Journal of International Islamic Law, Human Rights and Public Policy
(JISHUP Journal) | ISSN(e): 3031-2280
Website https://jishup.org/index.php/ojs/index | Volume 4 No. 4 (2026)

Indian Democracy in the Age of Artificial Intelligence: A Study of Its Impact on Political Public Opinion and Electoral Politics

Seema Choudhary
Assistant Professor, University of Technology, School of humanities, Arts and Social Science Department of Political Science Vatika, Jaipur
seemachoudhary21188@gmail.com

Received: 03/07/2026 | **Revised:** 10/07/2026 | **Accepted:** 29/07/2026 | **Published:**06/09/2026

Abstract

This paper examines how artificial intelligence, and generative AI in particular, is reshaping Indian democracy, with a focus on two closely related domains: the formation of political public opinion and the practical conduct of electoral politics. Drawing on documented developments from India's 2024 general election, subsequent state elections, and the evolving regulatory response of the Election Commission of India and the Ministry of Electronics and Information Technology, the paper argues that AI functions in the Indian context as a dual-use technology whose democratic consequences depend on who deploys it, for what purpose, and under what institutional constraints. The analysis is organized around three effects: enhancement, whereby AI has expanded the linguistic reach and inclusiveness of political communication in a highly diverse electorate; distortion, whereby deepfakes, voice cloning, and synthetic media have been used to mislead voters, fabricate endorsements, and inflame sensitive social cleavages; and asymmetry, whereby unequal access to AI infrastructure and resources has altered the competitive balance among political actors. The paper reviews India's regulatory trajectory, including Election Commission advisories issued between 2024 and 2025 and the Ministry of Electronics and Information Technology's October 2025 draft amendments to the Intermediary Guidelines, situating these developments within comparative international approaches from the European Union, the United States, and China. It concludes with policy recommendations addressing detection capacity, closed-network messaging platforms, campaign finance transparency, and voter digital literacy, and argues that India's experience offers an instructive case study for other large, diverse democracies confronting the same technological transformation.

Keywords: *Artificial intelligence, Indian democracy, electoral politics, deepfakes, disinformation, political communication, Election Commission of India, generative AI, public opinion*

1. Introduction

India's general election of 2024 was conducted across seven phases, involved nearly 970 million eligible voters, and was administered by more than one and a half million polling stations. It was also the first Indian election in which artificial intelligence, particularly generative AI, became a visible and consequential feature of political campaigning rather than a peripheral curiosity. Political parties, candidates, consultants, and independent operators used AI-generated audio, video, images, and translation tools at a scale that had no precedent in Indian electoral history. Commentators in India and abroad quickly labelled the contest the country's first 'AI election', a description that captured both the enthusiasm and the anxiety surrounding the technology's arrival in the world's largest democratic exercise. This paper examines how artificial intelligence is reshaping Indian democracy, with particular attention to two interlinked domains: political public opinion formation and electoral politics. The first domain concerns how AI-generated content, personalized messaging, and algorithmically curated information environments influence what citizens believe, whom they trust, and how they interpret political events. The second domain concerns how AI has altered the practical conduct

of election campaigns, the strategies of political parties, the administration of elections by the Election Commission of India (ECI), and the emerging regulatory apparatus designed to govern these new practices.

The central argument advanced here is that AI in the Indian context is not a single, unified force but a dual-use technology whose democratic consequences depend heavily on who deploys it, for what purpose, and under what regulatory constraints. On one hand, tools such as the government's Bhashini translation platform, AI-powered chatbots, and multilingual voice synthesis have measurably expanded the reach of political communication in a country with twenty-two scheduled languages and hundreds of dialects, arguably deepening the inclusiveness of democratic participation for non-Hindi-speaking and rural populations. On the other hand, the same underlying technologies, especially generative video, voice cloning, and synthetic image creation, have been used to fabricate statements by political leaders, resurrect deceased public figures for endorsement, and spread communally or socially sensitive misinformation capable of inflaming existing fault lines of caste, religion, and region.

India offers an unusually rich and important case study for this inquiry for several reasons. It is the largest democracy in the world by population and electorate, making the scale of any AI-enabled distortion or enhancement of the democratic process globally significant. Its linguistic and social diversity creates conditions in which AI translation and personalization tools provide genuine utility, but the same diversity also creates innumerable fault lines that bad-faith actors can exploit through targeted synthetic content. Its media ecosystem is dominated by messaging applications, particularly WhatsApp, which allow content to spread through closed, encrypted networks that are extraordinarily difficult for fact-checkers, platforms, or regulators to monitor. And its regulatory response, evolving through Election Commission advisories and a new Ministry of Electronics and Information Technology (MeitY) rule-making process launched in 2025, offers a live test case of how a large democracy attempts to govern generative AI in the electoral sphere without a dedicated AI statute.

This paper proceeds in the following manner. Section 2 situates the study within existing scholarship on AI, disinformation, and democratic theory, and outlines the conceptual framework used throughout the paper. Section 3 documents how AI has transformed the practice of political communication in India, from multilingual outreach to hyper-personalized voter messaging. Section 4 examines the darker side of this transformation: deepfakes, synthetic media, and disinformation campaigns, illustrated through documented cases from the 2024 general election and subsequent state elections. Section 5 considers the use of AI and data analytics for voter segmentation and microtargeting. Section 6 turns to the administrative dimension, examining how the Election Commission itself uses technology and how it has responded to AI-related risks. Section 7 analyses the regulatory and institutional response in detail, including the October 2025 MeitY draft amendments and comparative international approaches. Sections 8 and 9 offer an analytical assessment of AI's impact on political public opinion and on electoral politics respectively, drawing on the empirical material presented earlier. Section 10 discusses the specific risks and vulnerabilities that AI introduces into Indian democracy, while Section 11 considers the opportunities AI presents for strengthening democratic participation. Section 12 offers policy recommendations, and Section 13 concludes.

