# Characterization of User Engagement in Electronic News Media: A Case Study for India

**Abhishek Kumar, Abhik Jana, Debi Prosad Dogra**
School of Electrical and Computer Sciences, IIT Bhubaneswar
Argul, Khurda, Odisha, PIN 752050, India
{21ee01051, abhikjana, dpdogra}@iitbbs.ac.in

**Abstract**

Online news platforms have become central spaces for political discourse, playing a critical role in shaping public opinion and democratic participation. In large-scale democracies such as India, the combination of extensive user engagement, ideological polarization, and automated participation raises significant concerns regarding trust, transparency, and manipulation. This paper presents a large-scale empirical study of political discourse on a prominent Indian news platform, analyzing over 21,000 news articles and more than 1.5 million user comments. We investigate how ideological bias, sentiment dynamics, and non-organic user behavior interact to shape engagement patterns. Our methodology integrates hybrid article bias classification, large-scale sentiment analysis, heuristic-based bot detection, coordinated behavior analysis, and a focused examination of super-active users. In addition, we compare rule-based stance inference with large language model (LLM)-based stance classification to assess trade-offs between computational efficiency and contextual accuracy. The results reveal systematic sentiment skew, disproportionate influence by super-active and bot-like users, and coordinated campaigns aligned with specific political narratives. We conclude by discussing the implications of these findings for trust in online political discourse and reflecting on the dual role of generative AI as both an analytical tool and a potential vector for manipulation.

**Keywords:** Political NLP, Bias Analysis, Bot Detection, Political Discourse

## 1. Introduction

India's rapid growth in online news consumption has positioned digital media platforms at the centre of socio-political discourse. Automated analysis of political news and user engagement has therefore become an important research direction, enabling structured interpretation of large-scale political communication and narrative formation (Cartier and Tanev, 2022; Adhya and Sanyal, 2022). Computational approaches have been increasingly used to study political discourse across parliamentary proceedings, news reporting, and online platforms, underscoring their potential to elucidate ideological patterns and public engagement at scale.

News organizations significantly influence public opinion by shaping how political issues are framed and discussed (Gentzkow and Shapiro, 2010; D'Alessio and Allen, 2000). Prior research has shown that newspaper coverage often aligns more closely with real-world political developments than social media discussions, suggesting that institutional media remains a powerful driver of political awareness and civic interpretation (Sanders and van den Bosch, 2021). However, the growing digitization of news consumption has also introduced new challenges, including misinformation, ideological framing, and coordinated manipulation of public discourse (Lazer et al., 2018; Ferrara et al., 2016).

The *Times of India (TOI)* is one of the most influential digital news platforms in India, offering extensive political reporting and large-scale public engagement through its comment sections (Athique, 2012). These discussion spaces serve as critical arenas for civic expression but also remain vulnerable to manipulation through coordinated activity, automated participation, and ideologically driven amplification. Understanding how article tone, user sentiment, and engagement behavior interact within such environments is therefore essential for assessing the health of online democratic discourse.

Beyond overt expressions of opinion, media bias often emerges through subtle framing choices that prioritize particular narratives, actors, or perspectives. Such framing influences audience interpretation and may contribute to polarization, reinforcing ideological echo chambers within digital public spheres. Simultaneously, comment sections have become key sites of political expression, where large volumes of user-generated content reflect both organic public sentiment and coordinated behavior. The interaction between article bias and user engagement thus constitutes a complex socio-technical system that requires multilayered analysis.

This paper investigates ideological bias and engagement dynamics in Indian online news media by examining large-scale article and comment data from TOI. We propose a comprehensive analytical framework integrating article bias detection, sentiment modeling, stance classification, and behavioral analysis of user activity. By jointly studying article tone and user participation, the work aims to uncover how political narratives propagate and how coordinated engagement may influence public discourse.

To illustrate the overall analytical pipeline used



*Proceedings of The 3rd Workshop on Natural Language Processing for Political Sciences (PoliticalNLP 2026) @ LREC 2026*, pages 48–57
16 May 2026, Palma, Mallorca (Spain). 

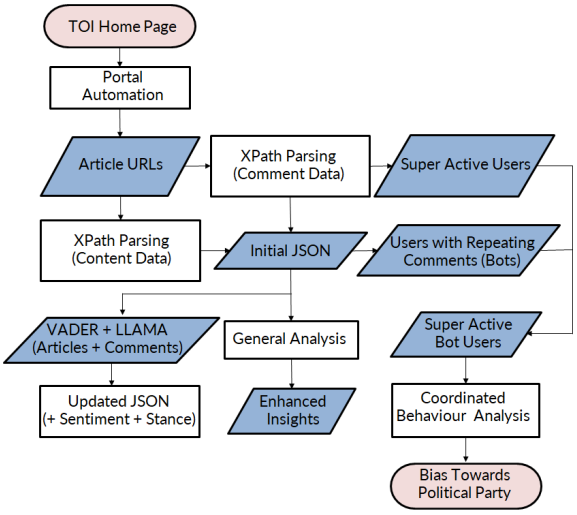


Figure 1: The process flow-diagram of the proposed analysis framework integrating article bias detection, sentiment modeling, stance classification, and bot identification.

