

# Fine-grained Narrative Classification in Biased News Articles

Zeba Afroz<sup>1</sup>, Harsh Vardhan<sup>1</sup>, Pawan Bhakuni<sup>2</sup>, Aanchal Punia<sup>2</sup>,  
Rajdeep Kumar<sup>2</sup>, Md. Shad Akhtar<sup>1</sup>

<sup>1</sup>Indraprastha Institute of Information Technology Delhi, <sup>2</sup>Bharat Electronics Limited

{zebaa, harsh25001, shad.akhtar}@iiitd.ac.in

{pawanbhakuni, aanchalpunia, rajdeepkumar}@bel.co.in

## Abstract

Narratives are the cognitive and emotional scaffolds of propaganda. They organize isolated persuasive techniques into coherent stories that justify actions, attribute blame, and evoke identification with ideological camps. In this paper, we propose a novel fine-grained narrative classification in biased news article. We also explore article-bias classification as the pre-cursor task to narrative classification and fine-grained persuasive technique identification. We develop *INDI-PROP*, the first ideologically grounded fine-grain narrative dataset with multi-level annotation for analyzing propaganda in Indian news media. Our dataset *INDI-PROP* comprises 1,266 articles focusing on two polarizing socio-political events in recent times: CAA/NRC and the Farmers' protest. Each article is annotated at three hierarchical levels: (i) ideological article-bias (*pro-government*, *pro-opposition*, *neutral*), (ii) event-specific fine-grained narrative frames anchored in ideological polarity and communicative intent, and (iii) persuasive techniques. We propose *FANTA* and *TPTC*, two GPT-4o-mini guided multi-hop prompt-based reasoning frameworks for the bias, narrative, and persuasive technique classification. *FANTA* leverages multi-layered communicative phenomenon by integrating information extraction and contextual framing for hierarchical reasoning. On the other hand, *TPTC* adopts systematic decomposition of persuasive cues via a two-stage approach. Our evaluation suggest substantial improvement over underlying baselines in each case.

## 1. Introduction

Propaganda refers to the systematic use of language, images, and framing to shape public perception, influence opinion, and advance ideological agendas. The biased connotation operates through selective presentation, emotional manipulation, and repetition to reinforce obscure objectivity and polarize discourse. By distorting factual reporting and embedding ideological bias in everyday news consumption, it can erode trust in institutions, normalize misinformation, and reinforce echo chambers. Therefore, tackling it is not only a question of fact-checking but also of understanding how language constructs different versions of reality.

Detecting propaganda at scale is central to ensuring media transparency, civic literacy, and informed decision-making. Early research in computational propaganda focused on identifying persuasive techniques that constitute the linguistic building blocks of manipulative discourse. This line of work was systematized through several shared tasks, including a series of SemEval workshops during 2020 – 2024 (Da San Martino et al., 2020; Dimitrov et al., 2021; Piskorski et al., 2023; Dimitrov et al., 2024). The scope of these works varies from detecting persuasive techniques at the span level to incorporation of multimodality and multilinguality. Although these datasets advanced the understanding of how propaganda operates linguistically, they remain predominantly in English (Maarouf et al., 2023; Gruppi et al., 2021) and Arabic (Al-Henaki et al., 2025; Alam et al., 2022), overlooking the socio-political

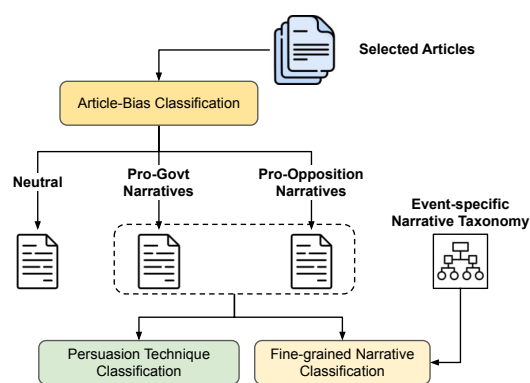


Figure 1: Multi-level schema of *INDI-PROP*.

and cultural context of countries like India. Moreover, these efforts primarily explore surface layer of persuasion to identify how manipulation is expressed without revealing what storyline, views, or narratives it promotes.

Narratives are the cognitive and emotional scaffolds of propaganda. They organize isolated persuasive techniques into coherent stories that justify actions, attribute blame, and evoke identification with ideological camps. Understanding narratives enables us to go beyond token-level detection and reveal how entire discourse ecosystems construct different realities. By modeling narrative explicitly, we can trace the mechanism of polarization, measure ideological symmetry between opposing sides, and build systems capable of narrative-aware propaganda analysis. Recently, SemEval-2025 Task 10 (Piskorski et al., 2025) studied narratives in propaganda in a global setting. However, these tax-

onomies treat narratives as thematic clusters without embedding ideological bias of the article.

To address the problem, we propose a multi-stage interpretive task to identify underlying storyline that links individual persuasive acts into a coherent ideological message. In particular, we explore three subtasks within the propaganda framework: a) **article-bias classification** to identify the ideological persuasion as *pro-government*, *pro-opposition*, or *neutral*; b) a novel **taxonomy-driven narrative classification** to reveal the underlying storyline; and c) **persuasive technique classification** used for propagating the agenda.

To support the tasks, we develop an ideologically grounded dataset, *INDI-PROP*, for two recent socio-political events in India: Farmers' protest and CAA/NRC. Inspired by the hierarchical narrative taxonomy of [Dash et al. \(2022\)](#), for each event, we propose a 2-level deep taxonomy to cover pro-government and pro-opposition aligned narratives. Each narrative in our framework is explicitly anchored to an ideological side, a communicative intent (legitimization, vilification, glorification, etc.), and a discursive focus that reveals how one side represents events, institutions, or social groups. In total, *INDI-PROP* comprises of 1266 articles with three bias labels, 20 narrative labels across two events, and 20 persuasion techniques.

We also propose two frameworks employing multi-hop prompt-based reasoning for the three classification tasks. We propose *FANTA* (**F**ine-grained **A**rticle **N**arrative **C**lassification) that leverages *entity-relation dynamics* and *context framing* for the article-bias and fine-grained narrative classification. In persuasive technique classification, conventionally, fine-tuned Pre-trained Language Models (PLMs) usually outperforms LLMs ([Jose and Greenstadt, 2025](#); [Hasanain et al., 2024](#)). However, we argue that due to subtle persuasive language LLMs struggles to appropriately decompose the article for underlying persuasive techniques. To counter this challenge, we propose *TPTC*, a novel **Two-stage Persuasion Technique Classification** that leverages the coarse persuasive intent before mapping it to fine-grained techniques. We evaluate both frameworks on *INDI-PROP* using GPT-4o-mini as the grounding LLM and observe performance-gain against fine-tuned Pre-trained Language Models (PLMs).

**Contributions:** We summarize contributions as:

- We propose a novel task of fine-grained narrative classification to provide interpretability of the underlying stories.
- We present *INDI-PROP*, a multi-level ideologically aligned dataset with annotations for bias, narratives, and persuasive techniques.

- We benchmark *INDI-PROP* using a multi-hop prompt-based inference only model, *FANTA*, for the fine-grained narrative classification and a novel two-stage persuasive intent decomposition model, *TPTC*, for the technique identification.

**Reproducibility:** The dataset and code can be accessed at <https://github.com/flamenlp/INDI-Prop>.

## 2. Related Work

Most of the prior research on propaganda is focused on linguistically identifying persuasive techniques. [Da San Martino et al. \(2019, 2020\)](#) introduced the first large-scale benchmark for detecting 18 persuasive techniques in English news articles. Subsequent shared tasks in SemEval (2021 and 2024) ([Dimitrov et al., 2021, 2024](#)) extended it to multimodality, incorporating memes as visual elements. In SemEval-2023, [Piskorski et al. \(2023\)](#) further extended it to multilinguality. These datasets have been instrumental in modeling propaganda linguistically. However, these resources remain predominantly western-centric, reflecting US & European media ideologies, whereas, non-Western democracies, like India, with its diverse ideological and complex politics, remained under-represented. Indian efforts have primarily focused on fact-checking ([Singhal et al., 2022](#)) and hate-speech detection ([K and B, 2025](#)), with limited exploration in the interplay between persuasion, ideology, and narrative framing. Notably, [Dash et al. \(2022\)](#) proposed a hierarchical narrative tree structure over Indian Twitter data, offering early insights into narrative construction in political discourse.