2. Literature Review and Conceptual Framework

The scholarly literature on artificial intelligence and democracy has grown rapidly since roughly 2016, when concerns about computational propaganda, algorithmic amplification, and foreign influence operations entered mainstream political science and communication studies. Early work in this literature, much of it centred on the United States and Western Europe, focused on the role of social media algorithms in creating filter bubbles and echo chambers, and on the use of bots and coordinated inauthentic behaviour to manipulate online discourse. The emergence of generative AI models capable of producing convincing synthetic text, images, audio, and video from roughly 2022 onward shifted scholarly attention toward a qualitatively different concern: the capacity of AI not merely to amplify existing content but to fabricate wholly new, individually convincing artefacts depicting real people saying or doing things they never said or did.

Within democratic theory, several concepts are useful for analysing these developments. The first is the idea of an 'epistemic commons', the shared factual and informational basis on which democratic deliberation depends. Scholars have argued that the proliferation of synthetic media threatens this commons not only through the direct spread of false content but through what has been termed the 'liar's dividend': the phenomenon by which the mere existence of convincing deepfake technology allows genuine, damaging evidence to be dismissed as fabricated, thereby degrading public trust in authentic information as much as spreading false information does. This dual erosion, of belief in false things and disbelief in true things, is arguably a more corrosive threat to democratic accountability than either effect in isolation.

A second relevant concept is agenda-setting and framing theory, which examines how media, including AI-mediated media, shapes not what people think but what they think about, and how issues are framed for public consumption. AI-generated content, precisely because it can be produced cheaply, rapidly, and in emotionally charged forms tailored to specific audiences, has the potential to intensify agenda-setting effects, allowing political actors to flood the information environment with narratives favourable to themselves or damaging to opponents at a volume and velocity that traditional media production could not match.

A third concept, drawn from political communication research on India specifically, is the notion of 'WhatsApp politics' or closed-network campaigning. Scholars studying the 2014 and 2019 Indian general elections had already documented the centrality of WhatsApp groups to political mobilization, rumour circulation, and the construction of partisan realities in India, given the app's massive user base, its use of end-to-end encryption, and its prevalence in both urban and rural contexts. This paper builds on that literature by examining how generative AI content, including personalized AI avatars, voice-cloned messages, and synthetic videos, has been layered onto this pre-existing closed-network infrastructure, producing effects that are qualitatively different from earlier text-and-forward misinformation because the new content is audiovisual, individually addressed, and considerably more difficult to disprove even after the fact.

Finally, this paper draws on a normative framework distinguishing three broad effects that AI can have on democratic processes: an enhancement effect, whereby AI expands access, inclusion, and the efficiency of democratic participation; a distortion effect, whereby AI degrades the quality of public deliberation through manipulation, misinformation, or manufactured consensus; and an asymmetry effect, whereby AI capabilities are unevenly distributed among political actors, favouring those with greater financial resources, technical sophistication, or access to state infrastructure, thereby altering the competitive balance of electoral contests independent of voter preference. These three effects, enhancement, distortion, and asymmetry, structure much of the empirical analysis that follows.

3. AI and the Transformation of Political Communication in India

The 2024 general election is widely regarded as the moment at which AI-enabled political communication moved from experimental novelty to mainstream campaign practice in India. Political parties reportedly spent tens of millions of dollars on authorized, sanctioned AI-generated content during the campaign, with independent estimates of total campaign expenditure across all parties running into the billions of dollars, making it, by most measures, the most expensive election ever conducted anywhere in the world. Within this vast expenditure, AI tools were used for purposes ranging from multilingual translation and dubbing to the creation of individually addressed video messages, automated phone calls in the voice of candidates, and AI-generated memes and satirical content targeting rivals.

3.1 Language, Translation, and the Bhashini Initiative

Perhaps the most consequential legitimate use of AI in the 2024 election was in the domain of language translation and dubbing. The Ministry of Electronics and Information Technology's Bhashini platform, originally developed as a National Language Translation Mission to support automatic speech recognition, text-to-speech synthesis, and machine translation across Indian languages, was employed extensively to translate Prime Minister Narendra Modi's campaign speeches, delivered in Hindi, into Tamil, Telugu, Kannada, Malayalam, Bengali, Odia, Marathi, and other regional languages in near real time. This use of AI addressed a longstanding structural feature of Indian politics: the historical advantage enjoyed by Hindi-speaking political figures and parties in reaching audiences across the Hindi Belt states, and the corresponding difficulty such figures have faced in connecting authentically with voters in South Indian and other non-Hindi-speaking states.

By enabling a single speech to be delivered, translated, and dubbed into a voter's native language essentially instantaneously, Bhashini and similar tools plausibly expanded the substantive inclusiveness of political communication for millions of citizens who would otherwise have relied on subtitles, delayed translations, or intermediaries to access a leader's message. This represents a clear instance of the enhancement effect described in the conceptual framework above: AI functioning to widen the practical reach of democratic communication across India's immense linguistic diversity, a diversity that has historically been treated by scholars as an obstacle to national political integration.

At the same time, the deployment of a government-built AI translation infrastructure by the ruling party for campaign purposes raises a subtler concern about asymmetry. Opposition parties without comparable access to state-funded linguistic infrastructure, or without the technical capacity to build equivalent tools in-house, faced a structural disadvantage in matching this scale of multilingual outreach, illustrating how AI capability, even when used for a broadly

legitimate purpose such as translation, can still tilt the competitive playing field between political actors depending on differential access to underlying infrastructure and resources.

3.2 Hyper-Personalized Outreach and AI Avatars

Beyond translation, the 2024 campaign saw the widespread use of AI-generated avatars of political leaders and candidates for individually addressed voter outreach. Reporting from major international outlets described AI-generated video clips in which an avatar of Prime Minister Modi addressed voters by name, discussing government welfare schemes the specific voter had received, in the voter's own language and dialect. Similar tools were used further down the party hierarchy, with BJP workers reportedly generating large numbers of personalized video messages using their own AI avatars, automatically rendered in whichever of India's languages a given voter spoke. Estimates cited in press coverage suggested that tens of millions of AI-generated voice calls mimicking politicians were placed to voters before and during the campaign period.

