://www.darioamodei.com/essay/machines-of-loving-grace> (suggesting that “powerful AI” “could come as early as 2026, though there are also ways it could take much longer”); Sam Altman, *The Intelligence Age* (Sept. 23, 2024), <https://ia.samaltman.com/> (“It is possible that we will have superintelligence in a few thousand days”).

understanding of the physical world.¹² What is fair to say is that the gap has narrowed faster than almost anyone predicted ten years ago,¹³ and serious people no longer treat AGI as science fiction.

D. And then ASI

Artificial Superintelligence is the step beyond. The worry, first set out by thinkers like I.J. Good in the 1960s¹⁴ and developed later by philosophers such as Nick Bostrom,¹⁵ is that a machine at human level could help design an even better machine, which designs a better one still. This “intelligence explosion” could, in theory, produce a mind that stands to us as we stand to a sparrow. Nobody knows if this is possible or how fast it would happen. But the consequences if it did would be so large that governments cannot responsibly ignore the question.

The honest summary is this: narrow AI is here, powerful general-purpose AI is here in early form, AGI is plausible within the working lives of today’s lawyers and legislators, and ASI is uncertain but not impossible. Law written today must survive into that world.

III. THE THREATS: WHAT COULD GO WRONG FOR PEOPLE

It helps to sort the dangers into two groups. The first group is already visible in Indian police stations, courtrooms and offices. The second group is more speculative but far more serious if it comes true. A good legal framework has to deal with both, without letting the loud, dramatic risks distract from the quiet ones hurting people today.

A. Harms we are already living with

Deepfakes and the collapse of “seeing is believing.” Anyone with a phone can now produce a convincing video of a politician saying something they never said, or a morphed intimate image of a college student. Women have borne the worst of this. Many victims hesitate to approach the police because the process is slow and humiliating, and by the time a complaint is registered the content has spread across dozens of platforms. Indian courts have already had to

¹² See, e.g., ASS’N FOR THE ADVANCEMENT OF A.I., *AAAI 2025 Presidential Panel on the Future of AI Research* 66 (Mar. 2025), <https://aaai.org/wp-content/uploads/2025/03/AAAI-2025-PresPanel-Report-FINAL.pdf> (reporting that 76% of respondents to its community survey considered it “unlikely” or “very unlikely” that “scaling up current AI approaches” would yield AGI).

¹³ See, e.g., Katja Grace et al., *Thousands of AI Authors on the Future of AI* (arXiv:2401.02843, 2024), <https://arxiv.org/abs/2401.02843> (reporting that surveyed researchers put a 50% chance on machines outperforming humans in every possible task by 2047, an estimate thirteen years earlier than the one reached in the same team’s survey a year before).

¹⁴ I.J. Good, *Speculations Concerning the First Ultraintelligent Machine*, in 6 *ADVANCES IN COMPUTERS* 31 (Franz L. Alt & Morris Rubino ed., 1965), [https://doi.org/10.1016/S0065-2458\(08\)60418-0](https://doi.org/10.1016/S0065-2458(08)60418-0).

¹⁵ See generally NICK BOSTROM, *SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES* (2014).

grant injunctions against AI-generated misuse of the faces and voices of well-known public figures, which shows how quickly the problem reached the courtroom.¹⁶

Fraud at industrial scale. Voice-cloning scams, where a caller sounds exactly like a son or a boss asking for urgent money, have moved from rare stories to routine news. “Digital arrest” frauds, in which criminals pose as police or customs officers on video calls, are becoming more believable with AI-generated uniforms, backgrounds and documents. An AI agent never gets tired, so a single criminal gang can run thousands of personalised conversations at once.

Silent discrimination. When a bank, insurer, employer or government scheme uses an algorithm to decide who gets a loan, a job interview or a benefit, bias in the training data can quietly shut out women, rural applicants, people with regional-language names or certain castes and communities. The person rejected usually has no idea an algorithm was involved, let alone why it said no.

Privacy without borders. Large models are trained on data scraped from the open internet, which includes a great deal of personal information Indians shared on social media, forums and public records. Once absorbed into a model, that data is extremely hard to remove, and the model may repeat it or combine it in ways that reveal things about a person they never agreed to share.

Work and dignity. India’s economic story over three decades has leaned on IT services, back-office work, customer support and content jobs. These are exactly the tasks generative AI does well. The impact will not be evenly spread. A senior architect may become more productive; a fresh graduate hoping for an entry-level testing or support role may find that the role no longer exists. Without planning, the demographic dividend could turn into a demographic strain.

