

Artificial Intelligence in Indian Elections: Synthetic Media, Campaign Technology and the Regulatory Mandate of Election Management Bodies

Tushar Panchal
ORCID ID: 0009-0003-3284-2472

Abstract: Artificial intelligence moved from the margins of Indian electioneering to its operational core within a single electoral cycle. The 2024 General Election was the first Indian election in which synthetic media appeared at a national scale; the 2026 Assembly elections in Assam, Keralam, Tamil Nadu, West Bengal and Puducherry, covering approximately 17.4 crore (174 million) electors, were the first to be conducted under a dedicated regulatory framework for AI-generated content. This paper asks how that framework developed, how it performed under electoral conditions and where it remains incomplete. It does so through a doctrinal analysis of the governing instruments—constitutional, statutory, delegated and advisory—read against a corpus of documented incidents assembled from fact-checking outlets, civil-society monitors and the press and against three comparative models (the European Union, Brazil and the United States). The paper makes three contributions. First, it traces the Election Commission of India's progression from a single advisory in May 2024 to a quantified labelling standard and cycle-wide enforcement in 2026, during which the Commission reported acting on more than 11,000 social media posts. Second, it identifies a divergence at the centre of the framework: a numerically precise labelling standard sits within a non-statutory advisory body, while the binding delegated legislation adopts a weaker qualitative standard—so the framework's bright line lies exactly where enforcement authority does not. Third, it argues that the framework's defining weakness is a measurement deficit that forecloses both accountability and institutional learning and specifies a measurement protocol to remedy it. It concludes with proposals for statutory codification, a standing technical directorate, extension of pre-certification, silence-period reform and a post-election audit mandate—each assessed against the free-expression objections it must survive.

Keywords: artificial intelligence; synthetic media; deepfakes; electoral integrity; Election Commission of India; Model Code of Conduct; election management bodies; campaign technology; electoral regulation.

1. Introduction

Every generation of communication technology has eventually arrived in the Indian election. Print, cinema, radio, television, the mobile phone and the social-media platform each altered the conduct of campaigns and each, in turn, compelled the electoral machinery to adapt. Artificial intelligence differs from its predecessors in one decisive respect. Earlier technologies changed how political messages were distributed. Generative AI changes how they are produced and at a marginal cost approaching zero. A campaign that once required studios, crews and days of production can now generate persuasive audio and video in minutes, in dozens of languages, personalised to segments of the electorate and, if it chooses, in the voice and likeness of persons living or dead.

This capability arrived in Indian elections faster than the apparatus designed to govern it. The 2024 General Election provided the first national demonstration. The 2026 Assembly elections—held in April–May across Assam, Keralam, Tamil Nadu, West Bengal and Puducherry—provided the first test of a purpose-built regulatory response, comprising amendments

to the Information Technology rules, a sequence of Election Commission of India (ECI) advisories and enforcement machinery deployed at scale for the first time.

The scholarly literature on AI and elections has so far concentrated on two questions: whether synthetic media measurably affects electoral outcomes and how democracies ought to regulate it in the abstract. The first question remains empirically unresolved and this paper does not attempt to settle it. The second has generated a comparative literature on codes of conduct and platform obligations, but relatively little that examines a single election management body (EMB) as an administrative institution—how it builds instruments, deploys them under the compressed and adversarial conditions of a live election and learns (or fails to learn) from the result. That is the gap this paper addresses. It treats the regulation of AI in Indian elections not as a problem of technology policy but as a problem of election administration and asks what an EMB with the ECI's constitutional standing can and should do.

Three research questions organise the paper. First, how has the framework governing AI in Indian elections developed

across the 2024–2026 period and what is the legal character of each instrument in it? Second, how did the framework perform when applied across the 2026 cycle, as documented in the available record? Third, where is the framework structurally incomplete and what institutional remedies lie within the Commission's existing powers?

The paper proceeds as follows. Section 2 sets out the analytical framework and situates the study in the literature, including the sceptical literature on deepfake effects. Section 3 states the method and sources and is candid about their limits. Section 4 examines the comparative context. Section 5 analyses the Indian legal framework as a composite of instruments built for other purposes. Section 6 reconstructs the Commission's regulatory trajectory. Section 7 presents the documented record of 2024 and 2026. Section 8 offers the analysis: five structural gaps, of which the measurement deficit is the culminating one. Section 9 sets out recommendations, each tested against its free-expression objections. Section 10 concludes.

2. Analytical framework and previous literature

2.1 The election management body as user and regulator of AI

The organising frame for this paper is drawn from the comparative EMB literature. International IDEA, in its guidance issued ahead of the 2024 global election cycle, framed the challenge as dual: EMBs must simultaneously evaluate AI as a tool for their own administrative functions and prepare to protect the electoral process against uses of AI by political actors, including synthetic media and micro-targeted information campaigns (International IDEA, 2024). This dual mandate—the EMB as both user and regulator of AI—provides the analytical lens used throughout. It matters because the two roles pull in different directions: the administrative adoption of AI for surveillance or monitoring accrues capacity to the Commission, while the regulatory task requires the Commission to constrain a technology whose detection it has not yet mastered. Both roles were visible within the 2026 Indian cycle, as Section 7 shows.