This form of communication represents a novel and potentially transformative addition to India's already well-established tradition of hyper-local political mobilization through booth-level workers, but it does so at a scale and level of individual customization that would have been prohibitively expensive using human labour alone. Political consultants interviewed by journalists during this period noted that generative AI had enabled an approximately twentyfold increase in the volume of digital campaign content produced without a corresponding increase in campaign staff, suggesting a genuine step change in the economics of political communication.

The personalization of political messaging through AI avatars raises distinct concerns that differ from those associated with mass-broadcast propaganda. Because such messages are individually addressed and often delivered through private channels such as WhatsApp, they are less visible to journalists, fact-checkers, and civil society monitors than public social media posts or television advertisements, creating what might be termed a transparency gap in campaign communication. A message that a single voter receives, addressed to them by name, discussing their specific circumstances, is, by its nature, largely invisible to outside scrutiny, complicating efforts to assess the accuracy, tone, or manipulative potential of such content at scale.

4. Deepfakes, Synthetic Media, and the Manipulation of Public Opinion

Alongside these largely constructive uses of AI, the 2024 general election and subsequent state elections in India were marked by a substantial volume of deceptive synthetic content designed to mislead voters, damage political opponents, or fabricate the appearance of support and endorsement. This section documents several prominent, well-reported cases and analyses their implications for public opinion formation.

4.1 Documented Cases from the 2024 Campaign and Beyond

One widely reported case involved a fabricated video appearing to show Prime Minister Modi performing a garba dance, a traditional Gujarati dance associated with the Navratri festival, in a manner intended to embarrass or mock him; Modi himself referenced this video publicly, including in remarks to G20 leaders, as an example of the dangers AI-generated content poses to public trust and social stability. A second, more politically consequential case involved a doctored video appearing to show Union Home Minister Amit Shah stating that the BJP government intended to scrap reservations, or affirmative-action quotas, for Scheduled Castes, Scheduled Tribes, and Other Backward Classes, a position of extraordinary sensitivity in Indian politics given the centrality of reservation policy to social justice debates. The circulation of this video prompted formal complaints, an Election Commission directive requiring parties to remove flagged fake content within three hours of notification, and the arrest of individuals associated with opposition parties in connection with its distribution.

A third category of case involved the digital resurrection of deceased public figures for political endorsement purposes, including AI-generated content in which a politician who had died years earlier appeared to endorse a living candidate's campaign, raising novel questions about consent, dignity, and the boundaries of acceptable campaign practice that existing electoral law had not anticipated. A fourth category involved the fabrication of celebrity endorsements, including AI-manipulated videos in which well-known Bollywood actors appeared to criticize the Modi government and endorse the opposition Indian National Congress, content that leveraged the trust and cultural authority of prominent entertainment figures to lend false credibility to partisan messaging.

Beyond individual high-profile cases, both major national parties, the BJP and the Congress, accused each other throughout the campaign of producing and disseminating deepfake content, with satirical AI-generated material becoming a routine feature of the campaign's meme warfare. In one notable instance, the Congress released AI-generated satirical content portraying Modi in an unflattering light through a modified popular song, prompting a retaliatory

response from BJP-aligned accounts that included AI-cloned audio of a deceased playback singer. These exchanges illustrate how generative AI has lowered the cost and effort required to produce sophisticated, emotionally resonant attack content, enabling a volume and speed of partisan content production that traditional production methods could not sustain.

Independent research and journalistic investigation during this period identified a cottage industry of AI content producers, sometimes referred to informally as commercial 'deepfake' studios, willing to produce synthetic political content on request. Reporting described individual operators receiving requests from campaigns not only for legitimate personalized outreach content but for material explicitly intended to embarrass rivals, including fabricated gaffes and manipulated compromising imagery, indicating that the market for deceptive synthetic political content operated alongside, and was not fully separable from, the market for legitimate AI-assisted campaign communication.

4.2 Voice Cloning, Robocalls, and Audio Deepfakes

Audio deepfakes present distinct challenges relative to video, because voice cloning technology requires comparatively little source material to produce convincing results, audio content is often shared without accompanying visual cues that might prompt scepticism, and voice notes are a dominant mode of communication on Indian messaging platforms, particularly WhatsApp, where voice messages are commonly forwarded within family and community networks with high levels of implicit trust. Cybersecurity officials within Delhi Police's cyber crimes unit publicly flagged voice cloning as one of the two AI tools, alongside deepfake video, most likely to be exploited at scale during the 2024 campaign, and noted that no reliable automated technology existed at that time to distinguish authentic from AI-generated audio content in real time.

The scale of AI-generated voice content during the campaign appears to have been substantial: press estimates suggest that tens of millions of automated calls employing synthetic renditions of politicians' voices were placed to voters, a figure that, if accurate, would represent one of the largest deployments of AI voice technology in any democratic election globally to that point. This form of outreach blurs the line between legitimate automated campaign communication, of the kind long used in the form of pre-recorded political robocalls in many democracies, and deceptive impersonation, since the same underlying technology can produce an authorized message from a candidate or an unauthorized, fabricated statement falsely attributed to that candidate, with little to distinguish the two from the recipient's perspective absent explicit disclosure.

5. AI, Data Analytics, and Voter Microtargeting

Beyond content generation, AI and advanced data analytics have reshaped how Indian political parties identify, segment, and target voters. Building on foundations laid in earlier elections, when parties experimented with data-driven micro-level targeting following global controversies such as the Cambridge Analytica affair, the 2024 campaign saw parties deploy AI-based analytics to segment the electorate along dimensions including geography, age cohort, caste, religion, economic status, and observed online behaviour, subsequently delivering differentially tailored messaging to each segment.