A complementary line of research has explored framing and narrative analysis in political communication. The Policy Frames Codebook ([Boydston et al., 2014](#)) and the Media Frames Corpus ([Card et al., 2015](#)) examined how policy issues are framed. More recent studies, SemEval-2025 Task 10 ([Piskorski et al., 2025](#)), formalized the multilingual characterization of narratives. Yet, these works typically define narratives as thematic clusters without embedding ideological polarity. Methodologically, previous work has employed transformer-based BERT ([Devlin et al., 2019](#)), RoBERTa ([Liu et al., 2019](#)), and DeBERTa ([He et al., 2020](#)), etc., for persuasive technique detection, which perform well in token-level tasks. However, LLMs despite their strong linguistic understanding, struggle to detect persuasive techniques in zero or few-shot settings ([Jose and Greenstadt, 2025](#); [Hasanain et al., 2024](#)), as persuasion often depends on subtle, context-driven cues beyond single-step inference.

Event	Article	Bias	Technique	Narrative
CAA	The reality is that after a detailed analysis of the provisions of the Act, we found that the law has nothing to do with Indian Muslims", the Chief added....	Pro-Govt	Assertion Glittering Generalities (Virtue)	Glorification of CAA Opposition spreading misinformation and fear
CAA	Across the country, as protests against the contentious Citizenship (Amendment) Act (CAA), NRC and police violence on students escalate, ....	Pro-Opp	Loaded Language Flag Waving	Glorifying Anti-CAA Protesters Framing Anti-CAA Protesters as Victims
Farmer's Protest	The list has doctors, lawyers, sportspersons, farmers, youth, women, and former IAS," said Bharatiya Janata Party (BJP) General Secretary Tarun Chugh Amid hectic campaigning ....	Pro-Govt	Smears	Glorification of Central Government Vilification of Opposition
Farmer's Protest	Amid growing outrage over the Lakhimpur Kheri violence, Rahul Gandhi questioned PM Modi's silence on the alleged murders of farmers, .....	Pro-Opp	Name Calling & labelling Red Hearing	Vilification of Central government Depicting Farmers as Victims

Table 1: A few annotated samples from INDI-PROP.

**Novelty in our work:** While prior research has studied bias detection and persuasive technique identification, our principal contribution lies in the new task of fine-grained narrative classification. Moreover, unlike prior corpora, INDI-PROP presents ideologically grounded fine-grained narratives along with persuasive techniques and bias annotations. This joint representation enables the analysis of how persuasive cues accumulate into broader ideological storylines. Furthermore, our approach of decomposing the narrative classification and persuasive techniques identification tasks into a multi-hop strategy inspires LLMs to comprehend the conceptual category of persuasion in an efficient manner, against state-of-the-art fine-tuned Pre-trained Language Models (PLMs) in existing works.

### 3. Dataset

We develop INDI-PROP, an ideologically grounded fine-grained narrative dataset that includes annotations of article-bias labels, taxonomy-driven narrative labels, and 20 persuasive techniques. The dataset centers on two major socio-political events in India, i.e., Citizenship Amendment Act (CAA)/National Register of Citizens (NRC) and the Farmers' protest. As shown in Fig 1, each article undergoes multi-level annotation schema for covering bias, fine-grained narrative, and persuasive techniques for both events. This setup enables us to capture how persuasive language, ideological bias, and narrative framing co-occur and interact within Indian news media. A few examples of the INDI-PROP dataset are depicted in Table 1.

#### 3.1. Data Collection

We select events based on following criteria: (a) it must span a prolonged duration, allowing us to trace how narratives emerge and evolve; (b) it must be highly polarizing, ensuring the presence of clear ideological divisions and contrasting perspectives; and (c) have sustained public engagement which have the potential to influence everyday discourse rather than being isolated incidents. Based on these criteria, we select two major socio-

political events in India in recent times, i.e., the Citizenship Amendment Act<sup>1</sup>/National Register of Citizens<sup>2</sup>(CAA/NRC) protests (2019 to 2024) and the farmers' protest<sup>3</sup> (2020 to 2024).

For data sources, we refer to Media Bias/Fact Check (MBFC)<sup>4</sup>, to identify outlets that represent an ideological spectrum. We include OpIndia, Republic World, Swarajya, The Quint, and Hindustan Times, covering right to left leaning perspectives. We collect articles using BeautifulSoup through a keyword driven web scrapping, employing terms such as *CAA*, *NRC*, *Farm Laws*, *Farmers' protest*, *Kisan Andolan*, etc.

**Data Cleaning:** We perform a systematic data cleaning process to ensure the quality and relevance of the collected corpus. We first remove duplicate entries by comparing URLs, titles, and content similarity. Next, we filter out of scope articles published outside the defined timeframes.

#### 3.2. Annotation Guidelines

The objective of the annotation process was to construct a qualitative resource that captures article-level ideologically grounded fine-grained narrative framing along with persuasive techniques within Indian spectrum.

**Article-Bias Guidelines:** Bias refers to the systematic ideological inclination. The task is to determine the overall ideological leaning of an article, whether it aligns with *pro-government* (Right) or *pro-opposition* (Left), or *neutral* reporting:

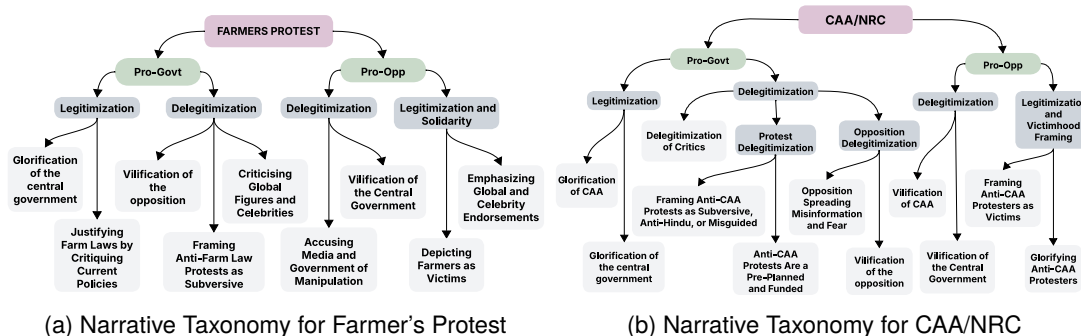
- **Pro-Govt:** Article aligns with pro-government, nationalist, or conservative narratives.  
E.g.: "The government's reform has empowered farmers to sell freely without a middleman" → *Pro-Govt*
- **Pro-Opp:** Article aligns with pro-opposition, liberal, or activist perspectives that question or critique government action.

<sup>1</sup>Citizenship Amendment Act, Wikipedia

<sup>2</sup>National Register of Citizens, Wikipedia

<sup>3</sup>Indian Farmers' Protest, Wikipedia

<sup>4</sup>mediabiasfactcheck.com



(a) Narrative Taxonomy for Farmer's Protest (b) Narrative Taxonomy for CAA/NRC  
 Figure 2: Narrative Taxonomy for two socio-political movements: Farmers' Protest and CAA/NRC. Taxonomy is organized along the ideological orientation of Pro-Govt and Pro-Opp.

E.g.: "The new farm laws are a betrayal of the farmer's trust and threaten livelihoods" → Pro-Opp

- **Neutral:** Article presents balanced reporting or lacks evaluative language.

E.g.: "The Parliament passed three new agricultural bills on Monday" → Neutral

Annotators were instructed to read the entire article and consider tone, framing, evaluative language, and attribution of responsibility before assigning the labels.

**Fine-Grain Narrative Guidelines:** After assigning the article-level bias label, annotators proceed to identify the dominant narrative expressed in each article. A narrative represents the storyline through which propaganda constructs legitimacy, blame, and moral position, and is unique for each event. It connects surface-level persuasive cues to deeper ideological framing, illustrating how each side legitimizes itself and delegitimizes its opponent. We utilize event-specific narrative taxonomies for fine-grained narrative classification.

**Narrative taxonomy:** We create separate narrative taxonomies for both events as depicted in Figure 2. We carefully analyze each event to understand its timeline and how it unfolded during its course. Each taxonomy was constructed after a detailed review of media coverage timelines, political speeches, protest statements, and counter narratives, to ensure a comprehensive mapping of the ideological space. The narratives for both events follow a hierarchical taxonomy that organizes ideological framing from broad to specific. The root node represents the overarching socio-political event. The second level divides the discourse into primary ideological clusters: *pro-government* (right or conservative) and *pro-opposition* (left or liberal). Subsequent branches captures broader narrative themes unique to each ideological side.

To finalize the labels, we conducted clustering to identify recurring ideological frames, storyline patterns, and rhetorical positioning strategies. These initial codes were iteratively grouped into higher-

level themes based on communicative intent (e.g., legitimization, delegitimization, vilification, victimhood framing, etc) and ideological alignment (pro-government vs. pro-opposition). Candidate leaf nodes were refined through multiple rounds of discussion to ensure conceptual distinctiveness, mutual exclusivity, and coverage of dominant discourse patterns observed in the corpus.