2.2 Advisory versus statutory regulation

The paper's second analytical commitment is a distinction between advisory and statutory regulation and a claim that this distinction does more work in the Indian case than is generally recognised. An advisory issued under a general power binds imperfectly, specifies consequences weakly and depends for enforcement on instruments designed for other purposes. A statutory rule, or a formal resolution of the kind Brazil's electoral authority employs, carries defined sanctions. The comparative literature on 2024 codes of conduct (see, e.g., the survey of election codes across jurisdictions in the digital-com-

munication literature) repeatedly finds that the binding-versus-normative character of an instrument shapes its effect. The Indian framework, this paper argues, has developed a precise standard in its least-binding instrument and a vague standard in its most-binding one—an inversion that the codification debate has yet to notice.

2.3 The sceptical literature on deepfake effects

A responsible framing must engage the counter-position. A substantial body of reporting and analysis concluded that the anticipated “deepfake apocalypse” of 2024 did not materialise: monitoring of AI use across the 2024 global cycle found that synthetic content functioned more often as trolling, satire and cheap engagement-bait than as decisive information warfare and that India's general election in particular “wasn't the deepfake doomsday many feared” (Rest of World, 2024b). Two implications follow for this paper. First, claims about the effect of synthetic media on outcomes should be treated as unproven in both directions; the paper accordingly makes no such claim. Second—and this is the paper's response to the sceptics—the case for regulation does not depend on proof of decisive effect. It depends on the administrative proposition that an EMB cannot demonstrate the effectiveness of its own interventions without measurement and that the absence of measurement is itself the problem. The sceptical literature, in other words, strengthens rather than weakens the measurement argument developed in Section 8.

3. Method and sources

This is a qualitative study combining doctrinal legal analysis, a documented-incident corpus and comparative institutional analysis. Each element and its limitations, is stated here.

Doctrinal analysis. The regulatory instruments are read as primary sources: the Constitution of India (Article 324); the Representation of the People Act, 1951 (Sections 123, 126 and 126A); the Information Technology Act, 2000 and the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021, as amended in February 2026; the Model Code of Conduct; and the three ECI advisories of May 2024, January 2025 and October 2025. Comparative instruments—the European Union's Artificial Intelligence Act and Digital Services Act, Brazil's Superior Electoral Court Resolution 23.732/2024 and the 2024 California statutes—are treated the same way. Where the text of a primary instrument was available, it is cited directly; where only a secondary report of an instrument was available, that limitation is flagged in the reference.

The incident corpus. Because no official dataset of synthetic-media incidents in Indian elections exists (a fact that is itself a finding, discussed in Section 8.5), the record in Section 7 is assembled from named public sources: fact-checking units (notably The Quint's WebQoof), civil-society monitors (nota-

bly the Netherlands-based Foundation Diaspora in Action for Human Rights and Democracy) and mainstream and trade press. An incident was included where (a) it was reported by at least one identifiable source with a stated basis for the AI attribution and (b) it related to campaign communication in the 2024 or 2026 Indian elections. Incidents are reported with their provenance attached and the evidentiary weight of each source is characterised rather than assumed. This is a purposive, not exhaustive, corpus; it documents the repertoire and scale of techniques, not their prevalence in any statistical sense.

Comparative case selection. The three comparators were selected on a most-different basis to span the design space available to Indian regulators: a horizontal legislative model (EU), an EMB-led resolution model (Brazil, the closest institutional analogue to the ECI) and a fragmented sub-national statutory model (United States). The purpose is to locate the Indian choice within a field of alternatives, not to rank jurisdictions.

Limitations. Three are material. First, the incident corpus depends on third-party detection using self-devised rubrics of varying transparency; it cannot support prevalence or effect estimates and is not offered as doing so. Second, the enforcement figures relied upon—most importantly the Commission’s statement that it acted on more than 11,000 posts—are the Commission’s own, unaccompanied by published methodology and cannot be independently verified. Third, the study is contemporaneous with the events it describes; some instruments were newly in force at the time of writing and remain subject to legal challenge. These limitations are not incidental to the argument. They are, in large part, the argument: the same opacity that constrains this study constrains the Commission’s own capacity to demonstrate that its framework works.

4. The comparative regulatory context

Three regulatory models had emerged internationally by the time India’s 2026 cycle began, each illustrating a different design choice.

4.1 The European Union: horizontal legislation

The European Union’s Artificial Intelligence Act, adopted in 2024, is the most comprehensive horizontal instrument. Its transparency provisions require that AI-generated or manipulated content constituting a deepfake be disclosed as such and it operates alongside the Digital Services Act, which obliges very large online platforms to identify and mitigate systemic risks to electoral processes arising from artificially generated or manipulated content. The European model regulates technology and platforms and leaves electoral application largely to member-state authorities. Its strength is comprehensiveness; its weakness, for electoral purposes, is that horizontal legislation moves slowly relative to electoral calendars and its staged application through 2026–2027 illustrates the lag.