This capacity for granular targeting is enabled in part by India's extensive digital public infrastructure, originally built to facilitate welfare delivery and financial inclusion, which holds vast databases of citizen contact information, including phone numbers linked to national identity and welfare-scheme enrolment systems. Reporting has documented instances in which millions of citizens received unsolicited WhatsApp messages from government-linked accounts seeking feedback on welfare schemes closely associated with the ruling party, raising serious concerns about the boundary between legitimate government communication and the use of state-built digital infrastructure for what critics characterize as ruling-party political outreach, a boundary problem that is likely to recur and intensify as AI lowers the cost of large-scale personalized messaging still further.

Political parties have also drawn on commercially available voter databases sold by private data brokers, combining these with AI-driven analytics to construct increasingly precise behavioural and demographic profiles of individual voters and households. When combined with generative AI content production capabilities, such targeting allows a party to deliver not merely a generic message calibrated to a demographic segment, as was common in earlier decades of targeted political advertising, but an individually rendered video or audio message addressed by name, in the recipient's own language, referencing their specific circumstances, substantially increasing the psychological salience and persuasive potential of political communication relative to earlier, less personalized forms of campaigning.

6. Election Administration, the Election Commission of India, and Technology

It is important to distinguish the use of AI in political campaigning, the primary focus of this paper, from the separate question of AI's role, or absence of a role, in the core administrative machinery of Indian elections, particularly Electronic Voting Machines (EVMs). India's EVMs are simple, standalone electronic devices without wireless connectivity, and the Election Commission of India has consistently maintained that they do not employ artificial intelligence or networked software of the kind that could be remotely manipulated. Persistent public debates and litigation concerning EVM integrity in India have generally centred on questions of hardware tampering, chain-of-custody procedures, and the demand for voter-verified paper audit trail (VVPAT) verification, rather than on AI-specific vulnerabilities, and this paper does not treat EVM controversies as part of the AI question proper, though public confusion conflating AI-related anxieties with pre-existing EVM scepticism has itself become a notable feature of India's contemporary political discourse.

Where AI does intersect meaningfully with election administration is in the Election Commission's own use of technology for voter education, misinformation monitoring, and content moderation coordination. The ECI has increasingly relied on social media platforms and digital channels to conduct voter awareness campaigns, and has, since 2024, issued a series of advisories and standard operating procedures instructing political parties and their representatives on the permissible use of AI-generated content during campaigns, discussed in detail in the following section. The Commission has also directed that flagged fake or doctored content be taken down within a specified window, typically three hours, once brought to a party's notice, an enforcement mechanism that depends heavily on the cooperation of political parties themselves and on the capacity of civil society fact-checking organizations to identify problematic content in the first instance.

7. The Regulatory and Institutional Response

India's regulatory response to AI in the electoral domain has evolved incrementally, through a combination of Election Commission advisories addressed to political parties and, more recently, formal rule-making by the Ministry of Electronics and Information Technology addressed to online intermediaries. As of early 2026, India does not possess a dedicated statute governing artificial intelligence generally, meaning that AI-related electoral harms have been addressed through amendments to existing information technology regulations and through the Commission's administrative advisory powers, rather than through purpose-built AI legislation.

7.1 Election Commission Advisories, 2024–2025

The Election Commission's engagement with AI-related risks began in earnest in early 2024, when it issued guidance to political parties emphasizing the need to uphold electoral integrity and refrain from circulating content that is patently false, misleading, or defamatory, alongside broader instructions concerning the treatment of women, children, and animals in campaign material. As deepfake-related controversies multiplied during the campaign, including the Amit Shah reservation video controversy, the Commission moved to a more explicit stance, directing parties to remove flagged fake content within three hours and warning against the misuse of AI-based tools to create deepfakes that distort information or propagate misinformation.

Following the conclusion of the 2024 general election, the Commission continued to refine its guidance. A May 2024 set of guidelines addressed social media conduct and AI-generated content more specifically, and a January 2025 advisory reiterated these principles with additional obligations, requiring parties to label AI-generated content appropriately and to include disclaimers on synthetic campaign advertisements. Ahead of the Bihar state assembly elections in late 2025, the Commission issued a further advisory, on October 24, 2025, explicitly addressing the use of AI-generated content in campaigning, a step that followed a public interest litigation filed by a civil society organization seeking judicial direction requiring the Commission to act against the circulation of deepfakes, after which the matter was referred back to the Commission itself for administrative resolution rather than pursued through court-ordered remedies. This advisory specifically addressed concerns about hyper-realistic synthetic content depicting political leaders and required political parties, candidates, and campaign representatives to follow specific directions intended to preserve the integrity of the electoral process.

7.2 The MeitY Draft Amendments of October 2025

The most significant recent development in India's AI governance landscape occurred on October 22, 2025, when the Ministry of Electronics and Information Technology released draft amendments to the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021, specifically to address the harms posed by what the amendments term 'synthetically generated information'. These amendments represent India's first attempt at an

explicit statutory framework for regulating AI-generated content, moving beyond the Commission's advisory approach, which binds only political parties and lacks direct enforcement mechanisms against platforms, toward obligations placed directly on social media intermediaries and significant social media intermediaries.

The draft amendments introduce a formal definition of synthetically generated information, and require covered platforms to implement labelling and traceability mechanisms, including visible markers or embedded metadata identifying AI-generated content, alongside a requirement that platforms take reasonable technical measures to detect such content and control its dissemination when it threatens public order or electoral integrity. The proposed framework establishes clearer lines of accountability for intermediaries that host or facilitate synthetically generated deepfake content, addressing a gap in the pre-existing Intermediary Guidelines framework, under which Rule 3(1)(b) provided only a general due-diligence obligation not specifically tailored to AI-generated harms.