Misinformation and democracy. India holds some of the largest elections in human history. Synthetic audio of a candidate, AI-written propaganda in twenty languages, and fake “news” channels run entirely by software can influence voters faster than any fact-checker can respond.

B. Harms that grow with AGI and ASI

Loss of meaningful human control. As systems become agents that set sub-goals and take actions, it becomes harder to say who actually decided anything. If an AI agent managing a company’s finances moves money in a way that turns out to be illegal, was it the developer, the

¹⁶ See *Anil Kapoor v. Simply Life India*, CS(COMM) 652/2023 (Del. HC Sept. 20, 2023) (India) (restraining unauthorised use of the plaintiff’s name, likeness and voice, including through AI-generated deepfakes); *Arijit Singh v. Codible Ventures LLP*, 2024 SCC OnLine Bom 2445 (India) (granting an ad-interim injunction against AI cloning of the plaintiff’s voice).

company deploying it, the employee who clicked “approve,” or nobody at all? Our entire legal system depends on finding a responsible person. AGI threatens to dissolve that link.

Concentration of power. The most capable models require enormous computing power, data and capital. Only a handful of companies and states can build them. If AGI arrives, whoever controls it could hold economic and military advantages no nation or company has ever held. For India, the risk of becoming a mere consumer of foreign intelligence, dependent on systems whose rules are written elsewhere, is a real question of sovereignty.

Security and catastrophic misuse. A highly capable AI could help a small group design cyber attacks on power grids, hospitals or payment systems, or lower the barrier to creating dangerous biological or chemical agents. AI labs themselves now test their models for such dangerous capabilities before release, which tells us how seriously they view the risk.¹⁷

Misaligned goals. The deepest worry about ASI is not that it would be evil, but that it might pursue goals slightly different from what we intended, with enormous competence. A system told to maximise a target, without a real grasp of human values, could cause great harm while doing exactly what it was asked. Researchers call this the alignment problem, and it remains unsolved.¹⁸

The erosion of human agency. Even in the best case, where nothing explodes, there is a subtler threat. If machines make better decisions in medicine, law, finance and governance, people may simply stop deciding. A society that hands over its judgment piece by piece may wake up one day to find it no longer knows how to govern itself.

None of these outcomes is certain. But the law does not wait for certainty before acting on fire safety or drug approvals. It acts on reasonable risk. The same logic should apply here.

IV. WHERE INDIAN LAW STANDS TODAY

India has chosen, so far, not to write a single standalone AI law. The India AI Governance Guidelines released by the Ministry of Electronics and Information Technology (MeitY) on November 5, 2025 say this directly.¹⁹ They favour “targeted amendments” to existing laws, set

¹⁷ See, e.g., OPENAI, *Preparedness Framework (Version 2)* 1, 3–4 (Apr. 15, 2025), <https://cdn.openai.com/pdf/18a02b5d-6b67-4cec-ab64-68cdfbdebcdd/preparedness-framework-v2.pdf> (tracking biological and chemical, cybersecurity and AI self-improvement capabilities, and requiring a model that reaches a “High” capability threshold to have safeguards that sufficiently minimise the associated risk of severe harm before it is deployed).

¹⁸ See generally STUART RUSSELL, *HUMAN COMPATIBLE: ARTIFICIAL INTELLIGENCE AND THE PROBLEM OF CONTROL* (2019).

¹⁹ MINISTRY OF ELEC. & INFO. TECH., GOV'T OF INDIA, *India AI Governance Guidelines: Enabling Safe and Trusted AI Innovation* 10 (2025), <https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc2025115685601.pdf> [hereinafter AI Governance Guidelines] (“[A]t this stage, a separate law to regulate AI is

out seven guiding principles (including People First, Accountability and Understandable by Design), recommend an AI Governance Group supported by a Technology and Policy Expert Committee, and call for the recently established AI Safety Institute to be resourced to provide technical expertise. They also call for the IT Act to be amended to clarify the roles of developers and deployers, and they identify as open questions how the DPDP Act's exemption for publicly available data and its principle of purpose limitation apply to AI.

