

# COSMMIC: Comment-Sensitive Multimodal Multilingual Indian Corpus for Summarization and Headline Generation

Raghvendra Kumar<sup>1</sup>, S. A. Mohammed Salman<sup>2</sup>, Aryan Sahu<sup>3</sup>, Tridib Nandi<sup>4</sup>,  
Pragathi Y. P.<sup>5</sup>, Sriparna Saha<sup>1</sup>, Jose G. Moreno<sup>6</sup>

<sup>1</sup>Department of Computer Science and Engineering, IIT Patna, India

<sup>2</sup>Department of Metallurgical and Materials Engineering, NIT Tiruchirappalli, India

<sup>3</sup>Department of Computer Science and Information Systems, BITS Pilani Goa, India

<sup>4</sup>Department of Computer Science and Engineering, IIIT Vadodara, India

<sup>5</sup>Department of Computer Science and Engineering, BMSCE Bangalore, India

<sup>6</sup>Université de Toulouse, IRIT UMR 5505 CNRS, France

**For inquiries:** {raghvendra.kumar1004, samohammedsalman, ariansahu010103}@gmail.com

**For inquiries:** {tridib.nandi25062002, pragathiyp2352, sriparna.saha}@gmail.com

**For inquiries:** jose.moreno@irit.fr

## Abstract

Despite progress in comment-aware multimodal and multilingual summarization for English and Chinese, research in Indian languages remains limited. This study addresses this gap by introducing COSMMIC, a pioneering comment-sensitive multimodal, multilingual dataset featuring nine major Indian languages. COSMMIC comprises 4,959 article-image pairs and 24,484 reader comments, with ground-truth summaries available in all included languages. Our approach enhances summaries by integrating reader insights and feedback. We explore summarization and headline generation across four configurations: (1) using article text alone, (2) incorporating user comments, (3) utilizing images, and (4) combining text, comments, and images. To assess the dataset's effectiveness, we employ state-of-the-art language models such as LLaMA3 and GPT-4. We conduct a comprehensive study to evaluate different component combinations, including identifying supportive comments, filtering out noise using a dedicated comment classifier using IndicBERT, and extracting valuable insights from images with a multilingual CLIP-based classifier. This helps determine the most effective configurations for natural language generation (NLG) tasks. Unlike many existing datasets that are either text-only or lack user comments in multimodal settings, COSMMIC uniquely integrates text, images, and user feedback. This holistic approach bridges gaps in Indian language resources, advancing NLP research and fostering inclusivity. The resources are available at: <https://github.com/AaryanSahu/COSMMIC>.

## 1 Introduction

In today's digital era, online content seamlessly integrates text and visuals, with user comments enriching discussions and providing diverse perspectives. **Reader-Aware Comment-Based Summarization (RACBS)** harnesses these interactions to generate more relevant and comprehensive summaries by incorporating user feedback (Kumar et al., 2023). However, the development of multilingual, comment-based, multimodal datasets for Indian languages remains largely unexplored. For instance, a news article on the Indian general elections<sup>1</sup> often features images alongside a flood of reader comments in multiple languages. While the text delivers factual information and images offer visual context, the comments reflect public sentiment, debates, and interpretations, making them a valuable yet underutilized resource for enhancing summarization.

**Motivation for the dataset:** The challenge lies in analyzing such multifaceted content comprehensively. Current datasets often fail to integrate all these elements—text, images, and comments—particularly in the context of Indian languages. A dataset limited to text or images alone cannot capture the full spectrum of user engagement and content richness. Our dataset COSMMIC addresses this shortcoming by combining article text, images, and user comments in **nine major Indian languages: Bengali, Hindi, Gujarati, Marathi, Malayalam, Odia, Tamil, Telugu, and Kannada**. This integrated approach offers a more

<sup>1</sup>[https://en.wikipedia.org/wiki/2024\\_Indian\\_general\\_election](https://en.wikipedia.org/wiki/2024_Indian_general_election)

complete view of online content. It is essential for developing advanced NLP models that can process and understand the full scope of information in today’s digital media landscape.

**NLG tasks driven by the dataset presented:** Summarization and headline generation are key NLP challenges with applications in news aggregation, content curation, and information retrieval (Tan et al., 2017; Singh et al., 2021; Verma et al., 2023a; Hamborg et al., 2017; Khatter and Ahlawat, 2022; Gaci et al., 2023; Mahalakshmi and Fatima, 2022; Esteva et al., 2021). Summarization condenses text into concise, coherent outputs, while headline generation crafts engaging titles, both leveraging advanced language models (Zhang et al., 2024) and multimodal data (Krubiński and Pecina, 2024b) for improved accuracy and relevance.

Integrating images and user comments enriches these tasks, adding context to summaries and enhancing headline relevance (Li et al., 2017; Gao et al., 2019). Multimodal approaches (He et al., 2023; Qiao et al., 2022) combining text, visuals, and user interactions produce more informative outputs. Building on these advancements, we assess the impact of filtered comments and categorized images on NLG tasks. By integrating only informative inputs, our approach enhances coherence and relevance, with evaluations confirming its effectiveness in multilingual and multimodal RACBS. Our **contributions** enhance context-aware summaries and headlines, as detailed below.

*(1) We introduce “COSMMIC”, an innovative comment-sensitive multimodal, multilingual dataset having article-image pairs with reader comments covering nine major Indian languages: Bengali, Hindi, Gujarati, Marathi, Malayalam, Odia, Tamil, Telugu, and Kannada. This dataset represents a pioneering effort in this field.*

*(2) Human-written ground truth summaries are provided for all multilingual news articles within the dataset.*

*(3) We establish benchmarks for our dataset by evaluating two NLG tasks—headline generation and summarization—using state-of-the-art language models such as GPT-4 and LLaMA3.*

*(4) An extensive study is conducted to explore various configurations: article text alone, article text with comments, article text with relevant comments, article text with images, article text with images and comments, and article text with images and relevant comments.*

*(5) To advance this task, we propose methods for filtering noisy comments and classifying images in the Hindi dataset. Using IndicBERT (Kakwani et al., 2020), we develop a classifier to retain only informative comments and filter out noisy comments. Additionally, a threshold-based multilingual CLIP (Carlsson et al., 2022; Reimers and Gurevych, 2019) classifier categorizes images as supplementary (reinforcing the text) or complementary (providing additional context). To assess linguistic preservation in RACBS, we perform a back-translation study on the Hindi dataset, highlighting quality degradation and the need to retain the original language.*

## 2 Related Works

We explore multilingual, multimodal datasets with reader comments for NLG tasks, as Table 1 details their linguistic, modal, and structural differences, excluding English from Indian languages.

**Multilingual Datasets:** Scialom et al. (2020) introduced MLSUM, a large-scale dataset containing 1.5 million article-summary pairs in five languages, emphasizing cross-lingual biases. Hasan et al. (2021) developed XL-Sum, featuring 1 million pairs across 44 languages for abstractive summarization. Ladhak et al. (2020) presented WikiLingua, which includes 18 languages and image-based alignments for cross-lingual summarization. Hewapathirana et al. (2024) introduced M2DS, the first multilingual dataset for multi-document summarization (MDS) using BBC articles. Aralikatte et al. (2023) developed Vārta, a dataset designed for headline generation in 14 Indic languages and English. Kumar et al. (2022) introduced the Indic-NLG (abbreviated as INLG in Table 1) benchmark for NLG tasks across 11 Indic languages. Readers are also referred to notable works such as Urlana et al. (2023), Cao et al. (2020), Varab and Schluter (2021), Feng et al. (2022), and Gliwa et al. (2019).