This regulatory move followed an earlier, widely criticized MeitY advisory from March 2024, which had attempted to require platforms to seek government permission before publicly releasing AI models deemed unreliable or under-tested, a requirement that was seen by many technology policy analysts as both impractical and constitutionally fraught, given its resemblance to a prior-restraint mechanism on the deployment of software tools. The October 2025 amendments reflect a partial retreat from that earlier, more interventionist posture toward a narrower focus specifically on the labelling, traceability, and platform accountability dimensions of synthetic content, while leaving the broader question of AI model governance comparatively under-addressed, an approach one prominent technology policy commentator characterized as combining a broadly 'hands-off' posture toward AI governance in general with unusually specific intervention on the narrower question of synthetic media disclosure.

The draft amendments remain, as of this writing, subject to public consultation, and unresolved questions persist regarding the technical reliability of AI-content detection and labelling systems, the practical feasibility of enforcing traceability requirements against sophisticated or foreign-hosted actors, and the cross-border enforcement challenges inherent in regulating content that may originate outside India's jurisdiction. Comparative experience is not encouraging on this point: AI-generated content continued to circulate widely during the 2024 European Parliament elections notwithstanding the disclosure obligations imposed by the European Union's AI Act, suggesting that formal labelling requirements alone are unlikely to fully resolve the underlying problem of synthetic political content, even where such requirements are backed by a more comprehensive statutory framework than currently exists in India.

7.3 Comparative Regulatory Approaches

India's evolving approach can usefully be situated alongside contemporaneous regulatory efforts elsewhere. The European Union's AI Act requires that AI-generated content, particularly deepfakes, be clearly disclosed as such, embedding this obligation within a broader, risk-tiered statutory framework for AI systems generally, a more comprehensive approach than India's current intermediary-rules-based strategy, though one that has not, as noted above, proven fully effective at preventing the circulation of synthetic election content in practice. In the United States, regulation has proceeded largely at the state level in the absence of comprehensive federal AI legislation, exemplified by statutes such as Texas's law criminalizing the creation and dissemination of deepfake videos intended to deceive voters or alter the electoral process within a defined window before an election, an approach that is narrower in scope but potentially more directly enforceable against domestic bad actors than broader disclosure mandates. China, meanwhoriginously moving in a different direction, has since September 2025 implemented a mandatory, standardized system for identifying AI-generated synthetic content, reflecting that jurisdiction's generally more centralized and interventionist approach to platform governance.

Set against this comparative backdrop, India's regulatory trajectory, combining Election-Commission advisories specific to political actors with newly proposed, narrower intermediary obligations focused on labelling and traceability, occupies a middle position: more developed than jurisdictions with no dedicated response at all, but considerably less comprehensive than the European Union's statutory model, and dependent for its practical effectiveness on the cooperation of political parties, the technical capacity of platforms to implement detection and labelling at scale, and the continued vigilance of an already strained ecosystem of independent fact-checking organizations.

8. Analytical Assessment: AI's Impact on Political Public Opinion

Drawing on the empirical material presented above, this section offers an analytical assessment of how AI is reshaping the formation of political public opinion in India, organized around the three effects introduced in the conceptual framework: enhancement, distortion, and asymmetry.

8.1 Enhancement: Expanding Access and Inclusion

The clearest enhancement effect documented in this paper concerns linguistic access. Tools such as Bhashini demonstrably reduced the practical barrier that language has historically posed to substantive political communication between national leaders and non-Hindi-speaking electorates, a genuinely significant development in a country where linguistic identity is closely bound up with regional political identity and historical grievances concerning central government dominance by Hindi-speaking elites. To the extent that AI translation allows citizens to receive political messages directly, in their own language, rather than through mediated or delayed translation, it plausibly strengthens the quality of political communication available to previously underserved linguistic communities, an outcome broadly consistent with democratic values of inclusion and equal access to information.

AI-powered chatbots and interactive tools, such as those integrated into the ruling party's NaMo mobile application, have also been used to disseminate information about government policies and welfare programmes at a scale and interactivity that traditional print or broadcast media could not match, potentially improving citizens' awareness of, and ability to access, entitlements to which they are eligible, though this benefit is naturally concentrated in favour of whichever party controls the state apparatus and its associated communication infrastructure at a given time.

8.2 Distortion: Manipulation, Misinformation, and the Liar's Dividend

The distortion effect is most visible in the documented cases of deepfake and synthetic media described in Section 4. Beyond the direct harm caused by any single fabricated video or audio clip, misleading a subset of voters who encounter and believe it, the cumulative effect of repeated, high-profile deepfake controversies during the 2024 campaign plausibly generated a broader corrosion of epistemic confidence among the Indian electorate: a growing uncertainty about whether any given piece of political audiovisual content, however genuine it might appear, can be trusted at face value. This uncertainty operates in both directions. It allows false content to mislead voters who accept it as authentic, and, equally corrosively, it allows genuine, damaging content, an authentic gaffe, an authentic admission, an authentic scandal, to be dismissed by its subject as an AI fabrication, the liar's dividend effect noted in Section 2. Indian political discourse has already exhibited instances of this dynamic, as accusations and counter-accusations of deepfake fabrication became a routine feature of disputes between the BJP and Congress during and after the 2024 campaign, with each side alleging the other had produced or benefited from synthetic disinformation.

A further distortion concern relates to the exploitation of India's deep social cleavages, particularly caste and religion, through AI-generated content. The Amit Shah reservation-policy deepfake illustrates how synthetic media can be weaponized to target the single most sensitive fault line in Indian electoral politics, the politics of caste-based affirmative action, in a manner calibrated for maximum social and political disruption. Because such content can be produced rapidly, cheaply, and in a form calibrated for virality on closed messaging platforms, AI arguably increases both the frequency and the precision with which such socially destabilizing content can be deployed relative to earlier, more labour-intensive forms of political misinformation.