4.2 Brazil: the electoral authority as regulator

Brazil offers the closest institutional analogue to India, because regulation there was driven by the election management body itself. The Superior Electoral Court (Tribunal Superior Eleitoral, TSE), exercising its regulatory powers under the Electoral Code, approved Resolution 23.732/2024 on 27 February 2024, ahead of that year’s municipal elections. The resolution prohibited synthetic audio or video that creates, replaces or alters the image or voice of a living, deceased or fictitious person to harm or favour a candidacy—even where the person depicted has consented. It required clear labelling of AI-assisted campaign material, barred chatbots and avatars from simulating conversation with a real candidate and established platform takedown obligations. The TSE supplemented the resolution with memoranda of understanding signed with Meta, Google, TikTok, X, Telegram, Kwai and LinkedIn and routed complaints through an inter-institutional coordination centre (TSE, 2024; DFRLab, 2024a).

The Brazilian model demonstrates that an electoral authority with strong plenary powers can construct a purpose-built AI regime without waiting for parliamentary legislation. It also demonstrates the model’s limits: independent researchers documented persistent difficulty in detecting violations at scale and self-declaration proved an unreliable foundation for labelling compliance (NetLab UFRJ, 2024). For the 2026 general elections the framework was extended to include a prohibition on publishing or boosting AI-altered content from 72 hours before polling until 24 hours after its close, with a candidate’s use of deepfakes punishable by annulment of registration or mandate.

4.3 The United States: sub-national statutes

In the absence of federal legislation, regulation in the United States has proceeded at state level. California enacted statutes in September 2024 requiring platforms to remove or label materially deceptive AI-generated election content, prohibiting advertisements using a candidate’s cloned voice or likeness to create false narratives and extending disclosure obligations to AI-generated political advertising. The American experience is chiefly a cautionary account of fragmentation: obligations that vary by jurisdiction impose compliance burdens without delivering uniform voter protection and the federal-level regulator has remained deadlocked on the question.

5. The Indian legal and regulatory framework

India entered the AI era with an electoral statute book written for earlier technologies. The framework nonetheless contains provisions of continuing application and the regulatory response of 2024–2026 has been constructed largely by interpretation and delegated rule-making rather than by amendment of electoral law.

The constitutional foundation is Article 324, which vests the superintendence, direction and control of elections in the ECI and has been consistently interpreted as a reservoir of plenary power where enacted law is silent. The Representation of the People Act, 1951, supplies three relevant instruments. Section 123 defines corrupt practices, including the publication of false statements of fact regarding a candidate's personal character or conduct—a provision that maps onto defamatory deepfakes. Section 126 prohibits the display of election matter by means of cinematograph, television or other similar apparatus during the 48 hours ending with the close of polling, the statutory silence period. Section 126A restricts exit polls. The Information Technology Act, 2000 and the 2021 Intermediary Rules govern platform obligations and provisions of the *Bharatiya Nyaya Sanhita* relating to forgery, impersonation and statements conducing to public mischief supply criminal-law backstops. The Model Code of Conduct (MCC), though not statutory, binds parties and candidates from the announcement of elections and prohibits unverified allegations and appeals to communal sentiment—both of which synthetic media can manufacture at scale.

The significant recent addition sits in delegated legislation. On 10 February 2026 the Ministry of Electronics and Information Technology notified amendments to the Intermediary Rules (Gazette notification G.S.R. 120(E)) directed specifically at synthetically generated information (SGI), in force from 20 February 2026. For the first time in Indian law, the amendments define SGI as audio, visual or audio-visual information created or altered by a computer resource in a manner that appears authentic or indistinguishable from a natural person or real-world event, excluding routine edits and accessibility features. They require intermediaries to label SGI, oblige significant social-media intermediaries to obtain a user declaration of synthetic origin and to deploy technical measures to verify it and require removal of unlawful content within three hours of a court order or authorised government notice—a substantial reduction from the earlier 36-hour window. The amendments took effect just over three weeks before the schedule for the 2026 Assembly elections was announced on 15 March 2026 and thereby became part of the operative framework for the cycle (MeitY, 2026).

A divergence at the centre of the framework. One feature of the delegated legislation deserves particular emphasis, because it goes to this paper's central argument. The draft rules circulated in October 2025 had proposed a quantified labelling standard—a visible label occupying not less than ten per cent of the display surface. The final Amendment Rules abandoned that bright line in favour of a qualitative standard: visual SGI must carry a label that is “prominent, easily noticeable and adequately perceivable,” with a prefixed audio disclosure for audio content (MeitY, 2026; Khaitan & Co., 2026).