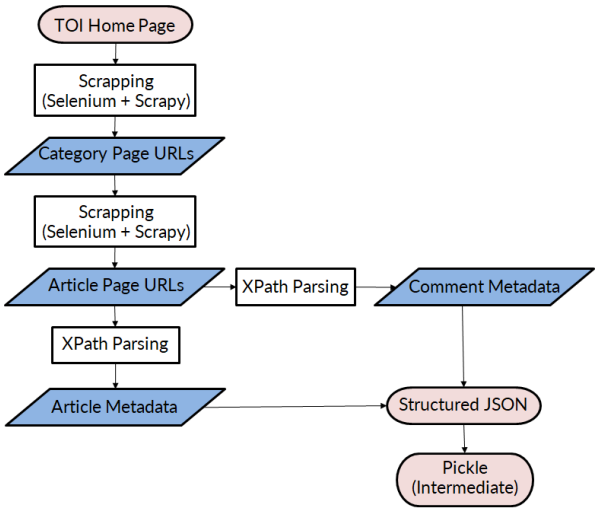


Figure 2: The proposed framework used for large-scale data collection and preprocessing using Selenium and Scrapy.

in this study, Figure 1 presents the process flow diagram of the proposed framework. We specifically analyze TOI user engagement to understand the intersection of article bias, sentiment trends, and bot activity. The proposed multi-stage pipeline performs sentiment scoring, stance classification, large-scale scraping, and heuristic behavior analysis. Bot-like users are detected using features such as bursty activity patterns, comment similarity, and unusually high posting frequency. The dataset preparation and scraping pipeline, implemented using Selenium and Scrapy, is illustrated in Figure 2. This architecture enables structured extraction of article content, comment threads, and metadata at scale.

This study is particularly timely given the increasing role of online discourse in shaping democratic participation and policy perception (Sunstein, 2001; Papacharissi, 2010). Our findings demonstrate that politicized narratives and coordinated behavior can disrupt organic discourse and skew engagement toward specific ideological positions.

**Objective:** In this paper, we aim to analyze user engagement patterns and ideological bias in Indian online news media by collecting large-scale comment and article data. We focus on detecting (i) ideological skew in article tone, (ii) sentiment patterns in user comments, and (iii) non-organic engagement driven by super-active or bot-like users. To this end, we propose a multi-stage framework that combines web scraping, sentiment analysis using VADER, stance detection using LLaMA 3.2, and heuristic-based bot detection. Our goal is to uncover the scale and structure of coordinated user behavior that could influence public opinion and democratic discourse.

In summary, our contributions are fivefold: (i) the construction of a dataset comprising over 21,000 political articles and 1.5 million comments from TOI across central and state government domains, (ii) a hybrid bias classification approach combining VADER-based sentiment scoring with LLaMA 3.2 stance detection, (iii) a heuristic-driven bot detection framework based on comment similarity, repetition, and activity frequency, (iv) a methodology to identify coordinated commenting patterns and their impact on sentiment dynamics, and (v) an empirical analysis of super-active and bot-like users in shaping political engagement and ideological discourse.

## 2. Related Work

**News Bias and Media Framing:** Ideological bias in news media has long been studied through content analysis, framing theory, and computational modeling. Gentzkow and Shapiro (Gentzkow and Shapiro, 2010) modeled media bias based on audience ideology, while Dalessio and Allen (D'Alessio and Allen, 2000) examined structural sources of bias in political reporting. More recent computational studies quantify bias in digital information environments, including search engine outputs and online news distributions (Kulshrestha et al., 2017; Gilani et al., 2017). Media framing further influences how audiences interpret political issues, and unsupervised approaches have been proposed to automatically identify latent frames in news narratives (Lai et al., 2017). These works establish the importance of computational approaches for analyzing bias but often focus primarily on content rather than user engagement dynamics.

**User Engagement and Political Bots:** Online

political discussions are increasingly shaped by automated and coordinated actors. Prior work demonstrates how bot networks can amplify ideological narratives, distort sentiment trends, and manipulate public discourse (Ferrara et al., 2016; Ratkiewicz et al., 2011; Cresci et al., 2017). Shao et al. (Shao et al., 2018) analyze the spread of misinformation through automated accounts, while Subrahmanian et al. (Subrahmanian et al., 2016) present large-scale frameworks for detecting influence operations on social platforms. However, most studies focus on social media ecosystems, leaving structured news comment environments comparatively underexplored. Our work extends this line by examining coordinated behavior within news platforms and linking it directly to article bias and sentiment patterns.

**Stance and Sentiment Modeling in Political Text:** Sentiment analysis tools such as VADER are widely used for analyzing social and political text due to their effectiveness on informal language and short-form content (Hutto and Gilbert, 2014). Stance detection further advances this by identifying political positioning toward entities, policies, or events (Augenstein et al., 2016; Mohammad et al., 2017). Recent research highlights the role of emotional and rhetorical language in shaping political stance, emphasizing the importance of pathos and affective framing in persuasive discourse (Evgrafova et al., 2024). Large language models have recently demonstrated strong performance in zero-shot and few-shot political sentiment and stance tasks, particularly in entity-centric political news analysis using contextual prompting and reasoning strategies (Kuila and Sarkar, 2024). These advances motivate hybrid approaches that combine rule-based interpretability with LLM-based reasoning.