**Multimodal Datasets:** In the realm of multimodal datasets, Zhu et al. (2018) introduced MSMO, a dataset for multimodal summarization that includes both text and image outputs and features a multimodal attention model. Verma et al. (2023b) developed M3LS, the largest multilingual multimodal summarization dataset, comprising over a million document-image pairs across 20 languages. Krubiński and Pecina (2024a) presented a unified encoder-decoder model for headline generation, which enhances M3LS by incorporating

Datasets:	MMSUM	M2DS	Varta	XLHT	INLG	MSMO	M3LS	RAMDS	RASG	COSMMIC
Multilingual	✓	✓	✓	✓	✓	✗	✓	✗	✗	✓
Multimodal	✓	✗	✗	✓	✗	✓	✓	✗	✗	✓
Comments	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓
Summaries	✓	✓	✗	✗	✓	✓	✓	✓	✓	✓
Headlines	✗	✗	✓	✓	✓	✓	✓	✓	✗	✓
# Indian Lang.	8	1	14	8	11	0	8	0	0	9

Table 1: Comparison of Datasets Based on Multilingualism, Multimodality, Comment Inclusion, Presence of Summaries and Headlines, and Coverage of Indian Languages.

additional images. Wang et al. (2022) released N24News, a dataset from the New York Times with both text and image information organized into 24 categories. Shohan et al. (2024) developed XL-HeadTags (shortened as XLHT in Table 1), a multilingual and multimodal dataset for headline and tag generation, using images and captions in 20 languages. Gotmare et al. (2024) proposed a machine-learning algorithm designed to generate news articles from geo-tagged images. Liang et al. (2023) introduced MM-Sum, a multilingual multimodal abstractive summarization dataset enhancing summary quality across 44 languages. Other noteworthy works include Kumar et al. (2024), Mita et al. (2023), Fu et al. (2021), Patil et al. (2024), Qiu et al. (2024), Nguyen et al. (2023), and Ghosh et al. (2024). **Limitations:** *Despite their significance, these datasets are text-only or multimodal without user comments.*

**Reader Comments Inclusive Datasets:** Zaidan and Callison-Burch (2011) pioneered the Topic-driven Reader Comments Summarization (Torcs) problem by creating a dataset of 1,005 Yahoo! News articles with over one million comments. Li et al. (2017) extended a variational autoencoder-based framework for reader-aware multi-document summarization (RA-MDS), presenting a new dataset with a detailed collection, annotation, and expert review process, demonstrating improved summarization performance with reader comments. Risch and Krestel (2020) introduced a dataset of journalist-reader dialogues from The Guardian, proposing a ranking task for comments. Gao et al. (2019) developed the Reader-Aware Summary Generator (RASG) using reader comments from Weibo to enhance summary relevance. Other noteworthy works include Zaidan and Callison-Burch (2011), Meng et al. (2024), Fujita et al. (2019), Roha et al. (2022), and Shi and Lam (2019). **Limitations:** *Despite their significance, these datasets lack focus on Indian languages, highlighting the need for our proposed dataset.*

### 3 COSMMIC Dataset

Our dataset is pioneering in the Indian context, featuring 4,959 article-image pairs and 24,484 reader comments, with comprehensive ground-truth summaries across all languages. It encompasses 500 to 650 articles for each of the nine major languages: *Bengali, Hindi, Gujarati, Marathi, Malayalam, Odia, Tamil, Telugu, and Kannada*. While the dataset might seem relatively small for each individual language, it offers a significant volume compared to the RAMDS benchmark dataset (Li et al., 2017), which includes 450 news articles and 10,000 comments across 45 topics. Our dataset provides broader coverage across nine languages, enhancing its utility. It also enabled us to manually verify each article to ensure that, after scraping and processing, the dataset is properly structured and free from any irrelevant information.

Additionally, with the rise of LLMs and their emphasis on zero-shot (Kojima et al., 2022) and few-shot learning (Brown et al., 2020), the size of our dataset is well-suited for fine-tuning these models. *Moreover, our dataset encompasses news exclusively from 2024. Extending further back proved impractical due to the **limited number of comments on older articles**. Even for articles from 2023 and 2022, there were only a few comments per 100 articles. Consequently, we concentrated on more recent news to ensure a robust dataset.*

#### 3.1 Dataset Curation

**Data Source:** We selected DailyHunt for data scraping due to its diverse, multilingual content and integration of text, images, and user comments, providing a rich multimodal dataset for tasks like summarization and headline generation. Its focus on regional languages and local news fills gaps in predominantly English-centric datasets, while its real-time updates enable applications such as trend analysis and fake news detection. Additionally, the platform’s vast articles and user interactions enable



Figure 1: Various components extracted from a news article on DailyHunt.

scalable research and model training, making it invaluable to our study.

**Scraping tools:** We used Selenium to scrape news articles from DailyHunt by navigating URLs, scrolling to load dynamic content, and collecting unique “Read more” links. For each article with comments, we extracted the headline, related images, main content (excluding disclaimers), and comments, saving the data to a CSV file. The process involved initializing the Chrome WebDriver, dynamically loading content, and using BeautifulSoup for HTML parsing. *Figure 1 illustrates the extracted components for a sample article.* The language-agnostic code was uniformly applied across nine languages. After scraping, we manually verified and cleaned the data, removing author and editor names (1-2% of cases) for consistency.

### 3.2 Statistics

We present a quantitative overview of our dataset. *Table 2* details key attributes per article, *Table 3* lists article and comment counts by language, and *Table 4* compares average word and sentence counts for headlines and summaries. Marathi has the highest comment volume and density, while Odia has the lowest. Bengali and Marathi lead in article counts, with Marathi, Bengali, and Hindi showing the most comment activity, reflecting higher user interaction.

### 3.3 Crafting Summaries

**Flexibility in Summary Composition:** Substantial variability in article features across different Indian languages is observed, as detailed in *Table 2*, underscoring the necessity of flexibility in summary composition. Imposing a rigid summary length, such as a fixed word count (Li et al., 2017), would not have accommodated the natural linguistic diversity and structural differences among these

languages. A one-size-fits-all approach would have risked compromising the quality and relevance of the summaries, failing to capture the true essence of the content. Therefore, flexibility in summary length and structure was crucial to accurately capture the unique characteristics and information density of each language, a principle central to our approach. **The summaries created were inclusive of user comments.** Noisy comments were manually removed (*in appendix A.8*).