The consequence is an inversion. India today has a numerically precise labelling standard—the ten-per-cent rule—but it lives in a non-statutory ECI advisory (see Section 6), while the binding delegated legislation carries only a qualitative standard whose application is contestable. The bright line and the enforcement authority have come apart. This is not a drafting curiosity; it is the codification problem in miniature and Section 8 returns to it.

Two characteristics of the framework as a whole follow. First, it is a composite—constitutional power, electoral statute, information-technology regulation and non-statutory code, assembled for a purpose none of its components was designed to serve. Second, the electoral-specific layer, examined next, rests almost entirely on the Commission's advisory jurisdiction rather than on law.

6. The Commission's regulatory trajectory, 2024–2026

The Commission's engagement with synthetic media has developed through three advisories and one cycle of enforcement—a progression from exhortation towards regulation.

The first advisory, dated 6 May 2024, was issued during the General Election, following a Delhi High Court direction that the Commission consider a representation on the electoral risks posed by deepfakes. Invoking the existing framework—the Representation of the People Act, the IT Act and the MCC—it directed parties not to publish or circulate deepfake audio or video through their social-media handles, required removal of such content within three hours of its being brought to notice and asked parties to identify and warn those responsible (ECI, 2024). Its character was corrective and reactive: it created no new obligation, prescribed no labelling standard and specified no mechanism for notification or verification.

The second advisory, dated 16 January 2025, moved towards disclosure. It directed parties to label AI-generated or synthetically altered images, audio and video used in campaigning and to comply with the IT Rules, thereby placing responsibility on parties, candidates and their campaign machinery (ECI, 2025a).

The third advisory, dated 24 October 2025 and issued ahead of the Bihar Assembly election, converted the disclosure principle into a measurable standard. AI-generated or AI-altered campaign material must carry a clear, prominent and legible label—“AI-Generated,” “Digitally Enhanced” or “Synthetic Content”—positioned near the top of the screen and covering not less than ten per cent of the visible display area, or the first ten per cent of the duration for audio. The advisory expressly invoked the Commission's plenary power under Article 324, required the responsible entity to be disclosed in metadata or caption, mandated takedown within three hours on official party handles and—significantly—stated that it

would take effect immediately and remain applicable during all general and bye-elections until further orders (ECI, 2025b). The quantified standard is the most consequential regulatory act in the sequence, because it created for the first time a bright-line test against which compliance could be audited.

The 2026 Assembly cycle supplied the first sustained enforcement test. The Commission convened senior representatives of the major platforms at Nirvachan Sadan on 11 March 2026, ahead of the announcement of polls, to establish a working framework for the cycle. During the elections it stated that it had acted on more than 11,000 social-media posts, a substantial proportion involving misleading or AI-generated content and it reiterated the mandatory labelling of synthetic campaign material (Storyboard18, 2026). The Commission's complaint infrastructure—the cVIGIL application and the Media Certification and Monitoring Committees (MCMCs) at district and state level—served as the machinery for receiving and processing complaints.

The trajectory is therefore legible: from a reactive advisory grounded in existing law (2024), to a disclosure obligation (2025), to a quantified labelling standard (late 2025), to cycle-wide enforcement supported by amended intermediary rules (2026). What it has not produced is codification. Each instrument remains an advisory, resting on the Commission's plenary authority and the MCC's moral force rather than on statute or statutory rule—and the “until further orders” language of the October 2025 advisory, while it answers the objection that advisories simply lapse with the cycle, does not convert an order under Article 324 into a rule with defined sanctions. The advisory specifies the size of the label; it does not specify the consequence of its absence.

7. The documented record

7.1 The General Election of 2024

The 2024 General Election established the repertoire of AI-enabled campaign techniques in India. Four categories are documented in the corpus.

Synthetic resurrection of deceased leaders. In Tamil Nadu, the Dravida Munnetra Kazhagam (DMK) circulated an AI-generated video of its former leader M. Karunanidhi, who died in 2018, delivering an address in his characteristic dress and voice praising the party's record under his son, Chief Minister M. K. Stalin (Harvard Political Review, 2024). The All India Anna Dravida Munnetra Kazhagam had earlier circulated an AI-generated audio message attributed to its deceased leader J. Jayalalithaa. Both were partisan communications by the sponsoring parties and neither carried disclosure adequate to inform an ordinary viewer of its synthetic character.

Malicious deepfakes targeting living figures. Manipulated videos of the Union Home Minister, Amit Shah, altered to suggest an announcement curtailing reservation entitlements,

circulated widely and led to criminal cases; deepfake videos of the actors Ranveer Singh and Aamir Khan, fabricated to depict political endorsement, spread across platforms in April 2024 despite public denials and police complaints by both actors (Boom, 2024). These cases are reported here on the basis of contemporaneous fact-checking and press coverage; where a specific detail (for example, the venue of a first-information report) rested on a single unverifiable source in the original manuscript, it has been generalised.

