Decoding Digital Campaigning with Deep Learning: Evidence from India

Amreeta Das*

September 15, 2025

Abstract

Right-wing parties have surged globally, yet explanations rooted in economic insecurity and cultural backlash cannot fully explain how their grievances become electoral victories. A critical but underexplored dimension is the dominance of the right in digital political spaces, shaped by opaque, data-driven strategies. This paper investigates this digital advantage by analyzing the Bharatiya Janata Party's (BJP) Facebook ad campaigns in India from 2020 to 2024. Using a novel deep learning approach, I uncover systematic differences in content and targeting strategies. Compared to their competitors, BJP ads feature distinct visual motifs and subtle symbolic cues that diverge from overt nationalism, and the party engages in more intensive microtargeting by tailoring content to specific geographic audiences. Format matters too: BJP's image ads emphasize saffron, the lotus, and Modi's portrait; its video ads are less distinctive, with fewer party, religious, and flag cues. A proposed survey experiment tests the persuasive power of these visual elements. Together, the findings shed light on how right-wing parties utilize opaque digital infrastructures and demonstrate how deep learning can scale the study of political campaigning in new ways.

*Ph.D. Candidate in Political Science, University of California Merced, adas13@ucmerced.edu

1 Introduction

The contemporary global political landscape is marked by the ascendance of right-wing movements. In the Global North, this trend is often attributed to broad structural forces such as economic insecurity and cultural backlash (Inglehart and Norris 2016; Norris and Inglehart 2019; Rodrik 2021). In the Global South, explanations differ, focusing more on uneven development, selective welfare policies, corruption, and identity-based appeals. Although the specific conditions differ across regions, what unites them is that they create fertile ground for right-wing actors to mobilize. A central question is how these actors convert grievances into durable electoral success.

This paper posits that a critical mechanism is the superior execution of digital political campaigns. A pattern has emerged wherein right-wing parties demonstrate a disproportionate effectiveness in leveraging social media and digital tools to mobilize supporters and shape political narratives (Chen et al. 2021; González-Bailón et al. 2022; Huszár et al. 2022). The central piece, therefore, is not simply why the right is rising, but why it is uniquely adept at fighting and winning on the digital battlefield.

This brings us to the specific puzzle that motivates this paper. Right-wing parties consistently dominate the online political space. What gives these parties a comparative advantage in digital campaigning compared to their opposition? This puzzle cannot be fully addressed with existing approaches. Much of the literature on digital politics relies on textual analysis, network structures, or disinformation tracing to highlight the advantage of the right (Huszár et al. 2022; McDonnell and Ondelli 2025; Nikolov, Flammini, and Menczer 2020). These studies illuminate important dynamics but largely overlook the visual and symbolic dimensions of digital campaigning, precisely the realm where right-wing movements thrive. Images, symbols, and design choices circulate with high velocity online, but they remain methodologically difficult to study at scale. Unlocking this "black box" requires new tools. The challenge is compounded by the fact that parties' visual strategies are format-dependent: the same organization may have different strategies when it comes to video vs image ads.

This paper introduces such tools. I employ deep learning techniques to inductively analyze the visual content of online political advertisements. This approach enables large-scale pattern recognition of symbolic cues, such as national flags, religious motifs, and leader iconography, that are often invisible to traditional methods. To complement these observational analyses, I propose a survey experiment that directly tests the persuasive power of these visuals, providing causal leverage on their political effects.

Focusing on India as a critical case where digital campaigning has reached unprecedented scale and sophistication, both in terms of market size and organizational investment, this paper investigates the strategies of the Bharatiya Janata Party (BJP) in comparison to its competitors. The analysis identifies two mechanisms that distinguish the BJP's digital advantage. First, while symbolic motifs appear across Indian political advertising, BJP advertisements rely more systematically on visuals that signal Hindu identity and belonging, rather than on broadly secular appeals. Second, although all major parties make use of geographic targeting, the BJP engages in microtargeting with greater granularity and intensity, tailoring visual content to specific constituencies more aggressively than its rivals. The results therefore offer two levels of insight: they document which visual strategies are prevalent across Indian digital campaigns, and they identify what is distinctive about the BJP's use of imagery and targeting. Although India is distinctive in scale and context, the mechanisms identified here provide new grounds for theorizing digital advantage and offer a framework that can be tested in other settings.

This paper makes three contributions to the study of right-wing politics and political communication.

First, it revisits the long-standing debate about the impact of campaigns. Classic accounts advanced a "minimal effects" perspective, arguing that campaigns rarely change minds because partisan loyalties and structural conditions already shape electoral outcomes (Berelson 1954; Brady and Johnston 2009; Kalla and Broockman 2018). At the same time, newer work has stressed that campaigns may still matter in specific ways, including mobilizing supporters, increasing turnout, or shifting the salience of issues (Druckman 2001; Finkel 1993; Green and Gerber 2019; Hersh 2015). This paper does not claim to resolve this debate. Instead, it provides evidence that campaigns may function through mechanisms often overlooked in minimal-effects frameworks, particularly in fragmented media environments where exposure to political messages is uneven. The analysis suggests that strategic design choices and the visual grammar of advertisements can create opportunities for parties to amplify their messages and sharpen symbolic appeals.

Second, the paper highlights how the nature of campaigning has been transformed by the rise of digital platforms. Earlier modes of campaigning such as television, radio, and print advertising relied on broadcast communication aimed at mass audiences, with limited scope for tailoring messages to different groups (Iyengar and Simon 2000; Hersh 2015). These formats were also more easily observable and subject to clearer regulatory oversight compared to online spaces. By contrast, contemporary campaigns operate in an environment where digital platforms enable microtargeting in real time, often with little transparency. Social media relies heavily on images, videos, and interactive content that spread rapidly across networks. Research shows that audiovisual formats evoke stronger emotions, travel farther and faster, and often carry subtle or coded appeals that text cannot easily convey (Albertson 2015; Engesser et al. 2017; Schmid, Schulze, and Drexel 2025). These features provide parties with new tools to connect with audiences while obscuring meaning from outsiders or regulators. By examining visual and symbolic strategies in digital advertising, this paper draws attention to a critical but underexplored aspect of campaigning that could be central to how right-wing movements gain traction in the contemporary media environment.

Third, the paper makes a methodological contribution. Much of the existing literature on digital politics relies on text-based analysis, network structures, or targeting strategies to understand how parties communicate online (Huszár et al. 2022; Nikolov, Flammini, and Menczer 2020; Votta et al. 2024). While these approaches reveal important dynamics, they rarely capture the visual dimension of political communication, which is precisely where right-wing parties have been most innovative. To address this gap, I apply deep learning methods to systematically analyze the visual content of more than 54,000 Facebook ads fielded by the top 6 national parties in India between 2020 and 2024. This approach enables large-scale detection of recurring motifs such as national flags, religious symbols, and leader iconography that are often overlooked by conventional methods. To connect these patterns to their political impact, the paper also proposes a survey experiment that tests the persuasive power of the identified visual cues. Together, this design illustrates how deep learning tools can open new avenues for studying the symbolic and emotional elements of campaigns at scale.

2 The Demand-Side Drivers of Right-Wing Populism

To analyze the contemporary right, it is crucial to distinguish between its constituent parts. The political right, historically defined in opposition to a left that prioritizes equality, is broadly characterized by an emphasis on hierarchy, order, and tradition (Bobbio 1996; Heywood 2021). Within this space, traditional conservatism represents an ideology committed to preserving established institutions and a deferential view of authority (Heywood 2021). Populism, in contrast, is a "thin-centered ideology" that structures politics around a moral antagonism between "the pure people" and "the corrupt elite" (Mudde 2004; Stanley 2008). As a thin ideology, populism requires a "thick" host; when fused with nativism and authoritarianism, it creates right-wing populism (Inglehart and Norris 2016; March 2017; Mudde 2019). This adds a second, exclusionary antagonism of "us" (the native people) versus "them" (outsiders), who are framed as being illegitimately favored by the corrupt elite (Inglehart and Norris 2016; March 2017; Mudde 2019).

Within this context, to understand the populist right's digital success, one must first grasp the underlying societal conditions that have created a demand for its political offerings. The process often begins with objective conditions of economic insecurity (Algan et al. 2017; Bossert et al. 2019; Guiso et al. 2024; Rodrik 2021). Decades of globalization, technological change, and the erosion of labor market protections have hollowed out the middle class and increased economic insecurity, particularly for low-skilled workers (Algan et al. 2017; Autor, Dorn, and Hanson 2013; Barone and Kreuter 2021). Foundational economic theories demonstrate that trade liberalization, while potentially increasing overall efficiency, inevitably creates economic losers in advanced economies. This widespread economic grievance was acutely amplified by events like the 2008 global financial crisis and subsequent austerity measures, which are strongly correlated with increased support for right-wing parties (Fetzer 2019; Funke, Schularick, and Trebesch 2016). The psychological consequence of this insecurity is a profound loss of trust in established institutions, including mainstream political parties, financial systems, and the media, and the crystallization of a powerful anti-elite sentiment (Algan et al. 2017). This widespread disillusionment creates a political opportunity and a demand for an alternative to the perceived failed consensus of the political establishment (Rodrik 2017, 2021).

It is at this stage that the right-wing populist narrative proves so effective. Political entrepreneurs successfully capture this free-floating economic anger by channeling it through pre-existing cultural anxieties and salient ethno-nationalist cleavages. These cultural grievances are often described as a "backlash" from once-predominant sectors of the population, particularly older, less educated cohorts, who perceive progressive cultural shifts as a fundamental threat to their traditional values and social status (Inglehart and Norris 2016; Norris and Inglehart 2019). The increasing political emphasis on issues like environmentalism, multiculturalism, gender equality, and LGBTQ+ rights can foment a desire for a return to traditional morality (Eatwell and Goodwin 2018; Inglehart and Norris 2016; Rothwell and Diego-Rosell 2016). Right-wing actors provide a simple, emotionally resonant, and identity-based explanation for complex economic problems: the "corrupt elite" has betrayed the "true people" by prioritizing the interests of global capital and domestic "out-groups" such as immigrants and minorities (Mudde 2004). Economic insecurity provides the essential fuel for populist demand, while cultural anxieties provide the specific narrative framework that gives this demand its distinctly right-wing character.

It is important to acknowledge that Western-centric frameworks inadequately capture the right's rise in the Global South, where patterns such as deindustrialization or large-scale immigration are less relevant. In regions like Latin America, Asia, and Africa, right-wing populism often arises from uneven capitalist development, selective welfare policies, and strategic exploitation of crises and identity politics. Right-wing populism in the Global South often mobilizes support through appeals to religious, ethnic, or national identity, and by exploiting dissatisfaction with corruption, crime, or weak institutions. For example, in India, the BJP's rise combines cultural nationalism with economic liberalism, building broad coalitions across caste and class lines (Chhibber and Verma 2014; Jaffrelot 2017). This is propelled by disillusionment with government inefficiency and corruption, alongside promises of development and market reforms. Leaders blend pro-growth narratives with nationalist symbolism that appeals to both emerging middle classes and marginalized communities through targeted welfare and identity politics. Such cases underscore the role of charismatic leadership and the fusion of social conservatism and economic agendas in right-wing consolidation across the Global South.

3 The Right's Asymmetric Advantage in the Digital Campaigning Space

The academic literature on political campaigns is marked by a long-standing and divided debate over their electoral impact. One influential school of thought posits that campaigns have "minimal effects" on vote choice, arguing that durable factors like partisanship largely predetermine outcomes (Berelson 1954; Brady and Johnston 2009; Kalla and Broockman 2018). Conversely, a substantial body of work suggests that while campaigns may not frequently convert partisans, they are crucial for mobilizing supporters, informing the electorate, and setting the political agenda (Druckman 2001; Finkel 1993; Green and Gerber 2019; Hersh 2015; Hillygus and Jackman 2003; Jacobson 2015). Consequently, the scholarly focus has shifted from whether campaigns matter to exploring for whom and under what conditions they matter most.

The modern digital media environment represents a critical new condition, with emerging evidence of a structural advantage for the populist right. Social media platforms facilitate the spread of polarized and populist narratives by enabling direct appeals to targeted audiences while bypassing traditional gatekeepers (Huszár et al. 2022; Votta et al. 2024). This supports the "amplification of the right" thesis: engagement-driven algorithms privilege divisive, novel, and emotionally charged content, features characteristic of right-wing populist communication (Engesser et al. 2017; González-Bailón et al. 2022; Rathje, Van Bavel, and Van Der Linden 2021; Vosoughi, Roy, and Aral 2018).

These dynamics amplify affective polarization, reinforcing identity-based mobilization and providing fertile ground for right-wing actors to exploit cultural grievances (Engesser et al. 2017; Iyengar, Sood, and Lelkes 2012). Exposure to sensationalist and partisan content further fosters cynicism and distrust in institutions, heightening the appeal of anti-establishment rhetoric (Gentzkow and Shapiro 2010; Mutz 2007). Thus, evolving media environments not only disseminate right-wing ideas but also actively shape demand-side conditions for their rise by influencing voter perceptions and engagement, creating a feedback loop between political actors and media consumption patterns.

Right-wing movements also excel at cultivating online communities that function as "emotional refuges" for supporters, where shared grievances and anxieties are validated and transformed into collective energy (Törnberg and Törnberg 2025). These spaces serve both recruitment and retention functions, reinforcing partisan identity and sustaining long-term engagement. To broaden their appeal, right-wing actors employ strategic ambiguity, using coded language, dog-whistles, irony, and humor to circulate radical ideas while maintaining plausible deniability (Albertson 2015; Bonikowski and Zhang 2023; Schmid, Schulze, and Drexel 2025). This "double communication" enables them to normalize exclusionary narratives while preserving mainstream legitimacy (Maly 2019; Marwick and Lewis 2017).

Campaign strategies have evolved simultaneously. While campaigns have long relied on structured voter data to mobilize supporters (Hersh 2015; Iyengar and Simon 2000), advances in digital technology have enabled sophisticated microtargeting and real-time personalization (Votta et al. 2024). Increasingly, this communication takes visual form. Images and videos evoke stronger emotional responses than text, allowing rapid transmission of complex messages. Yet, despite the "visual turn" in campaigning, research remains heavily text-centric, leaving a critical gap in understanding the informational and mobilizational power of visual content (McDonnell and Ondelli 2025).

4 The Case of the BJP in India

The rise of India's BJP provides a critical case for understanding right-wing digital dominance. Its success reflects not only effective online campaigning but also the fusion of three elements: a disciplined hierarchical organization, a resonant ethno-nationalist ideology, and the charismatic-populist leadership of Narendra Modi (Mahapatra and Plagemann 2019). While this combination is distinctive to the Indian context, it sheds light on broader dynamics of digital politics that can be observed in other settings.

Central to the BJP's mobilization is Hindutva, or Hindu nationalism, which defines Indian identity in majoritarian religious terms while positioning minorities as the "other" (Kaul 2017; Khan et al. 2017; Sarkar 2021). Though long present, Hindutva was mainstreamed under Modi's leadership after 2014, with recurring themes of Hindu insecurity and Muslim threat woven into campaign narratives (Leidig 2020). Modi's carefully cultivated online persona strengthens this project, offering an unmediated populist connection to millions of followers (Jaffrelot 2015; Rai 2019; Sircar 2020).

The party's digital dominance has been facilitated by India's rapidly expanding online ecosystem, with smartphones reaching nearly half the of the 1.5 billion population and Facebook capturing more than 60% of the social media market by 2024 ("India App Market Statistics (2024)" 2024; Statista 2024). Alongside Facebook, Instagram, YouTube, and WhatsApp have become crucial arenas for mobilization, each exploited by the BJP's IT Cell to amplify Hindutva narratives. The IT Cell, a centralized digital operations wing of the party, coordinates thousands of volunteers and paid staff to create, disseminate, and monitor online content at scale. The IT Cell orchestrates social media campaigns that amplify ideological narratives rooted in Hindutva, caste-based appeals, and economic liberalism, while also deploying disinformation and aggressive messaging against opposition parties (Garimella and Chauchard 2024; Jaffrelot 2017). While digital media has transformed communication, in-person rallies remain vital, often generating content that extends their reach online, reflecting a complementary campaign strategy (Sheikh 2024).