Following instructions are given to annotators:

- Narrative labels are applicable only for Pro-Govt and Pro-Opp biased articles.
- Refer to the event-specific taxonomy and select the leaf node whose description best captures the overall framing of the article. The chosen narrative should reflect the article's dominant ideological storyline rather than isolated sentences or quotations. In particular, annotators should identify:
  - **Legitimization:** Article portrays a particular side as just, necessary, patriotic, or reformative: analyze the article for evaluative language, success framing, and moral justification.
  - **Delegitimization:** Article undermines or attacks the credibility of the other party using blame, ridicule, misinformation, or moral discredit: analyze the article for lexical cues implying corruption, chaos, betrayal, or threat.
  - Based on above dominant framing, annotate the fine-grained narrative using the event-specific taxonomy.

The complete annotation guideline is provided in Appendix B.

**Persuasive Technique Guidelines:** Persuasion techniques determine the common practice people use to propagate their agenda. In one of the very first work on propaganda analysis, Da San Martino et al. (2019) defined 18 unique persuasion techniques. We also adopt these in our work and extend them with two additional techniques: *assertion* (Jowett and O'donnell, 2018), presenting a statement as fact without supporting evidence, often

Stance	Narrative class	Train	Test	Total
Pro-Govt	C1: Glorification of the central government	100	27	127
	C2: Vilification of the opposition	155	48	203
	C3: Glorification of CAA	59	19	78
	C4: Delegitimization of Critics	155	42	197
	C5: Framing anti-CAA protests as Subversive, Anti-Hindu, or Misguided	122	36	158
	C6: Opposition spreading misinformation and fear	119	30	149
	C7: Anti-CAA protests are a pre-planned and funded	49	11	60
Pro-Opp	C8: Vilification of the central government	80	13	93
	C9: Vilification of CAA	68	9	77
	C10: Glorifying anti-CAA protesters	57	6	63
	C11: Framing anti-CAA protesters as victims	50	3	53
Total	1014	244	1258	

(a) CAA/NRC

Stance	Narrative class	Train	Test	Total
Pro-Govt	F1: Glorification of the central government	15	6	21
	F2: Vilification of the opposition	19	7	26
	F3: Justifying farm laws by critiquing current policies	15	5	20
	F4: Criticizing global figures and celebrities	2	2	4
	F5: Framing anti-farm law protests as subversive	12	6	18
Pro-Opp	F6: Vilification of the central government	115	19	134
	F7: Depicting farmers as victims	117	20	137
	F8: Accusing media and government of manipulation	9	1	10
	F9: Emphasizing global and celebrity endorsements	6	1	7
Total	310	67	377	

(b) Farmer’s Protest

Table 2: Dataset statistics for the fine-grained narrative classes as per the taxonomy.

to persuade through confidence rather than proof and *glittering generalities* (Virtue) (Bauer, 2024) using emotionally appealing but vague and positive words to evoke approval without conveying specific meaning. We follow the guidelines of Da San Martino et al. (2019) for persuasive technique annotation. The complete annotation guideline is provided in Appendix B

### 3.3. Annotation

We employ nine annotators, aged between 20 - 40 (including five males and four females), for the annotation. All annotators are proficient in English and familiar with Indian politics. They represent a diverse background of research assistants, graduate/postgraduate students, and industrial researchers. In addition, two moderators oversee the annotation process and ensure quality control. All annotators/moderators are compensated according to the institute’s norm.

**Phase 1 – Annotator Training:** All annotators underwent a unified, multi-stage training process covering bias labeling, fine-grained narrative framing, and persuasive technique identification.

- **Orientation Session:** Annotators first attended a structured orientation that outlined the objectives of the project. Each annotator received the detailed annotation manual, which described every category with definitions and examples. Trainers clarified distinctions such as factual criticism vs ideological bias and topic description vs narrative framing. This session ensured that all annotators shared a common conceptual foundation.
- **Pilot Annotation and Calibration:** Annotators independently labeled a subset (i.e., 30) of articles spanning both events. For 5 pilot rounds and after each round, we conducted calibration meetings in which all annotations were discussed and disagreements were resolved. Through these sessions, annotators refined shared interpretations and harmonized their application of the guidelines across all three levels. In each calibration cycle, IAA scores were computed for bias, narrative, and

Persuasive Techniques	CAA/NRC			FARMERS		
	Train	Test	Total	Train	Test	Total
	537	131	668	480	118	598
T1: Causal Oversimplification	51	12	63	30	4	34
T2: Black-and-white Fallacy	12	3	15	6	1	7
T3: Straw Man	30	7	37	6	4	10
T4: Whataboutism	23	6	29	3	2	5
T5: Reductio ad Hitlerum	11	3	14	2	1	3
T6: Red Herring	24	7	31	5	1	6
T7: Loaded Language	112	29	141	55	11	66
T8: Name Calling and Labeling	104	27	131	43	12	55
T9: Appeal to Emotion, Fear, Prejudice	144	36	180	70	17	87
T10: Slogans	64	12	76	27	8	35
T11: Flag-Waving	52	8	60	11	4	15
T12: Exaggeration and Minimization	21	5	26	20	5	25
T13: Thought-Terminating Clichés	46	11	57	10	2	12
T14: Bandwagon	19	4	23	13	2	15
T15: Smears	176	49	225	100	22	122
T16: Obfuscation, intentional vagueness	19	4	23	7	2	9
T17: Doubt	49	12	61	26	5	31
T18: Appeal to Authority	30	9	39	8	3	11
T19: Assertion	148	51	199	70	17	87
T20: Glittering Generalities	51	16	67	46	17	63
<b>Total</b>	<b>1186</b>	<b>311</b>	<b>1497</b>	<b>558</b>	<b>140</b>	<b>698</b>

Table 3: Statistics of Persuasive Techniques.

technique labels to quantify consistency and accuracy. We repeat the pilot and calibration step until stable annotator behavior is observed.

**Phase 2 – Final Annotation:** The final annotation phase begins after achieving stable annotation consistency during training. At this stage, annotators annotate the articles independently. To assess the annotation quality, we compute inter-rater agreement on the annotation. For multi-class bias annotations, we compute Fleiss’  $\kappa$  (Fleiss, 1971; Carletta, 1996) and obtain a score of  $\kappa = 0.611$ . In multi-label fine-grained narrative annotations, we observe mean Fleiss’  $\kappa = 0.605$  and  $\kappa = 0.613$  for Farmers’ protest and CAA/NRC, respectively. We also compute mean pairwise Jaccard similarity (Jaccard, 1901) as 0.728 for Farmers’ protests and 0.673 for CAA/NRC. For multi-label persuasive technique annotations, we obtain mean Fleiss’  $\kappa = 0.312$  & pairwise Jaccard similarity of 0.364.

### 3.4. Dataset statistics

Our INDI-PROP dataset comprises of 1,266 news articles, 668 in CAA/NRC and 598 in Farmers’ protest. Overall, 54% of articles exhibit bias (Pro-Govt or Pro-Opp) in CAA/NRC, while the Farmers’ protest contains 30% biased content. Ideology-wise narrative class distribution for both events

is presented in Table 2. We note that observe right-leaning narratives dominate CAA/NRC, while Farmers’ protest is dominated by left leaning narratives. In total, INDI-PROP contains 1258 and 377 narrative labels in CAA/NRC and Farmers’ protest events, respectively. Finally, we report statistics of persuasive techniques in Table 3. We observe that a few techniques (*smears, assertion, appeal to emotion, name calling, loaded language, etc.*) dominate others in both events.

## 4. Methodology

We propose two distinct novel frameworks, employing multi-hop prompting strategy, for the three classification tasks. The first framework, FANTA, leverages the entity-relation dynamics in the article and its contextual framing for the article-bias classification. Subsequently, we employ an event-specific taxonomy for the fine-grained narrative classification. The second framework, TPTC, is a two-stage decomposition approach for the fine-grained persuasion techniques classification. It, at first, decomposes the article into identifying the conceptual category of persuasion, and subsequently, maps them into fine-grained persuasion techniques. A high-level overview of the proposed frameworks is presented in Figure 3.

### 4.1. Bias and Narrative Classification

We propose FANTA an inference-only prompt-based classification model via multi-hop reasoning process (Wei et al., 2022). Each hop incrementally aids the model’s interpretation of the article, from factual foundation to ideological abstraction. We group these hops into two functional stages: **information extraction** and **context framing**. The outcomes of these two steps are utilized in the prompt for article-bias classification. Subsequently, for biased articles, we perform narrative classification following the event-specific taxonomy.