The commercial supply chain. Reporting during the campaign documented a functioning commercial market for political synthetic media: consultancies openly producing consent-based content for candidates across parties and states, including cloned-voice messages personalised to individual voters at scale (Rest of World, 2024a; Harvard Political Review, 2024). Indian political parties were estimated to have spent over US\$50 million on AI-generated campaign material in 2024. Reported pricing varied by service—of the order of tens of thousands of rupees for voice cloning and around one to two lakh rupees for avatar creation and messaging integration—rather than the single figure cited in some accounts. The structurally important fact is that synthetic campaign content was not an exotic capability but a routine commercial service.

The institutional response in real time. The Commission's 6 May 2024 advisory, its collaboration with platforms and its “Myth vs Reality” register for rebutting circulating falsehoods together constituted the first live administrative response, is discussed in Section 6.

7.2 The Assembly elections of 2026

The 2026 cycle, conducted under the amended IT Rules and the labelling regime, demonstrated both the maturation of the techniques and the limits of the framework.

Synthetic resurrection recurred and crossed party lines. In Tamil Nadu, an AI-generated video of C. N. Annadurai, the founder of the DMK and Chief Minister of Madras State (1967–1969) and who died in 1969, was circulated online, purportedly soliciting support for the Tamilaga Vettri Kazhagam (TVK)—a party founded more than five decades after his death. Hologram projection allowed the TVK's leader to address multiple constituencies simultaneously and parties across the state deployed AI platforms for content generation, speech drafting and targeted outreach (Storyboard18, 2026).

Candidate-targeted deepfakes were reported across states. In Assam, a candidate of the Asom Jatiya Parishad filed a complaint about a fabricated AI-generated video of her; in Kerala, a cyber case was registered over a deepfake video circulated during the campaign (Storyboard18, 2026).

Systematic civil-society monitoring produced the first structured dataset. A report by the Netherlands-based Foundation Diaspora in Action for Human Rights and Democracy (DAHARD), examining the Assam campaign, classified 432

posts across Facebook and Instagram as likely AI-generated using a multi-signal detection rubric, recording 45.4 million views; it documented 31 deepfakes targeting the principal opposition chief-ministerial candidate, Gaurav Gogoi—fabricating his identity as a foreign agent—distributed in part through official party and ministerial handles, as well as fabricated videos targeting his spouse, a private individual (DAHRD, 2026; The Federal, 2026). The report's central finding was of a coordinated campaign of communal targeting against Bengali-origin Muslims and it documented 119 alleged breaches of the Model Code of Conduct on which, it stated, no enforcement action was taken. These are the findings of a civil-society organisation and carry the evidentiary weight of that provenance; they are cited here as documentation of scale and method and the paper takes no position on the contested political characterisations within the report. Their inclusion is nonetheless necessary, because they bear directly on the enforcement claim examined below.

Documented non-compliance postdated the labelling regime. The fact-checking unit of The Quint identified at least 19 AI-generated election videos circulated between January and April 2026 across Assam, West Bengal, Kerala and Tamil Nadu, none labelled by the posting account or the hosting platform, several shared by official handles of major national and regional parties (The Quint, 2026). This finding is significant precisely because it postdates both the ECI's quantified labelling standard and the amended IT Rules: a mandatory disclosure regime existed and routine non-compliance was documented across the party spectrum.

The Commission as user of AI. The 2026 cycle also recorded the first substantial use of AI by the electoral machinery itself. In West Bengal, the Commission deployed AI-enabled cameras for real-time monitoring of polling areas, supplementing body-worn cameras and physical perimeter controls (Storyboard18, 2026). The dual character anticipated by International IDEA—the EMB as user and as regulator of AI—was thus visible within a single election.

The enforcement claim against the compliance record. Two figures from this record must be read together. The Commission stated that it acted on more than 11,000 posts. An independent monitor of a single state's campaign recorded a substantial body of alleged breaches on which, it said, no action was taken and two fact-checking bodies documented systematic non-compliance with the labelling standard across the cycle. Neither the Commission's figure nor the monitors' figures can be reconciled against the other, because the Commission has published no methodology, no baseline and no disposition data for its 11,000 actions. The gap between the two accounts is not resolved in this paper; it cannot be, on the available evidence. That it cannot be is the subject of Section 8.5.

8. Analysis: five structural gaps

Five structural gaps emerge from the record.

8.1 The framework is advisory where it needs to be statutory

Every election-specific AI obligation in India today rests on ECI advisories backed by the MCC and Article 324. As Section 6 noted, the October 2025 advisory's "until further orders" clause answers the narrow objection that advisories lapse with the cycle; but an order under Article 324 remains an administrative instrument without the defined sanctions of a statutory rule. Brazil's contrast is instructive: its electoral AI rules are formal resolutions with graduated penalties up to annulment of candidature. India's labelling standard specifies the size of the label but not the consequence of its absence—and, as Section 5 showed, the one numerically precise standard India possesses sits in the advisory rather than in the binding Rules. Codification is required not merely to add sanctions but to reunite the bright line with the enforcement power.