The BJP's digital operation illustrates how right-wing populists adapt to a visual and interactive media environment. Targeted advertising and micro-messaging allow tailored appeals across caste, class, and regional divides, reinforcing Hindu majoritarianism while integrating pro-business governance claims (Chhibber and Verma 2014; Jaffrelot and Verniers 2020). The party's digital outreach deepens these appeals, allowing finely targeted communication that resonates with diverse constituencies. Recent research has revealed how WhatsApp groups in India exhibit right-wing dominance through dense, geographically dispersed networks sharing multimedia content (Bursztyn and Birnbaum 2019; Garimella and Eckles 2020).

Focusing on the BJP as a critical case where digital campaigning has reached unprecedented levels of scale and sophistication can help uncover mechanisms of the populist right's digital campaigning that can be tested in other settings.

5 Deep Learning Methodology

Despite cross-national variation, right-wing movements consistently dominate the online political space. What explains this comparative advantage? Do their advertisements differ from those of competitors in symbolic content or targeting sophistication, or is their dominance simply a matter of scale? This paper investigates these questions through the case of India's BJP, examining whether its success stems from distinctive visual strategies, aggressive targeting, or the scope of its digital operations.

The analysis proceeds in two steps. First, I employ deep learning techniques to inductively analyze Facebook advertisements disseminated by the BJP between 2020 and 2024. This computational approach enables large-scale detection of systematic patterns in the symbolic and visual appeals of right-wing digital campaigning, patterns that text-based analyses often overlook. Second, I field a survey experiment to assess whether these symbolic appeals influence political attitudes and engagement, thereby evaluating whether the strategies identified in the inductive analysis are not only distinctive but also persuasive.

Together, this design links supply (how appeals are constructed) with potential effectiveness (whether they resonate with audiences). The present section focuses on the first step, the deep learning analyses, which provide the foundation for the survey experiment introduced later.

This research addresses a critical gap by investigating how the visual design of political advertisements reveals distinctive features of right-wing digital campaigns. While prior work in digital politics has emphasized text-based cues such as partisan language or source credibility (Carnahan et al. 2022), the visual dimension remains understudied despite its salience and complexity (see Figure 1).

Visuals embed multiple cues simultaneously: local language markers, culturally significant colors, religious symbols, campaign promises, and leader iconography. These elements may serve as precise indicators of targeted demographic and electoral appeals. To capture such patterns, this study employs deep learning and explainable AI to inductively identify

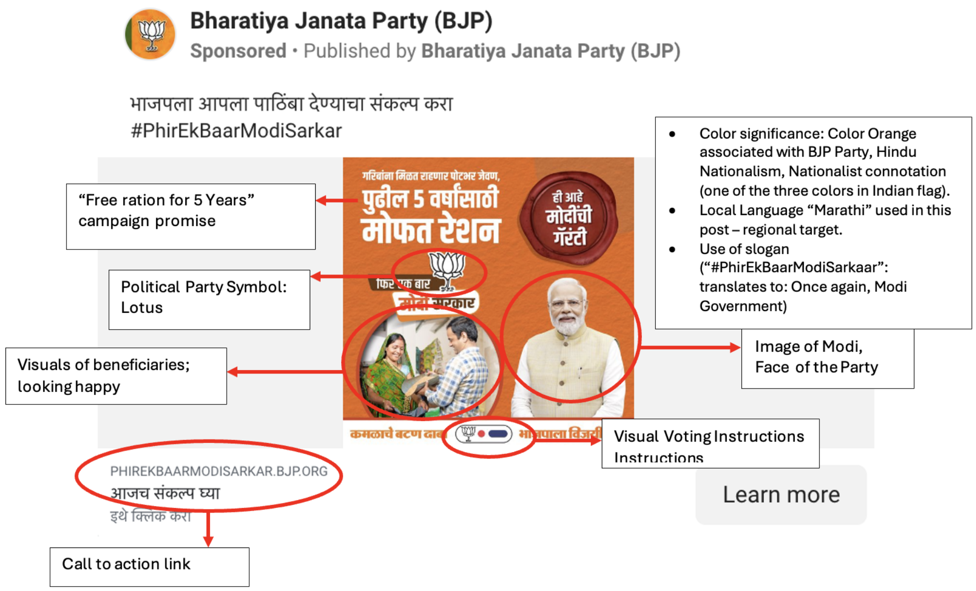

Figure 1: Example of an image advertisement by BJP on Facebook

the symbolic features of BJP advertisements between 2020 and 2024. Rather than treating images as ancillary, this approach recognizes them as structured political data, aligning with recent work in political science that applies computer vision for inductive theory-building (Breuer et al. 2025; Rizzo and Dasgupta 2024; Tarr, Hwang, and Imai 2023). The features identified computationally are then tested in a survey experiment to evaluate their persuasive effects.

Analyzing images at scale is challenging due to their high-dimensional structure, but advances in deep learning have made it possible to uncover systematic visual patterns. Within political science, scholars have increasingly turned toward advanced computational techniques, particularly deep learning and computer vision approaches, to systematically analyze these rich, multidimensional datasets. Recent studies leveraging such methodologies have contributed to diverse analytical tasks including protest detection (Won, Steinert-Threlkeld, and Joo 2017), ideological estimation through visual signaling (Xi et al. 2020), characterization of politicians' "home style" (Anastasopoulos et al. 2017), analysis of the built environment's influence on voting behaviors (Rizzo and Dasgupta 2024), automated detection of infrastructure and geographic features (Dasgupta and Ramirez 2025), and socioeconomic census estimation (Gebru et al. 2017), among numerous other applications.

Yet most prior work relies on supervised classification of labeled images. This study diverges by analyzing a complete corpus of ads linked to party and targeting metadata, enabling inductive discovery of how visual features, such as local language, cultural symbols, or stylistic design choices, function as strategic signals in digital campaigning. The following section details the data sources that make this analysis possible, including the corpus of Facebook advertisements and the auxiliary datasets used to contextualize visual content.

6 Data

6.1 Meta Ads Dataset

The primary data source is the Meta Ads Targeting Dataset, which includes all Social Issue, Electoral, and Political (SIEP) advertisements published on Facebook and Instagram in India between 2020 and 2024. Beginning August 3, 2020, the dataset records every SIEP ad run on these platforms, providing information on advertiser identity, expenditure, and targeting criteria, along with access to the associated ad creatives (images and videos).

6.2 Party Coverage

The analysis focuses on the six parties with the largest representation in the most recent general election: Bharatiya Janata Party (BJP), Indian National Congress (INC), All India Trinamool Congress (AITC), Dravida Munnetra Kazhagam (DMK), Samajwadi Party (SP), and Telugu Desam Party (TDP). These parties were selected to ensure comparability across major national actors.

6.3 Image and Video Advertisements

Between 2020 and 2024, the dataset contained approximately 54,000 image and video advertisements from the six parties. After removing duplicates to avoid data leakage, party-level distributions of ads are reported in Table 1. Data leakage refers here to the risk that identical or near-identical advertisements could appear in both the training and test sets, which would artificially inflate model performance by allowing the algorithm to "memorize" rather than genuinely learn underlying patterns.

Table 1: Dataset Composition by Party

Party	**Total image ads**	**Unique image ads**	**Total video ads**	**Unique video ads**	**Frames extracted**
BJP	14,289	226	32,389	880	6,417
INC	2,759	302	2,285	128	1,097
AITC	32	9	936	279	3,012
DMK	631	115	280	133	626
SP	115	16	69	34	207
TDP	47	25	304	118	1,477
Total	**17,873**	**693**	**36,263**	**1,572**	**12,836**

Although modest relative to the total volume of social media content, this corpus is significant because it consists entirely of paid political advertisements, each reflecting deliberate resource allocation and microtargeting decisions. Even after deduplication, many BJP ads remain highly similar, which could raise concerns about information leakage. Yet this similarity could reflect a deliberate strategy. The repetition of a narrow symbolic template may not only be a methodological artifact but also a substantive feature of how the party seeks to build coherence and recognition across contexts.

6.4 Auxiliary Datasets

To supplement the Meta Ads dataset and enable a fuller analysis of the visual, textual, and geographic dimensions of political advertising, I incorporate several external datasets and computational tools.

To classify targeted geographies as urban, peri-urban, or rural, I use the Degree of Urbanization dataset developed by the European Commission's Joint Research Centre (Schiavina, Melchiorri, and Pesaresi 2023). GHS-SMOD provides a globally consistent classification of settlement patterns at 1 km^2 resolution, combining satellite imagery with population data. This allows for systematic mapping of Facebook's PIN code and radius-based location targeting onto meaningful measures of urbanization.

In addition to geographic targeting, it is also important to capture the linguistic dimension of political advertising. Political video ads often rely heavily on spoken content in multiple Indian languages. To incorporate this dimension, I employ Whisper, a state-of-the-art automatic speech recognition (ASR) model trained on multilingual and multitask data (Radford et al. 2023). Whisper provides transcription, language identification, and English translation, thereby enabling standardized analysis of linguistic content across India's diverse language environment.

For video frame analysis, I use CLIP (Contrastive Language–Image Pretraining) (Radford et al. 2021), a vision–language model trained on 400 million image–text pairs. CLIP encodes images into 512-dimensional embeddings aligned with natural language descriptions, offering semantically rich visual features without requiring domain-specific training data. This is especially useful for political ads, where symbol recognition and thematic interpretation are highly variable.

To capture the affective tone of political messaging, I employ two lexicons. The NRC Emotion Lexicon (Mohammad and Turney 2013) maps words to basic emotional categories (e.g., joy, fear, anger), enabling a broad assessment of emotional appeals. In addition, the Lexicoder Sentiment Dictionary (Young and Soroka 2012) specifically designed for political communication, provides binary positive/negative classifications optimized for political text. Using both resources allows for validation across general-purpose and domain-specific sentiment measures.

7 Results

This section presents the empirical findings in four parts. I begin by comparing BJP and non-BJP advertisements on Facebook to establish differences in scale, spending, and efficiency, as well as the degree of creative diversity after deduplication. I then assess whether BJP ads exhibit a distinctive visual style using deep learning models trained on both image and video content, supplemented with feature coding and explainable AI analyses. Next, I examine the symbolic and emotional tone of advertisements, focusing on differences in imagery, color use, and sentiment between parties. Finally, I turn to targeting strategies, comparing BJP and INC across demographic and geographic criteria, before analyzing whether BJP adapts its visual strategies to rural and urban constituencies.

7.1 Are there differences between BJP and non-BJP ads on Facebook?

7.1.1 Descriptive statistics: Scale, Spending and Efficiency

Analysis of digital advertising activity from 2020-2024 reveals substantial asymmetry in campaign strategies across India's major political parties. BJP published 40,247 advertisements during the study period, constituting 94.8% of total advertising volume among the six parties examined (Figure 2, Panel A). INC, as the second-largest advertiser, published 1,847 advertisements, representing a 22:1 ratio between BJP and INC activity, while regional parties demonstrated considerably lower volumes.

This volume dominance corresponds with substantial financial investment, as BJP's minimum spending estimates reached \$2.35 million USD compared to INC's \$1.3 million USD, with regional parties investing between \$47,000-\$168,000 (Figure 2, Panel B). Notably, this volume-based approach trades individual content engagement for sustained presence, as BJP recorded substantially lower impressions per post (142,857) compared to regional parties like AITC (486,522) and TDP (441,739), indicating distinct strategic approaches where BJP em-

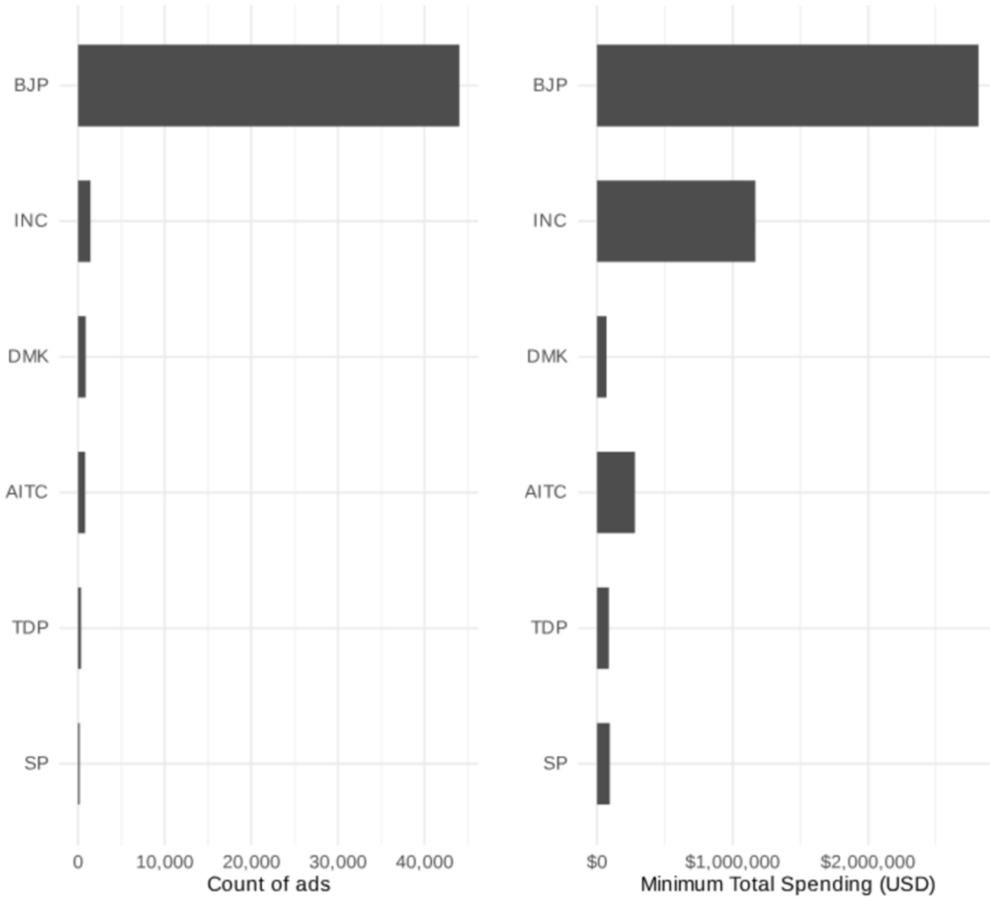

(A) **(B)**

Note: Dataset reports spending ranges. Graph shows lower bounds (minimum confirmed spending). Actual spending is higher. Converted from INR at 1 USD = 83 INR.

Figure 2: Advertisements on Facebook by top 6 National Parties (2020-2024)

ploys high-frequency distribution while competitors optimize for high-engagement content (Figure 3, Panel A). However, examination of cost-effectiveness metrics reveals that BJP achieved superior efficiency with the lowest average CPM (cost per mille, or the cost of one thousand ad impressions) at \$0.31, compared to \$2.51 for SP, suggesting optimized targeting strategies (Figure 3, Panel B). In practical terms, this means the BJP spent significantly less money for every additional thousand impressions its ads received, allowing greater reach at lower cost. These patterns collectively demonstrate that BJP has implemented an industrialized digital advertising model that prioritizes scale and cost efficiency fundamentally differentiating their approach from traditional campaign strategies employed by competing parties.