**Information Extraction:** Entity-relation dynamics in an article is one of the foremost steps in high-level reasoning comprehension. To achieve this, we begin with extracting structural information from each article in terms of NER, co-reference resolution, and entity-relation pair. Through NER and co-reference resolution, we identify all key entities as individuals, organizations, laws, political institutions, etc. along with their linking pronouns, aliases, and mentions within the article. Finally, we obtain relationship among entities through the entity-relation extraction module. By formalizing these relationships, we establish the necessary foundation for the analysis of the next hop.

Article	Info Extraction	Context Framing
NEW DELHI: <name>, the leader of <party1>, has criticized the central government for implementing the CAA, calling it unconstitutional. <name> said, <party2> has made unconstitutional law. It has been made based on religion, which goes against the right to equality, therefore I tore a copy of the bill in the Parliament. We were, are and continue to be against this." The <party1> chief expressed optimism about his party's candidates performing well in the upcoming Lok Sabha elections. He clarified that <party1> is not in alliance with any...	<ul style="list-style-type: none"> <li>• &lt;name&gt; is a leader of &lt;party1&gt; and criticizes the CAA, indicating a direct opposition to &lt;party2&gt; and its policies.</li> <li>• &lt;name&gt;'s remarks about the &lt;party2&gt; abusing the Election Commission suggest a conflict between &lt;party1&gt; and &lt;party2&gt;, focusing on election integrity.</li> <li>• ...</li> </ul>	The article frames <name> as a vocal opponent of the CAA and the policies of the <party2>-led central government, emphasizing his legal and constitutional arguments against what he deems unconstitutional actions. Further, <name> is depicted as actively campaigning for <party1> and being involved in electoral strategies, signifying a proactive political approach. The mentioning of specific accusations against the <party2> indicates a charged political atmosphere and reflects an ongoing rivalry. The article presents <name> statements on electoral integrity and the political alliances within different states, maintaining a focus on <party1> position in the political landscape.

Table 4: An instance of information extraction and context framing in FANTA.

**Context Framing:** Context framing refers to the process of determining how different actors, events, or institutions are positioned within the discourse. Building upon the fact that media bias rarely stems from isolated statements but from the interplay of actors/institutions relative to one another, we infer how these entities are framed within the same article. We hypothesize that bias is not always explicit in factual statements but it is often encoded in the tone, emphasis, and moral perspective adopted by the writer. By identifying framing patterns, we can capture how ideology manifests through language and thus, can infer the ideological alignment of an article’s portrayal. This step signifies the core novelty of FANTA by providing interpretive foundation for bias reasoning, as depicted in Table 4.

**Article-Bias Classification:** Utilizing the contextual grounding, we perform a three-way bias classification for each article into *pro-govt*, *pro-opp*, or *neutral* labels. This layered reasoning allows bias classification to function as an interpretive conclusion rather than a direct classification task.

**Narrative Classification:** Following the article-bias classification, we proceed to fine-grained narrative classification for biased articles – i.e., article identified as non-neutral in the previous stage. We utilize the biased label (pro-govt/pro-opp) to limit the underlying narrative classes in the taxonomy hierarchy for our input prompts in fine-grained narrative classification. The predicted narrative class, together with the context framing, explains the underlying storyline of the biased article with appropriate justification in term of contextual reasoning. The prompts corresponding to each stage of the framework are detailed in the appendix A.

### 4.2. Persuasive Technique Classification

Prior research has demonstrated that LLMs, despite their strong linguistic reasoning capabilities, struggle to accurately detect persuasive techniques in zero- or few-shot settings (Jose and Greenstadt, 2025; Hasanain et al., 2024). We argue

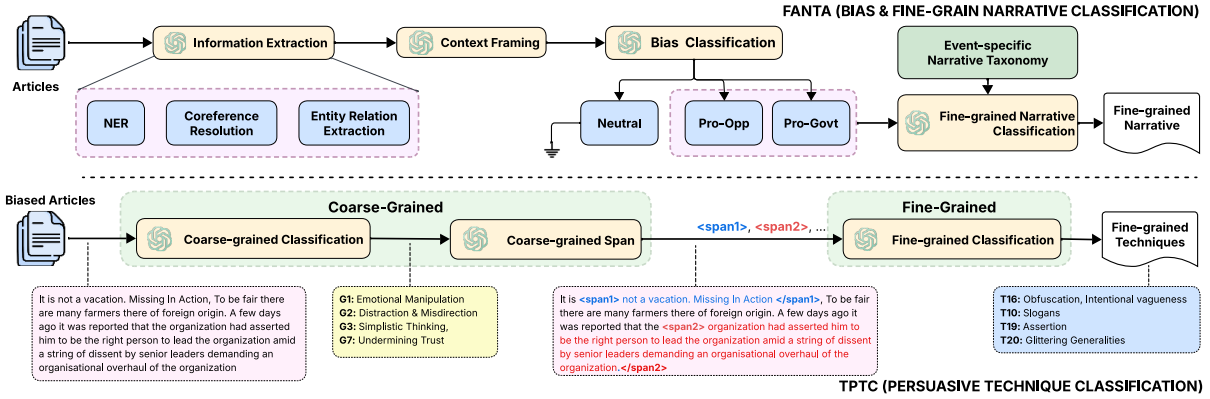


Figure 3: Overview of the proposed models architecture: The upper pipeline illustrates *FANTA* for the Bias and Fine-Grain Narrative Classification process. The lower pipeline illustrates *TPTC* for the Persuasive Technique Classification process.

Coarse-Grain	Fine-Grain
G1 Emotional Manipulation	Appeal to emotion/fear/prejudice; Loaded Language; Flag-Waving; Name Calling and Labeling; Glittering Generalities; Slogans.
G2 Distraction and Misdirection	Red Herring; Whataboutism; Obfuscation; Intentional Vagueness; Straw Man.
G3 Simplistic Thinking	Causal Oversimplification; Black-and-white Fallacy; Thought-Terminating Clichés; Assertion.
G4 Attack and Defamation	Name Calling and Labeling; Smears; Reductio ad Hitlerum.
G5 Manipulation through Popularity	Appeal to Authority; Bandwagon.
G6 Misrepresentation and Distortion	Straw Man; Exaggeration and Minimization.
G7 Undermining Trust	Doubt; Obfuscation; Intentional Vagueness.

Table 5: Coarse-grained to Fine-grained mapping for persuasion techniques.

that persuasion is a multi-layered communicative phenomenon and often realized through subtle contextual cues rather than being explicit in nature. As a result, identifying persuasive technique in a text becomes challenging in a single inference step for LLMs. To address this challenge, we propose *TPTC* that decomposes the problem into coarse-grained and fine-grained analysis. This decomposition allows the model to progressively reason about the conceptual category of persuasion at the coarse-level prior to mapping it to fine-level techniques. Such a design mimics the natural cognitive process of human annotators, who first recognize the general intent and then identify the precise persuasion device.

In the first stage, we organize the existing 20 fine-grained persuasive techniques into seven coarse-grained categories based on their abstract-level relatedness and functional similarity in terms of persuasion, as depicted in Table 5. Each category represents a shared communicative goal and serves as a conceptual scaffold that enables the model to establish a broader understanding of persuasion across fine-grained techniques. Utilizing this framework, at first, we extract the coarse-level persuasive categories and their collective textual grounding as spans. Subsequently, we utilize the category labels to obtain the subset of fine-grained

techniques belonging to that category to construct targeted prompts and spans to locate the potential regions-of-interest. We argue that by limiting the contextual region-of-interests, we allow the model to analyze the persuasive tone at a much finer-level for the technique identification. The prompts corresponding to each stage of the framework are detailed in the appendix A.

## 5. Experiments and Results

**Baselines and Comparative Systems:** Existing works (Da San Martino et al., 2020; Dimitrov et al., 2024; Piskorski et al., 2025; Dimitrov et al., 2021) have relied heavily on the fine-tuned Pre-trained Language Models (PLMs) due to their superior performance on propaganda-related tasks. We also comply and utilize several fine-tuned PLMs (BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), DistillBERT (Sanh et al., 2019), and DeBERTa (He et al., 2020)) for comparison. For PLM fine-tuning, we used Adam optimizer and trained for 20 epochs. Multi-label classification used sigmoid activation with a threshold of 0.5.

In addition, we also explore LLMs paradigms with GPT-4o-mini in zero- and few-shot settings as well. Moreover, we also ablate our proposed frameworks (i.e, *FANTA* and *TPTC*) with Gemma-3-12b (Team, 2025) and Mistral-8x7B-v0.1 (Mistral AI, 2024). Further, for the article-bias and narrative classification, we try a concise variant, *FANTA* (Concise), where we perform NER, Co-reference Resolution and entity-relation pair extraction in a single-hop.