8.3 Asymmetry: Uneven Capability Across Political Actors

A recurring theme across the empirical material reviewed in this paper is that AI capability in Indian politics is not evenly distributed. The ruling party's access to government-built digital infrastructure, including Bhashini and welfare-linked citizen databases, provides it with tools for AI-enabled political communication that opposition parties cannot easily replicate through private resources alone, at least not without incurring costs that smaller or less well-resourced parties may struggle to bear. Similarly, campaign expenditure data suggesting that authorized AI content spending ran into tens of millions of dollars during the 2024 campaign indicates that AI-enabled political communication at scale remains, for now, primarily accessible to well-funded national parties rather than smaller regional or local political formations, potentially reinforcing rather than disrupting existing patterns of political and financial power within the Indian party system.

This asymmetry effect has implications for electoral competitiveness that are distinct from, though related to, the distortion effect discussed above. Even where AI is used entirely for legitimate, transparent campaign communication rather than deception, its uneven availability across political actors can shift the competitive balance of an election in ways unrelated to underlying voter preferences, a concern that echoes longstanding debates in Indian electoral law concerning campaign finance disparities, but that now extends into the domain of information technology and communications infrastructure as well.

9. Analytical Assessment: AI's Impact on Electoral Politics

This section turns from public opinion formation to the broader question of how AI has reshaped the practice of electoral politics in India: campaign strategy, the competitive dynamics between parties, and the institutional processes that govern elections.

9.1 The Transformation of Campaign Strategy

AI has altered the underlying economics of political campaigning in India in ways that extend well beyond any single controversial incident. The capacity to generate large volumes of personalized, multilingual content without a proportional increase in campaign staffing represents a structural shift in how parties can approach voter outreach, moving from an emphasis on mass, undifferentiated messaging toward an emphasis on individually calibrated communication delivered through closed, high-trust channels such as WhatsApp. This shift builds on, and intensifies, a longer-running trend in Indian political communication toward decentralized, booth-level, WhatsApp-mediated campaigning that had already been well documented in scholarship on the 2014 and 2019 elections, but AI substantially lowers the marginal cost of producing genuinely individualized content at this scale, a capability that did not previously exist even within the pre-existing WhatsApp-centric campaign model.

This transformation has also altered the profile of campaign personnel and expertise that parties require, with AI content production specialists, data analytics teams, and third-party AI vendors, including small, specialized firms such as those described in journalistic coverage of individual AI content producers serving multiple campaigns, becoming an increasingly integral part of the electoral ecosystem alongside more traditional campaign roles such as booth-level organizers and media strategists.

9.2 Electoral Integrity and the Erosion of a Level Playing Field

The proliferation of deepfake controversies during the 2024 campaign, and the subsequent regulatory response, together suggest that AI poses a genuine and evolving threat to the perceived legitimacy of Indian electoral contests, independent of whether any single piece of synthetic content actually changed the outcome of a specific race. Democratic legitimacy depends not only on elections being free of manipulation in some objective sense, but on broad public confidence that they are so; the recurring cycle of deepfake allegations, counter-allegations, and disputed authenticity claims documented throughout the 2024 campaign, and continuing into subsequent state elections such as the 2025 Bihar assembly polls, risks normalizing a climate of generalized suspicion surrounding campaign communication that could, over time, undermine confidence in electoral outcomes regardless of their underlying integrity.

The Election Commission's response, requiring takedown of flagged content within three hours and mandating disclosure labels on AI-generated campaign material, represents a meaningful, if still evolving, institutional effort to preserve electoral integrity in the face of these pressures. However, the practical effectiveness of this response depends heavily on detection capacity that, as cybersecurity officials themselves acknowledged during the 2024 campaign, remains limited: no fully reliable automated technology yet exists to distinguish AI-generated from authentic audiovisual content at the speed and scale required to intercept viral misinformation before it achieves broad circulation on closed messaging platforms.

9.3 Implications for Political Parties, Candidates, and Voters

For political parties, the arrival of AI in Indian electoral politics presents both a strategic opportunity and a reputational risk. Parties that develop early, sophisticated AI content capabilities gain a communications advantage, but they simultaneously expose themselves to accusations of unethical or deceptive practice, as both major national parties experienced during mutual recriminations over deepfake content in 2024. For individual candidates, particularly those below the national tier who may lack the resources of a national party's central AI infrastructure, the technology may be primarily experienced as a source of vulnerability, since a well-produced deepfake targeting a local candidate, especially in the final days of a campaign when there may be insufficient time for rebuttal or independent verification, could materially affect a close local contest with limited institutional recourse available before polling concludes.

For voters, the cumulative effect of these developments is a political information environment that is simultaneously richer, in the sense of receiving more direct, personalized, and linguistically accessible communication from political leaders and parties, and more precarious, in the sense of facing a growing burden to critically evaluate the authenticity of the audiovisual political content they encounter, particularly on closed messaging platforms where the traditional cues that might prompt scepticism, an unfamiliar sender, an implausible source, are often absent because such content typically arrives via trusted personal or community contacts rather than from unknown senders.

10. Risks and Vulnerabilities

This section synthesizes the specific risks and vulnerabilities that emerge from the analysis above, several of which are distinctive to the Indian context rather than generic features of AI-enabled political communication observed elsewhere.

10.1 Detection and Enforcement Limitations

As multiple sources cited in this paper indicate, reliable automated detection of AI-generated audiovisual content remains an unsolved technical problem, particularly for detection that must operate at the speed and scale required to intercept content before it achieves viral circulation. This limitation constrains the practical effectiveness of both the Election Commission's takedown-based enforcement approach and MeitY's proposed labelling and traceability requirements, since labelling regimes depend on either voluntary compliance by content producers, who by definition have an incentive to evade detection when producing deceptive content, or on downstream platform detection capabilities that remain technically imperfect.

10.2 Linguistic and Cultural Diversity as a Vulnerability

The same linguistic diversity that AI translation tools can serve constructively also constitutes a vulnerability, since content moderation, fact-checking, and detection systems trained primarily on English-language or Hindi-language data may perform less reliably when confronted with synthetic content produced in less-resourced regional languages and dialects, potentially leaving voters in linguistically underserved regions more exposed to undetected AI-generated misinformation than voters in better-resourced linguistic markets.