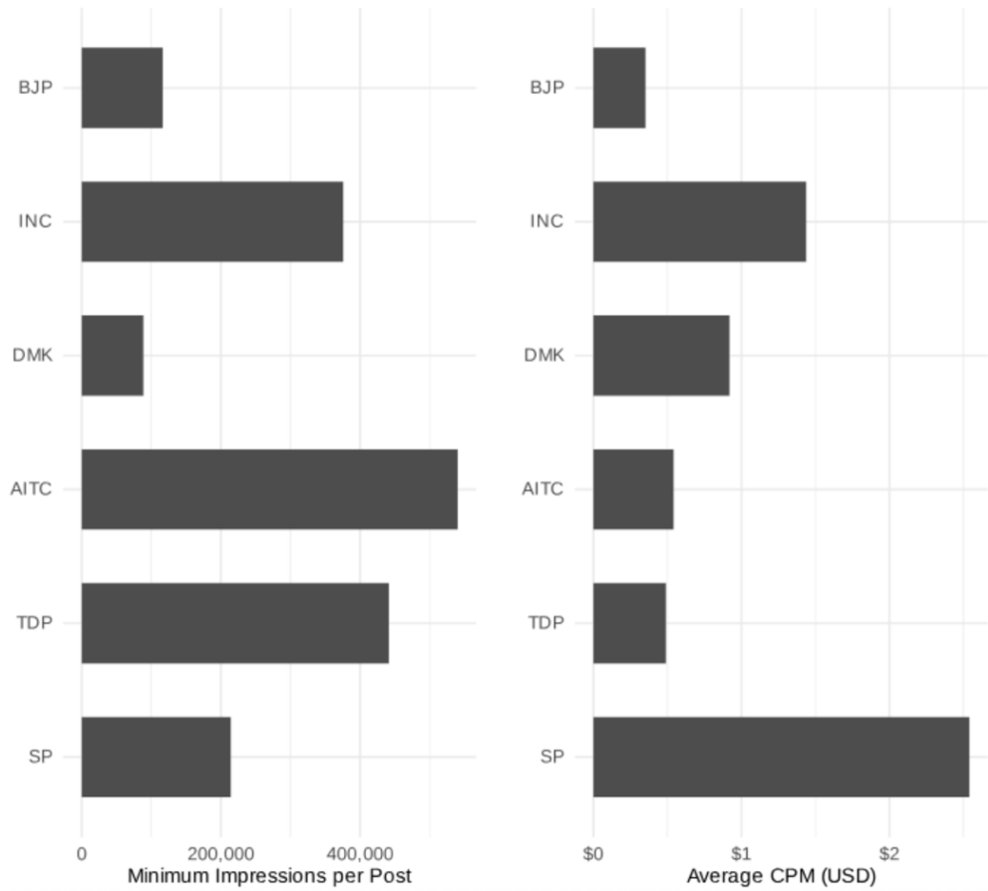

(A) **(B)**

Note: Dataset reports spending and impression ranges. Graph shows lower bounds (minimum confirmed spending and impressions). Actual spending and impressions are higher. Spending is converted from INR at 1 USD = 83 INR.

Figure 3: Advertising Performance Indicators by Political Party (2020-2024)

Note: Dataset reports spending and impression ranges. Graph shows lower bounds (minimum confirmed spending and impressions). Actual spending and impressions are higher. Spending is converted from INR at 1 USD = 83 INR.

BJP's digital operation is unmatched in both scale and efficiency, but raw counts alone do not tell us whether this activity reflects genuine creative diversity or the repeated circulation of the same messages. As seen in Table 1, once duplicates are removed, the corpus of unique ads shrinks dramatically across all parties. For example, BJP's 14,289 image ads reduce to just 226 distinct creatives, while INC's 2,759 collapse to 302, and smaller parties are left with only a handful of unique items. This pattern highlights that most parties, including BJP, rely heavily on recycling the same core visuals.

This raises a crucial question: are these few unique creatives themselves distinctive in ways that can be systematically recognized? Put differently, is BJP's advantage only a matter of industrialized distribution, or do its visuals carry a consistent branding style that sets them apart from competitors? To answer this, I turn to a deep-learning analysis of the de-duplicated ads.

7.1.2 Deep learning with Image ads

To assess visual distinctiveness, I use a convolutional neural network (CNN), a class of deep learning models that has become a dominant tool for analyzing visual imagery. The development of modern CNN architectures, such as AlexNet, VGGNet, GoogLeNet, and ResNet, has led to breakthrough performance in a wide range of tasks (He et al. 2016; Krizhevsky, Sutskever, and Hinton 2012; Simonyan and Zisserman 2014; Szegedy et al. 2015). In simple terms, the model takes an image as its input and, after processing it through multiple layers, produces a classification label as its output (e.g., identifying an object in the image).

The architecture is inspired by the human visual cortex and is designed to automatically learn a hierarchy of features from an image. The process begins when the model receives an input image, which it sees as a grid of numbers representing pixels. The core of a CNN is its 'convolutional layers,' which act like a series of digital scanners. Each scanner uses a 'filter' to look for a specific, simple feature, such as an edge, a corner, or a patch of color. As these filters slide across the image, they create 'feature maps' that highlight where these basic patterns appear.

This process is repeated through multiple layers. The initial layers detect simple features, and subsequent layers combine these to recognize more complex patterns, for instance, combining lines and curves to identify an eye, and then combining eyes and a nose to recognize a face. Between these layers, 'pooling layers' often simplify the information by summarizing features in a region, which makes the model more efficient and robust to small variations in the image. After passing through this hierarchy, the high-level features are analyzed by a final set of 'fully connected layers' which produce the output: a probability for each possible class (e.g., 95% 'dog', 5% 'cat'). To see details of the model architecture used for the party detection task, see Appendix A.

The dataset consisted of 691 unique, deduplicated advertisements, with 552 used for training and 139 for testing. The CNN achieved an overall test accuracy of 93.5 percent (see Appendices B.1 and B.2 for detailed results, party-level metrics and training history). This performance demonstrates that Indian political parties use systematic and machine-detectable visual styles in their advertisements. For context, prior research attempting to classify the ideology of U.S. politicians using social media photographs achieved only 59.3% accuracy from single images and 82.4% when aggregating across multiple photographs of the same individual (Xi et al. 2020). My results show that party advertisements in India, contain highly regularized visual cues, making them easier to detect and classify compared to the less systematic imagery studied in prior work.

Model performance varied across parties. The model was most accurate for BJP and Congress, the two largest national parties, whose ads consistently display strong and distinctive visual identities. BJP ads frequently use saffron coloring, the lotus party symbol, and a frontal image of Narendra Modi paired with campaign text. Because of this formulaic style, the model identified most BJP ads correctly with very high confidence (see Appendix B.3). However, this homogeneity creates the possibility of information leakage, since even deduplicated ads are often visually similar. The repetition is itself informative: BJP's use of a highly standardized template builds symbolic coherence and strengthens digital brand recognition across contexts. Errors in BJP classification occurred when this template was disrupted.

Misclassified BJP ads provide insight into the limits of formulaic branding. Figure 4 shows examples of such errors. In one case, a BJP ad was predicted as TDP because its color scheme overlapped with regional motifs. In another case, an event photo featuring Modi in a crowded background was predicted as Congress, indicating that contextual cues sometimes outweighed party symbols. These cases illustrate that while BJP's formulaic style provides a strong and easily recognizable visual fingerprint, deviations from the template blur partisan distinctions.

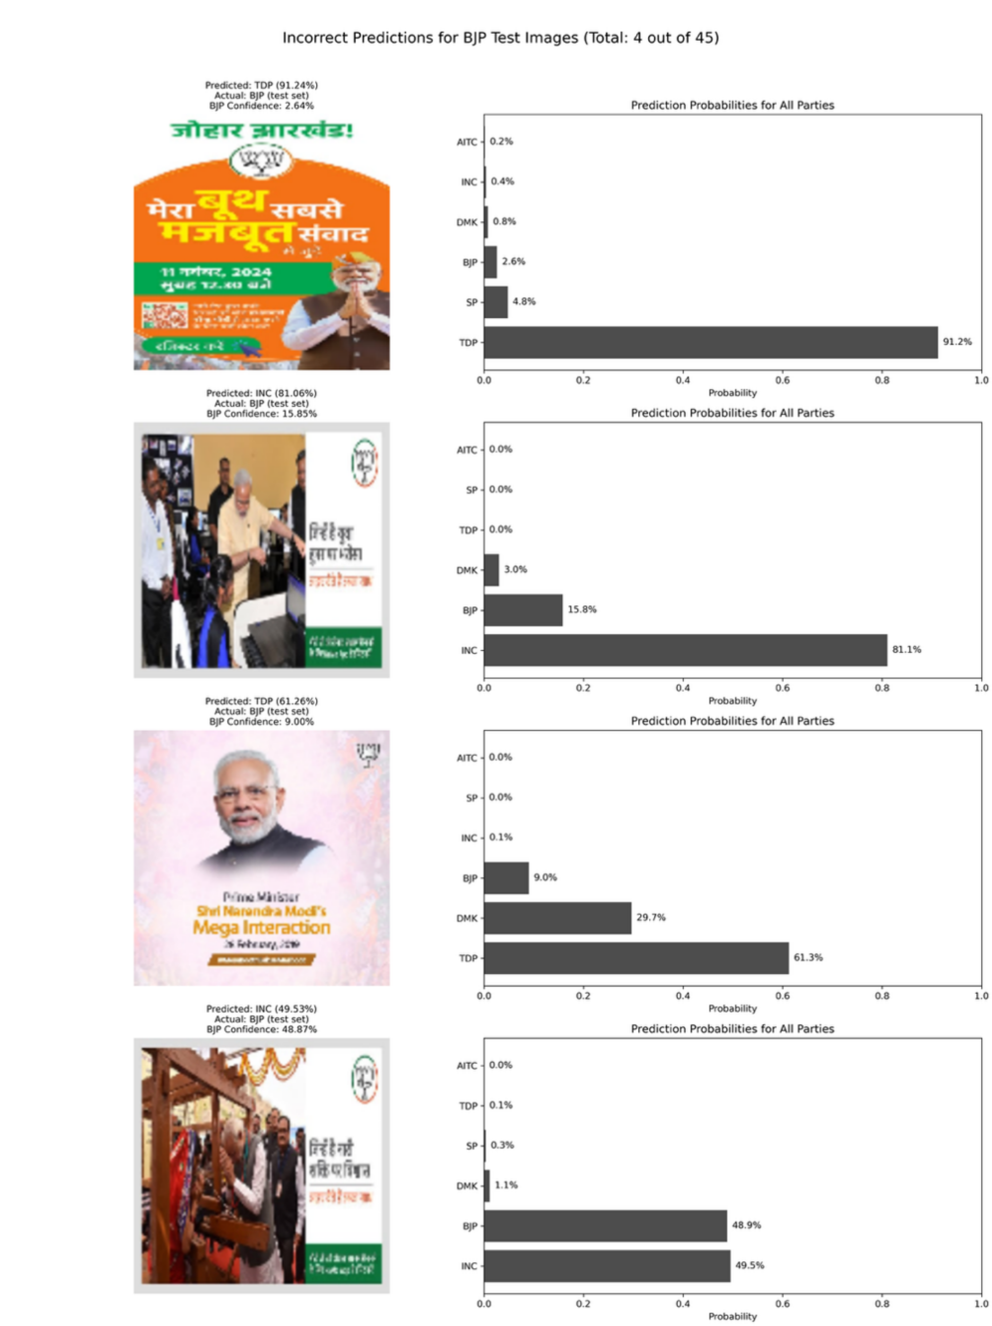

Figure 4: Misclassified BJP Advertisements

7.1.3 Deep Learning with Video ads

Video classification poses unique challenges compared to static image analysis. Unlike single images, which often contain clear visual cues such as text message, party symbols or political figures within one frame, videos consist of dynamic sequences where relevant political indicators may appear sporadically across multiple frames rather than concentrated in any single moment. Without additional context, like text, audio or the ability to view extended sequences, it is often difficult for human observers to identify party affiliation from individual frames alone. I extracted frames from each deduplicated video and employed several steps to ensure that these frames contain informative signals (Appendix C details the steps).

My approach employed a two-step methodology: first, classifying each video frame independently, and second, aggregating these frame-level predictions to classify the entire video using majority voting. The model was trained on 936 unique videos comprising 7,721 frames, and tested on 312 videos with 2,606 frames (see Appendix D for model architecture). Since the individual N at the party level is small for some parties, the output for the video model is a binary indicator of whether an ad is BJP or not. At the frame level, the model achieved an accuracy of 86.45%. When aggregating predictions across frames for video-level classification, accuracy improved significantly to 94.87% (see Appendix E for confusion matrix, training history and examples of frames with predicted probabilities). These results demonstrate that despite the temporal dispersion of visual signals, video frames alone provides strong predictive power for identifying BJP campaign material, even without contextual audio.

These findings underscore the extent to which BJP has operationalized Facebook advertising into an industrialized system, achieving unparalleled scale and cost efficiency relative to competitors. Yet the dominance reflected in counts and spending raises a deeper question about substance: does this vast activity represent genuine creative diversity, or the repeated circulation of a narrow set of messages? The sharp contraction of unique creatives once duplicates are removed suggests the latter, pointing to a campaign model built less on variety than on repetition. This observation motivates the next stage of analysis, using deep learning to assess whether these limited but highly standardized visuals carry a distinctive partisan style that sets BJP apart from other parties.

7.2 What differs visually between BJP and non-BJP ads?

7.2.1 Content of ads

To systematically compare symbolic and thematic patterns, all image and video ads were coded across a set of binary indicators. These captured whether an ad featured party symbols; rural settings (e.g., villages or farms) or urban settings (e.g., cityscapes, tall buildings); references to healthcare (e.g., hospitals or clinics), social welfare (e.g., benefit distribution), the military (e.g., soldiers or equipment), religion (e.g., temples or icons), or infrastructure development (e.g., bridges, construction sites). Additional indicators noted the display of the Indian flag, explicit mentions of political opposition, and the inclusion of men, women, or children. Finally, ads were coded for economic orientation, distinguishing between primary sector imagery (e.g., farmers in fields) and secondary sector imagery (e.g., industrial or construction workers). These categories are not mutually exclusive: a single ad could, for example, depict both rural settings and infrastructure development, or include party symbols alongside religious imagery.

For coding these features, I employed Pixtral-12B-2409 (Mistral AI), a multimodal large language model with strong OCR, translation, and visual classification capabilities. This approach enabled accurate identification of political figures, slogans, and symbols in a zero-shot setting, drawing on the model's general pretraining knowledge rather than task-specific fine-tuning, in line with recent advances in automated political media analysis (Breuer et al. 2025; Tarr, Hwang, and Imai 2023). In this context, a zero-shot setting means the model was applied directly to Indian political ads without any additional training on labeled examples from this dataset.

A two-sample test of proportions was used to statistically compare the frequency of each messaging element across party groups. The results reveal both clear differences and important areas of similarity. As shown in Table 2, BJP advertisements are significantly more likely to feature religious symbolism, party symbols, rural settings, and male figures.

Non-BJP ads, by contrast, more often highlight secondary sector themes, display the national flag, and mention opposition parties.

Table 2: Differences in Political Messaging in Image Ads: BJP vs Non-BJP Parties

Feature	BJP (Mean)	Non-BJP (Mean)	Difference in Proportions	χ^2 statistic	p-value	Significance
Panel A: Significant Differences (p < 0.05)						
Party Symbols	0.881	0.785	+0.096	9.26	0.002	**
Rural Focus	0.305	0.211	+0.095	7.42	0.006	**
Religious Symbolism	0.115	0.019	+0.096	28.96	<0.001	***
Secondary Sector	0.053	0.125	−0.072	8.57	0.003	**
Infrastructure Development	0.204	0.065	+0.139	30.03	<0.001	***
Indian Flag	0.018	0.178	−0.161	35.73	<0.001	***
Mentions Opposition	0.088	0.248	−0.160	26.96	<0.001	***
Has Man	0.854	0.699	+0.155	24.93	<0.001	***
Panel B: Non-Significant Differences (p ≥ 0.05)						
Urban Focus	0.217	0.178	+0.039	1.45	0.229	
Healthcare	0.093	0.080	+0.013	0.35	0.553	
Social Welfare	0.438	0.486	−0.048	1.40	0.236	
Defense & Military	0.044	0.026	+0.018	1.68	0.195	
Primary Sector	0.133	0.135	−0.003	0.01	0.921	
Has Woman	0.376	0.400	−0.024	0.36	0.546	
Has Children	0.053	0.075	−0.022	1.18	0.277	

Note: BJP N = 226, Non-BJP N = 465. *** p<0.001, ** p<0.01, * p<0.05. Panel A shows features with statistically significant differences. Panel B shows features with no statistically significant differences. Positive differences indicate higher proportion in BJP. Negative differences indicate higher proportion in Non-BJP parties. Difference in proportions = BJP proportion − Non-BJP proportion. Chi-square tests used for proportion comparisons.