**LLMs, Bias, and Political Communication:** The emergence of large language models has introduced new dimensions of bias and influence in political discourse. Studies show that conversational LLMs can exhibit partisan tendencies that affect users' political attitudes and decision-making (Fisher et al., 2025). Other work demonstrates bias in AI-generated news articles, where framing and content selection may distort neutrality and representation (Yoo and Shin, 2025). Generative search systems further introduce bias through citation selection, favoring particular media outlets and shaping information exposure (Dai et al., 2025). These findings highlight the importance of transparency and accountability in computational political analysis.

**Research Gaps:** Existing research largely treats article bias, sentiment dynamics, and automated engagement as separate problems. Social media platforms have been extensively studied, but structured news ecosystems with nested comment hierarchies remain comparatively underexplored. Moreover, few works jointly analyze article-level bias, user sentiment, stance behavior, and coordinated participation within a unified analytical framework. This study addresses this gap by integrating content analysis with behavioral modeling to provide a comprehensive view of ideological influence in digital news discourse.

## 3. Dataset Preparation

We present a hybrid web scraping framework depicted in Figure 2 that uses **Scrapy** (Scrapy Developers, 2023) and textbfSelenium (Silver, 2019) to obtain political news articles and the user engagement data from the TOI website (Scrapy Developers, 2023; Silver, 2019). The framework has been used to scrape more than **21,000 articles** and over **1.5 million comments** published during **October 2023 to April 2024** across various categories, including *Politics*, *India*, *Education*, *Tech-News*, *Good-News*, and *Sports*. We have used Scrapy to crawl category landing pages and efficiently extract article URLs. Selenium has been used to automate a Chrome WebDriver instance to load dynamic elements, such as nested-threaded comment boxes, user reactions, and Pagination buttons. All structured data (articles + comments) are stored in a nested `.json` format. Intermediate artifacts are serialized using Pickle (Python Software Foundation, 2023) for faster downstream loading.

The article metadata fields include `title`, `author`, `tag`, `date`, `time`, `content`, and `comments`: a list of nested dictionaries. Comment metadata fields include `user_id`, `text`, `upvotes`, `downvotes`, `replies`. A total of 27,000+ articles have been collected, out of which, **21,351 articles** have been categorized and distributed as follows: *India (3,077)*, *Education (1,477)*, *Politics (2,252)*, *Tech News (590)*, *Good News (104)*, and *Sports (13,851)*. We shall make this dataset publicly available upon acceptance.

## 4. Analysis & Findings

### 4.1. Genre-based User Engagement

We analyze user interaction based on article topical tags, focusing on **Politics** (2,187 articles with comments) and **India** (1,793 articles). 'Politics' shows higher volume, while 'India' exhibits deeper conversational engagement, measured by the ratio of replies to top-level comments. This ratio indicates how likely a comment is to receive a response, signaling the depth of discussion beyond surface-level reactions (Wang et al., 2016; Zhang et al.,

2019). Figure 4 shows article counts and reply-to-comment ratios, with *India* articles having a higher average reply ratio (**0.42**) than *Politics* (**0.34**), indicating more interactive discourse.

### 4.2. Ideological Bias Classification

In this context, we define *ideological bias* as the overall stance of the article content toward the government or any certain political party — categorized as pro-government, anti-government, or neutral. We use a hybrid two-stage pipeline to classify articles into *pro-government*, *anti-government*, or *neutral* classes. **(a) VADER-Based Sentiment Scoring** (Hutto and Gilbert, 2014) provides lexicon-based sentiment scores. VADER's method lacks the ability to capture subtle ideological cues and complex political narratives. To address this, we use **(b) LLaMA 3.2 Stance Detection:** (Touvron et al., 2024) (via Ollama API), which enhances detection by interpreting nuanced political framing. Figure 4 (left) shows article stance distribution, revealing dominance of anti-government articles, especially at the state level.

### 4.3. Sentiment Analysis of User Comments

Applying VADER (Hutto and Gilbert, 2014) on comments, sentiment is labeled as Positive (compound $>$ 0.05), Negative (compound $<$ -0.05), or Neutral otherwise. A predominance of negative sentiment is observed, consistent with trends in political discourse (Soroka and McAdams, 2015; Blitzer et al., 2007). Negative sentiment dominates political discussions (Soroka and McAdams, 2015; Blitzer et al., 2007). Emotional polarity aligns with trends in political discourse, where negative framing tends to provoke greater engagement. These sentiment dynamics contextualize bias and inform subsequent **bot behaviour analysis**.