10.3 The Closed-Network Problem

India's political communication ecosystem is unusually dependent on closed, encrypted messaging networks, particularly WhatsApp, relative to open social media platforms such as X or public Facebook pages. This structural feature, well documented in earlier scholarship on Indian political communication, substantially complicates AI-related content moderation and fact-checking, since content circulating within private WhatsApp groups and individual chats is largely invisible to platform-level detection systems, independent fact-checkers, and the Election Commission's own monitoring mechanisms until it has already achieved significant circulation, at which point takedown or correction is considerably less effective at limiting the content's ultimate influence on public opinion.

10.4 Digital Divides and Differential Exposure

India's substantial urban-rural and socioeconomic digital divides mean that exposure to, and capacity to critically evaluate, AI-generated political content is unevenly distributed across the electorate. Voters with lower digital literacy, including many older and rural voters who form a significant share of the Indian electorate, may be simultaneously more receptive to the benefits of AI-enabled multilingual outreach and less equipped to critically assess the authenticity of synthetic audiovisual content, creating a compounded vulnerability that existing regulatory frameworks have not yet directly addressed.

11. Opportunities: AI as a Tool for Democratic Strengthening

Notwithstanding the risks documented above, this paper's framework of enhancement, distortion, and asymmetry effects implies that AI is not inherently or exclusively a threat to Indian democracy, and several genuine opportunities for democratic strengthening merit consideration alongside the risks.

First, AI-powered translation and accessibility tools, exemplified by Bhashini, could be extended beyond partisan campaign use to serve broader civic purposes, including multilingual voter education materials produced and disseminated by the Election Commission itself, multilingual access to legislative proceedings and government documents, and improved accessibility of civic information for citizens with disabilities, applications that would extend the enhancement effect identified in this paper beyond its current, largely partisan deployment toward more genuinely public and non-partisan ends.

Second, AI-based analytics could, in principle, assist the Election Commission's own administrative functions, including the detection of coordinated inauthentic behaviour and misinformation campaigns, more efficient processing of the millions of grievances and complaints the Commission receives during a campaign period, and improved voter list maintenance and electoral roll accuracy, functions that are distinct from, and less contested than, AI's use in political campaign messaging itself.

Third, civil society fact-checking organizations, which have played an important role in identifying and flagging deepfake content during recent Indian elections, could themselves benefit from AI-assisted detection tools, provided such tools receive adequate institutional support and are deployed with sufficient technical sophistication to keep pace with

the generative capabilities of the AI content they are attempting to detect, an ongoing technical race that currently appears to favour content producers over detectors.

Realizing these opportunities while containing the risks documented earlier will require deliberate institutional choices, discussed in the following section, rather than being an automatic byproduct of the technology's continued diffusion through India's political system.

12. Policy Recommendations

Based on the analysis presented in this paper, several policy directions merit consideration by Indian regulators, the Election Commission, political parties, and civil society organizations.

First, the Election Commission's advisory-based approach to regulating political parties' use of AI should be strengthened with clearer, more specific, and more consistently enforced disclosure requirements, moving beyond general exhortations toward specific technical standards for labelling AI-generated content, potentially including standardized metadata formats that would be more resistant to removal or manipulation than visible on-screen labels alone.

Second, the MeitY draft amendments of October 2025 should be finalized with particular attention to the practical enforceability of traceability and labelling requirements against sophisticated and cross-border actors, informed by the comparative experience of the European Union, where similar disclosure requirements have not fully prevented the circulation of synthetic election content, suggesting that labelling requirements should be understood as one component of a broader strategy rather than a complete solution in themselves.

Third, given the centrality of closed messaging platforms, particularly WhatsApp, to the spread of both authentic and synthetic political content in India, regulators should explore mechanisms, developed in consultation with platform operators and with due regard for the privacy and encryption protections that make such platforms valuable to users in the first place, to improve the traceability of virally forwarded content without undermining the end-to-end encryption on which user privacy and security depend, an admittedly difficult balance that has proven contentious in previous Indian intermediary-liability debates.

Fourth, investment in independent, multilingual fact-checking capacity should be expanded, with particular attention to regional and non-Hindi languages that have historically received less attention from both platform-level content moderation systems and civil society fact-checking organizations, in order to close the detection gap identified in Section 10.2 of this paper.

Fifth, campaign finance and disclosure requirements should be updated to require political parties to report AI-related campaign expenditure as a distinct category, improving transparency around the asymmetry effect documented in this paper and enabling more informed public and regulatory scrutiny of the scale and distribution of AI-enabled campaign communication across parties of different sizes and resource levels.

Sixth, voter education initiatives, ideally coordinated by the Election Commission in partnership with civil society and educational institutions, should be expanded to improve public digital literacy specifically with respect to AI-generated content, building on the Commission's existing voter education mandate but extending it to address the particular epistemic challenges that generative AI introduces, including practical guidance on recognizing common signs of synthetic media and the importance of verifying sensational political content, particularly content received through closed messaging channels, before sharing it further.

13. Conclusion

Artificial intelligence has become, within the span of a single general election cycle, an integral feature of Indian democratic practice, reshaping how political leaders communicate with citizens, how parties compete for votes, and how the Election Commission and regulators attempt to preserve the integrity of the electoral process. The evidence reviewed in this paper suggests that AI's impact on Indian democracy cannot be reduced to a single verdict of harm or benefit. The technology has demonstrably enhanced the reach and linguistic inclusiveness of political communication in a country whose scale and diversity have long posed genuine challenges to national political integration, while simultaneously introducing new and serious risks of manipulation, misinformation, and the erosion of shared factual ground on which democratic deliberation depends.