8.2 Detection capacity is external to the Commission

The 2026 enforcement effort depended on complaints, platform cooperation and third-party fact-checking. The Commission lacks the standing technical capability to identify synthetic content at scale, verify provenance or audit compliance and the MCMC architecture—built for the certification of conventional advertisements—was not designed for content that is continuously generated and mutated. That the external capacity exists is not in doubt: civil-society detection efforts, from the Deepfakes Analysis Unit convened by the Misinformation Combat Alliance during the 2024 cycle to the monitoring bodies cited in Section 7, demonstrated real capability. But reliance on external detection has a cost the 2026 record makes plain: it leaves the Commission unable to originate, rather than merely receive, findings of non-compliance—and therefore unable to answer the monitors' figures with its own.

8.3 Enforcement speed is mismatched to distribution speed

The operative windows—three hours for party takedowns under the 2024 advisory, three hours for platform action on authorised notice under the amended IT Rules—are short by regulatory standards and long by virality standards. Independent documentation of the 2026 Assam campaign found the synthetic-content operation front-loaded months before the MCC took effect, with narratives established before the regulatory window opened (DAHRD, 2026). A framework that activates with the electoral schedule cannot reach content seeded before it. This is a structural limit of election-triggered regulation, not a failure of diligence and it bears on the design of any remedy.

8.4 The statutory silence period is technologically obsolete

Section 126 restricts electioneering through cinematograph, television and similar apparatus in the final 48 hours. It does not, on its face, reach social media or automated distribution. The Commission's own committee on Section 126, chaired by Senior Deputy Election Commissioner Umesh Sinha, which reported on 10 January 2019, recommended extending the provision's logic to print, digital and social media and building a notification-and-takedown mechanism with intermediaries; the core statutory recommendation remains unimplemented (ECI, 2019). Generative AI sharpens the anomaly: content can be produced and distributed autonomously, at scale, inside the silence period, through channels the section does not name.

8.5 Measurement is absent—the culminating gap

The most consequential gap is what the record does not contain. There is no official dataset of synthetic-media incidents in Indian elections, no published audit of labelling compliance, no systematic assessment of the reach of AI-generated campaign content and no post-election technical report equivalent to the statistical publications the Commission has produced on turnout and participation since 1952. The best available incident data for 2026 were produced by a foreign civil-society organisation and a domestic fact-checking newsroom, both working with partial visibility and self-devised rubrics.

This mirrors a pattern documented across sectors by the Stanford Institute for Human-Centered AI, whose 2026 AI Index recorded documented AI incidents rising to 362 in 2025 from 233 in 2024 while noting that capability reporting is near-universal and responsible-use reporting sparse (Stanford HAI, 2026). The Index figure is an all-domain count and cannot be read as an electoral measure—a caution the present paper observes; it is cited only for the structural point that adoption has outrun evaluation. In the electoral domain the deficit is more consequential, because regulation without measurement cannot demonstrate its own effectiveness. The Commission's statement that it acted on more than 11,000 posts and the monitors' findings of unremedied non-compliance, are simply not commensurable in the absence of published baselines, detection methodologies and disposition data. The first casualty of the measurement deficit is not transparency but institutional learning.

What “measurement” would require. To specify the remedy rather than merely name the gap, a minimum protocol would define: (i) an incident as a discrete item of campaign communication attributed to synthetic generation on a stated basis; (ii) a denominator—the population of campaign communications sampled, so that incidents can be expressed as a rate rather than a raw count; (iii) a detection-and-classification

rubric published in advance, with an assessment of inter-rater reliability where human coding is involved; (iv) a sampling frame for compliance auditing, drawn from official party and candidate handles against the labelling standard; and (v) a disposition record for enforcement actions, stating what was acted upon, how quickly and to what effect. None of this is beyond an institution that has measured electoral participation at the scale of a billion electors for seven decades; its absence is a choice, not a constraint.

9. Recommendations

The gaps admit institutional answers, several within the Commission's existing powers. Five are proposed and each is tested against the free-expression objection it must survive—because a regime governing political speech that ignores Article 19(1)(a) will neither withstand challenge nor deserve to.

9.1 Codify the synthetic-media framework

The advisory sequence of 2024–2025 should be consolidated into a formal, permanent instrument: either dedicated rules under the Conduct of Elections Rules, 1961, framed in consultation with the Law Ministry, or a standing ECI order under Article 324. The instrument should carry forward the quantified labelling standard—reuniting the bright line with the enforcement power, as Section 8.1 argued—define prohibited categories on the Brazilian pattern (including synthetic depiction of any person, living or deceased, in a manner capable of misleading electors as to authenticity) and specify graduated consequences, from censure through MCC proceedings to reference under Section 123 where fabricated content defames a candidate.