Some movement has already happened. The IT Amendment Rules, 2026, in force from February 20, 2026, defined “synthetically generated information,” required clear labels on AI-made audio and video, and cut takedown times to three hours for government or court orders and two hours for non-consensual intimate imagery and impersonation.²⁰ A further draft in March 2026 proposed continuous on-screen labels and permanent traceable metadata.²¹ The DPDP Rules were notified in November 2025 with an eighteen-month phased rollout, so most obligations take effect only in May 2027.²² And in *Pune Bar Association v. Union of India*, decided on May 22, 2026, the Supreme Court upheld the hash-value and expert-certificate requirement for electronic evidence under section 63(4) of the BSA.²³

These are useful patches. But they are mostly delegated rules and court rulings layered over statutes that still assume a human author behind every act. Table 1 sums up the gap.²⁴

Statute	What it does well today	Where AI slips through
IT Act, 2000	Cyber offences (§§ 66C, 66D, 67), blocking (§ 69A), intermediary safe harbour (§ 79)	No category for AI developers or deployers; safe harbour built for passive hosts, not systems that generate content
DPDP Act, 2023	Consent, notice, children's data, penalties up to ₹250 crore	Publicly available data largely excluded; no right against purely automated decisions; no rule on training data
BNS, 2023	Cheating by personation, defamation, voyeurism, organised crime	Offences hinge on a human's intention; no specific offence for making or spreading harmful deepfakes

not needed given the current assessment of risks.”); *see also id.* at 5 (seven guiding principles), 6 (“targeted amendments”; AI Governance Group, Technology & Policy Expert Committee and AI Safety Institute), 19 (IT Act and DPDP Act), 36 (the “recently established” AI Safety Institute).

²⁰ The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, G.S.R. 120(E) (Feb. 10, 2026) (India) [hereinafter IT Amendment Rules, 2026].

²¹ Rohit Singh, *MeitY Tightens AI Label Rules, Mandates Continuous Disclosure*, MEdIANAMA (Apr. 23, 2026), <https://www.medianama.com/2026/04/223-meity-ai-label-rules-mandates-continuous-disclosure/> (reporting a draft amendment of March 30, 2026).

²² The Digital Personal Data Protection Rules, 2025, G.S.R. 846(E), r. 1(2)–(4) (Nov. 13, 2025) (India).

²³ *Pune Bar Ass’n v. Union of India*, 2026 SCC OnLine SC 1297 (India).

²⁴ *See* IT Act §§ 66C, 66D, 67, 69A, 79; DPDP Act §§ 3(c)(ii), 9, 33, sch.; BNS §§ 77, 111, 319, 356; BSA §§ 57, 61, 63; BNSS §§ 176(3), 530.

Statute	What it does well today	Where AI slips through
BSA, 2023	Electronic records admissible as primary evidence with a § 63 certificate	Certificate proves a file was not altered after capture, not that it was real when created
BNSS, 2023	Mandatory forensics for serious offences; electronic summons and trials	No standards for AI forensics, AI-assisted policing or AI tools used by courts

Table 1: The Five Statutes and the AI Gap

V. AMENDMENTS NEEDED IN THE INFORMATION TECHNOLOGY ACT, 2000

The IT Act is twenty-six years old. It was written when the big worry was hacking a website and the word “intermediary” meant a company that passively carried other people’s messages. It needs to be updated in the parent statute itself, not only through rules, so that the obligations rest on a firm legislative base and survive court challenge.

1. **Define the AI value chain.** Section 2 should add definitions for an “AI system,” a “general-purpose AI model,” a “developer” (who builds or trains the model), a “deployer” (who puts it to use for customers or citizens) and an “autonomous agent” (a system that takes actions, not just produces content).²⁵ Right now, a company that trains a frontier model and a small startup that plugs it into a chatbot are treated the same, or not at all.
2. **Rethink safe harbour for generative systems.** Section 79 protects intermediaries who do not initiate a transmission or select or modify the information it contains.²⁶ A platform whose own model writes the defamatory paragraph or renders the fake video is not a neutral carrier. The Act should say clearly that immunity is not available for content generated by the platform’s own AI system, while keeping safe harbour for genuine user uploads. A middle path is conditional immunity for providers who follow published safety standards, label outputs and act on complaints quickly.
3. **Put labelling and provenance into the statute.** The 2026 rules on synthetically generated information are a good start, but they live in delegated legislation that can be challenged as going beyond the parent Act.²⁷ A new section should require providers of generative tools above a certain scale to embed tamper-resistant

²⁵ See IT Act § 2(1). Cf. Regulation (EU) 2024/1689, of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence, art. 3, 2024 O.J. (L 1689) 1 (defining “AI system,” “provider,” “deployer” and “general-purpose AI model”) [hereinafter EU AI Act].