**Annotator Guidelines:** A diligently selected team of ten annotators, each holding a master’s in linguistics for their respective languages, was engaged to ensure the highest standards in creating ground truth summaries. Each annotator was proficient in at least three languages, allowing for a profound understanding of linguistic sharpness and cultural contexts across the nine languages in the dataset. The process involved four annotators per language: two generated the summaries, with each summarizing half of the sample set, such as 50 out of 100 samples, while the remaining two conducted thorough reviews of all summaries. A summary was deemed final only when it received approval from both reviewing annotators, ensuring a robust quality assurance process. The guidelines mandated that summaries be concise yet comprehensive, capturing essential ideas and significant details while accommodating natural variability in length. Annotators were instructed to faithfully reflect the meaning and context of the original text, avoiding biases and preserving the tone and intent of the source material. Importantly, reader comments were excluded from consideration to prevent external influences. Adhering to these standards and employing expert annotators ensured the summaries were of high quality and preserved the dataset’s integrity. We compensated annotators at USD 1.20 per 10 summaries.

## 4 Summarization and Headline Generation

### 4.1 Methodology

While prior studies Ahuja et al. (2023, 2024); Kumar et al. (2022); Singh et al. (2024); Bhat et al. (2023) have benchmarked NLG for Indic languages using LLMs, none have explored reader-aware summarization or headline generation, a gap our work addresses.

**Configurations Used:** Our comprehensive studies explore various configurations to assess the



Languages →	Hindi			Bengali			Marathi		
Features ↓	Min	Max	Average	Min	Max	Average	Min	Max	Average
Token	118	4160	518	28	1962	339	36	2652	353
Words	100	3712	466	24	1661	288	31	2229	303
Unique Words	69	966	214	23	834	193	24	1079	202
Sentences	2	228	25	2	119	26	2	210	24
Comments	1	148	6	1	107	6	1	153	6
Languages →	Gujarati			Malayalam			Odia		
Features ↓	Min	Max	Average	Min	Max	Average	Min	Max	Average
Token	45	2254	410	31	2005	259	46	3058	281
Words	43	2027	362	27	1666	218	41	2511	242
Unique Words	40	1010	212	24	1186	172	34	1156	161
Sentences	1	137	23	1	186	23	2	297	20
Comments	1	34	3	1	82	6	1	25	2
Languages →	Tamil			Telugu			Kannada		
Features ↓	Min	Max	Average	Min	Max	Average	Min	Max	Average
Token	37	1882	289	48	1918	327	53	1641	258
Words	28	1403	239	40	1651	268	43	1402	215
Unique Words	24	866	171	38	1113	191	39	820	155
Sentences	2	137	20	1	245	28	2	189	19
Comments	1	93	6	1	103	4	1	62	5

Table 2: Statistical Summary of Feature Counts across Various Indian Languages in the COSMMIC Dataset. Note that the average values have been rounded to the nearest integer.

Language	# Articles	# Comments
Hindi	543	3372
Bengali	627	3742
Marathi	650	3870
Gujarati	537	1790
Malayalam	505	2833
Odia	577	1241
Tamil	505	3204
Telugu	505	1774
Kannada	510	2658
Total	4959	24484

Table 3: Distribution of Articles and Comments Across Indian Languages in the COSMMIC Dataset.

Attribute	Headlines		Summaries	
	# Words	# Sent.	# Words	# Sent.
Hindi	18	1	117	6
Bengali	12	1	86	6
Marathi	13	1	110	7
Gujarati	13	1	90	6
Malayalam	12	1	59	6
Odia	10	1	83	6
Tamil	11	2	77	6
Telugu	8	1	90	9
Kannada	10	1	66	5

Table 4: Average Word and Sentence Counts for Headlines and Summaries (rounded to the nearest integer).

COSMMIC dataset for headline generation and summarization, utilizing both GPT-4 and LLAMA-3 models. The configurations included (1) using only the article text, (2) combining the article text with comments, (3) integrating the article text with relevant comments, (4) incorporating images along with the article text, (5) merging article text with both images and comments, and (6) combining article text, images, and relevant comments. These configurations were rigorously tested to assess how different data combinations influence the models’ performance in generating accurate and contextually rich headlines and summaries across the various languages.

**Models Used:** We employed GPT-4 (Achiam et al., 2023) and LLAMA-3 (Dubey et al., 2024), state-of-the-art NLP models. GPT-4, a paid multimodal model, excels in processing text and images, making it ideal for our dataset. LLAMA-3, an open-source alternative with 8 billion parameters, offers strong language generation capabilities. Recent studies Katz et al. (2024); Waisberg et al. (2023) highlight GPT-4’s superiority in complex tasks, while our evaluation of LLAMA-3 provides insights into the performance of a lighter, cost-effective model for reader-aware summarization and headline generation.

**Evaluation Metrics:** We assessed NLG tasks using ROUGE (Lin, 2004) for n-gram overlap and BERTScore (Zhang et al., 2019) for semantic similarity, ensuring precise n-gram segmentation in Indic languages with the Polyglot tokenizer (Al-Rfou’ et al., 2013). *While these metrics may appear limited, they are the standard for state-of-the-art datasets (Table 1) and provide a rigorous benchmark. Their widespread adoption reinforces their reliability, making them well-suited for our dataset evaluation.*

## 4.2 Experiments, Results and Analysis

**Key Prompting Strategies** To enhance headline generation and summarization, we employ structured zero-shot prompts tailored for different multi-modal scenarios. For headline generation, prompts ensure factual, concise, and engaging outputs, whether integrating images with article text or incorporating user comments as supplementary context. Similarly, summarization prompts maintain linguistic consistency while prioritizing core information, ensuring balanced and insightful summaries. *Complete prompt structures and rationale are provided in the appendix A.2.*

### Analysis of Summarization with GPT-4

Table 5 shows that “Article + Images” generally provides the best summarization results across most Indian languages, improving both ROUGE and BERTScore metrics. Images offer valuable contextual support, especially in languages with complex scripts and structures, enhancing the understanding and interpretation of content. However, Telugu is an exception where “Article + Images + Rel. Com.” yields the highest scores, highlighting the importance of integrating relevant comments for improved summarization. This variation underscores the need for tailored summarization strategies that account for each language’s unique linguistic and contextual features. *The table further shows that incorporating **relevant comments enhances summarization quality**, as reflected in improved ROUGE and BERTScore metrics across most languages.* This reinforces the effectiveness of reader-aware summarization, where comments provide crucial context and sentiment, leading to more accurate summaries.

### Analysis of Headline Generation with GPT-4

While article-only inputs yield strong headlines, incorporating relevant comments enhances performance in certain languages by providing crucial

context. *Languages like Telugu and Kannada benefit significantly, demonstrating that **comments help bridge informational gaps and refine headlines**.* The minimal drop in BERTScore suggests that while images may not improve exact textual matches, they contribute valuable contextual insights. Overall, GPT-4 effectively optimizes NLG tasks across languages with diverse inputs.

### Summarization with LLama-3

Table 6 shows that “Only Article” input generally yields the best summarization scores with LLAMA-3. However, incorporating relevant comments significantly boosts performance for languages like Tamil and Telugu, as seen in improved BERTScore metrics. For Bengali and Kannada, including all comments enhances contextual richness, leading to superior BERTScore results. This demonstrates that while minimal input is effective, relevant comments improve summary quality, even for lightweight models.