### 4.4. Probable Bot Detection

We detect probable bots using two heuristics, (i) **Repetition Heuristic**, where the users posting identical or near-identical comments across various articles (e.g., *eknath* with 128 repeats on 48 articles), and (ii) **Semantic Similarity**, where average cosine similarity on TF-IDF vectors can identify repetitive messaging. Thresholds have been empirically tuned to classify **461 users** as probable bots, indicating *scripted activity*, i.e., automated or template-based posting of similar content across articles, aimed at *narrative amplification*, where repeated messaging artificially boosts the visibility and apparent support for certain political stances, rather than genuine discussion. We formulate a multi-feature detection algorithm to identify probable bots based on behavioral heuristics and textual patterns. Let $U$ be the set of all users. For each user $u \in U$:

1. Extract comment set $C_u = \{c_1^u, c_2^u, ..., c_n^u\}$
2. Compute:
   - **Length Variance** ($V_u$): Standard deviation of character lengths of all comments in $C_u$
   - **Average Cosine Similarity** ($S_u$): Mean of pairwise cosine similarities on TF-IDF vectors of comments
   - **Repetition Score** ($R_u$): Fraction of repeated comments across different articles
3. Normalize all scores to $[0, 1]$ and compute:

   $$\text{BotLikelihood}_u = \alpha \cdot V_u + \beta \cdot S_u + \gamma \cdot R_u$$

   where $\alpha, \beta, \gamma$ are empirically set weights.

   Additionally, we introduce a similarity decision threshold $t_s$ applied to the average cosine similarity score $S_u$. The threshold is selected from:

   $$t_s \in \{0.10, 0.15, 0.20, 0.25, 0.30, 0.40, 0.50, 0.60\}$$

   If $S_u \geq t_s$, the user's comments are considered semantically repetitive.
4. If $\text{BotLikelihood}_u \geq \tau$ (e.g., 0.7), classify $u$ as a probable bot.

### 4.5. Threshold Tuning and Justification

We conducted a preliminary threshold analysis on the top 50 users by total comment count (25 labeled as human and 25 as bot-like based on independent review by two annotators; Cohen's $\kappa = 0.76$, indicating substantial agreement). The aim was not to establish definitive operating points, but to explore how heuristic thresholds influence detection performance. For each candidate pair $(t_s, t_c)$ - where $t_s$ is the cosine similarity threshold and $t_c$ is the minimum comment count - we qualitatively evaluated precision and recall.

Although this subset was too small for statistically significant results, it provided insight into selecting operating points. We observed that $t_s \approx 0.20$ and $t_c \approx 6$ offered a balance: stricter thresholds (e.g., $t_s = 0.30$) missed coordinated behavior, whereas looser ones (e.g., $t_s = 0.10$) admitted many false positives. These choices were therefore adopted as pragmatic defaults.

We emphasize that this tuning is dataset-specific and exploratory. A larger annotated validation set will be required in future work to confirm the robustness of these thresholds and to calibrate them more rigorously.


Figure 3: (Left) Articles with Comments by Category; (Right) Reply-to-Comment Ratio.

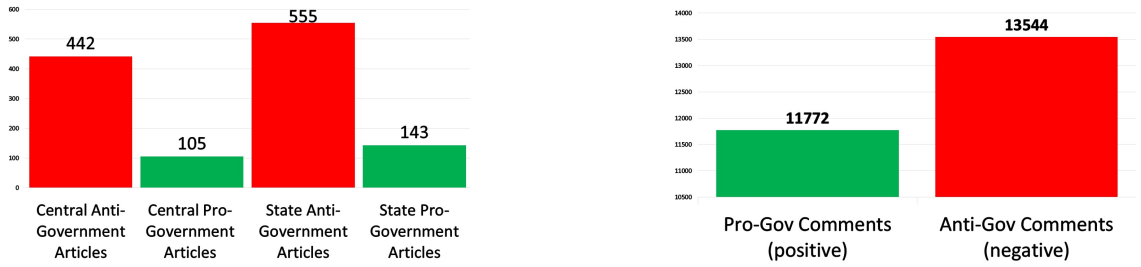


Figure 4: (Left) Article Stance: State vs National; (Right) Sentiment Distribution Across All Comments.

### 4.6. Coordinated Behavior Detection

We cluster comments using **cosine similarity** to identify similar messages across users, then map clusters to party-aligned articles (e.g., BJP, INC, AAP). This reveals organized groups likely engaged in influence campaigns or mobilization, shaping political discourse. To detect coordinated messaging, we cluster similar comments across users using cosine similarity. To detect reused text by user $u$, we use the following heuristics. Let $C_u = \{c_1, c_2, ..., c_n\}$ be the set of comments posted by $u$. We define,

$$R_u = \frac{|\{c \in C_u : \text{count}(c) \geq T\}|}{|C_u|}$$

where count$(c)$ is the number of distinct articles where $c$ appears and $T$ is a repetition threshold (e.g., 2). A high $R_u$ value (e.g., $\geq 0.3$) indicates repeated comment use and suggests non-organic behavior.

**Coordinated Behavior Beyond Detected Bots:** A natural question is whether coordinated patterns persist after removing flagged bot users. To assess this, we re-examined the repeated-comment clusters after excluding all 461 detected bots from the analysis. We find that reused comment sequences continue to appear among the remaining users: for instance, accounts not flagged as bots still post identical text across 3–10 distinct articles. This suggests that coordinated behavior — while concentrated among bot-flagged accounts — is not exclusively limited to them. A subset of human users appears to engage in coordinated or semi-coordinated activity (e.g., organised groups sharing and reposting content), though at a lower frequency and with greater textual variety than detected bots. This observation aligns with prior work on hybrid influence operations (Subrahmanian et al., 2016), where automated and manual actors operate in tandem to amplify narratives.