²⁶ IT Act § 79(2)(b).

²⁷ See *Kunal Kamra v. Union of India*, 2024 SCC OnLine Bom 3025 (India) (striking down the 2023 amendment to the intermediary rules that created a government fact-check unit as *ultra vires* the IT Act and violative of arts. 14, 19(1)(a) and 19(1)(g) of the Constitution).

watermarks or cryptographic provenance data, and make stripping these markers with intent to deceive an offence.

4. **A tiered duty for high-capability and frontier models.** For the most powerful models, especially those approaching general capability, the Act should require pre-deployment safety testing, red-teaming for cyber, biological and chemical misuse, incident reporting to CERT-In or the AI Safety Institute within a fixed time, and a documented ability to shut the system down.²⁸ This mirrors ideas now common in international discussions and is the most direct legal response to AGI-level risk.²⁹
5. **New cyber offences for AI-enabled crime.** Sections 66C and 66D (identity theft and cheating by personation using a computer resource) carry modest punishments that do not reflect the scale of AI-driven fraud.³⁰ Parliament should add an aggravated form where AI is used to impersonate a real person or to run fraud at scale, with higher punishment and the power to attach proceeds quickly.
6. **Guard against overreach.** Blocking powers under section 69A and the growing list of MeitY “advisories” must come with reasoned orders and a right to be heard, the safeguards on which the Supreme Court upheld section 69A in *Shreya Singhal*, and with greater transparency.³¹ AI regulation that quietly becomes a censorship tool would damage the very trust it is meant to protect.

VI. AMENDMENTS NEEDED IN THE DIGITAL PERSONAL DATA PROTECTION ACT, 2023

The DPDP Act is India’s first real privacy law, and its simple, consent-based design suits a country of first-time internet users. But data is the fuel of AI, and a privacy law that does not speak to AI will quickly be bypassed.

1. **Close the “publicly available” loophole for AI training.** Section 3(c)(ii) excludes personal data that a person has made publicly available.³² In practice this lets AI developers argue that anything posted on social media is free to scrape and train on.

²⁸ See IT Act § 70B (Indian Computer Emergency Response Team).

²⁹ Cf. EU AI Act art. 55 (requiring providers of general-purpose AI models with systemic risk to perform model evaluations and adversarial testing, report serious incidents and ensure cybersecurity protection).

³⁰ IT Act §§ 66C, 66D (each punishable with imprisonment of up to three years and a fine of up to one lakh rupees).

³¹ *Shreya Singhal v. Union of India*, (2015) 5 SCC 1 (India) (upholding § 69A of the IT Act in light of its procedural safeguards, including reasoned orders, while striking down § 66A); see also *Kunal Kamra*, 2024 SCC OnLine Bom 3025.

³² DPDP Act § 3(c)(ii); see also AI Governance Guidelines, *supra* note 19, at 19 (listing among open questions “the scope and applicability of exemptions available for the training of AI models on publicly available personal data” and whether the principles of collection and purpose limitation are compatible with how modern AI systems operate).

The Act should say that large-scale collection of public personal data to train AI models is still “processing,” subject at least to notice, an easy opt-out, and a ban on using it to profile or identify individuals.

2. **A right against purely automated decisions.** Section 11 gives a right to a summary of processing, but nothing about decisions made by algorithms.³³ Indians should have a clear right to know when a significant decision about them (a loan, a job, insurance, a welfare benefit, bail) was made mainly by AI, to receive a plain-language explanation, and to ask for human review.
3. **Make erasure meaningful for trained models.** The right to erasure under section 12 works for a database row but not for knowledge absorbed into a model’s weights.³⁴ The law should require fiduciaries to use reasonable technical measures, such as output filtering and retraining at set intervals, so that a person’s erased data is not reproduced by the model.
4. **Treat frontier AI companies as Significant Data Fiduciaries.** Section 10 already lets the government notify Significant Data Fiduciaries, who must conduct impact assessments and audits.³⁵ Developers of large general-purpose models, and deployers using AI in health, finance, policing or elections, should be notified by default, with a mandatory algorithmic impact assessment covering bias and discrimination, not only data security.
5. **Protect children from AI companions and profiling.** Section 9 bars tracking and targeted advertising aimed at children.³⁶ That protection should expressly extend to emotionally persuasive AI chatbots and “companion” apps, which can shape a young mind far more than an advertisement ever could.
6. **Narrow the government exemption.** Section 17 allows broad exemptions for state agencies.³⁷ When the state uses AI for surveillance, facial recognition or predictive policing, the exemption should be subject to the tests of legality, necessity and

³³ DPDP Act § 11. *Cf.* Regulation (EU) 2016/679, of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), art. 22, 2016 O.J. (L 119) 1 (granting a right not to be subject to a decision based solely on automated processing).