### Headline Generation with LLama-3

While “Only Article” inputs typically perform best for headline generation, Table 6 shows that adding all comments improves BERTScore across languages. This suggests that while simpler inputs work well, richer contextual data enhances semantic relevance and alignment. Thus, LLama-3 benefits from minimal inputs but achieves better headline quality with comments. **Human analysis is included in the appendix A.7 due to space constraints.**

## 5 Classifying Comments & Images

As an initial benchmark, we focus on the Hindi section of our dataset due to annotation costs and feasibility. Tasks below aim to enhance data quality and assess the impact of linguistic transformations.

**Comment Classification:** To refine our dataset, we propose an automated approach to retain only informative comments while filtering out noisy ones in the Hindi section. Using IndicBERT (Kakwani et al., 2020), we develop a classifier to enhance the quality of reader-aware summarization. For training, our annotator manually labelled 250 comments into three categories: **Supporting** (aligned with the article), **Enriching** (adding new insights), and **Disconnected** (irrelevant to the article). We fine-tuned IndicBERT on this labelled data and then applied the trained model to classify the remaining comments. Once categorized, we conducted further experiments to analyze how comment clas-

Reader Aware Summarization with GPT-4												
Lang	Only Article				Article + Comments				Article + Rel. Com.			
	R1	R2	RL	BS	R1	R2	RL	BS	R1	R2	RL	BS
Hin.	44.18	18.94	33.47	76.61	50.98	24.82	38.34	78.22	52.13	25.52	39.43	78.71
Ben.	29.40	9.32	21.52	74.69	36.64	13.38	26.15	76.94	37.19	13.73	26.82	77.34
Mar.	27.24	10.49	25.79	74.28	40.38	18.56	37.97	78.02	41.50	19.47	39.03	78.56
Guj.	26.97	8.38	25.13	75.00	32.00	12.00	29.46	76.26	32.89	12.41	30.29	76.63
Mal.	17.38	5.45	16.55	72.78	18.88	6.16	17.90	73.49	20.18	<b>6.71</b>	19.11	74.01
Odia.	25.45	7.98	18.90	92.47	34.99	13.73	26.47	92.91	36.36	14.11	27.68	92.83
Tam.	25.07	9.52	23.60	74.05	32.63	14.32	30.67	77.20	34.94	15.52	32.73	77.83
Tel.	21.91	6.83	20.79	73.05	34.32	13.89	32.35	76.80	<b>36.19</b>	<b>15.07</b>	<b>33.93</b>	<b>77.75</b>
Kan.	23.94	8.04	22.70	74.83	28.85	10.68	27.08	75.89	29.10	10.79	28.01	76.01
Lang	Article + Images				Article + Images + Comments				Article + Images + Rel. Com.			
	R1	R2	RL	BS	R1	R2	RL	BS	R1	R2	RL	BS
Hin.	<b>53.41</b>	<b>26.10</b>	<b>40.41</b>	<b>79.38</b>	49.63	23.67	36.66	77.69	50.90	24.42	37.64	78.15
Ben.	<b>38.16</b>	<b>14.16</b>	<b>27.62</b>	<b>77.87</b>	34.90	12.53	24.34	76.49	36.12	13.20	25.63	76.89
Mar.	<b>42.83</b>	<b>20.11</b>	<b>40.38</b>	<b>79.03</b>	39.50	18.02	37.43	77.85	40.62	18.57	38.27	78.28
Guj.	<b>33.74</b>	<b>12.60</b>	<b>31.15</b>	<b>77.20</b>	31.06	11.15	28.66	76.17	31.78	11.52	29.38	76.30
Mal.	<b>20.27</b>	6.53	<b>19.22</b>	<b>74.15</b>	18.47	5.73	17.60	73.64	19.00	5.90	18.04	73.92
Odia.	<b>36.31</b>	<b>14.22</b>	<b>27.81</b>	<b>93.40</b>	33.06	12.36	24.59	91.86	34.60	13.21	25.82	92.29
Tam.	<b>35.87</b>	<b>15.94</b>	<b>33.78</b>	<b>78.32</b>	31.85	13.34	29.89	76.83	32.96	14.05	31.01	77.28
Tel.	35.33	14.03	33.17	77.54	32.61	12.51	30.61	76.52	33.45	13.15	31.56	76.82
Kan.	<b>30.15</b>	<b>11.12</b>	<b>28.28</b>	<b>76.55</b>	27.81	9.89	26.28	75.43	28.60	10.50	26.88	75.68
Reader Aware Headline Generation with GPT-4												
Lang	Only Article				Article + Comments				Article + Rel. Com.			
	R1	R2	RL	BS	R1	R2	RL	BS	R1	R2	RL	BS
Hin.	<b>32.23</b>	<b>9.58</b>	<b>24.54</b>	<b>72.97</b>	30.27	8.99	23.61	72.89	30.49	9.12	23.94	72.87
Ben.	<b>22.07</b>	4.96	17.27	72.63	21.01	4.80	17.14	72.66	21.27	<b>5.14</b>	<b>17.28</b>	<b>72.72</b>
Mar.	<b>21.18</b>	<b>6.51</b>	<b>18.33</b>	<b>72.83</b>	20.19	5.46	17.47	72.63	20.01	5.51	17.48	72.63
Guj.	<b>20.47</b>	<b>5.37</b>	<b>17.49</b>	73.67	19.81	4.86	17.13	<b>73.89</b>	20.10	5.26	17.42	73.84
Mal.	<b>20.41</b>	<b>5.53</b>	<b>17.48</b>	<b>72.19</b>	18.99	4.87	16.38	71.75	19.19	4.89	16.09	71.87
Odia.	18.23	4.80	15.55	90.03	19.09	<b>4.83</b>	16.14	<b>91.35</b>	<b>19.38</b>	4.74	<b>16.32</b>	91.27
Tam.	<b>24.63</b>	<b>6.66</b>	<b>22.13</b>	<b>72.59</b>	22.03	6.45	19.77	72.37	22.71	6.62	20.45	72.55
Tel.	<b>18.82</b>	4.51	<b>16.33</b>	70.61	17.97	<b>4.56</b>	15.90	70.63	18.11	4.55	16.10	<b>70.86</b>
Kan.	20.88	5.04	17.97	73.22	22.12	6.32	19.16	74.07	<b>22.87</b>	<b>6.54</b>	<b>19.78</b>	<b>74.12</b>
Lang	Article + Images				Article + Images + Comments				Article + Images + Rel. Com.			
	R1	R2	RL	BS	R1	R2	RL	BS	R1	R2	RL	BS
Hin.	28.56	8.31	22.44	72.34	29.14	8.53	22.88	72.29	28.69	8.37	22.64	72.38
Ben.	18.21	3.66	14.71	72.24	18.67	3.63	14.85	72.11	19.45	4.04	15.81	72.22
Mar.	18.18	4.49	15.59	72.26	19.03	5.02	16.34	72.25	19.07	4.67	16.45	72.30
Guj.	15.98	3.52	13.83	72.51	16.42	3.43	14.25	72.47	16.86	3.56	14.51	72.65
Mal.	14.80	2.78	12.92	70.40	15.18	3.45	13.02	70.37	15.41	3.53	13.25	70.43
Odia.	15.89	3.43	13.32	90.11	16.76	3.84	14.18	90.23	16.94	3.89	14.53	90.69
Tam.	18.93	4.57	17.18	71.81	20.03	4.36	17.99	71.56	20.31	4.74	18.42	71.80
Tel.	15.44	3.31	14.01	70.28	16.96	3.84	15.21	70.46	17.85	4.36	16.15	70.64
Kan.	19.46	4.35	16.95	72.92	19.27	4.63	17.04	73.08	19.62	5.05	17.37	73.28

Table 5: Performance of Summarization and Headline Generation using GPT-4 across different languages and input combinations. The table shows ROUGE-1 (R1), ROUGE-2 (R2), ROUGE-L (RL), and BERTScore (BS) metrics. Bold values indicate the highest performance metrics achieved for each language and input type, highlighting the most effective combination.

sification improves summarization performance, as shown in Table 7.