Some of these patterns are expected and less surprising. For instance, opposition mentions are more common in non-BJP ads, which aligns with literature showing that challengers are more likely to frame campaigns against incumbents. Similarly, the prevalence of male figures in BJP ads is partly mechanical, given that Prime Minister Narendra Modi is the party's central campaign face and appears in virtually all their ads.

Other findings are more revealing. The relative absence of an urban focus in BJP ads stands out. Despite the party's documented electoral strength in cities, its visual strategy does not emphasize urban backdrops but instead leans heavily on rural representation, suggesting a deliberate effort to mobilize rural constituencies. Likewise, the use of religious symbolism and avoidance of national symbols such as the flag point to a strategy of consolidating a Hindu majoritarian identity rather than appealing to broad, secular nationalism (Charnysh, Lucas, and Singh 2015). Indeed, while 88% of BJP images analyzed contained party symbols, only 2% contained the national flag (see Appendix F for full list of feature prevalence in BJP vs non-BJP ads). These choices highlight how BJP's advertising departs from inclusive tropes and reinforces a selective ideological frame.

At the same time, several dimensions show no significant differences. Urban versus non-urban aside, healthcare, social welfare, and the inclusion of women or children appear at comparable rates across party groups. These null results are important in themselves, reflecting domains where electoral competition does not produce symbolic divergence, either because they are shared campaign tropes or because the ad format encourages convergence.

Video advertisements present a different picture from image ads (Table 3). Non-BJP parties are significantly more likely than BJP to use party symbols, rural settings, healthcare and social welfare themes, religious symbolism, primary sector imagery, and the Indian flag. They also feature women and children at higher rates. BJP video ads, by contrast, rely far less on these elements, and the only areas of convergence are urban imagery, defense and military themes, secondary sector references, and infrastructure development, where no significant differences emerge.

The contrast with image-based results is instructive. In image ads, BJP was more likely to employ party symbols, rural imagery, and religious motifs, while non-BJP parties leaned on the flag, the opposition, and the secondary sector. In video ads, however, the balance flips: non-BJP parties dominate across most symbolic categories, while BJP's distinctiveness largely disappears. This divergence suggests that the two media are not interchangeable. Still images, which are cheaper and easier to circulate, appear to be used by BJP for strong symbolic signaling, foregrounding Modi, religion, and partisan identity. Video ads, by contrast, often take the form of narrative skits or short explanatory clips, where it is harder to isolate single cues from frames and where challengers may have more to gain by layering multiple policy and representational appeals. This media-specific divergence highlights the

Table 3: Differences in Political Messaging in Video Ads: BJP vs Non-BJP Parties

Feature	BJP (Mean)	Non-BJP (Mean)	Difference in Proportions	χ^2 statistic	p-value	Significance
Panel A: Significant Differences ($p < 0.05$)						
Party Symbols	0.541	0.631	−0.091	14.10	<0.001	***
Rural Focus	0.206	0.422	−0.217	86.45	<0.001	***
Healthcare	0.070	0.114	−0.044	10.63	0.001	**
Social Welfare	0.078	0.192	−0.114	49.31	<0.001	***
Religious Symbolism	0.143	0.227	−0.084	19.79	<0.001	***
Primary Sector	0.041	0.185	−0.145	78.20	<0.001	***
Indian Flag	0.074	0.347	−0.273	189.37	<0.001	***
Has Woman	0.390	0.610	−0.220	82.44	<0.001	***
Has Man	0.771	0.816	−0.046	5.25	0.022	*
Has Children	0.078	0.109	−0.031	4.79	0.029	*
Panel B: Non-Significant Differences ($p \geq 0.05$)						
Urban Focus	0.362	0.364	−0.003	0.01	0.907	
Defense & Military	0.023	0.035	−0.013	2.43	0.119	
Secondary Sector	0.040	0.030	+0.010	1.19	0.276	
Infrastructure Development	0.112	0.099	+0.013	0.72	0.397	

Note: BJP N = 1,012, Non-BJP N = 708. *** $p<0.001$, ** $p<0.01$, * $p<0.05$. Panel A shows features with statistically significant differences ($p < 0.05$). Panel B shows features with no statistically significant differences ($p \geq 0.05$). Positive differences indicate higher proportion in BJP. Negative differences indicate higher proportion in Non-BJP parties. Difference in proportions = BJP proportion − Non-BJP proportion. Chi-square tests used for proportion comparisons.

importance of considering format in studies of political communication: the same party's visual strategy can look very different depending on whether we analyze images or videos.

Analysis of color choice (Table 4) reveals systematic differences between BJP and non-BJP advertisements. BJP ads employ orange and red at much higher rates, while non-BJP ads rely more on white and blue. These differences are consistent with both color psychology and party branding.

Table 4: Significant Differences in Color Usage: BJP vs Non-BJP Parties

Feature	BJP (Mean)	Non-BJP (Mean)	Difference in Means	t-statistic	p-value	Significance
Orange	20.157	1.313	+18.844	14.48	<0.001	***
Red	11.299	2.301	+8.998	9.29	<0.001	***
White	27.986	32.978	−4.992	2.86	0.004	**
Blue	1.176	3.900	−2.724	4.59	<0.001	***
Green	0.358	0.279	+0.079	0.87	0.384	
Yellow	0.508	0.580	−0.072	0.47	0.637	

Note: BJP N = 226, Non-BJP N = 302. *** $p<0.001$, ** $p<0.01$, * $p<0.05$. Positive differences indicate higher mean in BJP. Negative differences indicate higher mean in Non-BJP parties. Difference in means = BJP mean − Non-BJP mean. Values represent percentages of total pixels. t-tests used for mean comparisons.

Orange and saffron tones are strongly associated with Hindu identity and have long served as symbolic colors of the BJP. Their prominence in BJP ads reflects a deliberate effort to embed Hindu majoritarian symbolism into campaign visuals, reinforcing the party's ideological brand. Red, often linked to urgency, energy, and emotional intensity in psychological research, further reinforces a sense of mobilization and passion in BJP's imagery. Multiple studies have shown that red can increase physiological arousal, such as heart rate and blood pressure, and is frequently used to capture attention and convey a sense of urgency (Elliot and Maier 2014). The color is known to attract attention, particularly in emotionally charged contexts, making it a powerful tool for political messaging intended to evoke strong feelings and prompt action. Within the campaign context, these effects are not neutral: BJP's use of red and saffron could be an attempt to energize its base, signal partisan loyalty, and portray the stakes of the election as existential.

In contrast, non-BJP parties make greater use of white and blue. In many cultures, white is strongly associated with concepts like peace, purity, and simplicity, often used to create a sense of calm and cleanliness (Aslam 2006). Blue is one of the most extensively studied colors and is broadly linked to feelings of stability, competence, and trust, making it a strategic choice for parties wishing to project an image of dependable governance and peace (Labrecque and Milne 2012; Su, Cui, and Walsh 2019).

The key takeaway from this section is that BJP's advertisements consistently use saffron coloring, religious motifs, Modi's face, and the absence of the Indian flag as part of a broader symbolic strategy. These choices could be symbolic of efforts to consolidate a Hindu nationalist base by emphasizing ethno-religious identity over inclusive nationalism (Charnysh, Lucas, and Singh 2015). Non-BJP parties, in contrast, rely more on the national flag and calmer colors like white and blue to project inclusiveness and trust. These findings show that visual choices are a central tool of political communication.

7.2.2 Explainable AI Analysis

In the previous section, I manually identified common visual features in campaign advertisements and compared their distribution across parties. While this approach established clear contrasts, it was constrained by pre-defined categories. To complement this analysis, I use explainable AI techniques on the party classification CNN trained earlier. This allows systematic identification of the visual elements the model attends to when distinguishing BJP from non-BJP ads, thereby providing evidence of whether the classifier is leveraging politically salient cues rather than incidental background details.

I applied Integrated Gradients (IG), a widely used attribution method that highlights the regions of an image most influential in shaping a model's prediction. IG compares each ad frame to a fully blank baseline and then incrementally reconstructs the original image, calculating the contribution of each region to the prediction. In the resulting visualizations, green areas indicate features that increased the probability of a BJP classification, while red areas indicate features that decreased it. The analysis was conducted exclusively on the test set. Figure 5 presents randomly sampled BJP advertisements stratified by model confidence, illustrating how the classifier's attributions differ between high-, medium-, and low-confidence predictions.

Note: Green overlays indicate visual regions that increased the model's confidence in predicting the frame as BJP (pro-BJP features), while red overlays show regions that decreased it (anti-BJP features). The model's explanations are based on comparisons to a completely blank baseline image.

Figure 5: Integrated Gradients Analysis with Confidence levels

Across all panels, the classifier relies most on three families of cues: (i) explicit party branding (lotus, Modi's portrait, saffron-orange bands), (ii) stereotyped layout grammar (wide lower-third banners with dense text and icons, leader on the right, symbol near lower corners), and (iii) "development" visuals (roads, skylines, infrastructure, schools). Where these cues are clean and placed in the expected locations, IG shows strong positive attribution (green) and confidence is very high. Confidence falls when ads depart from that grammar.

The contrast between the middle-right panel (64.8% confidence) and the bottom-right panel (11.4% confidence) is especially revealing. Although both ads share the same basic design, Modi's portrait on the right, a lotus symbol and slogan banner at the bottom, the model treats them very differently. IG shows that in both frames Modi's portrait is overlaid in red, indicating it reduces BJP confidence, while the bottom banner consistently provides the strongest pro-BJP signal. The higher-confidence example highlights the banner and top-left text in green, conforming to the model's learned BJP layout, whereas in the lower-confidence example much of the text and portrait are marked red, weakening the partisan signal.

The classifier does not appear to equate the presence of Modi's portrait with a definitive BJP cue. Instead, the model has internalized a broader visual grammar characteristic of BJP advertising, in which partisan identification is driven primarily by bottom banners, lotus symbols, and bold text blocks. This result suggests that BJP's the routinization of partisan symbols and layouts into a standardized visual grammar that can be scaled and adapted at low cost is an important part of their digital strategy.

These results matter because they validate and extend the manual coding from the previous section. While earlier analysis showed BJP's reliance on symbolic branding and development motifs, the IG results demonstrate that the model has internalized a broader "visual grammar" of BJP ads, with bottom banners, lotus symbols, and bold text as the strongest cues. Crucially, Modi's portrait alone does not function as a reliable partisan marker, underscoring that it is the structured composition of these elements, rather than any single feature, that defines BJP's distinctive advertising style.

7.2.3 Sentiment of image ads

To assess the emotional tone of political advertisements, I employed two complementary lexicon-based approaches. The NRC sentiment lexicon provides a broad mapping of words to basic emotions and evaluative categories such as joy, anger, sadness, or trust. While this lexicon captures the emotional palette of language, it is not tailored to political texts and can at times misclassify policy-oriented terms. To address this limitation, I also applied the Lexicoder Sentiment Dictionary (Young and Soroka 2012) which was developed specifically for the analysis of political communication. The LSD offers a binary classification of positive versus negative sentiment and includes extensive preprocessing rules for negations and contractions, thereby aligning more closely with the kinds of language used in political discourse. Using both dictionaries provides a robustness check: NRC highlights variation across discrete emotions, while LSD benchmarks the overall valence of political communication in a way validated for this domain.

The results are shown in Figure 6. Panel A, based on the NRC lexicon, indicates that BJP advertisements rely more heavily on positive emotions, especially joy and anticipation, and also draw on fear to some degree. Non-BJP advertisements, by contrast, exhibit higher proportions of sadness, anger, and disgust, alongside somewhat greater use of trust-related language. Panel B, based on the LSD, confirms this divergence in more systematic terms. Both BJP and non-BJP campaigns use predominantly positive language, but BJP's messaging is markedly more positive (approximately 87 percent of sentiment words) and correspondingly less negative (13 percent), while non-BJP campaigns are less positive (78 percent) and more negative (22 percent). Together, these patterns suggest that BJP emphasizes optimistic and mobilizing tones, while non-BJP parties rely more on critical and grievance-oriented rhetoric.

These findings are consistent with established campaign strategy research. Challengers frequently use negative campaigning to overcome the inherent advantages of incumbents, who tend to focus on positive achievements. This aligns with the Indian case, where the

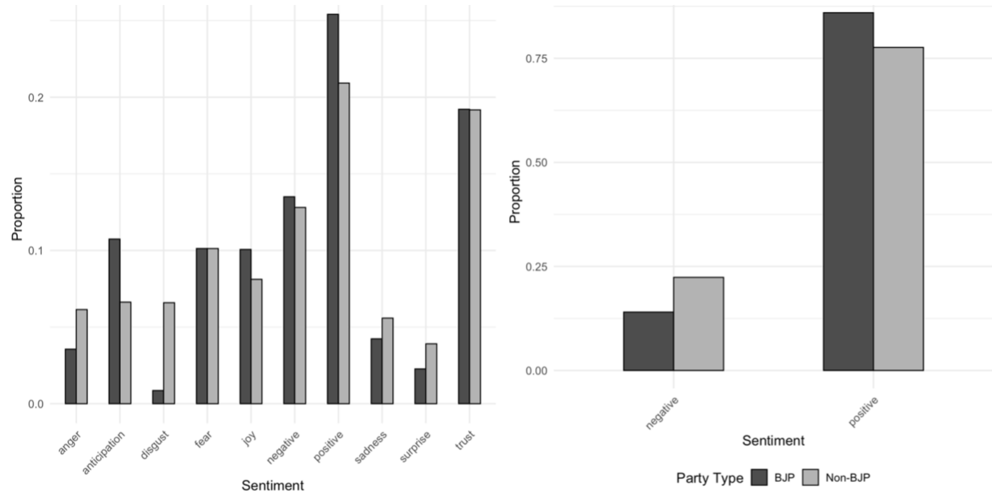

(A) Using NRC Lexicon **(B)** Using LSD

Figure 6: Sentiment Distribution by Party Type

incumbent BJP uses positive messaging while its rivals employ more negative appeals. This negativity can be a strategic device. Negative ads can be more memorable, contain more substantive policy information, and successfully mobilize dissatisfied voters (Kahn and Kenney 1999; Valentino, Hutchings, and Williams 2004). The most striking finding is that BJP campaigns overwhelmingly center on Modi and party identifiers such as BJP and the lotus, reflecting a strategy of leader-centric personalization and partisan branding, whereas non-BJP campaigns foreground institutional references like Congress alongside issue-oriented terms such as justice, women, and voice, suggesting a more programmatic and oppositional style of communication that contrasts sharply with the BJP's emphasis on charismatic leadership (see Appendix G for list of most used campaign words).

One way to interpret these patterns is through the dual lens of incumbency and right-wing communication strategy. On the one hand, the BJP's reliance on overwhelmingly positive sentiment and its avoidance of overtly negative messaging is consistent with classic findings that incumbents emphasize achievements and optimism when they are in power, while challengers resort to negativity to unsettle the status quo. On the other hand, the combination of optimism with symbolic cues and personalization reflects a broader right-wing communication style that emphasizes identity, belonging, and destiny, rather than programmatic achievements alone. In this sense, the BJP exemplifies what Maly (2019) describes as the "metapolitical" strategy of the New Right: reframing political discourse through emotionally charged but affirmative messages centered on leadership, party symbols, and national destiny. Thus, while opposition parties draw on negative tones to highlight grievances and demand accountability, the BJP uses the advantages of incumbency to project a confident, emotionally mobilizing vision that both consolidates its base and normalizes right-wing ideological frames.

7.3 What targeting strategies are used by the BJP?

Comparing the BJP with the leader of the opposition INC (Table 5), we notice that zip-code level geographical targeting appears is the dominant version of targeting done by BJP. On social media, geographic targeting represents a central feature of microtargeting, distinguishing it from traditional campaign strategies. From an advertiser's perspective, location targeting can be done in several ways on Facebook, including country (eg. India), region (eg. Punjab region), state (eg. Odisha), city (eg. Bhubaneshwar), district (eg. Rayagadha), zip code/ pin code (eg. Ambodala 765021).

For the BJP, geographic targeting makes sense from a practical standpoint, since the party has access to past election records at the polling booth level and can draw on local officials to support such efforts. By contrast, since the INC spends considerably less on digital advertising overall, one possibility is that it prefers to allocate its more limited resources to broader categories such as age or gender, aiming for reach rather than precision. The evidence does not allow us to determine whether this reflects constraints or strategic choice, but the pattern suggests that both capacity and campaign priorities shape the granularity of targeting strategies across parties.