³⁴ DPDP Act § 12.

³⁵ *Id.* § 10(2) (requiring Significant Data Fiduciaries to appoint a Data Protection Officer and an independent data auditor and to undertake periodic Data Protection Impact Assessments).

³⁶ *Id.* § 9(3).

³⁷ *Id.* § 17(2)(a).

proportionality that the *Puttaswamy* judgment requires,³⁸ and to independent oversight.

VII. AMENDMENTS NEEDED IN THE BHARATIYA NYAYA SANHITA, 2023

The BNS replaced the Indian Penal Code in July 2024 and did modernise several provisions, for example by recognising organised crime and cyber-crime in section 111.³⁹ Still, its core logic is the classic one: a human being, with a guilty mind, does a guilty act. AI stretches both halves of that sentence.

1. **A specific deepfake offence.** Today a victim of a morphed video must fit their case into cheating by personation (section 319), defamation (section 356), voyeurism (section 77) or forgery of an electronic record (sections 335 and 336).⁴⁰ In the meantime, victims who can afford to litigate have turned to civil courts for injunctions protecting their personality rights, a remedy out of reach for most ordinary people.⁴¹ A standalone offence should punish creating or knowingly spreading synthetic media of a real, identifiable person without consent, where it causes or is likely to cause harm to reputation, dignity, finances or public order. Sexual deepfakes should carry the heaviest punishment and be treated on par with the most serious offences against women's dignity.
2. **AI as an aggravating factor.** Using AI to commit cheating, extortion, stalking or election-related offences should attract a higher sentence, much as using a weapon does for hurt. The reason is scale: one person with an AI agent can harm thousands of victims in a day.
3. **Clear rules on who is liable when an AI acts.** The Sanhita should state that a person who deploys or directs an AI system is treated as having done what the system does, if they intended the result, knew it was likely, or were reckless about it. For companies, a failure-to-prevent offence, similar to the model used for corporate bribery in some countries, would make organisations liable when their AI causes serious harm and they did not have reasonable safeguards in place.⁴² This keeps responsibility with humans and avoids the dangerous idea of treating the machine itself as the criminal.

³⁸ *Justice K.S. Puttaswamy (Retd.) v. Union of India*, (2017) 10 SCC 1 (India).

³⁹ BNS § 111(1); *see also id.* § 358 (repealing the Indian Penal Code, 1860).

⁴⁰ *Id.* §§ 77, 319, 335, 336, 356.

⁴¹ *See, e.g., Anil Kapoor*, CS(COMM) 652/2023; *Arijit Singh*, 2024 SCC OnLine Bom 2445.

⁴² *Cf. Bribery Act 2010*, c. 23, § 7 (UK) (creating a corporate offence of failure to prevent bribery, subject to a defence of adequate procedures).

4. **Offences for dangerous AI development.** For the AGI horizon, Parliament should consider an offence of knowingly developing, releasing or modifying an AI system to assist in making weapons of mass destruction or in attacks on critical infrastructure, and of deliberately disabling safety controls on a frontier system.
5. **Handle misinformation with care.** Section 197(1)(d) already punishes false or misleading information that jeopardises the sovereignty, unity and integrity or security of India, and section 353 covers statements conducing to public mischief.⁴³ Any AI-specific extension must be narrow and tied to real harm, so that satire, criticism and honest mistakes are not criminalised.⁴⁴

VIII. AMENDMENTS NEEDED IN THE BHARATIYA SAKSHYA ADHINIYAM, 2023

Evidence law is where the deepfake problem bites hardest, because a trial is ultimately a search for what really happened. The BSA made electronic records primary evidence⁴⁵ and set out a certificate procedure in section 63.⁴⁶ This continues a line of authority developed under section 65B of the old Indian Evidence Act.⁴⁷ The Supreme Court's decision in *Pune Bar Association* confirmed that a hash value works like an electronic fingerprint.⁴⁸ But a hash only proves that a file has not changed since it was hashed. It says nothing about whether the video was real when it was first created. A perfect deepfake, faithfully hashed, passes the test.