**Insight:** The results in Table 7 indicate that integrating enriching comments (OA+EC) yields

the highest improvement in summarization. Interestingly, supporting comments (OA+SC) also enhance performance but slightly less than enriching ones. Including disconnected comments (OA+DC)

Reader Aware Summarization with Llama-3												
Lang	Only Article				Article + Comments				Article + Rel. Com.			
	R1	R2	RL	BS	R1	R2	RL	BS	R1	R2	RL	BS
<b>Hin.</b>	<b>36.49</b>	<b>14.88</b>	<b>34.97</b>	<b>84.45</b>	31.27	11.95	29.89	83.63	34.35	12.76	33.13	84.26
<b>Ben.</b>	2.97	1.04	<b>2.91</b>	82.10	<b>2.98</b>	<b>1.14</b>	2.89	<b>82.12</b>	2.64	0.76	2.64	82.08
<b>Mar.</b>	19.96	<b>8.36</b>	<b>19.41</b>	<b>85.29</b>	18.13	4.19	11.09	79.38	<b>20.05</b>	8.01	19.39	85.21
<b>Guj.</b>	<b>29.71</b>	<b>11.54</b>	<b>28.70</b>	<b>93.48</b>	29.44	11.53	28.45	93.69	27.89	9.82	26.92	93.46
<b>Mal.</b>	<b>22.04</b>	<b>7.60</b>	<b>21.65</b>	<b>93.60</b>	20.60	6.66	20.10	93.38	20.76	7.36	20.28	93.50
<b>Odia.</b>	<b>1.96</b>	<b>0.96</b>	<b>1.96</b>	<b>92.18</b>	2.03	0.70	2.03	91.77	0.12	0.03	0.12	91.50
<b>Tam.</b>	33.61	12.35	32.51	86.91	32.41	13.11	31.23	86.93	<b>34.01</b>	<b>13.72</b>	<b>33.19</b>	<b>87.01</b>
<b>Tel.</b>	<b>26.21</b>	<b>11.37</b>	<b>25.49</b>	92.01	22.11	10.3	21.59	91.89	21.94	9.88	21.43	<b>92.25</b>
<b>Kan.</b>	<b>25.34</b>	<b>8.81</b>	<b>24.36</b>	93.29	23.98	8.71	23.47	<b>93.46</b>	23.57	8.36	22.86	93.38

Reader Aware Headline Generation with Llama-3												
Lang	Only Article				Article + Comments				Article + Rel. Com.			
	R1	R2	RL	BS	R1	R2	RL	BS	R1	R2	RL	BS
<b>Hin.</b>	<b>7.76</b>	<b>0.86</b>	<b>7.59</b>	<b>76.60</b>	5.29	0.42	5.19	76.35	4.20	0.33	4.08	76.11
<b>Ben.</b>	<b>0.59</b>	<b>0.23</b>	<b>0.59</b>	<b>73.51</b>	0.58	0.16	0.58	72.87	0.25	0.07	0.25	72.50
<b>Mar.</b>	<b>4.25</b>	<b>0.41</b>	<b>4.15</b>	<b>76.03</b>	2.97	0.15	2.89	75.47	2.57	0.23	2.57	75.03
<b>Guj.</b>	<b>4.57</b>	<b>1.08</b>	<b>4.51</b>	82.94	4.56	0.91	4.50	<b>82.95</b>	4.09	0.59	4.09	82.49
<b>Mal.</b>	<b>3.56</b>	<b>0.43</b>	<b>3.56</b>	81.89	2.79	0.20	2.79	<b>82.34</b>	2.62	0.06	2.52	76.31
<b>Odia.</b>	<b>0.76</b>	0.23	<b>0.76</b>	82.63	0.65	0.03	0.65	<b>82.91</b>	0.62	<b>0.24</b>	0.62	81.93
<b>Tam.</b>	5.34	<b>0.99</b>	5.26	75.85	<b>6.36</b>	0.63	<b>6.36</b>	<b>75.88</b>	4.24	0.34	4.14	75.71
<b>Tel.</b>	<b>3.94</b>	<b>0.25</b>	<b>3.84</b>	75.40	3.73	<b>0.25</b>	3.63	75.84	2.62	0.06	2.52	<b>76.31</b>
<b>Kan.</b>	<b>4.83</b>	<b>0.76</b>	<b>4.83</b>	79.37	3.36	0.47	3.36	<b>79.96</b>	0.13	0.04	0.13	79.20

Table 6: Performance of Summarization and Headline Generation using LLAMA-3 across languages and input types, showing ROUGE-1 (R1), ROUGE-2 (R2), ROUGE-L (RL), and BERTScore (BS) metrics. Bold values highlight the best results.

significantly deteriorates performance, highlighting the importance of filtering relevant contextual information. This trend continues for headline generation, where **OA+EC** outperforms other configurations, reinforcing the benefit of additional insights over mere agreement. *A detailed discussion is furthermore provided in the appendix A.1 due to space constraints.*

**Image Classification:** We use a threshold-based classifier with multilingual CLIP (sentence-transformers/clip-ViT-B-32-multilingual-v1) (Carls-son et al., 2022; Reimers and Gurevych, 2019) to assess image relevance in multimodal summarization (Jangra et al., 2021). Since images accompany articles from a reputable media house, classification into unrelated categories is unnecessary. Instead, we classify images as *supplementary* (reinforcing the text) or *complementary* (providing additional context). CLIPScore is computed using both the ground truth summary and the image, as well as the headline and the image. The average score determines classification: above 0.25 indicates reinforcement, while lower values suggest added context.

The 0.25 threshold is empirically set based on CLIPScore distributions in our dataset, ensuring balanced classification. Prior work Jangra et al. (2021) shows real-world text-image pairs often have low CLIPScores, making higher thresholds (e.g., 0.5) impractical. This approach effectively differentiates between reinforcing and contextual images. We observed that 36.84% of images reinforced the article (*supplementary*), while most (63.16%) added extra information (*complementary*). This suggests that Hindi news articles are often accompanied by images that provide background information rather than directly supporting the text. As a result, multimodal summarization models must focus on interpreting contextual visuals rather than relying solely on direct text-image alignment (**further strengthening the findings of Table 5**). Likewise, headline generation may gain from complementary images for a broader context beyond the main points.