Taking a closer look at geographical targeting, Panel A and B in Figure 7 show parliamentary constituencies and pin codes/zip codes in India shaded by count of targeting by the BJP during the 2024 Indian general elections campaign, respectively. In Panel A we can

Table 5: Targeting Criteria used by BJP and INC between April – June 2024

Targeting Criteria	**BJP (% of posts)**	**INC (% of posts)**
Zip-code level targeting	78%	9%
Location-based exclusion	1.2%	19%
Gender	2%	22%
Age	19%	42%
Interests	0%	15%

Note: Based on BJP N = 40,247 & INC N = 1,847. Targeting criteria used by the national level pages of the current ruling party in 2020-2024.

see some spatial clustering, especially in the Western parts of the country. These patterns become much more evident as we zoom into the pin code level in Panel B. We can see some clear spatial clustering here, especially in the North-West and Western parts of the country. These regions are BJP's stronghold and have played a key role in their rise to national prominence since the 1980s.

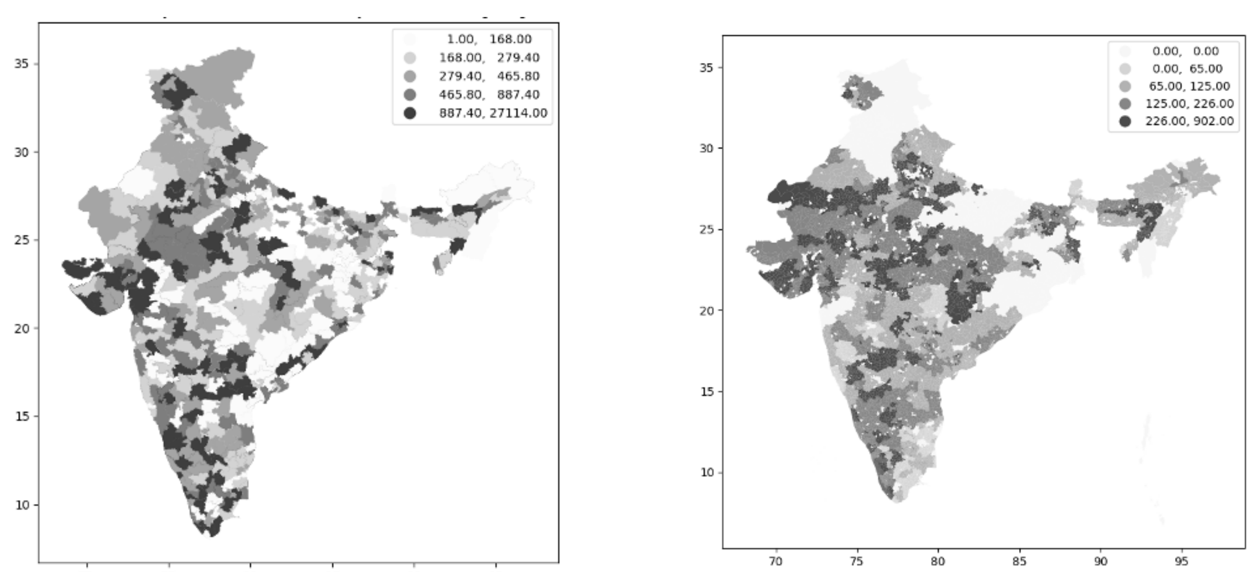

(A) **(B)**

Figure 7: Parliamentary Constituencies and Pincodes shaded by Quartiles of FB Targeting Count by BJP (April – July 2024)

An examination of co-targeting patterns between states in the six months before the 2024 election (see Appendix H) reveals three strategies. First, reinforcing strategies, which refers to concentrating resources on states where the BJP already has high support, aiming to secure turnout and prevent slippage. Second, building regional linkages, which involves coordinated targeting of neighboring or culturally connected states with stronghold states, suggesting an effort to leverage shared cultural identities, linguistic ties, or regional issues to expand BJP's campaigning. Third, independent targeting in pair of states in isolation, pointing to a localized, issue-specific approach.

These strategies align with campaign research showing that reinforcing strongholds is a classic and effective mobilization tactic, as it ensures core supporters turn out on election day, often more cost-effective than attempting to convert undecided voters (Druckman 2001; Finkel 1993; Green and Gerber 2019; Hersh 2015; Hillygus and Jackman 2003; Jacobson 2015; Kalla and Broockman 2018). By contrast, the regional linkage and independent targeting strategies represent persuasion-oriented efforts designed to extend the party's reach into new constituencies. This approach resonates with findings in the literature on right-wing populist communication. Such movements cultivate a durable bond with their base by constructing an "in-group" identity, often defined in opposition to perceived external or internal threats (Jaffrelot 2024; Mudde 2004). Yet, electoral success requires more than mobilization of the base. Populist parties supplement it by adapting their broad anti-elite message to local contexts, reframing it around regionally specific issues or grievances (Bos, Van der Brug, and De Vreese 2011). If the BJP employs this strategy, we should expect to see variation in the content of its advertisements across different settings. The next section explores this possibility by comparing rural and urban targeting, building on the previous section's finding that rural symbols appeared more prominently in BJP ads than in those of other parties.

7.4 Is there urban vs. rural visual differentiation in targeted ads by BJP?

To examine whether BJP advertisements are visually adapted to different local contexts, I focus on urban versus rural targeting. Each advertisement's geographic target was drawn from Facebook metadata and then classified as predominantly urban or rural using the Global Human Settlement Layer's Degree of Urbanization (GHS-SMOD 2023) (Schiavina, Melchiorri, and Pesaresi 2023). The analysis is restricted to video ads, which constitute the bulk of the dataset. To analyze their content, I use CLIP (Contrastive Language–Image Pretraining), a vision–language model trained on millions of image–text pairs that is well-suited for extracting semantically meaningful features from visual and textual inputs (Radford et al. 2021). CLIP enables me to capture the thematic and symbolic content of ads without training a bespoke model from scratch, providing a scalable way to compare video frames and accompanying audio across urban and rural constituencies. I combined this with transcriptions from the audio using OpenAI's Whisper model (see Appendix I for details on data preparation steps).

Using both the CLIP visual embeddings and the Whisper transcriptions, I trained a multi-modal deep learning model to predict whether an ad was targeted to an urban or rural audience. After preprocessing and deduplication, the dataset contained 385,352 frames (272,062 frames targeted at rural zipcodes and 113,290 frames targeted at urban zipcodes), with a class imbalance ratio of about 2:1. The data was split into a training set (229,970 frames, 297 unique ads), a validation set (89,223 samples, 99 unique ads), and a test set (66,159 frames, 100 unique ads), with no overlap in ads across splits.

The final model achieved an overall accuracy of 75.55% on the test set. Performance varied by class: it achieved 83.4% accuracy on rural samples, but only 46.7% on urban ones. This asymmetry likely reflects both the class imbalance favoring rural ads and the greater visual and thematic diversity of urban-targeted advertisements, which makes them harder for the model to classify consistently.

This analysis demonstrates that rural and urban targeting is not just reflected in ad placement but can be inferred from ad content itself. The model's stronger performance on rural ads suggests more consistent visual and linguistic patterns in messaging to rural audiences perhaps due to centralized party messaging or a narrower range of symbols and themes. In contrast, the lower accuracy for urban ads may reflect the greater diversity

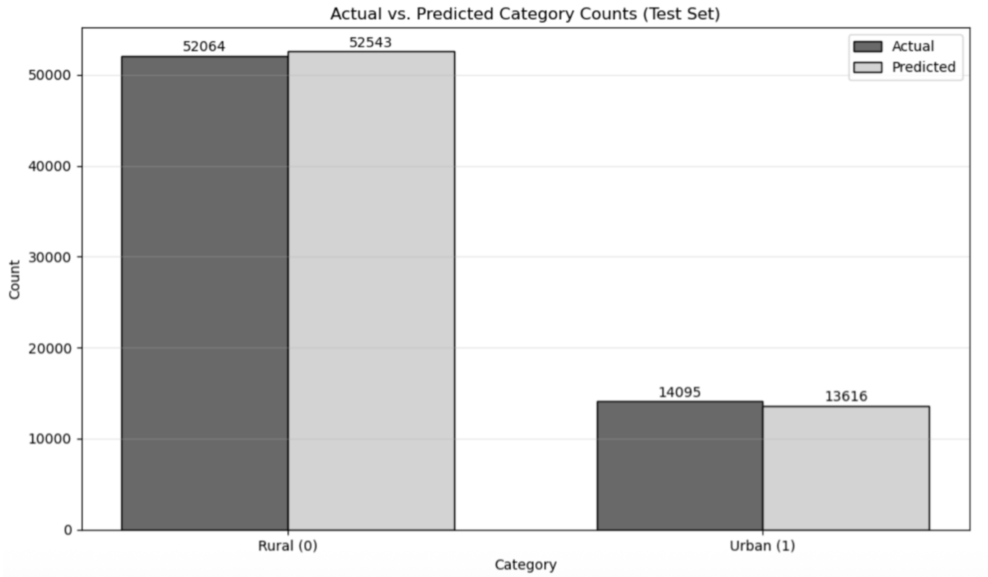

Figure 8: Actual vs. Predicted Category Counts in the Test Set

and complexity of appeals in urban constituencies, where audiences are more heterogeneous. These findings highlight the value of multi-modal machine learning approaches for uncovering subtle forms of political microtargeting in visual media.

8 Discussion and Next Steps

The results of this study suggest that right-wing parties' digital strategies are both distinctive and internally varied, with image and video ads reflecting different forms of communication. BJP's advertising strategy is built around a small set of symbolic cues that are repeated across a very large volume of ads, creating a consistent partisan and brand identity. At the same time, the party varies its use of these cues across audiences, such as with greater uniformity in rural constituencies and more visual diversity in urban ones.

These findings contribute to a broader understanding of why right-wing parties have been particularly effective in digital spaces. Structural explanations often emphasize that digital platforms reward divisive or emotive content. The evidence here suggests that BJP's digital strategy combines standardized visual templates with targeted distribution. The party relies on a consistent symbolic grammar, like saffron coloring, the lotus symbol, and Modi's portrait, that makes its ads recognizable and coherent across campaigns. At the same time, microtargeting ensures that these templates are delivered selectively across constituencies, allowing a uniform brand to be strategically deployed in different electoral contexts. While these dynamics are specific to the Indian case, the mechanisms identified here provide a framework that can be tested in other contexts to assess whether similar patterns underpin right-wing digital campaigning elsewhere.

The observational findings establish that BJP's ads are visually distinctive, but they cannot show whether distinctiveness is persuasive. The key question is whether symbolic elements that make BJP ads look different from competitors actually change attitudes and behaviors when voters encounter them.

To test this, I propose a survey experiment that employs a randomized factorial vignette design in which participants view realistic ad images generated by systematically varying visual features such as leader presence, saffron color scheme, the national flag, the party symbol, and slogan focus. This design allows for clean estimation of the causal effects of each visual element on vote intention, willingness to share, and perceived credibility.

The findings of this paper, together with the proposed experiment, suggest possible mechanisms behind right-wing digital dominance. The evidence indicates that right-wing parties may succeed online not only because of how platforms are structured but also because they seem to turn ideological symbols into repeatable and targeted campaign tools. This study shows how these strategies work in the Indian case and sets up a way to test their effects on voters. While the results are specific to this context, they provide a starting point for examining whether similar patterns shape the digital advantage of right-wing movements elsewhere.

9 Conclusion

This paper asked why right-wing parties hold a comparative advantage in digital campaigning. Existing explanations often point to algorithms that amplify divisive content or to social conditions that create demand for right-wing appeals. While important, these accounts do not explain how parties turn these conditions into concrete campaign strategies. To explore this, the paper examined the BJP, which dominates India's digital advertising landscape, and found evidence that its advantage may lie in the combination of standardized symbolic branding and more intensive microtargeting. These patterns suggest that BJP ads are both more distinctive and more precisely delivered than those of its competitors, although further work is needed to test whether similar dynamics operate beyond this case.

The analysis used large-scale Facebook advertising data together using deep learning techniques. This approach allowed me to identify systematic patterns in visual content that are not accessible through text analysis alone. The study also linked these patterns to targeting strategies, showing how symbolic design choices interacts with audience segmentation.

The findings highlight three main mechanisms. First, BJP campaigns operate at scale, producing tens of thousands of ads at lower costs than competitors. Second, the repeated templates carry consistent symbolic features. Ads prominently use saffron tones, religious motifs, and the party's lotus symbol, while the national flag is used less often. These features form a stable visual grammar that distinguishes BJP ads from those of other parties. Deep learning analyses showed that this symbolic grammar, more than leader images alone, explains the distinctiveness of BJP's ads. Third, the BJP combines repetition with adaptation. Ads are targeted down to the pin-code level, and their content varies across audiences. Rural ads are more uniform, while urban ads are more diverse. These results show that BJP's digital advantage rests on the combination of repetition and adaptation. The party projects a recognizable symbolic brand while adjusting its appeals to specific audiences. The findings suggest that right-wing digital dominance stems not only from structural features of platforms but from a strategic capacity to transform ideological symbols into repeatable and adaptable campaign infrastructures. This helps explain why right-wing parties often dominate digital campaigning: they turn ideological symbols into scalable, repeatable, and targetable assets.

This study also makes a methodological contribution. It demonstrates how deep learning models, generative AI and explainable AI can be used to study political communication at scale. These tools make it possible to move beyond text and capture the symbolic and affective dimensions of campaigns. By integrating large-scale observational data with a planned survey experiment, the project shows how inductive methodologies can be linked with causal discovery.

References

Albertson, Bethany L. 2015. "Dog-whistle politics: Multivocal communication and religious appeals." *Political Behavior* 37 (1): 3–26.

Algan, Yann, Sergei Guriev, Elias Papaioannou, and Evgenia Passari. 2017. "The European trust crisis and the rise of populism." *Brookings papers on economic activity* 2017 (2): 309–400.

Anastasopoulos, L Jason, Dhruvil Badani, Crystal Lee, Shiry Ginosar, and Jake Ryland Williams. 2017. "Political image analysis with deep neural networks." *University of Georgia, United States.*

Aslam, Mubeen M. 2006. "Are you selling the right colour? A cross-cultural review of colour as a marketing cue." *Journal of marketing communications* 12 (1): 15–30.

Autor, David H, David Dorn, and Gordon H Hanson. 2013. "The China syndrome: Local labor market effects of import competition in the United States." *American economic review* 103 (6): 2121–2168.

Barone, Guglielmo, and Helena Kreuter. 2021. "Low-wage import competition and populist backlash: The case of Italy." *European Journal of Political Economy* 67:101970.

Berelson, Bernard. 1954. *Voting: A study of opinion formation in a presidential campaign.* University of Chicago Press.

Bobbio, Norberto. 1996. *Left and right: The significance of a political distinction.* University of Chicago Press.

Bonikowski, Bart, and Yueran Zhang. 2023. "Populism as dog-whistle politics: Anti-elite discourse and sentiments toward minority groups." *Social Forces* 102 (1): 180–201.

Bos, Linda, Wouter Van der Brug, and Claes De Vreese. 2011. "How the media shape perceptions of right-wing populist leaders." *Political Communication* 28 (2): 182–206.

Bossert, Walter, Andrew Eric Clark, Conchita d'Ambrosio, and Anthony Lepinteur. 2019. "Economic insecurity and the rise of the right."

Brady, Henry E, and Richard GC Johnston. 2009. *Capturing campaign effects.* University of Michigan Press.

Breuer, Adam, Bryce J Dietrich, Michael H Crespin, Matthew Butler, JA Pryse, and Kosuke Imai. 2025. "Using AI to Summarize US Presidential Campaign TV Advertisement Videos, 1952-2012." *arXiv preprint arXiv:2503.22589.*

Bursztyn, Victor S, and Larry Birnbaum. 2019. "Thousands of small, constant rallies: A large-scale analysis of partisan WhatsApp groups." In *Proceedings of the 2019 IEEE/ACM international conference on advances in social networks analysis and mining,* 484–488.