1. **A presumption-shifting rule for contested audio and video.** Where a party raises a credible, specific challenge that an electronic record is synthetic, the burden should move to the party relying on it to prove authenticity through forensic analysis or provenance data. Without this, courts risk two failures at once: convicting on fake evidence, and letting guilty people escape by simply shouting "deepfake," the so-called liar's dividend.⁴⁹
2. **Recognise provenance and watermark data.** The BSA should expressly allow courts to rely on embedded provenance metadata and AI watermarks, and to draw an adverse inference where such markers have been deliberately stripped.

⁴³ BNS §§ 197(1)(d), 353.

⁴⁴ See *Shreya Singhal*, (2015) 5 SCC 1 (striking down § 66A of the IT Act for vagueness and overbreadth).

⁴⁵ BSA § 57 expls. 4–7; see also *id.* § 61.

⁴⁶ *Id.* § 63(4) & sch.

⁴⁷ See *Anvar P.V. v. P.K. Basheer*, (2014) 10 SCC 473 (India); *Arjun Panditrao Khotkar v. Kailash Kushanrao Gorantyal*, (2020) 7 SCC 1 (India).

⁴⁸ *Pune Bar Ass'n*, 2026 SCC OnLine SC 1297.

⁴⁹ Bobby Chesney & Danielle Citron, *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security*, 107 CALIF. L. REV. 1753, 1785–86 (2019), <https://doi.org/10.15779/Z38RV0D15J>.

3. **Rules for AI-generated evidence.** Outputs of AI tools, such as an enhanced CCTV image, an AI transcription of a phone call or a machine-translated statement, are increasingly placed before courts. The Adhiniyam should require disclosure that AI was used, the tool and version, and a human expert willing to be cross-examined on its reliability. Enhanced or reconstructed images should never be presented as if they were original recordings.
4. **Accredited AI forensic experts.** Section 39 on expert opinion should be read with a statutory scheme of certified examiners able to detect synthetic media, building on the existing examiner scheme under section 79A of the IT Act.⁵⁰ Every district needs access to such experts, not only metro courts.
5. **Machine statements are not confessions.** As AI agents act on behalf of people, the law must clarify that an AI's log or output is a record to be proved like any other document, and cannot be treated as an admission or confession of the person who used it unless independently linked to that person's knowledge and intent.

IX. AMENDMENTS NEEDED IN THE BHARATIYA NAGARIK SURAKSHA SANHITA, 2023

The BNSS governs how crimes are reported, investigated and tried. It already embraced technology: e-FIRs under section 173, video-recorded searches under section 105, mandatory forensic visits for offences punishable with seven years or more under section 176(3), and trials in electronic mode under section 530.⁵¹ The next step is to make this procedure fit for AI-era crime and AI-era policing.

1. **Fast-track procedure for synthetic sexual content.** A victim of a sexual deepfake should be able to file an e-FIR from home, and the police should be under a duty to send a takedown request to platforms within hours, not days. This would link the BNSS to the two-hour takedown window now in the IT Rules, so that the criminal process and the platform process move together.⁵²
2. **Preserve evidence before it disappears.** Section 94 should allow police to issue urgent preservation orders to AI providers for prompts, outputs and account logs connected to a crime, with a clear retention period and judicial oversight, so that crucial records are not deleted before a formal request arrives.⁵³

⁵⁰ BSA § 39(2); IT Act § 79A.

⁵¹ BNSS §§ 105, 173, 176(3), 530.

⁵² See IT Amendment Rules, 2026, *supra* note 20.

⁵³ BNSS § 94.

3. **Rules for AI in policing.** Facial recognition, predictive policing and automated suspect profiling are spreading across states without any statutory basis. The Sanhita should require that such tools be authorised by law, tested for accuracy and bias across gender, region and community, used only for serious offences, and never be the sole ground for arrest.⁵⁴ An arrest memo should record whether an AI tool contributed to identifying the accused.
4. **AI in courts, with humans in charge.** The Supreme Court already uses AI tools to transcribe oral arguments and to translate judgments, and is developing a research tool, SUPACE.⁵⁵ That is welcome in a system with a large backlog. But the BNSS should state plainly that bail, conviction and sentencing are human judicial decisions, that any AI assistance is disclosed to the parties, and that an accused may challenge it.
5. **Cross-border cooperation.** Most frontier AI providers are based abroad. Procedures for obtaining evidence from foreign providers need to be faster, through updated mutual legal assistance arrangements and a clear duty on large AI services operating in India to appoint a local officer who responds to lawful requests, similar to the duty already placed on significant social media intermediaries.⁵⁶

X. BEYOND THE FIVE STATUTES: WHAT TIES IT TOGETHER

Amending five laws separately risks creating five half-solutions that do not talk to each other. A few cross-cutting reforms would give the whole effort a spine.