*Remark: For a deeper understanding, additional discussion and a detailed back-translation study focusing on preserving linguistic authenticity are provided in the appendix (see A.9 and A.10). These sections explore the nuances of language preserva-*



Config	R1	R2	RL	BS	Config	R1	R2	RL	BS
<b>OA</b>	44.18	18.94	33.47	76.61	<b>OA</b>	<b>32.23</b>	<b>9.58</b>	<b>24.54</b>	72.97
<b>OA+SC</b>	56.64	28.33	42.59	82.32	<b>OA+SC</b>	31.68	8.98	24.02	71.53
<b>OA+EC</b>	<b>58.22</b>	<b>29.19</b>	<b>43.93</b>	<b>84.15</b>	<b>OA+EC</b>	32.06	9.52	23.89	<b>73.44</b>
<b>OA+DC</b>	35.62	20.44	29.85	70.44	<b>OA+DC</b>	26.54	7.11	20.48	69.85
<b>OA+AC</b>	50.98	24.82	38.34	78.22	<b>OA+AC</b>	30.27	8.99	23.61	72.89
<b>OA+RC (GPT)</b>	52.13	25.52	39.43	78.71	<b>OA+RC (GPT)</b>	30.49	9.12	23.94	72.87

Table 7: Performance of different configurations for **Summarization (left)** and **Headline Generation (right)** using GPT-4. Configurations: **OA** = Only Article, **SC** = Supporting Comments, **EC** = Enriching Comments, **DC** = Disconnected Comments, **AC** = All Comments, **RC (GPT)** = Relevant Comments selected by GPT.

tion and offer insights into the effectiveness of our approach.

## 6 Conclusion

In summary, the **COSMMIC** dataset is a notable advancement in multimodal, multilingual, and comment-sensitive resources for Indian languages. It integrates text, images, and reader comments across nine major languages, addressing gaps in existing datasets. This dataset enhances headline generation and summarization while offering insights into regional language dynamics and user interactions. Our ablation study demonstrates the effectiveness of different component combinations, making **COSMMIC** a valuable resource.

## Limitations

Although the dataset encompasses nine major Indian languages, it does not represent all regional dialects and languages within India. Also, the dataset’s reliance on publicly available sources from DailyHunt means that the coverage may be skewed towards specific topics or viewpoints prevalent on the platform.

**Ethics Statement** Our dataset, including URLs to original articles and images, content, ground truth summaries, and headlines, is publicly available on GitHub: <https://github.com/AaryanSahu/COSMMIC>, along with the scraping and experimental code. The scraping process deliberately excludes any user information to protect privacy.

## Acknowledgments

Raghvendra Kumar extends appreciation to the Prime Minister’s Research Fellowship (PMRF) for its assistance and successful completion of this research. Sriparna Saha recognizes the support received from the “Technology Innovation Hub, Vishlesan I-Hub Foundation IIT Patna” (Project No.

TIH/CSE/ASMO/05) for the achievement of this research. Aryan Sahu sincerely thanks the India Chapter of ACM SIGKDD for the opportunity to participate in the Uplink Internship program, which played a pivotal role in the successful completion of this research work.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Kabir Ahuja, Harshita Diddee, Rishav Hada, Millicent Ochieng, Krithika Ramesh, Prachi Jain, Akshay Nambi, Tanuja Ganu, Sameer Segal, Mohamed Ahmed, et al. 2023. Mega: Multilingual evaluation of generative ai. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4232–4267.
- Sanchit Ahuja, Divyanshu Aggarwal, Varun Gumma, Ishaan Watts, Ashutosh Sathe, Millicent Ochieng, Rishav Hada, Prachi Jain, Mohamed Ahmed, Kalika Bali, et al. 2024. Megaverse: Benchmarking large language models across languages, modalities, models and tasks. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2598–2637.
- Rami Al-Rfou’, Bryan Perozzi, and Steven Skiena. 2013. *Polyglot: Distributed word representations for multilingual NLP*. In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 183–192, Sofia, Bulgaria. Association for Computational Linguistics.
- Rahul Aralikkatte, Ziling Cheng, Sumanth Doddapaneni, and Jackie Chi Kit Cheung. 2023. Varta: A large-scale headline-generation dataset for indic languages. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3468–3492.
- Savita Bhat, Vasudeva Varma, and Niranjan Pedanekar. 2023. *Generative models for Indic languages: Evaluating content generation capabilities*. In *Proceed-*

- ings of the 14th International Conference on Recent Advances in Natural Language Processing*, pages 187–195, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Yue Cao, Xiaojun Wan, Jinge Yao, and Dian Yu. 2020. Multisumm: towards a unified model for multi-lingual abstractive summarization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11–18.
- Fredrik Carlsson, Philipp Eisen, Faton Rekathati, and Magnus Sahlgren. 2022. Cross-lingual and multilingual clip. In *Proceedings of the thirteenth language resources and evaluation conference*, pages 6848–6854.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Andre Esteva, Anuprit Kale, Romain Paulus, Kazuma Hashimoto, Wenpeng Yin, Dragomir Radev, and Richard Socher. 2021. Covid-19 information retrieval with deep-learning based semantic search, question answering, and abstractive summarization. *NPJ digital medicine*, 4(1):68.
- Xiachong Feng, Xiaocheng Feng, and Bing Qin. 2022. Msamsum: Towards benchmarking multi-lingual dialogue summarization. In *Proceedings of the Second DialDoc Workshop on Document-grounded Dialogue and Conversational Question Answering*, pages 1–12.
- Xiyan Fu, Jun Wang, and Zhenglu Yang. 2021. **MM-AVS: A full-scale dataset for multi-modal summarization**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5922–5926, Online. Association for Computational Linguistics.
- Soichiro Fujita, Hayato Kobayashi, and Manabu Okumura. 2019. **Dataset creation for ranking constructive news comments**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2619–2626, Florence, Italy. Association for Computational Linguistics.
- Yacine Gaci, Boualem Benatallah, Fabio Casati, and Khalid Benabdeslem. 2023. Targeting the source: Selective data curation for debiasing nlp models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 276–294. Springer.
- Shen Gao, Xiuying Chen, Piji Li, Zhaochun Ren, Li-dong Bing, Dongyan Zhao, and Rui Yan. 2019. Abstractive text summarization by incorporating reader comments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6399–6406.
- Akash Ghosh, Arkadeep Acharya, Raghav Jain, Sriparna Saha, Aman Chadha, and Setu Sinha. 2024. Clipsyntel: clip and llm synergy for multimodal question summarization in healthcare. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 22031–22039.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. Samsam corpus: A human-annotated dialogue dataset for abstractive summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79.
- Abhay Gotmare, Gandharva Thite, and Laxmi Bewoor. 2024. A multimodal machine learning approach to generate news articles from geo-tagged images. *International Journal of Electrical & Computer Engineering* (2088-8708), 14(3).
- Felix Hamborg, Norman Meuschke, and Bela Gipp. 2017. Matrix-based news aggregation: exploring different news perspectives. In *2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, pages 1–10. IEEE.
- Tahmid Hasan, Abhik Bhattacharjee, Md Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M Sohel Rahman, and Rifat Shahriyar. 2021. Xl-sum: Large-scale multilingual abstractive summarization for 44 languages. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703.
- Bo He, Jun Wang, Jieli Qiu, Trung Bui, Abhinav Shrivastava, and Zhaowen Wang. 2023. Align and attend: Multimodal summarization with dual contrastive losses. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14867–14878.
- Kushan Hewapathirana, Nisansa de Silva, and CD Athuraliya. 2024. M2ds: Multilingual dataset for multi-document summarisation. *arXiv preprint arXiv:2407.12336*.
- Anubhav Jangra, Sriparna Saha, Adam Jatowt, and Mohammed Hasanuzzaman. 2021. Multi-modal supplementary-complementary summarization using multi-objective optimization. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 818–828.
- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, NC Gokul, Avik Bhattacharyya, Mitesh M Khapra, and Pratyush Kumar. 2020. Indicnlp suite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for indian

- languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4948–4961.
- Daniel Martin Katz, Michael James Bommarito, Shang Gao, and Pablo Arredondo. 2024. Gpt-4 passes the bar exam. *Philosophical Transactions of the Royal Society A*, 382(2270):20230254.
- Harsh Khatter and Anil Kumar Ahlawat. 2022. Content curation algorithm on blog posts using hybrid computing. *Multimedia Tools and Applications*, 81(6):7589–7609.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Mateusz Krubiński and Pavel Pecina. 2024a. [Towards unified uni- and multi-modal news headline generation](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 437–450, St. Julian’s, Malta. Association for Computational Linguistics.
- Mateusz Krubiński and Pavel Pecina. 2024b. Towards unified uni-and multi-modal news headline generation. In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 437–450.
- Aman Kumar, Himani Shrotriya, Prachi Sahu, Amogh Mishra, Raj Dabre, Ratish Puduppully, Anoop Kunchukuttan, Mitesh M. Khapra, and Pratyush Kumar. 2022. [IndicNLG benchmark: Multilingual datasets for diverse NLG tasks in Indic languages](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5363–5394, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Raghvendra Kumar, Ratul Chakraborty, Abhishek Tiwari, Sriparna Saha, and Naveen Saini. 2023. [Diving into a sea of opinions: Multi-modal abstractive summarization with comment sensitivity](#). In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM ’23*, page 1117–1126, New York, NY, USA. Association for Computing Machinery.
- Ritesh Kumar, Shyam Ratan, Siddharth Singh, Enakshi Nandi, Laishram Niranjana Devi, Akash Bhagat, Yogesh Dawer, Bornini Lahiri, and Akanksha Bansal. 2024. A multilingual, multimodal dataset of aggression and bias: the comma dataset. *Language Resources and Evaluation*, 58(2):757–837.
- Faisal Ladhak, Esin Durmus, Claire Cardie, and Kathleen Mckeown. 2020. Wikilingua: A new benchmark dataset for cross-lingual abstractive summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4034–4048.
- Piji Li, Lidong Bing, and Wai Lam. 2017. Reader-aware multi-document summarization: An enhanced model and the first dataset. In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 91–99.
- Yunlong Liang, Fandong Meng, Jinan Xu, Jiaan Wang, Yufeng Chen, and Jie Zhou. 2023. Summary-oriented vision modeling for multimodal abstractive summarization. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- P Mahalakshmi and N Sabiyath Fatima. 2022. Summarization of text and image captioning in information retrieval using deep learning techniques. *IEEE Access*, 10:18289–18297.
- Zixiang Meng, Qiang Gao, Di Guo, Yunlong Li, Bobo Li, Hao Fei, Shengqiong Wu, Fei Li, Chong Teng, and Donghong Ji. 2024. [Mmlscu: A dataset for multi-modal multi-domain live streaming comment understanding](#). In *Proceedings of the ACM on Web Conference 2024, WWW ’24*, page 4395–4406, New York, NY, USA. Association for Computing Machinery.
- Masato Mita, Soichiro Murakami, Akihiko Kato, and Peinan Zhang. 2023. Camera: A multimodal dataset and benchmark for ad text generation. *arXiv preprint arXiv:2309.12030*.
- Laura Nguyen, Thomas Scialom, Benjamin Piwowarski, and Jacopo Staiano. 2023. Loralay: A multilingual and multimodal dataset for long range and layout-aware summarization. In *The 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2023)*.
- Vaidehi Patil, Leonardo Ribeiro, Mengwen Liu, Mohit Bansal, and Markus Dreyer. 2024. Refinesumm: Self-refining mllm for generating a multimodal summarization dataset. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13773–13786.
- Lingfeng Qiao, Chen Wu, Ye Liu, Haoyuan Peng, Di Yin, and Bo Ren. 2022. Grafting pre-trained models for multimodal headline generation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 244–253.
- Jielin Qiu, Jiacheng Zhu, William Han, Aditesh Kumar, Karthik Mittal, Claire Jin, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Ding Zhao, et al. 2024. Mm-sum: A dataset for multimodal summarization and thumbnail generation of videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21909–21921.



- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Julian Risch and Ralf Krestel. 2020. [A dataset of journalists’ interactions with their readership: When should article authors reply to reader comments?](#) In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management, CIKM ’20*, page 3117–3124, New York, NY, USA. Association for Computing Machinery.
- Vishal Singh Roha, Naveen Saini, Sriparna Saha, and Jose G Moreno. 2022. Unsupervised framework for comment-based multi-document extractive summarization. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 574–582.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2020. Mlsum: The multilingual summarization corpus. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8051–8067.
- Bei Shi and Wai Lam. 2019. [Reader comment digest through latent event facets and news specificity](#). *IEEE Transactions on Knowledge and Data Engineering*, 31(8):1581–1594.
- Faisal Tareque Shohan, Mir Tafseer Nayeem, Samsul Islam, Abu Ubaida Akash, and Shafiq Joty. 2024. Xl-headtags: Leveraging multimodal retrieval augmentation for the multilingual generation of news headlines and tags. *arXiv preprint arXiv:2406.03776*.
- Harman Singh, Nitish Gupta, Shikhar Bharadwaj, Dinesh Tewari, and Partha Talukdar. 2024. Indicgenbench: A multilingual benchmark to evaluate generation capabilities of llms in indic languages. *arXiv preprint arXiv:2404.16816*.
- Rajeev Kumar Singh, Sonia Khetarpaul, Rohan Gorantla, and Sai Giridhar Allada. 2021. Sheg: summarization and headline generation of news articles using deep learning. *Neural Computing and Applications*, 33:3251–3265.
- Jiwei Tan, Xiaojun Wan, and Jianguo Xiao. 2017. From neural sentence summarization to headline generation: A coarse-to-fine approach. In *IJCAI*, volume 17, pages 4109–4115.
- Ashok Urlana, Pinzhen Chen, Zheng Zhao, Shay Cohen, Manish Shrivastava, and Barry Haddow. 2023. [PMIndiaSum: Multilingual and cross-lingual headline summarization for languages in India](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11606–11628, Singapore. Association for Computational Linguistics.
- Daniel Varab and Natalie Schluter. 2021. Massivesumm: a very large-scale, very multilingual, news summarisation dataset. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10150–10161.
- Prince Verma, Kiran Jyoti, Varsha Sahni, S Jafar Ali Ibrahim, and NS Kalyan Chakravarthy. 2023a. Synthesized feature learning model on news aggregator for chatbot. *Soft Computing*, pages 1–14.
- Yash Verma, Anubhav Jangra, Raghvendra Kumar, and Sriparna Saha. 2023b. Large scale multi-lingual multi-modal summarization dataset. *arXiv preprint arXiv:2302.06560*.
- Ethan Waisberg, Joshua Ong, Mouayad Masalkhi, Sharif Amit Kamran, Nasif Zaman, Prithul Sarker, Andrew G Lee, and Alireza Tavakkoli. 2023. Gpt-4: a new era of artificial intelligence in medicine. *Irish Journal of Medical Science (1971-)*, 192(6):3197–3200.
- Zhen Wang, Xu Shan, Xiangxie Zhang, and Jie Yang. 2022. [N24News: A new dataset for multimodal news classification](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6768–6775, Marseille, France. European Language Resources Association.
- Omar F. Zaidan and Chris Callison-Burch. 2011. [The Arabic online commentary dataset: an annotated dataset of informal Arabic with high dialectal content](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 37–41, Portland, Oregon, USA. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen Mckeown, and Tatsunori B Hashimoto. 2024. Benchmarking large language models for news summarization. *Transactions of the Association for Computational Linguistics*, 11:39–57.
- Junnan Zhu, Haoran Li, Tianshang Liu, Yu Zhou, Jijun Zhang, and Chengqing Zong. 2018. [MSMO: Multimodal summarization with multimodal output](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4154–4164, Brussels, Belgium. Association for Computational Linguistics.