**Article-Bias Classification:** We report results of article-bias classification in Table 6. We observe that, in most of the cases, PLMs report better precision scores than LLMs based systems. Further-

Prompt Settings	CAA/NRC									Farmer's Protest								
	Micro			Macro			Weighted			Micro			Macro			Weighted		
	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1
RoBERTa	0.705	0.710	0.707	0.475	0.520	0.488	0.641	0.710	0.662	0.735	0.703	0.719	0.433	0.460	0.446	0.681	0.703	0.691
BERT	<b>0.766</b>	0.626	0.689	<b>0.733</b>	0.580	0.641	<b>0.777</b>	0.626	0.685	0.738	0.670	0.702	0.433	0.435	0.434	0.683	0.670	0.676
DistillBERT	<b>0.766</b>	0.626	0.689	0.532	0.457	0.475	0.715	0.626	0.645	0.761	0.729	0.745	0.462	0.428	0.432	0.678	0.729	0.691
DeBERTa	0.713	0.664	0.688	0.520	0.483	0.471	0.697	0.664	0.640	<b>0.771</b>	0.712	0.740	<b>0.750</b>	0.642	0.658	<b>0.831</b>	0.712	0.749
GPT-4o-mini (Zero-Shot)	0.473	0.473	0.473	0.158	0.333	0.214	0.224	0.473	0.304	0.695	0.695	0.695	0.232	0.333	0.273	0.483	0.695	0.570
GPT-4o-mini (Few-Shot)	0.374	0.374	0.374	0.194	0.344	0.245	0.231	0.374	0.285	0.542	0.542	0.542	0.248	0.278	0.262	0.477	0.542	0.507
FANTA – Mistral	0.730	0.730	0.730	0.660	0.640	0.640	0.720	0.730	0.720	<b>0.771</b>	0.771	<b>0.771</b>	0.743	0.654	0.652	0.793	0.771	0.751
FANTA – Gemma	0.530	0.530	0.530	0.630	0.530	0.420	0.730	0.530	0.420	0.474	0.474	0.474	0.531	0.704	0.467	0.759	0.474	0.475
FANTA – GPT-4o-mini	0.670	<b>0.830</b>	<b>0.740</b>	0.590	<b>0.790</b>	<b>0.660</b>	0.700	<b>0.830</b>	<b>0.740</b>	0.730	<b>0.810</b>	0.770	0.650	<b>0.830</b>	<b>0.700</b>	0.780	<b>0.810</b>	<b>0.780</b>
FANTA – GPT-4o-mini (Concise)	0.520	0.420	0.410	0.520	0.420	0.410	0.520	0.520	0.480	0.720	0.720	0.720	0.570	0.490	0.500	0.700	0.710	0.680

Table 6: Article-Bias Classification: Comparative Results

Prompt Settings	CAA/NRC									Farmer's Protest								
	Micro			Macro			Weighted			Micro			Macro			Weighted		
	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1
RoBERTa	<b>0.744</b>	0.379	0.502	0.518	0.234	0.299	0.708	0.379	0.457	0.714	0.671	0.692	0.410	0.343	0.345	0.726	0.671	0.684
BERT	0.646	0.399	0.493	0.502	0.289	0.337	0.618	0.399	0.448	0.712	0.664	0.688	0.183	0.226	0.201	0.604	0.664	0.629
DistillBERT	0.664	0.497	0.568	0.542	0.366	0.405	0.678	0.497	0.555	0.621	0.584	0.602	0.163	0.212	0.179	0.562	0.584	0.561
DeBERTa	0.680	0.444	0.538	0.350	0.251	0.269	0.566	0.444	0.471	0.681	0.658	0.669	0.211	0.236	0.221	0.621	0.658	0.637
GPT-4o-mini (Zero-Shot)	0.340	0.200	0.250	0.190	0.120	0.090	0.300	0.200	0.150	0.300	0.450	0.367	0.119	0.279	0.130	0.470	0.460	0.420
GPT-4o-mini (Few-Shot)	0.060	0.100	0.080	0.020	0.330	0.040	0.007	0.100	0.014	0.100	0.240	0.150	0.048	0.298	0.077	0.050	0.235	0.082
FANTA – Mistral	0.650	0.520	0.580	0.580	0.410	0.430	0.700	0.520	0.550	<b>0.757</b>	0.691	0.723	0.522	0.381	0.427	0.746	0.691	0.696
FANTA – Gemma	0.400	0.420	0.410	0.450	0.380	0.320	0.590	0.420	0.370	0.370	0.563	0.446	0.336	<b>0.792</b>	0.393	0.676	0.563	0.482
FANTA – GPT-4o-mini	0.654	0.810	0.724	0.564	0.768	0.632	<b>0.708</b>	0.810	0.732	0.732	0.805	0.767	<b>0.702</b>	<b>0.780</b>	<b>0.709</b>	0.775	0.805	0.772
FANTA – GPT-4o-mini (Concise)	0.670	<b>0.830</b>	<b>0.740</b>	<b>0.590</b>	<b>0.790</b>	<b>0.660</b>	0.700	<b>0.830</b>	<b>0.740</b>	0.730	<b>0.810</b>	<b>0.770</b>	0.650	0.830	0.700	<b>0.780</b>	<b>0.810</b>	<b>0.780</b>

Table 7: Fine-Grained Narrative Classification: Comparative Results.

Prompt Settings	CAA/NRC									Farmer's Protest								
	Micro			Macro			Weighted			Micro			Macro			Weighted		
	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1
RoBERTa	<b>0.749</b>	0.329	0.457	0.206	0.118	0.140	0.509	0.329	0.376	0.687	0.509	0.585	0.156	0.167	0.157	0.495	0.509	0.496
BERT	0.689	0.364	0.476	<b>0.294</b>	0.166	0.201	0.545	0.364	0.423	0.719	0.455	0.557	0.156	0.130	0.136	0.522	0.455	0.476
DistillBERT	0.699	0.375	0.488	0.221	0.157	0.181	0.502	0.375	0.423	0.669	0.405	0.504	0.134	0.119	0.124	0.491	0.405	0.439
DeBERTa	0.730	0.342	0.466	0.235	0.134	0.162	0.520	0.342	0.394	0.674	0.423	0.520	0.201	0.128	0.125	0.568	0.423	0.439
GPT-4o-mini (Zero-Shot)	0.208	0.517	0.297	0.167	0.385	0.190	0.272	0.517	0.298	0.098	0.398	0.157	0.147	0.342	0.106	0.293	0.398	0.159
GPT-4o-mini (Few-Shot)	0.210	0.500	0.300	0.198	<b>0.390</b>	0.200	0.270	0.500	0.297	0.090	0.345	0.143	0.739	0.300	0.084	0.129	0.345	0.134
TPTC – Mistral	0.250	0.292	0.269	0.246	0.236	0.201	0.451	0.292	0.301	0.188	0.252	0.215	0.202	0.256	0.171	0.487	0.252	0.266
TPTC – Gemma	0.440	0.270	0.330	0.259	0.169	0.174	0.401	0.268	0.283	0.600	0.418	0.493	0.294	0.191	0.208	0.513	0.418	0.421
TPTC – GPT-4o-mini	0.728	<b>0.728</b>	0.219	0.234	<b>0.214</b>	<b>0.840</b>	<b>0.766</b>	<b>0.784</b>	<b>0.847</b>	<b>0.847</b>	<b>0.847</b>	<b>0.847</b>	<b>0.814</b>	<b>0.887</b>	<b>0.834</b>	<b>0.820</b>	<b>0.910</b>	<b>0.853</b>

Table 8: Fine-Grained Persuasive Technique Classification: Comparative Results.

more, GPT-4o-mini in zero- and few-shot settings yield inferior results against all competition systems. This is in line with existing works, which showed poor adaptability of LLMs in propaganda settings. However, on the contrary, leveraging explicit information extraction and context framing, our proposed framework, FANTA, outperform PLMs on recall and F1-score across both events.

**Narrative Classification:** We observe similar trends in the narrative classification step as well in Table 7. Moreover, the concise framework further assists the model in improved prediction. This indicates that explicitly merging entity and relation information strengthens discourse-level coherence by enabling the model to reason about who interacts with whom within ideological frames. Overall, the structured reasoning approach integrating entity-relational analysis and context framing consistently outperforms all fine-tuned baselines and prompting variants across all metrics.

**Persuasive Technique Identification:** Table 8 presents the comparative results for the fine-grained persuasive technique classification for

both events. We observe moderate performance of PLMs – in the range of [0.457, 0.488] micro-F1, [0.140, 0.201] macro-F1, and [0.376, 0.423] weighted-F1 on the CAA dataset and in the range of [0.504, 0.585] micro-F1, [0.124, 0.157] macro-F1, and [0.439, 0.496] weighted-F1 on the Farmers’ dataset. In contrast, our TPTC achieves a substantial leap in performance, attaining 0.728 micro-F1, 0.214 macro-F1, and 0.784 weighted-F1 on CAA. Similarly, TPTC yields 0.847 micro-F1, 0.834 macro-F1, and 0.853 weighted-F1 on Farmers’ protest. We argue that the performance-gain is due to our two-stage hierarchical prompting design, which recognizes the coarse persuasive intent before mapping it to fine-grained persuasive techniques.