### 4.7. Super Active User Analysis

The top **5%** commenting users generate over **60%** of the comments, defined as *super active users*. We define super-active users as those in the top 5% of the total commenter population, ranked by the number of comments posted. This threshold was chosen based on the observed long-tail distribution of user activity, in which a small fraction of users accounted for a disproportionate share of total engagement. Specifically, out of all unique users in our dataset, the top 5% accounted for over 60% of the total comments posted. Their behaviour shows high engagement, ideological intensity, and comment repetition. Here, *ideological intensity* refers to the consistency and extremity of a user's political stance across multiple articles. It is measured based on the frequency of polarized (pro/anti) stances identified via LLaMA and rule-based methods, especially when repeated over diverse article topics, suggesting agenda-driven or semi-automated posting. Figure 5 visualizes their comment sentiment via VADER(Hutto and Gilbert, 2014). Bot-like behavior among these users is identified using (i) **Comment Similarity**, where average comment cosine similarity ($\geq 0.2$) on TF-IDF vectors has been used, (ii) **Comment Volume**, where users with minimum 6 comments are considered, and (iii) **Length Variance**, where variance in comment length is used to identify automation. Figure 6 and Figure 7 show sentiment for top 10 bot users and detected super active bots, respectively. Two users, *User A* and *User B*, are flagged as super active bot users.

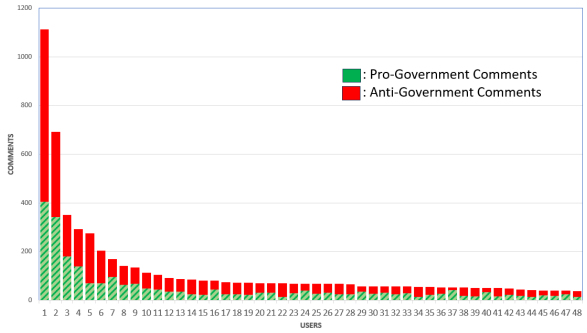


Figure 5: Sentiment Distribution of Super Active Users.

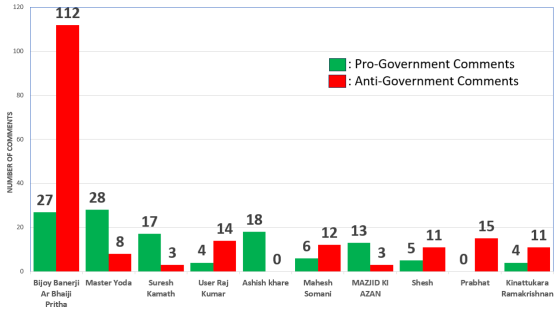


Figure 6: Sentiment Distribution of top 10 detected Bot Users.

## 4.8. Party Biasness Analysis

To study user engagement and behaviour, we have used two methods to classify **repeated comments** as pro-government or anti-government to certain political parties and analyzed how users, especially bots, align with these political biases. Two methods have been used. (i) **LLaMA-based Stance Classification:** Two-step querying on fine-tuned transformer architecture of LLaMA 3.2(Touvron et al., 2024) for party identification and stance classification. (ii) **Rule-based Sentiment and Bias Matching:** Combines VADER sentiment(Hutto and Gilbert, 2014) of article and comment with keyword bias tagging for stance mapping.

### 4.8.1. LLaMA-based Classification Results

Locally deployed LLaMA 3.2 model(Touvron et al., 2024) is queried using structured prompts to perform (i) **Party Identification:** for identifying which party the article supports and (ii) **Stance Classification:** for determining if the comment supports, criticizes, or is neutral toward that party. Table 1 shows bot user stance distribution by party, highlighting strong opposition towards AAP.

### 4.8.2. Sentiment and Bias Matching Rules

This method assigns stances based on article bias and comment sentiment following **Stance Mapping Logic:**

- If comment sentiment = $+1$:
  - Article bias = pro-gov $\Rightarrow$ stance = $+1$
  - Article bias = anti-gov $\Rightarrow$ stance = $-1$
  - No clear bias $\Rightarrow$ stance = $0$
- If comment sentiment = $-1$:
  - Article bias = pro-gov $\Rightarrow$ stance = $-1$
  - Article bias = anti-gov $\Rightarrow$ stance = $+1$
  - No clear bias $\Rightarrow$ stance = $0$
- If comment sentiment = $0$:
  - Stance = $0$ (neutral), regardless of article bias

Here, $+1$ indicates support and $-1$ indicates opposition.