The free-expression test. Codification must be drawn to survive Article 19(1)(a) scrutiny. This means targeting deception rather than AI as such (satire, parody and disclosed synthetic content must remain protected), preferring disclosure obligations to content prohibitions wherever disclosure suffices and building in expedited review of takedown so that the remedy does not become an instrument of pre-emptive censorship. The constitutional critique already levelled at the 2026 IT Rules—vagueness, overbreadth and the risk that mandatory pre-verification amounts to prior restraint—maps directly onto any electoral codification and should be met in the drafting, not after it.

9.2 Establish a standing Electoral Technology Directorate

The Commission requires a permanent technical unit, staffed across cycles, with three functions: detection and forensic verification of synthetic media using licensed tools and empanelled domestic expertise; maintenance of an authenticated repository of certified campaign material against which circulating content can be checked; and technical support to Chief

Electoral Officers and MCMCs during elections. The unit should publish its verification standards. The precedent is the Commission's own history: the Systematic Voters' Education and Electoral Participation programme institutionalised voter education that had previously been episodic and the same institutional logic applies to technology response. Publishing verification standards is also the answer to the free-expression concern that in-house detection could become opaque state adjudication of truth.

9.3 Extend pre-certification to synthetic content

The MCMC architecture already requires pre-certification of political advertisements on electronic media. Its mandate could be extended so that AI-generated or AI-altered campaign material issued by a party or candidate requires pre-certification, with the certification identifier embedded in metadata as the amended IT Rules already contemplate. This would convert the labelling obligation from self-declaration—which the Brazilian experience and the 2026 compliance audits both show to be unreliable—into a verifiable chain of provenance.

The free-expression test and a limit. This is the most constitutionally exposed of the five proposals, because pre-certification of political content is a form of prior restraint and prior restraint on political speech attracts the highest scrutiny. Two design constraints follow and they are conditions of the proposal rather than afterthoughts. First, pre-certification should attach to paid or amplified campaign material issued by parties and candidates—actors already subject to the MCC and to advertisement pre-certification—not to organic speech by citizens. Second, the scheme must carry a strict timeline and a fast appeal, so that certification cannot function as delay. A further hazard must be named: any labelling or certification regime creates a “liar’s dividend,” whereby authentic material can be dismissed as unlabelled fakery. Codification should therefore protect the status of verified-authentic content as much as it flags synthetic content, or it will hand a weapon to the very actors it aims to constrain. If these limits cannot be secured, this recommendation should yield to the disclosure-first approach of 9.1.

9.4 Implement the Section 126 reform

The Sinha Committee's 2019 recommendation to extend silence-period restrictions to digital media should be pursued legislatively, updated for present technology: a prohibition on the publication and paid amplification of new political content, including synthetic content, during the 48-hour window, on the model of Brazil's 72-hour restriction on AI-altered material. The free-expression calibration here is the distinction the Committee itself drew—between restricting fresh campaign solicitation during a defined tranquillity period, which the courts have long accepted and restricting political discussion as such, which they have not.

9.5 Institute a post-election technology audit

After each general and assembly election the Commission should publish a technology report documenting complaints received, content acted upon, platform response times, labelling-compliance rates and detection outcomes, on the model of its statistical reports on electors and turnout and using the measurement protocol specified in Section 8.5. The Commission's statistical publications since 1952 created the evidentiary foundation on which the study of Indian elections rests; an equivalent series for the technology dimension would do the same for the coming decades and convert the present measurement deficit into a public asset. This proposal alone carries no free-expression cost and it is the precondition for evaluating all the others.

These proposals are deliberately institutional rather than technological. Detection tools will change annually; the architecture that deploys them should not.

10. Conclusion

India's encounter with artificial intelligence in elections has moved through three phases in three years: demonstration in 2024, when the techniques entered the national campaign; regulation in 2025, when disclosure obligations were specified; and enforcement in 2026, when a labelling regime and amended intermediary rules were applied across a five-state cycle for the first time. The Commission has, on the evidence, done what a capable administrative institution does with a new challenge: it has built instruments quickly and deployed them at scale.

The record also counsels against declaring the problem managed. Independent audits of the 2026 cycle documented non-compliance with labelling obligations across the party spectrum, incident volumes in the hundreds by civil-society monitors and synthetic campaigns seeded before the regulatory window opened—and, most tellingly, an enforcement claim and a non-compliance finding that cannot be reconciled because the data to reconcile them do not exist. The framework that exists is advisory where it should be codified, dependent where it should be capable and—above all—unmeasured where it should be audited. The inversion identified in this paper, whereby the one bright-line standard India possesses sits outside its binding law, is a compact emblem of the whole.

The Commission's institutional history suggests the remedy is within reach. It built the world's largest electoral roll, pioneered voter education at the scale of a billion electors and has repeatedly used its plenary authority to fill statutory silence. The regulation of synthetic media is the same species of problem: a permanent feature of the electoral environment requiring permanent institutional machinery and—before anything else—a measurement regime capable of telling the institution whether its machinery works. The 2024 and 2026 elections

have provided the diagnosis. The years before the next general election are the window in which to build.