Carnahan, Dustin, Ezgi Ulusoy, Rachel Barry, Johnny McGraw, Isabel Virtue, and Daniel E Bergan. 2022. "What should I believe? A conjoint analysis of the influence of message characteristics on belief in, perceived credibility of, and intent to share political posts." *Journal of Communication* 72 (5): 592–603.

Charnysh, Volha, Christopher Lucas, and Prerna Singh. 2015. "The ties that bind: National identity salience and pro-social behavior toward the ethnic other." *Comparative Political Studies* 48 (3): 267–300.

Chen, Wen, Diogo Pacheco, Kai-Cheng Yang, and Filippo Menczer. 2021. "Neutral bots probe political bias on social media." *Nature communications* 12 (1): 5580.

Chhibber, Pradeep, and Rahul Verma. 2014. "The BJP's 2014'Modi Wave': An Ideological Consolidation of the Right." *Economic and Political Weekly,* 50–56.

Dasgupta, Aditya, and Elena Ramirez. 2025. "Explaining rural conservatism: political consequences of technological change in the Great Plains." *American Political Science Review* 119 (1): 277–299.

Druckman, James N. 2001. "The implications of framing effects for citizen competence." *Political behavior* 23 (3): 225–256.

Eatwell, Roger, and Matthew Goodwin. 2018. *National populism: The revolt against liberal democracy.* Penguin UK.

Elliot, Andrew J, and Markus A Maier. 2014. "Color psychology: Effects of perceiving color on psychological functioning in humans." *Annual review of psychology* 65 (1): 95–120.

Engesser, Sven, Nicole Ernst, Frank Esser, and Florin Büchel. 2017. "Populism and social media: How politicians spread a fragmented ideology." *Information, communication & society* 20 (8): 1109–1126.

Fetzer, Thiemo. 2019. "Did austerity cause Brexit?" *American Economic Review* 109 (11): 3849–3886.

Finkel, Steven E. 1993. "Reexamining the" minimal effects" model in recent presidential campaigns." *The Journal of Politics* 55 (1): 1–21.

Funke, Manuel, Moritz Schularick, and Christoph Trebesch. 2016. “Going to extremes: Politics after financial crises, 1870–2014.” *European Economic Review* 88:227–260.

Garimella, Kiran, and Simon Chauchard. 2024. “WhatsApp explorer: A data donation tool to facilitate research on WhatsApp.” *Mobile Media & Communication,* 20501579251326809.

Garimella, Kiran, and Dean Eckles. 2020. “Images and misinformation in political groups: Evidence from WhatsApp in India.” *arXiv preprint arXiv:2005.09784.*

Gebru, Timnit, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei. 2017. “Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States.” *Proceedings of the National Academy of Sciences* 114 (50): 13108–13113.

Gentzkow, Matthew, and Jesse M Shapiro. 2010. “What drives media slant? Evidence from US daily newspapers.” *Econometrica* 78 (1): 35–71.

González-Bailón, Sandra, Valeria d’Andrea, Deen Freelon, and Manlio De Domenico. 2022. “The advantage of the right in social media news sharing.” *PNAS nexus* 1 (3): pgac137.

Green, Donald P, and Alan S Gerber. 2019. *Get out the vote: How to increase voter turnout.* Brookings Institution Press.

Guiso, Luigi, Helios Herrera, Massimo Morelli, and Tommaso Sonno. 2024. “Economic insecurity and the demand for populism in Europe.” *Economica* 91 (362): 588–620.

He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. “Deep residual learning for image recognition.” In *Proceedings of the IEEE conference on computer vision and pattern recognition,* 770–778.

Hersh, Eitan D. 2015. *Hacking the electorate: How campaigns perceive voters.* Cambridge University Press.

Heywood, Andrew. 2021. *Political ideologies: An introduction.* Bloomsbury Publishing.

Hillygus, D Sunshine, and Simon Jackman. 2003. "Voter decision making in election 2000: Campaign effects, partisan activation, and the Clinton legacy." *American Journal of Political Science* 47 (4): 583–596.

Huszár, Ferenc, Sofia Ira Ktena, Conor O'Brien, Luca Belli, Andrew Schlaikjer, and Moritz Hardt. 2022. "Algorithmic amplification of politics on Twitter." *Proceedings of the national academy of sciences* 119 (1): e2025334119.

"India App Market Statistics (2024)." 2024. Business of Apps. Accessed November 14, 2024. https://www.businessofapps.com/data/india-app-market/.

Inglehart, Ronald F, and Pippa Norris. 2016. "Trump, Brexit, and the rise of populism: Economic have-nots and cultural backlash."

Iyengar, Shanto, and Adam F Simon. 2000. "New perspectives and evidence on political communication and campaign effects." *Annual review of psychology* 51 (1): 149–169.

Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes. 2012. "Affect, Not Ideology: A Social Identity Perspective on Polarization." *Public Opinion Quarterly* 76 (3): 405–431. https://doi.org/10.1093/poq/nfs038.

Jacobson, Gary C. 2015. "How Do Campaigns Matter?" *Annual Review of Political Science* 18:31–47. https://doi.org/10.1146/annurev-polisci-072012-113556.

Jaffrelot, Christophe. 2015. "The Modi-centric BJP 2014 election campaign: New techniques and old tactics." *Contemporary South Asia* 23 (2): 151–166.

———. 2017. "India's democracy at 70: Toward a Hindu state?" *Journal of Democracy* 28 (3): 52–63.

———. 2024. "Hindutva, Caste, and State Vigilantism." *The Troubling State of India's Democracy,* 300–320.

Jaffrelot, Christophe, and Gilles Verniers. 2020. "The BJP's 2019 election campaign: Not business as usual." *Contemporary South Asia* 28 (2): 155–177.

Kahn, Kim Fridkin, and Patrick J Kenney. 1999. "Do negative campaigns mobilize or suppress turnout? Clarifying the relationship between negativity and participation." *American political science review* 93 (4): 877–889.

Kalla, Joshua L, and David E Broockman. 2018. "The minimal persuasive effects of campaign contact in general elections: Evidence from 49 field experiments." *American Political Science Review* 112 (1): 148–166.

Kaul, Nitasha. 2017. "Rise of the political right in India: Hindutva-development mix, Modi myth, and dualities." *Journal of Labor and Society* 20 (4): 523–548.

Khan, Sammyh S, Ted Svensson, Yashpal A Jogdand, and James H Liu. 2017. "Lessons from the past for the future: The definition and mobilisation of Hindu nationhood by the Hindu nationalist movement of India." *Journal of Social and Political Psychology* 5 (2): 477–511.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E Hinton. 2012. "Imagenet classification with deep convolutional neural networks." *Advances in neural information processing systems* 25.

Labrecque, Lauren I, and George R Milne. 2012. "Exciting red and competent blue: the importance of color in marketing." *Journal of the Academy of Marketing Science* 40 (5): 711–727.

Leidig, Eviane. 2020. "Hindutva as a variant of right-wing extremism." *Patterns of Prejudice* 54 (3): 215–237.

Mahapatra, Sangeeta, and Johannes Plagemann. 2019. "Polarisation and politicisation: The social media strategies of Indian political parties."

Maly, Ico. 2019. "New right metapolitics and the algorithmic activism of Schild & Vrienden." *Social Media+ Society* 5 (2): 2056305119856700.

March, Luke. 2017. "Left and right populism compared: The British case." *The British Journal of Politics and International Relations* 19 (2): 282–303.

Marwick, Alice, and Rebecca Lewis. 2017. "Media manipulation and disinformation online." *New York: Data & Society Research Institute* 359:1146–1151.

McDonnell, Duncan, and Stefano Ondelli. 2025. "The distinctive vocabularies of right-wing populists." *Government and Opposition* 60 (2): 335–357.

Mohammad, Saif M, and Peter D Turney. 2013. "Nrc emotion lexicon." *National Research Council, Canada* 2:234.

Mudde, Cas. 2004. "The populist zeitgeist." *Government and opposition* 39 (4): 541–563.

———. 2019. *The far right today.* John Wiley & Sons.

Mutz, Diana C. 2007. "Effects of "in-your-face" television discourse on perceptions of a legitimate opposition." *American Political Science Review* 101 (4): 621–635.

Nikolov, Dimitar, Alessandro Flammini, and Filippo Menczer. 2020. "Right and left, partisanship predicts (asymmetric) vulnerability to misinformation." *arXiv preprint arXiv:2010.01462.*

Norris, Pippa, and Ronald Inglehart. 2019. *Cultural backlash: Trump, Brexit, and authoritarian populism.* Cambridge University Press.

Radford, Alec, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. "Learning transferable visual models from natural language supervision." In *International conference on machine learning,* 8748–8763. PmLR.

Radford, Alec, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. "Robust speech recognition via large-scale weak supervision." In *International conference on machine learning,* 28492–28518. PMLR.

Rai, Swapnil. 2019. ""May the force be with you": Narendra Modi and the celebritization of Indian politics." *Communication, Culture & Critique* 12 (3): 323–339.

Rathje, Steve, Jay J Van Bavel, and Sander Van Der Linden. 2021. "Out-group animosity drives engagement on social media." *Proceedings of the national academy of sciences* 118 (26): e2024292118.

Rizzo, Tesalia, and Adi Dasgupta. 2024. "Does the Built Environment Shape Voter Participation? Learning from Polling Place Imagery in Mexico." Conditionally accepted manuscript at APSR.

Rodrik, Dani. 2017. *Populism and the economics of globalization.* Technical report. National Bureau of Economic Research.

———. 2021. "Why does globalization fuel populism? Economics, culture, and the rise of right-wing populism." *Annual review of economics* 13 (1): 133–170.

Rothwell, Jonathan T, and Pablo Diego-Rosell. 2016. "Explaining nationalist political views: The case of Donald Trump." *Available at SSRN 2822059.*

Sarkar, Tanika. 2021. "The Dominant Face of Religious Nationalism in India." *When politics are sacralized: Comparative perspectives on religious claims and nationalism,* 161.

Schiavina, Marcello, Michele Melchiorri, and Martino Pesaresi. 2023. *GHS-SMOD R2023A - GHS Settlement Layers, Application of the Degree of Urbanisation Methodology (Stage I) to GHS-POP R2023A and GHS-BUILT-S R2023A, Multitemporal (1975-2030).* Dataset. https://doi.org/10.2905/A0DF7A6F-49DE-46EA-9BDE-563437A6E2BA.

Schmid, Ursula Kristin, Heidi Schulze, and Antonia Drexel. 2025. "Memes, humor, and the far right's strategic mainstreaming." *Information, Communication & Society* 28 (4): 537–556.

Sheikh, Shanana. 2024. "How Technology Is (and Isn't) Transforming Election Campaigns in India." Carnegie Endowment for International Peace, March. Accessed November 14, 2024. https://carnegieendowment.org/research/2024/03/how-technology-is-and-isnt-transforming-election-campaigns-in-india?lang=en.

Simonyan, Karen, and Andrew Zisserman. 2014. "Very deep convolutional networks for large-scale image recognition." *arXiv preprint arXiv:1409.1556.*

Sircar, Neelanjan. 2020. "The politics of vishwas: Political mobilization in the 2019 national election." *Contemporary South Asia* 28 (2): 178–194.

Stanley, Ben. 2008. "The thin ideology of populism." *Journal of political ideologies* 13 (1): 95–110.

Statista. 2024. "India: Leading Social Media Sites 2024." Accessed November 14, 2024. https://www.statista.com/statistics/1115648/india-leading-social-media-sites/.

Su, Lixun, Annie Peng Cui, and Michael F Walsh. 2019. "Trustworthy blue or untrustworthy red: The influence of colors on trust." *Journal of Marketing Theory and Practice* 27 (3): 269–281.

Szegedy, Christian, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. "Going deeper with convolutions." In *Proceedings of the IEEE conference on computer vision and pattern recognition,* 1–9.

Tarr, Alexander, June Hwang, and Kosuke Imai. 2023. "Automated coding of political campaign advertisement videos: An empirical validation study." *Political Analysis* 31 (4): 554–574.

Törnberg, Anton, and Petter Törnberg. 2025. "White supremacists anonymous: how digital media emotionally energize far-right movements." *Journal of Information Technology & Politics* 22 (1): 131–148.

Valentino, Nicholas A, Vincent L Hutchings, and Dmitri Williams. 2004. "The impact of political advertising on knowledge, Internet information seeking, and candidate preference." *Journal of communication* 54 (2): 337–354.

Vosoughi, Soroush, Deb Roy, and Sinan Aral. 2018. "The spread of true and false news online." *science* 359 (6380): 1146–1151.

Votta, Fabio, Simon Kruschinski, Mads Hove, Natali Helberger, Tom Dobber, and Claes de Vreese. 2024. "Who Does (n't) Target You? Mapping the Worldwide Usage of Online Political Microtargeting." *Journal of Quantitative Description: Digital Media* 4.

Won, Donghyeon, Zachary C Steinert-Threlkeld, and Jungseock Joo. 2017. "Protest activity detection and perceived violence estimation from social media images." In *Proceedings of the 25th ACM international conference on Multimedia,* 786–794.

Xi, Nan, Di Ma, Marcus Liou, Zachary C Steinert-Threlkeld, Jason Anastasopoulos, and Jungseock Joo. 2020. "Understanding the political ideology of legislators from social media images." In *Proceedings of the international AAAI conference on web and social media,* 14:726–737.

Young, Lori, and Stuart Soroka. 2012. "Affective news: The automated coding of sentiment in political texts." *Political communication* 29 (2): 205–231.

APPENDIX

A CNN Architecture for Political Advertisement Classification: Images

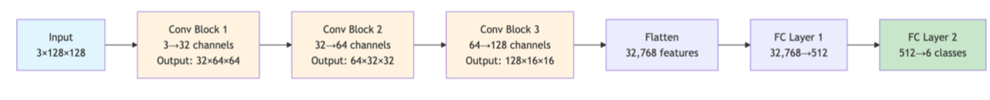

Figure A.1: Appendix Figure 1: CNN Architecture for Image Classification

The CNN architecture employed for predicting party labels from image ads follows a hierarchical feature extraction approach designed to classify political party advertisements across six Indian political parties (BJP, Congress, AITC, TDP, SP, and DMK). The network processes RGB input images normalized and resized to 128×128 pixels.

The architecture comprises three convolutional blocks that implement a systematic feature learning hierarchy. The first convolutional block applies two sequential 3×3 convolutions to extract low-level visual features such as edges, textures, and basic geometric patterns from the input images. This block increases the feature representation from the original 3 input channels to 32 feature maps, followed by batch normalization and ReLU activation functions to stabilize training dynamics and introduce non-linearity. Spatial dimensionality is reduced through 2×2 max pooling operations that downsample the feature maps from 128×128 to 64×64 pixels, while 25% dropout regularization prevents overfitting by randomly zeroing selected features during training.

The second and third convolutional blocks follow an identical architectural pattern but progressively double the number of feature channels from 32 to 64 and subsequently to 128, enabling the network to learn increasingly complex and abstract visual representations. Each block applies max pooling to further reduce spatial dimensions to 32×32 and finally 16×16 pixels, respectively. This progressive dimensionality reduction coupled with increased feature depth allows the network to capture mid-level features such as color patterns, shapes, and party-specific visual elements in the second block, while the third block learns high-level semantic features that distinguish between different political party visual identities.

The convolutional feature extraction stage concludes with a flattening operation that transforms the three-dimensional feature tensor (128×16×16) into a one-dimensional vector of 32,768 features. This representation is subsequently processed through two fully connected layers that perform the final classification task. The first fully connected layer compresses the high-dimensional feature space from 32,768 to 512 neurons, incorporating batch normalization, ReLU activation, and 50% dropout regularization to maintain training stability and prevent overfitting. The final fully connected layer maps these 512 compressed features to the six output classes, producing probability distributions across all political parties through a softmax activation function, with the highest probability indicating the predicted party affiliation of the input advertisement image.