**Error Analysis:** We report our error analyses in Figure 4 and Table 9. For the bias classification task, the confusion matrix in Figure 4 shows that the model performs well overall, particularly in identifying right-leaning articles, with a high precision of 64 *true-positives* and only 2 *false-negatives*. The main source of error stems from the model’s difficulty in distinguishing neutral articles from bi-

Task	Ground truth	Model	Prediction
Persuasive Technique Classification	[appeal to emotion/fear/prejudice; assertion; glittering generalities (virtue); loaded language; slogans; smears;]	GPT-4o-mini (FS)	[Assertion;] [Causal Oversimplification; Exaggeration and Minimization; Black-and-white Fallacy; Name Calling and Labeling; Appeal to Authority; Flag-Waving; appeal to emotion/fear/prejudice; glittering generalities (virtue); loaded language; slogans; smears;]
		TPTC – GPT-4o-mini	[appeal to authority; smears; bandwagon; exaggeration and minimization; name calling and labeling; reductio ad hitlerum; straw man; assertion; glittering generalities (virtue); loaded language; slogans;]
Narrative Classification	[glorification of CAA; vilification of the opposition; opposition spreading misinformation and fear; delegitimization of critics; glorification of the central government;]	GPT-4o-mini (FS)	[Vilification of CAA; Vilification of Central Government; Glorifying Anti-CAA Protesters; Framing Anti-CAA Protesters as Victims; Glorification of CAA; Vilification of the opposition; opposition spreading misinformation and fear; Delegitimization of critics; Glorification of the central government;]
		FANTA – GPT-4o-mini	[Vilification of the opposition; opposition spreading misinformation and fear; Glorification of the central government; Glorification of CAA; Delegitimization of critics;]

Table 9: Error analysis: True Positives, False Positives, False Negatives.

ased ones. This suggests that the model is highly sensitive to political keywords, sometimes over interpreting factual reporting as ideologically biased.

On the other hand, a qualitative analysis of the model’s predictions reveals challenges in fine-grained classification for both narrative and technique classification tasks. In most of the cases, we observe that models make partially correct predictions, capturing some narrative themes while missing the complete picture. However, we also notice that proposed model usually predicts within the ideological polarity.

True Labels	Pro-Opp	28	4	7
	Pro-Govt	0	64	2
	Neutral	25	12	107
		Pro-Opp	Pro-Govt	Neutral
		Predicted Labels		

Figure 4: Confusion matrix for Article-Bias Classification.

## 6. Conclusion

In this paper, we proposed a novel fine-grained narrative classification for biased articles to unveil the storyline at play. To support this, we developed the first ideologically grounded fine-grained narrative dataset for Indian news media, covering two recent socio-political events (CAA and the Farmers’ protest) and annotated them across three levels: bias, narrative, and persuasive techniques. Our proposed multi-stage hierarchical models, FANTA and TPTC leveraged contextual reasoning over entities and their relations to uncover how ideological bias manifests through structured narratives and rhetorical persuasion. We observed that, LLMs, when provided with decomposed persuasive intent adapt efficiently for the fine-grained classifications for both narrative and persuasive techniques. Overall, these findings validate that hierarchical reasoning enables LLMs to bridge micro-level persuasive cues and macro-level ideological framing, capturing propaganda as a multi-layered communicative phenomenon. We believe this work lays the groundwork for narrative-aware propaganda analysis in diverse, multilingual contexts, and pro-

vides a strong empirical benchmark for future research on media framing, ideology, and persuasion.

## 7. Ethical Considerations

The INDI-PROP dataset was created with careful attention to ethical issues related to data collection, annotation, and release. All articles included in the dataset were collected from publicly accessible news outlets, and no personally identifiable information (PII) beyond what was already published by the media was retained. The dataset does not include user-generated content such as social media posts, thereby avoiding concerns around individual consent or inadvertent disclosure of private data.

All news outlets were selected based on their rating on the Media Bias/Fact Check (MBFC)<sup>4</sup> platform to ensure a balanced representation of ideological perspectives. We emphasize that we do not glorify, diminish, support, or oppose any of these outlets, they are included solely for research purposes to study media bias, framing, and propaganda within a controlled, transparent framework.

Nevertheless, the dataset reflects the biases, inaccuracies, and manipulative strategies present in the sources. Some outlets were deliberately chosen to include those flagged for poor sourcing or failed fact checks, some articles may contain offensive language, or extreme ideological framing. These elements are integral to studying propaganda but may reproduce harmful discourse if taken out of context. To mitigate this risk, the dataset is intended strictly for academic and research purposes. Users are encouraged to apply critical thinking when interpreting model outputs trained on the data.

Another consideration arises from the subjectivity of annotations. Persuasive techniques and narrative frames are inherently interpretive categories, and annotators may bring their own perspectives into the labeling process. While we employed multiple rounds of training, pilot annotations, and consensus building to reduce subjectivity, some level of disagreement is unavoidable. Researchers us-

ing the dataset should therefore treat annotations as a curated but fallible representation rather than the absolute truth.

From a broader perspective, there is a risk that models trained on `INDI-PROP` could be misused for purposes such as automated censorship, political profiling, or targeted manipulation of readers. We strongly caution against such applications. The dataset is released to advance transparency in media analysis, to support work in computational social science and NLP, and to foster a deeper understanding of how propaganda functions in public discourse.

Finally, we recognize that the dataset focuses on two politically sensitive events in India. Our goal is not to endorse or delegitimize any political position but to provide a resource for studying persuasion in context. By making the data openly available, we hope to support a diversity of perspectives and to enable critical research on the media ecosystem in the Global South, which remains underrepresented in computational resources.

## 8. Limitations

While `INDI-PROP` dataset advances the study of propaganda and ideological narratives in Indian media, several limitations remain. The dataset focuses on two major socio-political events (CAA/NRC and the Farmers' Protest), which may not capture the full diversity of propaganda strategies across other domains or time periods. Our models, despite leveraging multi-hop reasoning, rely primarily on textual cues and do not incorporate multimodal evidence such as images or meta-data. Furthermore, the experiments are limited to English news sources, extending the framework to multilingual or code-mixed Indian media would further enhance its generalizability. Future work will address these limitations by expanding event coverage, integrating multimodal reasoning, and incorporating multilinguality.

## Acknowledgment

We acknowledge the partial support of Bharat Electronics Limited (BEL) and Infosys Foundation through Center of AI (CAI) at IIIT Delhi.

## 9. Bibliographical References

Lubna Al-Henaki, Hend Al-Khalifa, Abdulmalik Al-Salman, Hajar Alqubayshi, Hind Al-Twailay, Gheeda Alghamdi, and Hawra Aljasim. 2025. Multiprose: A multi-label arabic dataset for

propaganda, sentiment, and emotion detection. In *International Conference on Applications of Natural Language to Information Systems*, pages 156–172. Springer.

Firoj Alam, Hamdy Mubarak, Wajdi Zaghoulani, Giovanni Da San Martino, and Preslav Nakov. 2022. [Overview of the WANLP 2022 shared task on propaganda detection in Arabic](#). In *Proceedings of the Seventh Arabic Natural Language Processing Workshop (WANLP)*, pages 108–118, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

AJ Bauer. 2024. Glittering generalities: Reconsidering the institute for propaganda analysis. *International Journal of Communication*, 18:19.

Amber E Boydston, Dallas Card, Justin H Gross, Philip Resnik, and Noah A Smith. 2014. Tracking the development of media frames within and across policy issues. In *APSA 2014 annual meeting paper*, pages 1–25.

Dallas Card, Amber Boydston, Justin H Gross, Philip Resnik, and Noah A Smith. 2015. The media frames corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444.

Jean Carletta. 1996. Assessing agreement on classification tasks: the kappa statistic. *arXiv preprint cmp-lg/9602004*.

Giovanni Da San Martino, Alberto Barrón-Cedeño, Henning Wachsmuth, Rostislav Petrov, and Preslav Nakov. 2020. [SemEval-2020 task 11: Detection of propaganda techniques in news articles](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1377–1414, Barcelona (online). International Committee for Computational Linguistics.

Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. [Fine-grained analysis of propaganda in news articles](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5636–5646, Hong Kong, China. Association for Computational Linguistics.