Declarations

Conflict of interest. The author is Founder and Chief Executive of WarRoom Strategies, a political communication and government-affairs consultancy that has advised political parties and senior political leaders across Indian states. Several proposals in Section 9—in particular the pre-certification regime in 9.3—would regulate the industry in which the author works. The author discloses this interest in the interest of transparency and confirms that the analysis and recommendations reflect his professional judgement. It is also important to know that WarRoom Strategies did produce synthetic or AI-assisted campaign material in the 2024 or 2026 electoral cycles, but none of the client relationship material is related to the actors discussed in Section 7.

Funding. This research received no external funding.

Use of AI tools. The author did not use any generative-AI tool to prepare the manuscript, per journal policy and confirms that the author takes full responsibility for the content.

Data availability. This study draws on publicly available regulatory instruments, reports and press sources cited in the reference list. No proprietary dataset was generated. The incident corpus described in Section 3 can be reconstructed from the cited sources.

REFERENCES

- Boom. (2024). Fact-checks of deepfake and manipulated political videos during the 2024 General Election [fact-check archive]. BOOM. <https://www.boomlive.in>
- DAHRD. (2026, 6 April). AI-weaponised disinformation, systematic exclusion and institutional failure in India's 2026 Assam Assembly election. Foundation Diaspora in Action for Human Rights and Democracy. <https://dahrd.org>
- DFRLab. (2024a, 2 October). The challenges of identifying deepfakes ahead of the 2024 Brazil election. Atlantic Council Digital Forensic Research Lab. <https://dfrlab.org>
- DFRLab. (2024b, 29 May). Regulating the use of AI for Brazilian elections: what's at stake. Atlantic Council Digital Forensic Research Lab. <https://dfrlab.org>
- Election Commission of India. (2019, 10 January). Report of the Committee on Section 126 of the Representation of the People Act, 1951 [Umesh Sinha Committee]. <https://eci.gov.in>
- Election Commission of India. (2024, 6 May). Advisory on responsible and ethical use of social-media platforms by political parties. New Delhi: ECI.
- Election Commission of India. (2025a, 16 January). Advisory on labelling of AI-generated content in election campaigning. New Delhi: ECI.
- Election Commission of India. (2025b, 24 October). Advisory on the use of AI-generated content in election campaigning [invoking Article 324; effective until further orders]. New Delhi: ECI.
- European Union. (2024). Regulation (EU) 2024/1689 (Artificial Intelligence Act); Regulation (EU) 2022/2065 (Digital Services Act). Official Journal of the European Union.
- Government of India. (1951). The Representation of the People Act, 1951 (ss. 123, 126, 126A).
- Government of India. (2000). The Information Technology Act, 2000.
- Government of India. (2021). Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021.
- Harvard Political Review. (2024, 6 December). Shaping robust AI regulation: lessons from India's 'deepfake' election. <https://the-harvardpoliticalreview.com>
- International IDEA. (2024). Artificial intelligence for electoral management. International Institute for Democracy and Electoral Assistance. <https://www.idea.int>
- Khaitan & Co. (2026, 16 February). MeitY notifies the IT Amendment Rules 2026 [legal analysis of the qualitative labelling standard]. <https://www.khaitanco.com>
- Medianama. (2025, 28 October). ECI cracks down on deepfakes before Bihar polls. <https://www.medianama.com>
- Ministry of Electronics and Information Technology. (2026, 10 February). Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026 [G.S.R. 120(E); in force 20 February 2026]. Government of India.
- NetLab UFRJ. (2024, 11 December). Lessons from the pre-campaign: the challenge of combating deepfakes in the 2024 elections in Brazil. Federal University of Rio de Janeiro. <https://netlab.eco.ufrj.br>
- Rest of World. (2024a, 23 May). Indian politicians spend big on AI election content. <https://restofworld.org>
- Rest of World. (2024b, 5 June). India's election wasn't the deepfake doomsday many feared. <https://restofworld.org>
- Stanford Institute for Human-Centered Artificial Intelligence. (2026). The 2026 AI Index Report. Stanford University. <https://hai.stanford.edu>
- Storyboard18. (2026, 21 April). Holograms, robots and deepfakes: how AI shaped the 2026 Assembly elections. <https://www.storyboard18.com>
- The Federal. (2026, 8 April). Assam election saw industrial-scale AI disinformation campaign: report. <https://thefederal.com>
- The Quint. (2026, 8 May). WebQoof: AI videos surge amid state elections in India, defy labelling rules. <https://www.thequint.com>
- Tribunal Superior Eleitoral. (2024, 27 February). Resolution No. 23.732/2024 [amending Resolution 23.610/2019 on electoral propaganda]. Brasília: TSE.

About Author



Tushar Panchal

Tushar Panchal is the Founder and Chief Executive of WarRoom Strategies, a political communications and government affairs consultancy. He has advised various political parties and senior political leaders across Indian states for nearly three decades. He frequently writes on political communication, electoral technology and AI governance.

Email: tushar@warroom.live