B CNN Results: Images

B.1 Party-wise precision, recall and f-1 score

Table B.1: CNN Performance Metrics by Party (Image Classification)

Party	**Precision**	**Recall**	**F1-Score**	**Support**
BJP	1.00	0.91	0.95	45
CONGRESS	0.95	0.98	0.97	61
AITC	1.00	1.00	1.00	2
TDP	0.50	0.40	0.44	5
SP	1.00	1.00	1.00	3
DMK	0.85	0.96	0.90	23
Accuracy	0.94	0.94	0.94	139
Macro avg	0.88	0.88	0.88	139
Weighted avg	0.94	0.94	0.93	139

B.2 Training History

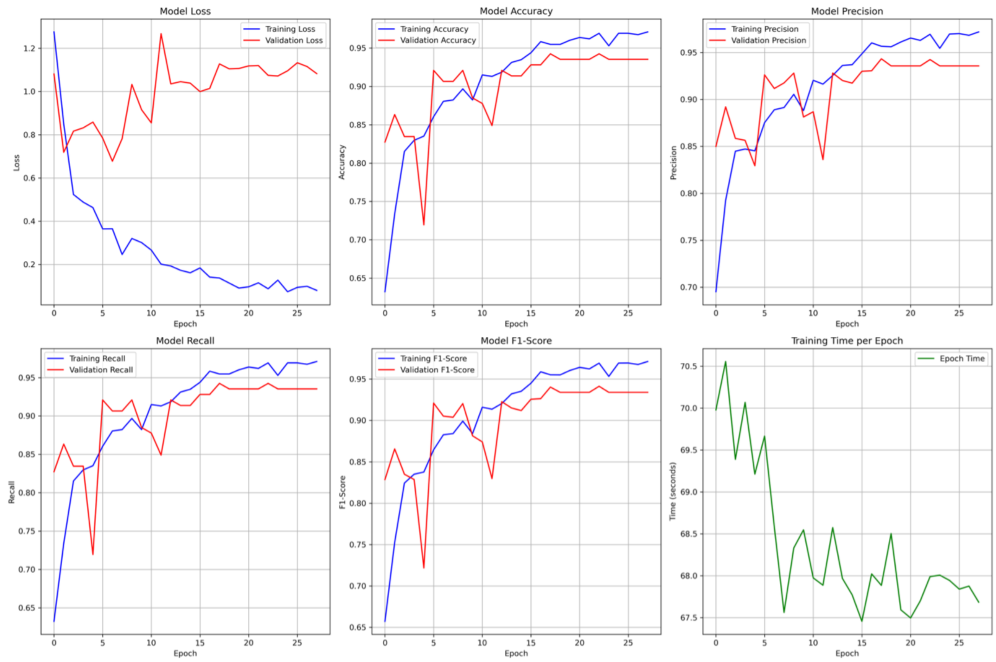

Figure B.1: Training History for Image Classification CNN

B.3 Randomly sampled BJP images from the test set with predicted probabilities

Figure B.2: BJP Test Images with Predicted Probabilities

C Video Preprocessing Pipeline

Several preprocessing steps were implemented on the political advertisement videos before analysis. First, to address the risk of data leakage, where identical ads may appear in both the training and validation sets, I deduplicated the videos by analyzing both visual and audio elements of each video. Once the duplicates were identified, the system created a new collection of unique videos, ensuring that each video in the dataset represented distinct content. This step was critical in eliminating redundancy and preventing data leakage during model training.

After this, I extracted frames from videos at regular one-second intervals. The problem with video frames is that it often contains transitional frames where the frame is shaky, all black or incoherent. Using all of these frames increases the noise to signal ratio. To overcome this, filtration mechanisms were essential to ensure the frames used for analysis were stable and free from transitional artifacts. Each set of resulting frames from a video underwent a quality filtration process, where brightness, contrast, and sharpness were evaluated. Specifically, brightness was checked for a minimum mean pixel value of 30, contrast was assessed with a standard deviation threshold of 20, and sharpness was measured using Laplacian variance, with a threshold set to 50 to discard frames with insufficient sharpness. Additionally, artifacts such as horizontal streaking, split-frame transitions, and black regions were detected and removed. Horizontal streaking was identified using Sobel operator ratios, which detect edge gradients in the horizontal direction. Split-frame transitions were flagged by comparing the brightness of the two halves of the frame, and black regions were detected if a frame contained more than 30% black pixels. Manual inspection revealed that this process improved noise, but some noise remained. For audio analysis, I used OpenAI's Whisper model which excels at multi-language speech recognition. I used it to extract audio, detect the language and content, and translate to English.

A critical methodological consideration in video frame analysis involves preventing data leakage that could artificially inflate model performance. Unlike static image datasets where individual images are independent, video frames exhibit high temporal correlation and visual similarity within the same advertisement. To address this challenge, the dataset splitting was implemented at the video level rather than the frame level.

Specifically, the train-validation-test split was performed by assigning entire video advertisements exclusively to one partition before frame extraction. This ensures that all frames extracted from a single political advertisement remain within the same data split, preventing the model from encountering visually similar frames during training that would later appear in validation or test sets. The video-level splitting process involved stratified sampling based on party labels to maintain balanced representation across splits while preserving the independence requirement for robust evaluation.

D CNN Architecture for Political Advertisement Classification: Videos

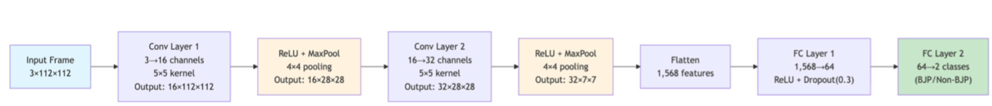

Figure D.1: CNN Architecture for Video Frame Classification

The CNN model processes RGB input frames normalized and resized to 112×112 pixels, a more compact resolution than the 128×128 pixels used for static images. The feature extraction component employs only two convolutional layers, representing a significantly streamlined approach compared to the three-block architecture used for image classification. The first convolutional layer applies sixteen 5×5 kernels to the three input channels, generating 16 feature maps that capture low-level visual patterns such as edges, textures, and color gradients. The adoption of 5×5 kernels, rather than the 3×3 kernels used in the image CNN, enables broader spatial context capture in each convolution operation, which proves particularly beneficial for detecting larger visual elements like party logos, text layouts, and color schemes that may span multiple pixels in lower-resolution video frames.

Following the first convolution, a 4×4 max pooling operation aggressively reduces spatial dimensions from 112×112 to 28×28 pixels, representing a 16-fold reduction in feature map size. The second convolutional layer doubles the feature representation from 16 to 32 channels through another set of 5×5 convolutions, enabling learning of more complex visual patterns specific to political party identification. These 32 feature maps capture mid-level features such as political symbols, candidate imagery, and text patterns crucial for accurate binary classification. A second 4×4 max pooling operation further reduces spatial dimensions to 7×7 pixels, yielding a compact 32×7×7 feature representation totaling 1,568 features—significantly fewer than the 32,768 features produced by the image classification architecture.

The classification component begins with flattening the three-dimensional tensor into a 1,568-dimensional vector. The first fully connected layer compresses this representation to merely 64 neurons, incorporating ReLU activation and 30% dropout regularization. The final layer maps these 64 features to two output classes, producing BJP versus non-BJP probability distributions through softmax activation.

The video frame classification methodology employs a dual-level evaluation framework that generates both frame-level predictions for individual video frames and video-level predictions through majority voting aggregation across all frames within each advertisement. Frame-level accuracy measures the model's ability to correctly classify individual frames based solely on single-frame visual content, while video-level accuracy reflects performance when leveraging accumulated evidence across multiple frames from the same political advertisement, typically achieving higher accuracy due to the error-correcting effect of majority voting where individual misclassified frames are outnumbered by correctly classified frames within the same video sequence.