The Indian case also illustrates a broader, more general dynamic likely to recur in other large, linguistically and socially diverse democracies as generative AI diffuses further into political life: that the same underlying technical capabilities, translation, personalization, synthetic media generation, can simultaneously serve democratically enhancing and democratically corrosive purposes, often within a single campaign and sometimes even within a single piece of content, depending on questions of consent, accuracy, and disclosure that existing electoral and information-technology

law was not originally designed to address. India's evolving regulatory response, combining Election Commission advisories specific to political actors with newly proposed intermediary obligations focused on labelling and traceability, represents a genuine, if still incomplete, institutional effort to grapple with this dynamic, and its ultimate effectiveness will depend on technical, political, and civil-society capacities that are, as of this writing, still very much in the process of development.

As India moves toward its next general election cycle, the trajectory of AI's role in the country's democratic life will likely be shaped by the interaction of several factors examined in this paper: the pace and comprehensiveness of regulatory development at MeitY and the Election Commission, the technical evolution of both AI content generation and AI content detection capabilities, the degree to which political parties across the spectrum choose to exercise restraint in deploying AI for deceptive rather than constructive purposes, and the capacity of India's civil society and media ecosystem to build the fact-checking and digital literacy infrastructure necessary to help citizens navigate an information environment of unprecedented technological sophistication. The stakes of getting this balance right are considerable, given both the sheer scale of the Indian electorate and the likelihood that India's experience will continue to serve, as it already has, as an instructive and closely watched case study for democracies worldwide confronting the same underlying technological transformation.

References

- Asia Pacific Foundation of Canada. (2024, June 5). In Indian election, AI amplifies political reach but raises new risks. https://www.asiapacific.ca/publication/indian-election-use-of-ai-political-campaigns-voter-engagement
- CoinGeek. (2025, November 11). India proposes law to label AI-generated social media content. https://coingeek.com/india-proposes-law-to-label-ai-generated-social-media-content/
- Deccan Herald. (2024, April 2). AI content, deep fake big headache above all during polls. https://www.deccanherald.com
- Deccan Herald. (2024). Keep 2024 campaign free of AI mischief. https://www.deccanherald.com/amp/story/opinion%2Fkeep-2024-campaign-free-of-ai-mischief-2784403
- Deccan Herald. (2024). Lok Sabha election 2024: Remove fake content within 3 hours, EC tells political parties.
- Global Network on Extremism and Technology (GNET). (2024, September 11). Deep fakes, deeper impacts: AI's role in the 2024 Indian general election and beyond. https://gnet-research.org/2024/09/11/deep-fakes-deeper-impacts-ais-role-in-the-2024-indian-general-election-and-beyond/
- Khurana and Khurana. (2025). Deepfake regulation India 2025: MeitY's comprehensive IT law. https://www.khuranaandkhurana.com/deepfake-regulation-india-2025-meity-s-comprehensive-it-rules-amendment
- Mondaq. (2025, December 18). Deepfake regulation India 2025: MeitY's comprehensive IT rules amendment. https://www.mondaq.com/india/new-technology/1720652/deepfake-regulation-india-2025-meitys-comprehensive-it-rules-amendment
- MediaNama. (2025, October 28). ECI cracks down on deepfakes before Bihar polls. https://www.medianama.com/2025/10/223-election-commission-deepfakes-advisory-bihar-elections/
- NetMission.Asia. (2025, February 26). The digital election: How technology redefined India's 2024 general election. https://netmission.asia/2025/02/26/the-digital-election-how-technology-redefined-indias-2024-general-election-ankita-rathi/
- New Lines Magazine. (2024, July 15). AI and deepfakes played a big role in India's elections. https://newlinesmag.com/spotlight/ai-and-deepfakes-played-a-big-role-in-indias-elections/
- NPR. (2024, December 21). How AI deepfakes polluted elections in 2024. https://www.npr.org/2024/12/21/nx-s1-5220301/deepfakes-memes-artificial-intelligence-elections
- Record of Law. (2025, December 28). AI in Indian elections: Deepfakes, disinformation and the challenge to electoral integrity. https://recordoflaw.in/ai-in-indian-elections-deepfakes-disinformation-and-the-challenge-to-electoral-integrity/
- Ronin Legal Consulting. (2025, October 29). India's draft deepfake rules: MeitY's first regulation move. https://roninlegalconsulting.com/india-finally-stands-tentatively-up-to-ai-deepfakes/
- Journalism University. (2026, April 11). The role of social media in Indian political communication. https://journalism.university/social-media-and-society/role-social-media-indian-politics/
- StartupArticle. (2026, April 23). 6 ways how AI changed Indian elections and the electoral process. https://startuparticle.com/technology/2026/04/how-is-ai-changing-indian-elections-and-the-electoral-process/
- TechPolicy.Press. (2024, May 25). The era of AI-generated election campaigning is underway in India. https://www.techpolicy.press/the-era-of-aigenerated-election-campaigning-is-underway-in-india/
- TechPolicy.Press. (2024, September 25). India's experiments with AI in the 2024 elections: The good, the bad and the in-between. https://www.techpolicy.press/indias-experiments-with-ai-in-the-2024-elections-the-good-the-bad-the-inbetween/
- TechPolicy.Press. (2026, January 14). Synthetic media in elections and politics in India. https://www.techpolicy.press/synthetic-media-in-elections-and-politics-in-india/
- Trace Law Partners. (2025, November 3). Regulating deepfakes in India: Understanding the draft amendments to the IT rules and their implications. https://www.tracelawpartners.com/post/regulating-deepfakes-in-india-understanding-the-draft-amendments-to-the-it-rules-and-their-implicat
- UC Berkeley School of Information. (2024, April 23). AI deepfakes are threatening to upend global elections. Hany Farid dives deep. https://www.ischool.berkeley.edu/news/2024/ai-deepfakes-are-threatening-upend-global-elections-hany-farid-dives-deep

- Press Reader / Santa Fe New Mexican. (2024, April 19). 'Wild West' for AI tools ahead of India's election.