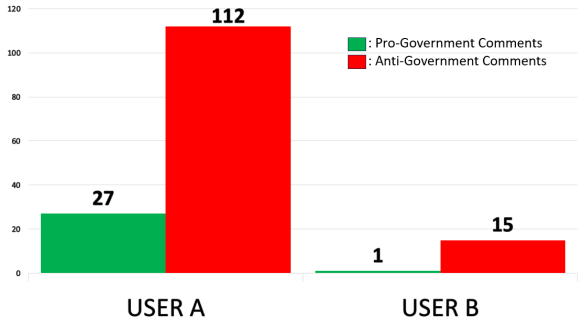


Figure 7: Sentiment Distribution of Detected Super Active Bot Users.

| ID | User | Party | % Support | % Oppose |
|---|---|---|---|---|
| 0 | User A | AAP | 0.0 | 100.0 |
| 1 | User A | BJP | 48.1 | 40.7 |
| 2 | User A | Congress | 0.0 | 90.0 |
| 3 | User B | AAP | 0.0 | 100.0 |
| 4 | User B | BJP | 0.0 | 100.0 |
| 5 | User B | Congress | 0.0 | 100.0 |

Table 1: Party-Wise Stance Distribution (LLaMA-based)

On a **volumetric analysis** rather than a percentage-based comparison, the following pattern is uncovered:

### 4.8.3. Comparative Analysis of the Two Approaches

Upon manual inspection of a part of the dataset, the **LLaMA-based** approach is found to offer higher accuracy and a deeper understanding, but it is computationally intensive. The **rule-based** method is faster but limited to surface cues. **Overall Accuracy:** Across 249 articles, rule-based accuracy compared to LLaMA was **88.35%** (Table 3).

| ID | User | Party | % Support | % Oppose |
|---|---|---|---|---|
| 0 | User A | AAP | 0.0 | 99.0 |
| 1 | User A | BJP | 29.6 | 25.9 |
| 2 | User A | Congress | 35.0 | 35.0 |
| 3 | User B | AAP | 1.1 | 98.9 |
| 4 | User B | BJP | 11.1 | 88.9 |
| 5 | User B | Congress | 0.0 | 100.0 |

Table 2: Party-Wise Stance Distribution (Rule-Based)

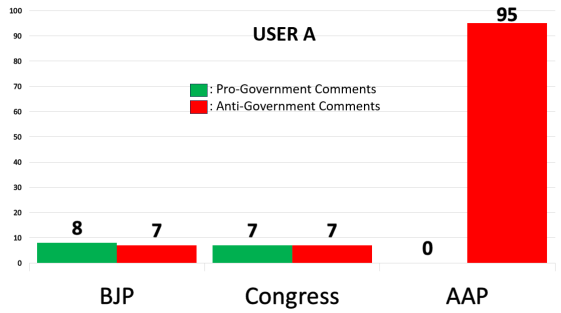


Figure 8: Party-Wise Stance (Rule-Based) for Bot 1 (User A).

#### 4.8.4. Qualitative Analysis of Mismatches

Out of the 249 articles analyzed, 29 cases (11.65%) showed disagreement between the rule-based and LLaMA-based stance classification. A closer manual inspection revealed several recurring patterns:

- **Sarcasm and Figurative Language:** Comments such as *"What a visionary move by the government to raise taxes again"* were labeled as positive by VADER due to lexical polarity, but LLaMA correctly classified them as anti-government given the sarcastic context.
- **Mixed Sentiment in a Single Comment:** Comments containing both praise and criticism (e.g., *"The government did well on healthcare, but utterly failed on jobs"*) were often misclassified by the rule-based method as neutral, whereas LLaMA placed them in the anti-government class by weighing the critical portion more heavily.
- **Implicit Bias Detection:** Articles with subtle framing (e.g., emphasis on corruption investigations) caused the rule-based approach to miss ideological leanings, while LLaMA captured these cues from the broader context.
- **Domain-Specific Vocabulary:** Terms like "jumla" (colloquial criticism) or "achhe din" (political slogan) were not captured by VADER lexicon, leading to misclassifications; LLaMA's contextual understanding handled these better.

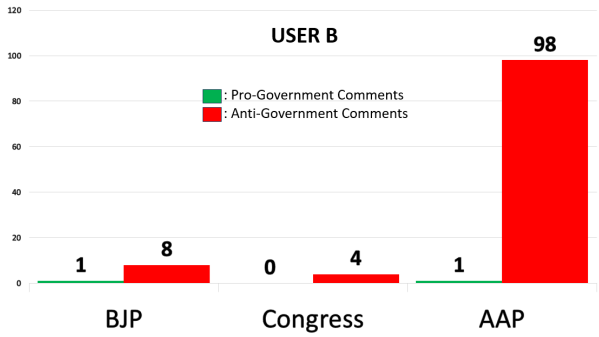


Figure 9: Party-Wise Stance (Rule-Based) for Bot 2 (User B).

| ID | User | Matched | Total | Mismatched | Accuracy (%) |
|---|---|---|---|---|---|
| 0 | User A | 115 | 142 | 27 | 80.99 |
| 1 | User B | 105 | 107 | 2 | 98.13 |

Table 3: Rule-Based Accuracy vs LLaMA

These examples suggest that the rule-based approach is effective for clear-cut sentiment but struggles with sarcasm, mixed sentiment, and culturally specific expressions. LLaMA, while computationally heavier, demonstrates greater robustness in such nuanced cases.

**Validation of LLaMA-Based Classification:** In addition to the rule-based comparison, we manually annotated a validation set of 100 political articles. LLaMA achieved an overall accuracy of **91.0%** with a macro-averaged F1 of **0.89**, demonstrating strong alignment with human labels.

#### 4.8.5. Key Observations

Both methods show that highly active bot users predominantly **oppose AAP-affiliated content**, with users such as *User A* and *User B* opposing over 99% of AAP articles, indicating likely coordinated negative engagement targeting a single party.