E CNN Results: Videos

E.1 Confusion Matrix

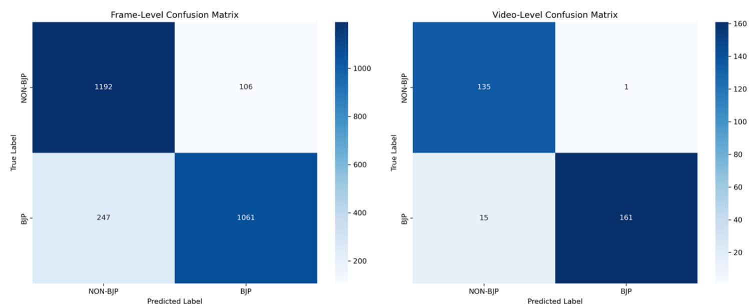

Figure E.1: Confusion Matrix for Video Frame Classification

E.2 Training History

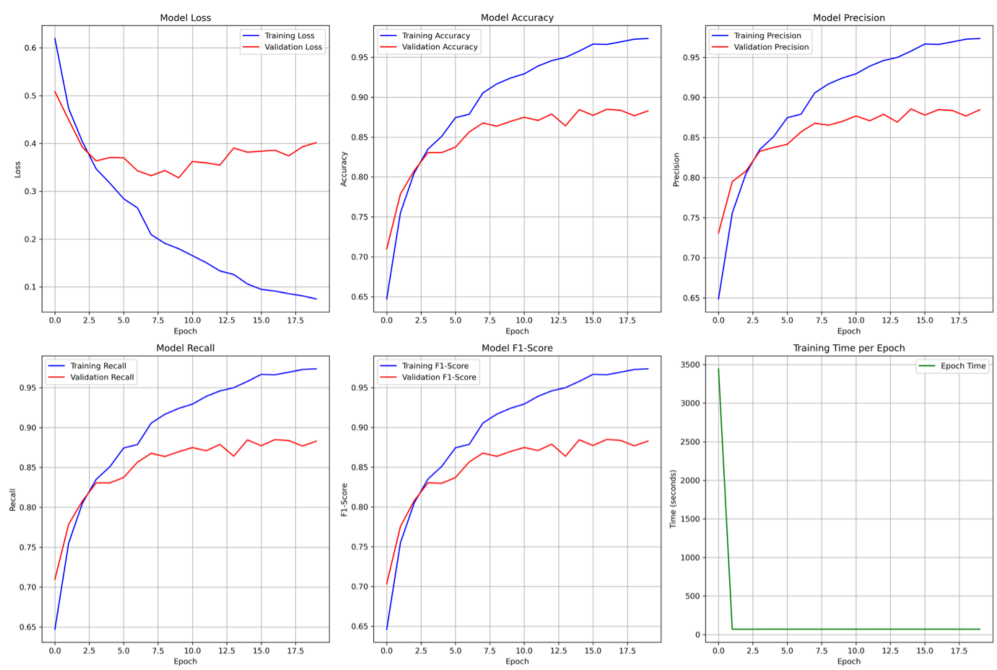

Figure E.2: Training History for Video Classification CNN

E.3 Randomly sampled frames from the test set with predicted probabilities

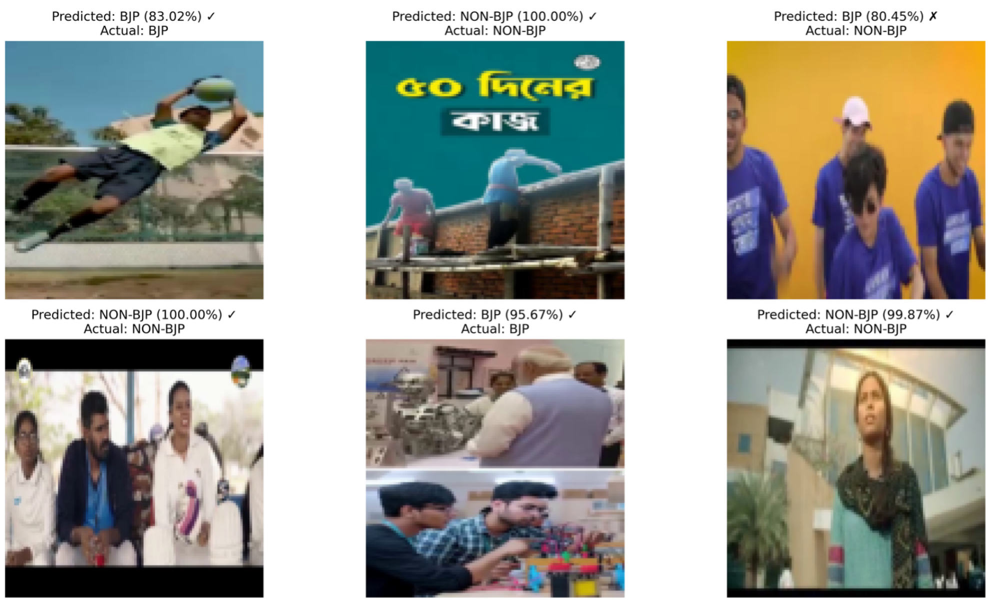

Figure E.3: Sample Video Frames with Predicted Probabilities

F Feature Prevalence between BJP and non-BJP ads

F.1 In Image Ads

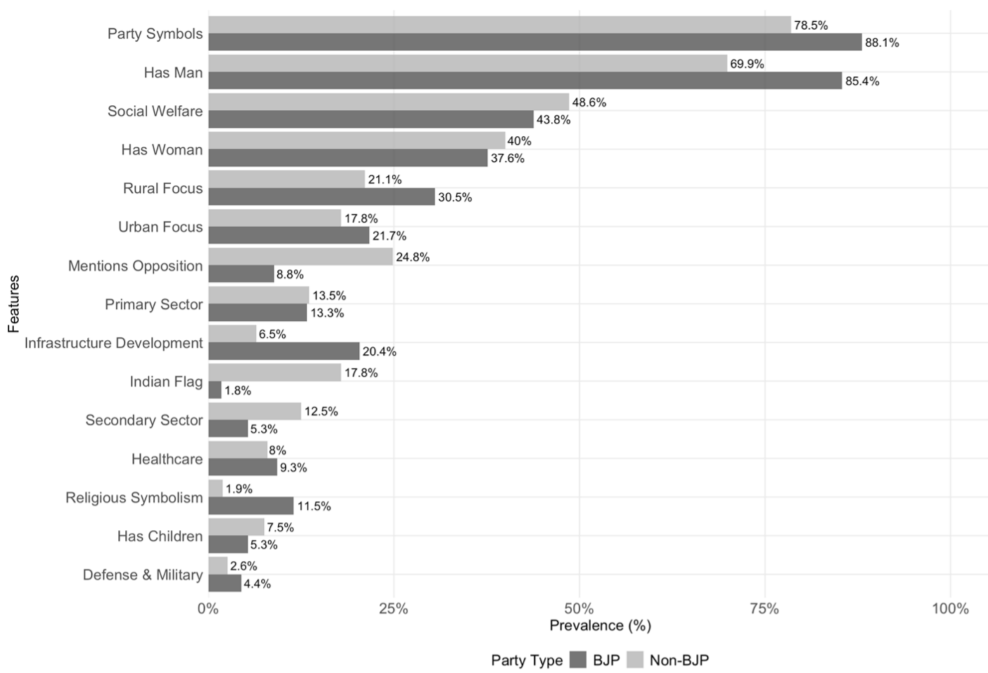

Figure F.1: Feature Prevalence in Image Ads by Party Type

F.2 In Video Ads

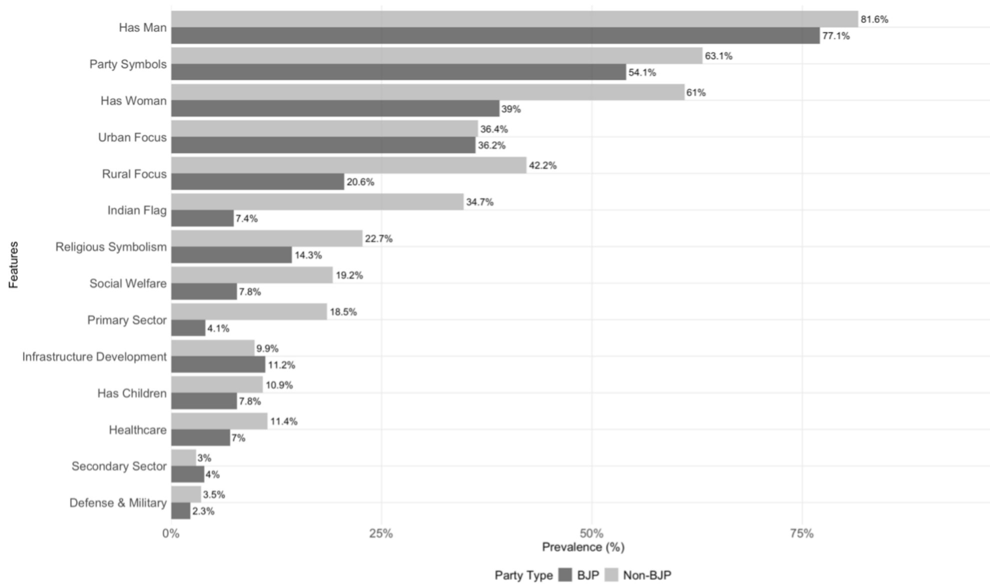

Figure F.2: Feature Prevalence in Video Ads by Party Type

G Most used words in campaign images

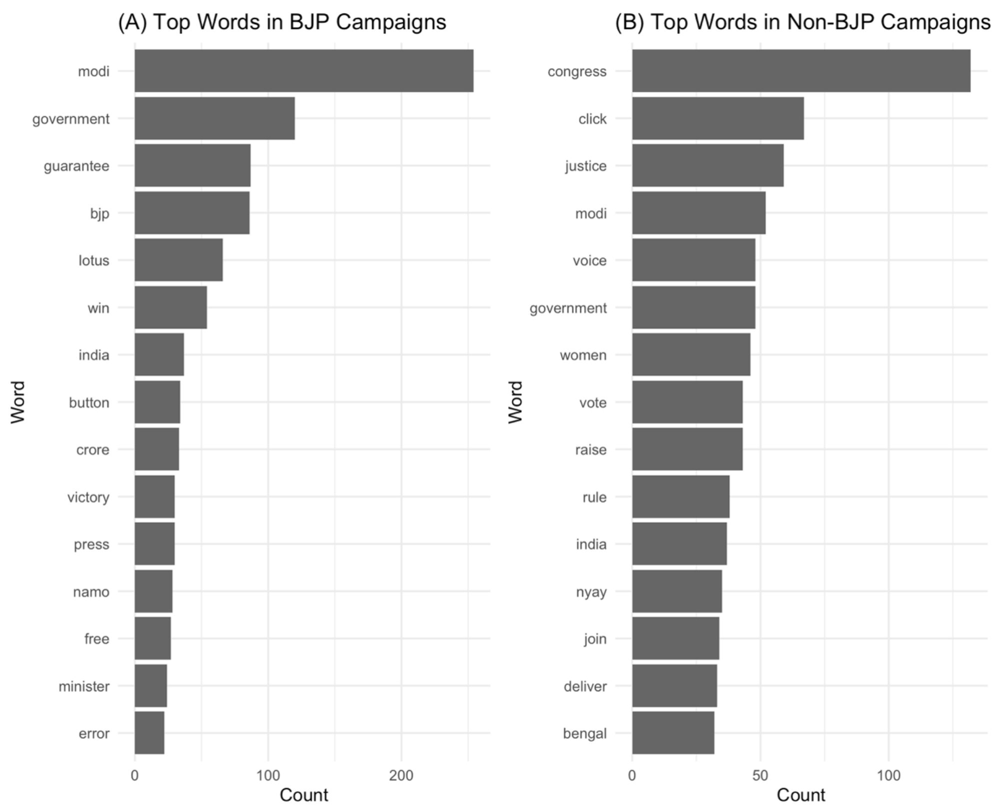

Figure G.1: Most Frequently Used Words in Campaign Images

H State-level Political Ad co-occurrence Network

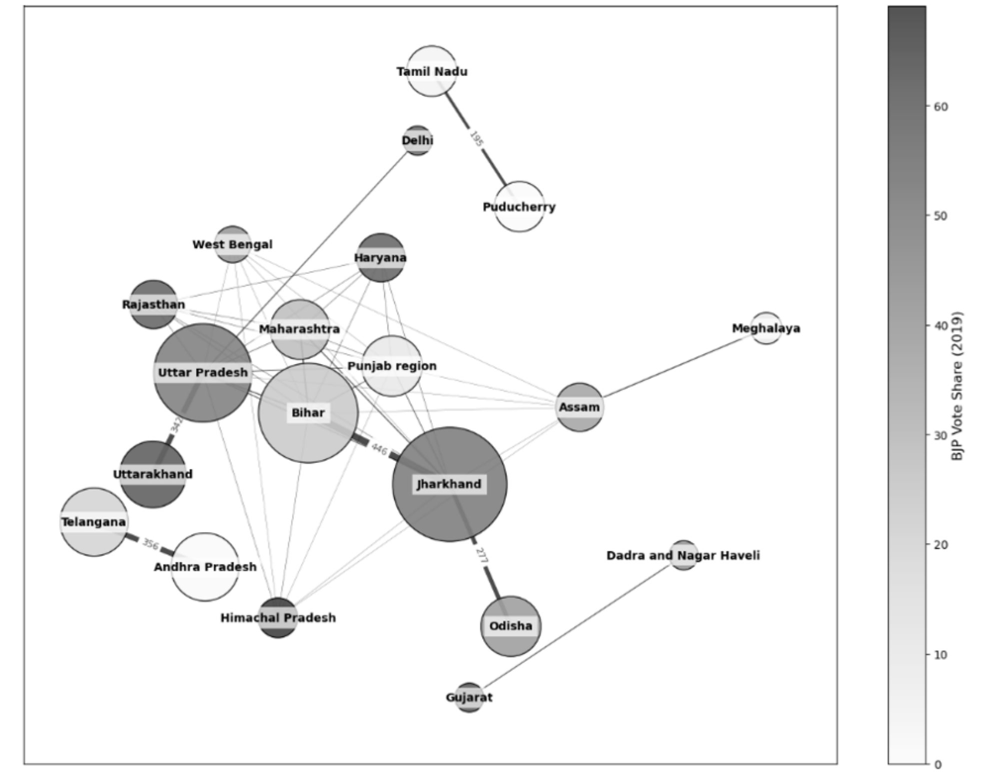

Note: Node size represents targeting frequency, connection thickness between two nodes indicates relationship strength, and node shading reflecting BJP's 2019 vote share.

Figure H.1: State-level Political Ad Co-occurrence Network (Shaded by BJP Vote Share in 2019 general elections)

I Data Preparation to test if rural or urban can be predicted using advertisements alone

I restrict the analysis to video ads only due to the limited number of unique image advertisements available after deduplication. Several preprocessing steps were necessary before analyzing the video advertisements. To address potential data leakage where identical ads might appear in both training and validation sets I deduplicated the dataset by analyzing both visual and audio components of each video, resulting in a set of approximately 900 unique ads.

For audio, I used OpenAI's Whisper model to extract transcriptions, detect language, and generate English translations. I extracted one frame per second from the videos and then used some filtration mechanisms (using mean contrast, brightness) to remove noisy frames. Each cleaned video frame was then passed through the CLIP model (Contrastive Language–Image Pre-training), a powerful vision-language network trained on millions of image–text pairs (Radford et al. 2021). Given the relatively small number of unique political advertisements after deduplication, CLIP's use of pre-trained image encoders enables us to extract semantically rich features without requiring a large training dataset. I extracted 512-dimensional embedding vectors per frame using CLIP's image encoder, producing high-dimensional numerical representations of visual content suitable for downstream classification tasks. These embedding features, combined with the transcribed audio text, form a multi-modal dataset at the frame-geography level.

J Assigning land categories to targeted constituencies

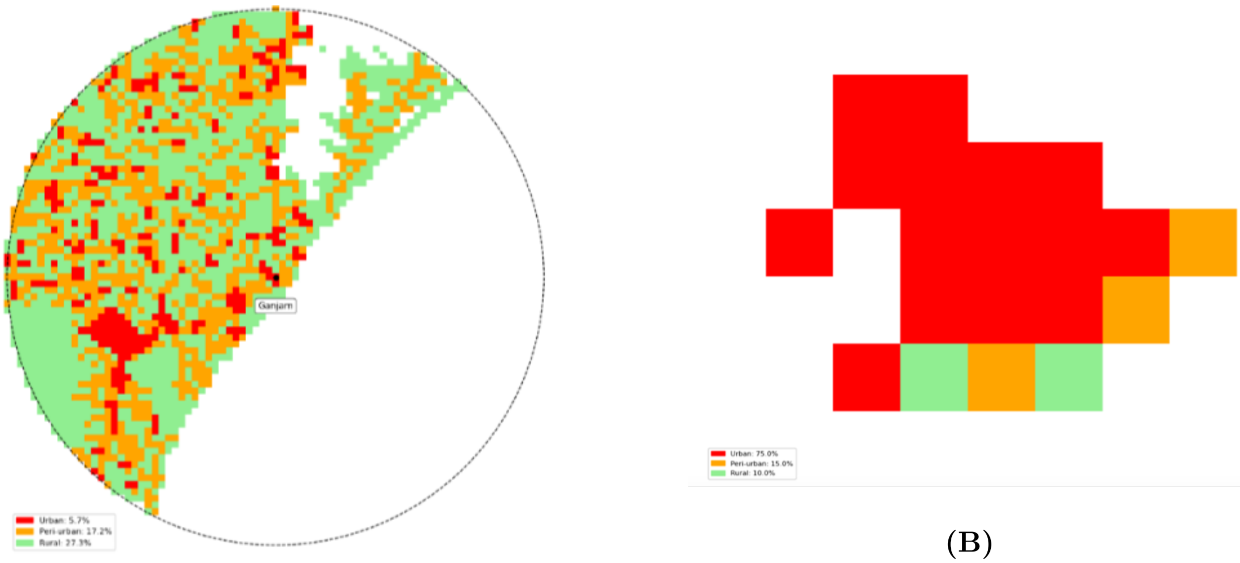

(B)

(A)

Note: Each grid corresponds to 1 sq km spatial resolution. Red pixels correspond to Urban, green correspond to rural, and orange correspond to peri-urban. Colors pixels on Panel B represent the boundaries of a pincode.

Figure J.1: Land categories in (A) a 40 km radius around Ganjam in Odisha and (B) pincode 560066 in Bangalore

K Multi-modal Model Architecture for Rural vs Urban Prediction

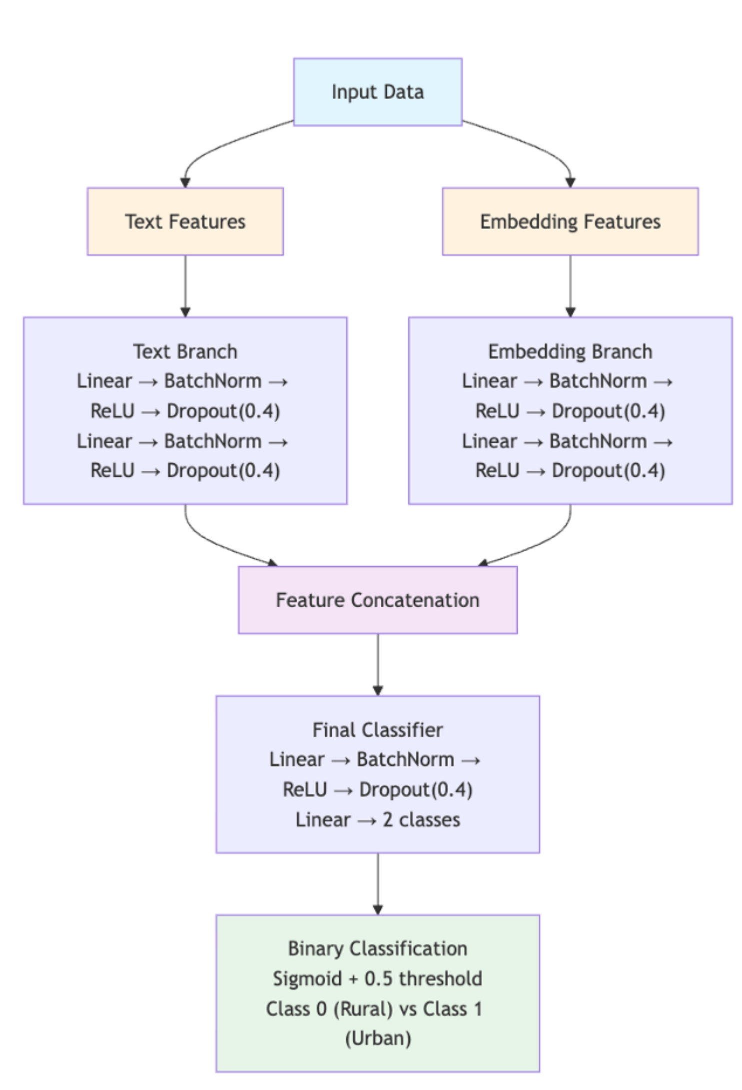

Figure K.1: Multi-modal Model Architecture for Rural vs Urban Prediction

L Multi-modal Model Results

L.1 Confusion Matrix (Test Set)

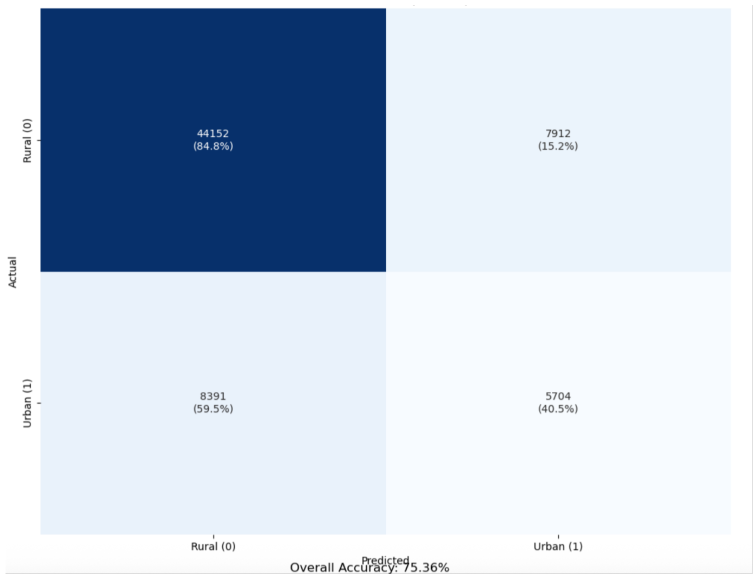

Figure L.1: Confusion Matrix for Rural vs Urban Classification (Test Set)

L.2 Probability Distribution by Actual Class (Test Set)

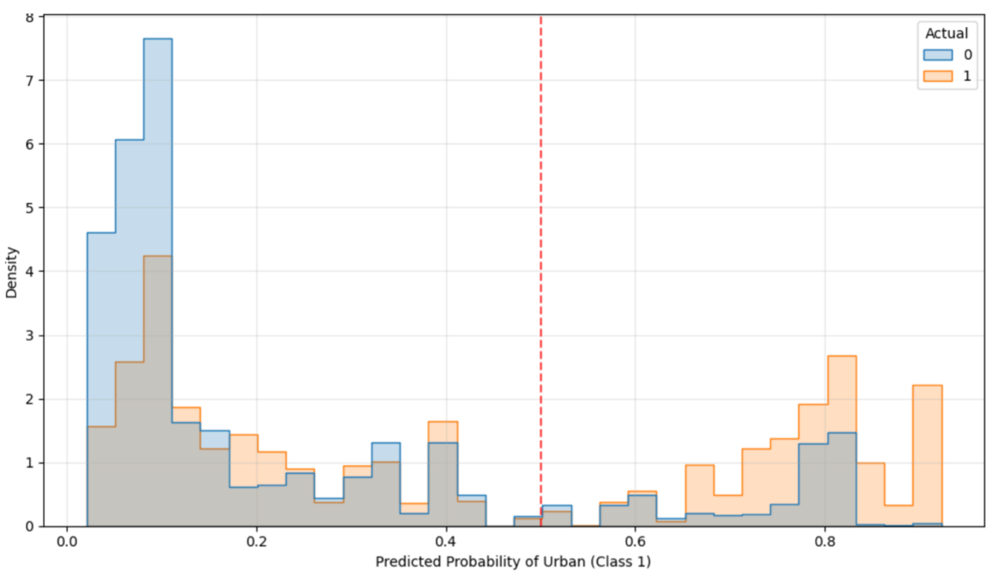

Figure L.2: Probability Distribution by Actual Class (Test Set)