A statutory AI Safety Institute. The AI Safety Institute, which the 2025 Guidelines make the source of technical expertise on AI safety, should be given a legal foundation, a budget and real technical staff.⁵⁷ It should test frontier models, set standards for watermarking and forensic detection, run a national AI incident database and advise courts and police. Guidelines are useful, but an institution without statutory backing can be ignored.

Risk-based, not one-size-fits-all. A farmer's crop-advice chatbot and a model capable of designing malware should not carry the same burden. Obligations should scale with the system's

⁵⁴ See *K.S. Puttaswamy*, (2017) 10 SCC 1 (requiring that any state intrusion on privacy satisfy the tests of legality, legitimate aim and proportionality).

⁵⁵ MINISTRY OF LAW & JUSTICE, *Use of Artificial Intelligence in Supreme Court* (Press Info. Bureau, July 25, 2025), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2148356> (written reply in the Lok Sabha: AI tools used to transcribe oral arguments in Constitution Bench matters, to translate judgments into eighteen Indian languages and to identify filing defects; SUPACE "in experimental stage of development").

⁵⁶ The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021, G.S.R. 139(E), r. 4(1) (Feb. 25, 2021) (India).

⁵⁷ AI Governance Guidelines, *supra* note 19, at 6, 36.

capability and the stakes of its use, so that Indian startups are not crushed by paperwork meant for global giants.

A compensation route for ordinary people. Criminal law punishes; it does not always heal. Victims of AI harm, whether a wrongful loan rejection or a viral deepfake, need a quick civil remedy, perhaps through the Data Protection Board or a dedicated tribunal, with compensation paid by the responsible developer or deployer.

Investment in human capacity. No law works without people who understand it. Judges, police officers, public prosecutors and lawyers need regular training in AI evidence and AI crime. Workers whose jobs are changing need reskilling support funded partly by the productivity gains AI brings.

India's voice in global rules. AGI and ASI are global problems. A superintelligent system built in one country does not stop at the border. India, as a large democracy, a signatory to the Bletchley Declaration⁵⁸ and a major tech talent base, should push for international agreements on frontier AI safety, shared testing, and a clear ban on AI that autonomously controls weapons of mass destruction. It should also invest in its own computing capacity and models, so that it helps write the rules rather than only following them.

XI. CONCLUSION: KEEPING HUMANS AT THE CENTRE

Every powerful technology has forced the law to grow. Railways brought accident law and compensation. The printing press forced societies to think about libel and free speech. The internet gave us cyber law. Artificial intelligence, and especially the prospect of AGI and ASI, is different in one important way: for the first time, the technology itself can make decisions, create evidence and take actions that look like the work of a human mind.

That is why patchwork will not be enough. India needs its IT Act to recognise who builds and who deploys AI, its DPDP Act to protect people whose data trains these systems and whose lives are decided by them, its BNS to punish those who use AI to cheat and humiliate, its BSA to help courts separate the real from the fabricated, and its BNSS to give police and judges the tools, and the limits, they need in an AI-driven world.

The goal is not to stop AI. India has too much to gain from it: better healthcare in villages, faster justice, services in every Indian language, and new kinds of work for a young population. The

⁵⁸ *The Bletchley Declaration by Countries Attending the AI Safety Summit, 1-2 November 2023*, GOV.UK (Nov. 1, 2023), <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023>.

goal is to make sure that as machines grow more capable, human beings remain the ones who set the direction, carry the responsibility and enjoy the benefits.

It is often said that a society is best judged by how it treats its weakest members. In the age of AI, the weakest may be the woman whose face is misused, the job-seeker rejected by a biased algorithm, or the elderly man deceived by a voice that sounds like his son. If our laws protect them, they will likely protect the rest of us too. And if we get the foundations right today, while the machines are still our tools, we give ourselves the best chance of staying their masters tomorrow.