### 4.9. Validation of Individual Modules

To ensure the robustness and reliability of our framework, we validated each core component—stance classification, sentiment analysis, and bot detection—on manually labeled data.

**1. LLaMA-Based Stance Classification:** We manually annotated a validation set of 100 randomly sampled political articles as *pro-government*, *anti-government*, or *neutral*. LLaMA 3.2, run locally via prompt-based querying, achieved **91.0% accuracy** and a macro-averaged **F1-score of 0.89**, indicating strong alignment with human-labeled stance.

**2. VADER Sentiment on Indian Political Comments:** To validate VADER's effectiveness in the Indian political context, we manually labeled 300 user comments. VADER matched the labels with

an overall accuracy of **87.6%**, especially in clearly positive or negative comments. Errors were primarily due to sarcasm or contextual ambiguity, which were mitigated in downstream tasks by article-bias-aware stance mapping.

**3. Bot Detection Heuristics:** We reviewed 50 super-active users (the top 50 by total comment count) for bot-like behavior. These users were selected deterministically based on the observed long-tail distribution of user activity, as they represent the highest-impact commenters in the dataset. Each user was independently assessed by two annotators who examined posting frequency, comment similarity, and textual repetition patterns; the resulting labels were reconciled to produce a ground-truth set. Of the 50 users, 25 were labeled as likely human and 25 as bot-like. Of users flagged by our heuristic model from this cohort, **46 (92%)** aligned with the manual labels, exhibiting strong non-human behavior patterns such as comment repetition across 10+ articles or average cosine similarity $\geq 0.5$. The detection thresholds (similarity $\geq 0.2$, minimum 6 comments) were fine-tuned to reduce false positives based on this validation. We note that the absence of a fully independent annotated benchmark is a limitation; the annotators share domain knowledge as co-authors, and the 50-user set is small relative to the 461 flagged bots across the full dataset.

These validation studies support the credibility of our framework and show that existing tools, when carefully calibrated, are effective for analyzing ideological bias and coordinated behavior in Indian news media.

## 5. Conclusion and Broader Impact

This study presents a scalable framework to detect ideological bias, sentiment patterns, and bot-like coordinated behavior in Indian electronic news media. By analyzing over 21K articles and 1.5M comments, we infer article-level bias, classify comment stance, and identify manipulation through repeated and bot-generated engagement. Consistent anti-party sentiment from highly active users points to coordinated, potentially automated behavior.

We apply two stance classification methods: a LLaMA 3.2-based model for deep contextual understanding and a rule-based method that combines sentiment polarity with article alignment. Both approaches showed strong agreement in detecting oppositional trends among super active users. This dual-method approach enables accurate and interpretable stance mapping at scale.

**Reproducibility:** To support transparency, we release a publicly available GitHub repository with key modules, including web scraping scripts, bot detection logic, stance classification code, and visualization templates[1].

**Broader Implications:** Our findings reveal that a small subset of active or automated users can disproportionately shape public discourse by artificially inflating certain ideological narratives. Such manipulation distorts user engagement metrics, reinforces echo chambers, and can erode trust in digital platforms—especially during sensitive political periods. The framework underscores the need for bot-aware sentiment analysis, transparent engagement metrics, and moderation strategies that account for coordinated or automated influence.

**Limitations:** Despite strong results, the framework has several constraints. (i) LLaMA 3.2 may carry pretraining biases that affect stance predictions. (ii) Bot detection relies on empirically tuned thresholds (e.g., cosine similarity, repetition rate), which may limit generalizability. (iii) The dataset is platform-specific (TOI), and findings may not fully generalize across India's diverse media ecosystem. (iv) Sentiment tools may miss sarcasm, satire, or deeper intent in human discourse. (v) The bot detection validation relies on a small manually annotated set of 50 super-active users labeled by co-authors; inter-annotator agreement was assessed, but independent external validation would further strengthen reliability.

**Future Work:** We aim to fine-tune LLMs for Indian political contexts, incorporate temporal and network-level features, expand to multilingual platforms, and integrate misinformation detection to explore its overlap with engagement bias.

## 6. Ethical Considerations

This study analyzes publicly available user comments without collecting any personally identifiable information (PII). While labeling users as bots poses ethical risks, we mitigate false positives through a conservative, rule-based detection framework using comment repetition, textual similarity, and volume. Thresholds were empirically tuned and validated on manually reviewed subsets. Usernames in all visualizations were anonymized or replaced with placeholders.

We do not assert definitive bot identity, but highlight patterns of non-organic behavior in aggregate. No individual profiling was performed, and all analysis adhered to ethical standards for public data use. The goal is to surface systemic risks in discourse, not moderate users.

[1] https://github.com/DiabolikArms/OnlineNewsMediaCharacterization

## References


Suman Adhya and Debarshi Kumar Sanyal. 2022. What does the indian parliament discuss? an exploratory analysis of the question hour in the lok sabha. In *Proceedings of the Language Resources and Evaluation Conference (LREC)*. European Language Resources Association (ELRA).

Adrian Athique. 2012. The global dynamics of indian media. *Media, Culture and Society*, 34(4):407–422.

Isabelle Augenstein, Sebastian Riedel, and Stefan Augenstein. 2016. Stance detection with bidirectional conditional encoding. In *Proceedings of EMNLP*, pages 876–885.