Saloni Dash, Arshia Arya, Sukhridh Kaur, and Joyojeet Pal. 2022. Narrative building in propaganda networks on indian twitter. In *Proceedings of*

- the 14th ACM Web Science Conference 2022, pages 239–244.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Dimitar Dimitrov, Firoj Alam, Maram Hasanain, Abul Hasnat, Fabrizio Silvestri, Preslav Nakov, and Giovanni Da San Martino. 2024. [SemEval-2024 task 4: Multilingual detection of persuasion techniques in memes](#). In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, pages 2009–2026, Mexico City, Mexico. Association for Computational Linguistics.
- Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021. [SemEval-2021 task 6: Detection of persuasion techniques in texts and images](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 70–98, Online. Association for Computational Linguistics.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- Maurício Gruppi, Benjamin D Horne, and Sibel Adali. 2021. Nela-gt-2020: A large multi-labelled news dataset for the study of misinformation in news articles. [arXiv preprint arXiv:2102.04567](#).
- Maram Hasanain, Fatema Ahmad, and Firoj Alam. 2024. [Can GPT-4 identify propaganda? annotation and detection of propaganda spans in news articles](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2724–2744, Torino, Italia. ELRA and ICCL.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. [arXiv preprint arXiv:2006.03654](#).
- Paul Jaccard. 1901. Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bull Soc Vaudoise Sci Nat*, 37:547–579.
- Julia Jose and Rachel Greenstadt. 2025. Are large language models good at detecting propaganda? [arXiv preprint arXiv:2505.13706](#).
- Garth S Jowett and Victoria O’donnell. 2018. *Propaganda & persuasion*. Sage publications.
- Sreeja K and Bharathi B. 2025. [SS-NCSE@DravidianLangTech 2025: Multimodal hate speech detection in Dravidian languages](#). In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, pages 98–102, Acoma, The Albuquerque Convention Center, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. [arXiv preprint arXiv:1907.11692](#).
- Abdurahman Maarouf, Dominik Bär, Dominique Geissler, and Stefan Feuerriegel. 2023. Hqp: a human-annotated dataset for detecting online propaganda. [arXiv preprint arXiv:2304.14931](#).
- Mistral AI. 2024. [Mixtral 8x7B](#).
- Hygeinus Nwokwu and Chinasa Johson. 2019. Propaganda and electioneering campaigns in 2019 general elections. *South-East Journal of Public Relations*, 1(1).
- Jakub Piskorski, Tarek Mahmoud, Nikolaos Nikolaidis, Ricardo Campos, Alipio Mario Jorge, Dimitar Dimitrov, Purificação Silvano, Roman Yangarber, Shivam Sharma, Tanmoy Chakraborty, Nuno Guimaraes, Elisa Sartori, Nicolas Stefanovitch, Zhuohan Xie, Preslav Nakov, and Giovanni Da San Martino. 2025. [SemEval 2025 task 10: Multilingual characterization and extraction of narratives from online news](#). In *Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025)*, pages 2610–2643, Vienna, Austria. Association for Computational Linguistics.
- Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023. [SemEval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multi-lingual setup](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2343–2361, Toronto, Canada. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. [arXiv preprint arXiv:1910.01108](#).

Shivangi Singhal, Rajiv Ratn Shah, and Pon-nurangam Kumaraguru. 2022. Factdrill: A data repository of fact-checked social media content to study fake news incidents in india. In Proceedings of the international AAAI conference on web and social media, volume 16, pages 1322–1331.

Gemma Team. 2025. [Gemma 3](#).

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems, 35:24824–24837.

## A. Prompt templates

This section presents the complete set of prompts used in our frameworks. The prompts for bias classification, fine-grained narrative classification, and persuasive technique identification are provided in tabular form. Table 10 presents the prompts for bias and narrative classification, while Tables 11 present the prompts for coarse-grained and fine-grained persuasive technique detection.

## B. Annotation Guidelines

In this section, we present the annotation guidelines for fine-grained narrative labels, including two new persuasive techniques introduced in addition to the 18 techniques proposed by (Da San Martino et al., 2019). The corresponding guidelines are provided in Table 12.

### Pro-Govt Aligned Narratives (Farmer’s protest):

1. Glorification of the Central Government: Articles portraying the central government as visionary, reformist, or acting in farmers’ best interest.
  - Presents farm laws as historic, bold, or progressive.
  - Credits leadership for modernization or economic reform.
  - Highlights efficiency, decisiveness, or long-term vision.
  - Eg: *"The PM’s reform vision will finally free farmers from middlemen."*, *"These laws mark a new era of prosperity for rural India."*
2. Justifying Farm Laws by Critiquing Existing Policies: Articles defending new farm laws by emphasizing flaws or inefficiencies of the older system.
  - Frames previous market structures as exploitative or outdated.

- Argues that reforms are necessary corrections.
  - Eg: *"For decades, farmers were trapped in the APMC monopoly."*, *"The old system benefited brokers, not farmers."*
3. Vilification of the Opposition: Articles blaming opposition parties or leaders for politicizing or misleading farmers.
    - Claims protests are driven by opposition propaganda.
    - Describes leaders as hypocritical or self-serving.
    - Eg: *"Opposition leaders are inciting unrest for political mileage."*, *"Those who failed farmers for 60 years now pretend to care."*
  4. Framing Anti-Farm Law protests as Subversive: Articles depicting the protests as anti-national or disruptive to peace.
    - References funding or conspiracy.
    - Describes protests as violent, separatist, or extremist.
    - Eg: *"These so-called protests are an attempt to destabilize the nation."*
  5. Criticizing Global Figures and Celebrities: Articles discrediting international figures or celebrities who expressed solidarity with farmers.
    - Dismisses foreign commentary as ignorant or biased.
    - Emphasizes domestic sovereignty over external opinions.
    - Eg: *"Foreign celebrities should stop lecturing India without understanding facts."*, *"These global influencers are tools of anti-India campaigns."*

### Pro-Opp Aligned Narratives (Farmer’s protest):

1. Vilification of the Central Government: Articles accusing the government of arrogance, repression, or lack of empathy toward farmers.
  - Frames government as authoritarian or indifferent.
  - Uses moral condemnation or emotional tone.
  - Eg: *"The government silences farmers instead of hearing their pain."*, *"Farmers’ pleas are ignored by those in power."*
2. Depicting Farmers as Victims: Articles portraying farmers as suffering under unjust laws or state apathy.
  - Focuses on hardships, deaths, or sacrifices of protesters.
  - Evokes empathy and moral outrage.
  - Eg: *"Elderly farmers spend nights on the streets for their rights."*, *"protesters continue to die, but their voices remain unheard."*
3. Emphasizing Global and Celebrity Endorsements: Articles highlighting international attention or celebrity support for protests to validate legitimacy.

<p><b>Step: Prompt</b></p> <p><b>NER:</b> {article} Identify key entities mentioned in the text to understand the focal points. NOTE: Only include individuals, Organizations/ Groups, Locations, and Laws/Policies. Just name the entities, no need to explain.</p> <hr/> <p><b>Coreference Resolution:</b> {article} Now perform Coreference Resolution to resolve pronouns and indirect references to maintain consistency in analysis. Note: Just include the sentences where the entity is coreferred.</p> <hr/> <p><b>Entity-Relation Analysis:</b> {article} Now perform relation extraction analysis for the text and identify the relationship between entities and their contexts to detect IF THERE IS potential biases or associations, OR the article has unbiased reporting.</p> <hr/> <p><b>Context and Framing Analysis:</b> {article} Analyze how the article frames entities, events, and legislation based on the analysis from previous steps.</p> <hr/> <p><b>Bias Detection:</b> {article} Based on the analysis done so far, identify whether the article is Neutral, Left-biased, or Right-biased. Do not give explanation just do Classification and Classify among – <b>Bias Classes</b> – Neutral, Right-biased, Left-biased <b>Instructions:</b> Instead of numbering the classes, provide the full name of the selected Bias class. <b>Output Format:</b> Bias class: [Bias class Name]</p> <hr/> <p><b>Bias Check:</b> {step 5 result} If the article is 'Neutral', terminate the analysis. If it is 'Left-biased' or 'Right-biased', proceed to Step 7.</p> <hr/> <p><b>Narrative Classification:</b> {article} Since the article is {step_5_result} biased, it is likely to support a {step_5_result} narrative. Based on the analysis so far, write a one-line narrative that best represents the article's perspective. Narrative Classes – {Event-wise Narrative class} <b>Instructions:</b></p> <ul style="list-style-type: none"> <li>• Instead of numbering the classes, provide the full names of the selected narratives.</li> <li>• Only include the <b>subclasses</b> of Left, Right, and Both-Side Narratives in the output.</li> </ul> <p><b>Output Format:</b> Narrative Classes: [list of narrative names separated by commas]</p>
--