John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 440–447.

Emmanuel Cartier and Hristo Tanev. 2022. Event detection in the socio-political domain. In *Proceedings of the Language Resources and Evaluation Conference (LREC)*, Marseille, France. European Language Resources Association (ELRA).

Stefano Cresci, Roberto Di Pietro, Marinella Petrocchi, Angelo Spognardi, and Maurizio Tesconi. 2017. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 963–972.

Sunhao Dai, Zhanshuo Cao, Wenjie Wang, Liang Pang, Jun Xu, See-Kiong Ng, and Tat-Seng Chua. 2025. Media source matters more than content: Unveiling political bias in llm-generated citations. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.

Dave D'Alessio and Mike Allen. 2000. Media bias in presidential elections: A meta-analysis. *Journal of Communication*, 50(4):133–156.

Natalia Evgrafova, Véronique Hoste, and Els Lefever. 2024. Analysing pathos in user-generated argumentative text. In *Proceedings of the Second Workshop on Natural Language Processing for Political Sciences (PoliticalNLP 2024)*, pages 39–44, Malta. ELRA Language Resource Association.

Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. The rise of social bots. *Communications of the ACM*, 59(7):96–104.

Jillian Fisher, Shangbin Feng, Robert Aron, Thomas Richardson, Yejin Choi, Daniel W. Fisher, Jennifer Pan, Yulia Tsvetkov, and Katharina Reinecke. 2025. Biased LLMs can influence political decision-making. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6559–6607. Association for Computational Linguistics.

Matthew Gentzkow and Jesse M. Shapiro. 2010. What drives media slant? evidence from u.s. newspapers. *Econometrica*, 78(1):35–71.

Zafar Gilani, Reza Farahbakhsh, Gareth Tyson, and Jon Crowcroft. 2017. Classification of political ideology on twitter: A comparative study. In *International Conference on Social Informatics*, pages 282–300. Springer.

C. J. Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 8, pages 216–225.

Alapan Kuila and Sudeshna Sarkar. 2024. Deciphering political entity sentiment in news with large language models: Zero-shot and few-shot strategies. In *Proceedings of the Second Workshop on Natural Language Processing for Political Sciences (PoliticalNLP 2024)*, pages 1–11, Malta. ELRA Language Resource Association.

Juhi Kulshrestha, Muhammad Bilal Zafar, Motahhare Eslami, Saptarshi Ghosh, and Krishna P. Gummadi. 2017. Quantifying search bias: Investigating sources of bias for political searches in social media. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pages 417–432.

Sha Lai, Yanru Jiang, Lei Guo, Margrit Betke, Prakash Ishwar, and Derry T. Wijaya. 2017. An unsupervised approach to discover media frames. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.

David M. J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, et al. 2018. The science of fake news. *Science*, 359(6380):1094–1096.

Saif M. Mohammad, Parinaz Sobhani, and Svetlana Kiritchenko. 2017. Stance and sentiment in tweets. In *ACM Transactions on Internet Technology*, volume 17, pages 1–23. ACM.

Zizi Papacharissi. 2010. *A Private Sphere: Democracy in a Digital Age*. Polity Press.

Python Software Foundation. 2023. pickle – python object serialization. https://docs.python.org/3/library/pickle.html.

Jacob Ratkiewicz, Michael Conover, Mark Meiss, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. 2011. Detecting and tracking political abuse in social media. In *Proceedings of the International AAAI Conference on Web and Social Media*.

Eric Sanders and Antal van den Bosch. 2021. Correlating political party names in tweets, newspapers and election results. In *Proceedings of the Workshop on Natural Language Processing for Political Sciences (PoliticalNLP)*. Association for Computational Linguistics.

Scrapy Developers. 2023. Scrapy: A fast and powerful scraping and web crawling framework. https://scrapy.org.

Chengcheng Shao, Giovanni Luca Ciampaglia, Onur Varol, Kai-Cheng Yang, Alessandro Flammini, and Filippo Menczer. 2018. Anatomy of an online misinformation network. *PLOS One*, 13(4):e0196087.

Simon Silver. 2019. Selenium webdriver. https://www.selenium.dev/documentation/.

Stuart N. Soroka and Stephen McAdams. 2015. News, politics, and negativity. *Political Communication*, 32(1):1–22.

V. S. Subrahmanian, A. Azaria, S. Durst, V. Kagan, A. Galstyan, K. Lerman, et al. 2016. The darpa twitter bot challenge. *Computer*, 49(6):38–46.

Cass R. Sunstein. 2001. *Republic.com*. Princeton University Press.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, et al. 2024. Llama 3: Open foundation and instruction models. https://ai.meta.com/llama/.

Xinyue Wang, Robert E. Kraut, and John M. Levine. 2016. Understanding conversational engagement in online news comments. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, pages 1379–1391. ACM.

Junho Yoo and Youhyun Shin. 2025. Fair or framed? political bias in news articles generated by llms. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.

Jian Zhang, Subhajit Ghosh, and Catherine Tucker. 2019. User engagement in online news comments: Behavioral patterns and the role of sentiment. In *Proceedings of the 2019 World Wide Web Conference*, pages 375–385.