Table 10: Bias and Narrative Classification Prompts

- Cites solidarity tweets, global rallies, or NGO statements.
  - Frames global support as moral validation.
  - Eg: "Global icons from across continents have supported India's farmers.", "The world is watching as India's farmers fight for justice."
4. Accusing Media and Government of Manipulation: Articles claiming media and government misrepresent or suppress protest realities.
    - Mentions biased reporting or censorship.
    - Suggests narrative control or propaganda.
    - Eg: "The media is running government-fed lies about farmers' violence." "Independent journalists face harassment for showing ground truth."
- Pro-Govt Aligned Narratives (CAA):**
1. Glorification of CAA: Articles portraying the Citizenship Amendment Act as patriotic, humanitarian, or constitutionally sound.
    - Emphasizes moral duty or compassion.
    - Links CAA to India's civilizational ethos.
    - Eg: "CAA offers refuge to persecuted minorities - it's an act of compassion.", "This law strengthens India's humanitarian spirit."
  2. Glorification of the central government: Articles presenting the government as protector of India's sovereignty or integrity.
    - Frames CAA/NRC as defense against illegal immigration.
3. Vilification of the Opposition: Articles that morally discredit or delegitimize opposition parties, portraying them as politically motivated, anti-national, or obstructive to government initiatives.
    - Appeals to national unity or security.
    - Eg: "The government stands firm to protect India's borders and identity."
  4. Opposition Spreading Misinformation and Fear: Articles that cognitively discredit the opposition or media by accusing them of spreading falsehoods, misinterpretations, or panic regarding the CAA/NRC.
    - Frames the opposition as divisive or acting against national interest.
    - Uses morally charged language to question credibility or intent.
    - Eg: "Opposition leaders are opposing CAA only to appease vote banks.", "They are obstructing reforms meant for the welfare of persecuted minorities."
- Mentions deliberate fear-mongering or fake news campaigns.
  - Highlights ignorance or distortion of the law's provisions.
  - Frames unrest as a result of misinformation rather than a legitimate grievance.
  - Eg: "False propaganda is misleading citizens about CAA.", "Unfounded rumors are fueling

---

**Step: Prompt**

**Coarse Grain Classification:** You are a classification expert for Propaganda. You have to classify the given article into one or more of the following classes. Each class has a definition given.

- 'G1': Emotional manipulation – Statements that rely heavily on triggering emotional responses rather than presenting factual or logical arguments.
- 'G2': Distraction and Misdirection – Statements that aim to divert attention away from the central issue or argument by introducing irrelevant, distorted, or ambiguous information.
- 'G3': Simplistic thinking – Statements that simplify complex issues into overly simplistic narratives, either by reducing causes, forcing binary choices, or using authoritative statements without evidence.
- 'G4': Attack and Defamation – Statements that undermine credibility by targeting individuals/groups with negative portrayals.
- 'G5': Manipulation through Popularity and Authority – Persuasion relying on expert credibility or peer pressure.
- 'G6': Misrepresentation and Distortion – Altering perception of reality through straw man, exaggeration, minimization.
- 'G7': Undermining Trust – Introducing uncertainty, confusion, skepticism, eroding confidence in claims or sources.
- 'G0': No Narrative.

Read the article below:

**Article:** {article}

**Output Format:** Coarse-Grain Classes: [list of classes separated by commas]

**Coarse Grain Span Identification:** Read the article: {article}

My model predicted these classes: prediction. For all the classes present in this prediction, their definitions are given below. Read carefully.

{ Coarse Grain Class definition }

**NOTE:** Now extract spans ONLY for the following predicted classes: {prediction}

**Output Format:** Coarse-Grain spans: [list of spans]

**Fine Grain Classification:** You are an expert annotator.

**Article:** {article}

From the list fine grain list, determine which fine-grain propaganda techniques are present in the spans below and assign spans accordingly.

**Spans:** {spans}

**Instructions:**

- Match each span with the most appropriate fine-grain technique(s). All spans must be matched only with techniques from the list: fine grain list.
- If a span clearly matches multiple techniques, assign it to all relevant ones.
- Output strictly in the following format.
- Only if list have 'No Technique', then output: NaN

**Output format** – Techniques: [list of spans]

---

Table 11: Persuasive Technique Classification Prompts

- |  |   |
|--|---|
| <p><i>unnecessary fear among minorities."</i></p> <p>5. <u>Anti-CAA protests as Pre-Planned and Funded:</u> Articles alleging protests are orchestrated or foreign-funded to defame India.</p> <ul style="list-style-type: none"><li>• Mentions NGOs, conspiracies, or hidden motives.</li><li>• Frames unrest as part of an international agenda.</li><li>• Eg: <i>"The protests were planned months before the Act passed."</i>, <i>"Foreign agencies are behind the chaos."</i></li></ul> <p>6. <u>Framing Anti-CAA protests as Subversive or Misguided:</u> Articles portraying protesters as misinformed, anti-national, or extremist.</p> <ul style="list-style-type: none"><li>• Frames protests as threats to national unity.</li><li>• Suggests participants are misled or ignorant.</li><li>• Eg: <i>"protesters have no idea what the law even says."</i></li></ul> <p>7. <u>Delegitimization of Critics:</u> Articles suggesting critics of CAA are biased or communal.</p> <ul style="list-style-type: none"><li>• Attacks left-wing commentators or state-aligned voices.</li><li>• Calls out intolerance disguised as patriotism.</li><li>• Eg: <i>"Those protesters against CAA expose</i></li></ul> | <p><i>their deep-rooted prejudice."</i></p> <p><b>Pro-Opp Aligned Narratives (CAA):</b></p> <ol style="list-style-type: none"><li>1. <u>Vilification of CAA:</u> Articles framing the law as discriminatory, unconstitutional, or anti-minority.<ul style="list-style-type: none"><li>• Highlights exclusion of Muslims or erosion of secularism.</li><li>• Uses moral or legal condemnation.</li><li>• Eg: <i>"CAA legitimizes religious discrimination under the guise of compassion."</i>, <i>"This Act dismantles India's secular foundation."</i></li></ul></li><li>2. <u>Vilification of the Central Government:</u> Articles accusing the government of authoritarianism, repression, or communal politics.<ul style="list-style-type: none"><li>• Focuses on police brutality or suppression of dissent.</li><li>• Frames leadership as divisive or intolerant.</li><li>• Eg: <i>"The government treats peaceful protesters as enemies."</i></li></ul></li><li>3. <u>Glorifying Anti-CAA protesters:</u> Articles portraying protesters as brave defenders of democracy and justice.<ul style="list-style-type: none"><li>• Describes protests as peaceful or youth-led.</li><li>• Frames protesters as voices of conscience.</li></ul></li></ol> |
|--|---|

Persuasion Technique	Existing	Ours	guidelines
Assertion (Nwokwu and Johson, 2019)	✗	✓	<p>Involves presenting a confident statement as fact, often without supporting evidence. It aims to persuade through authority and certainty. The following are the Characteristics:</p> <ul style="list-style-type: none"> <li>• Confident Presentation: Delivers statements with conviction.</li> <li>• Lack of Evidence: Omits detailed support or reasoning.</li> <li>• Persuasive: Aims to convince through apparent certainty.</li> </ul> <p>Example: <i>"This new policy will solve all problems of farmers."</i></p>
Glittering Generalities (Virtue) (Bauer, 2024)	✗	✓	<p>Involve using vague, positive-sounding words or phrases to promote an idea or product. These terms are emotionally appealing but lack concrete details. The following are the Characteristics:</p> <ul style="list-style-type: none"> <li>• Positive Language: Uses appealing but imprecise terms.</li> <li>• Emotional Appeal: Aims to evoke positive feelings.</li> <li>• Lack of Substance: Offers no specific information or evidence.</li> </ul> <p>Example: <i>"He is our hope for a brighter future"</i></p>
Causal Oversimplification, Black-and-white Fallacy, Straw Man, Whataboutism, Reductio ad Hitlerum, Red Herring, Loaded Language, Name Calling and Labeling, Appeal to emotion/fear/prejudice, Slogans, Flag-Waving, Exaggeration and Minimization, Thought-terminating Cliches, Bandwagon, Smears, Obfuscation/intentional vagueness/confusion, Doubt, Appeal to Authority	✓	✓	<p>guidelines as per (Da San Martino et al., 2019)</p>

Table 12: Annotation guidelines for persuasive techniques.

- Eg: *"Women of Shaheen Bagh became the conscience of the nation."*
4. Framing Anti-CAA protesters as Victims: Articles emphasizing state violence or persecution against protesters.
- Focuses on casualties, injuries, or arrests.
  - Highlights moral injustice or fear.
  - Eg: *"protesters face jail for exercising free speech."*