## A Appendix

### A.1 Detailed Analysis of Comment Classification Impact

The results in Table 7 provide key insights into how different types of comments influence summariza-



tion and headline generation quality. Below, we analyze these findings in detail.

### A.1.1 Impact on Summarization

1. **Only Article (OA)** serves as the baseline, achieving an **R1 score of 44.18** and a **BERTScore (BS) of 76.61**. This reflects the model’s ability to generate summaries without any external context.

2. **Adding Supporting Comments (OA+SC)** boosts performance significantly (**R1: 56.64, BS: 82.32**). These comments align with the article’s content and reinforce key information, leading to improved coherence and factual consistency.

3. **Enriching Comments (OA+EC)** yield the highest scores, achieving **R1: 58.22** and **BS: 84.15**. Unlike supporting comments, these introduce new insights that enhance the richness of generated summaries.

4. **Disconnected Comments (OA+DC)** degrade performance dramatically (**R1: 35.62, BS: 70.44**). The presence of unrelated information confuses the model, resulting in incoherent or misleading summaries.

5. **All Comments (OA+AC)** moderately improve results (**R1: 50.98, BS: 78.22**) but underperform compared to selective filtering. This suggests that a mix of useful and irrelevant comments dilutes the benefits of contextual enrichment.

6. **Relevant Comments (OA+RC) by GPT** provide a slight edge over OA+AC (**R1: 52.13, BS: 78.71**), indicating that LLM-based filtering can help, but our classifier-based selection performs better in identifying enriching content.

### A.1.2 Impact on Headline Generation

The trends observed in summarization largely hold for headline generation:

1. **OA achieves R1: 32.23, BS: 72.97**, serving as the baseline.

2. **OA+EC provides the highest BERTScore (BS: 73.44)**, showing that adding new insights aids headline generation.

3. **OA+SC performs slightly worse than OA+EC (BS: 71.53)** but still improves over the baseline.

4. **OA+DC leads to a major drop (BS: 69.85)**, reaffirming the negative impact of irrelevant content.

5. **OA+AC and OA+RC** show marginal improvements over the baseline, but the classifier-filtered enriching comments outperform all.

### A.1.3 Key Takeaways

These findings demonstrate that **not all comments are equally beneficial**. Enriching comments contribute the most to both summarization and headline generation, while disconnected comments negatively impact performance. Automated classification, as performed in this study, proves essential for optimizing results.

## A.2 Detailed Prompting Strategies

**Prompts Used:** We present the zero-shot prompts used in two headline generation scenarios: (i) integrating images with article text and (ii) incorporating user comments alongside the article, as illustrated in Figure 2. The first prompt ensures that headlines effectively capture key textual and visual content, maintaining clarity, conciseness, and factual accuracy. The second prompt integrates relevant user comments while preserving the article’s primary message, subtly reflecting public sentiment without distorting factual integrity.

For summarization, the prompt shown in Figure 3 ensures concise and coherent summaries by distilling essential information. When incorporating images and comments, the prompt balances factual accuracy with supplementary insights from comments, maintaining a focus on the article’s core message. All prompts follow expert guidelines to optimize headline generation and summarization.

Figures 2 and 3 illustrate these prompt designs in detail.

### A.3 Summarization using GPT

#### Utilization of Comments and Images for summarization:

In the provided example summary as shown in figure 4 generated using GPT-4-turbo, the first paragraph primarily draws information from the article. In contrast, the second paragraph incorporates details about the weather in Florida impacting the Pakistan match, which can be inferred from the accompanying image. This highlights the context and background information effectively. The final paragraph captures user opinions from comments expressed in Hindi, suggesting that teams should be awarded total points if the match is cancelled, thereby adding valuable insights. These additional elements significantly enhance the overall quality of the summary.

Article + Images	Article + Relevant Comments
<p><i>"role": "system", "content": "You are an expert news editor. Your primary task is to create headlines that accurately reflect the main content of news articles and images."</i></p> <p><i>"role": "user", "content": f""Craft a precise and engaging news headline that primarily captures the main point of the article and image. Follow these guidelines:</i></p> <ol style="list-style-type: none"> <li><i>1. Focus predominantly (about 80-90%) on the key information and main message from the article and image.</i></li> <li><i>2. Keep the headline concise, ideally under 15 words.</i></li> <li><i>3. Use active voice and clear, impactful language.</i></li> <li><i>4. Ensure the headline is factual and directly related to the article's main content and the key elements of the image.</i></li> </ol>	<p><i>"role": "system", "content": "You are an expert news editor. Your primary task is to create headlines that accurately reflect the main content of news articles. User comments should be considered as supplementary information only."</i></p> <p><i>"role": "user", "content": f""Craft a precise and engaging news headline in {l} that primarily summarizes the main point of the article. Use insights from user comments only as secondary, supplementary information. Follow these guidelines:</i></p> <ol style="list-style-type: none"> <li><i>1. Focus predominantly (about 80-90%) on the key information and main message from the article.</i></li> <li><i>2. If relevant, subtly incorporate a minor element (about 10-20%) from the user comments to reflect public sentiment or add context.</i></li> <li><i>3. Keep the headline concise, ideally under 15 words.</i></li> <li><i>4. Use active voice and clear, impactful language.</i></li> <li><i>5. Ensure the headline is factual and directly related to the article's main content.</i></li> <li><i>6. Avoid overemphasizing or sensationalizing elements from user comments.</i></li> </ol>

Figure 2: Headline generation prompts: (a) integrating article text with images and (b) using relevant comments with the article text.

Only Article	Article + Image + Comments
<p><i>"role": "system", "content": "Your task is to summarize the text in the same language as the text."</i></p> <p><i>"role": "user", "content": f""Generate a summary using the contents. The summary should be concise and cover all the important key points as conveyed by the original content and should have minimal noise. The generated summary should be in the same language as the original content."</i></p>	<p><i>"role": "system", "content": "You are an expert news editor. Your primary task is to create summaries that accurately reflect the main content of news articles and images. User comments should be considered as supplementary information only."</i></p> <p><i>"role": "user", "content": f""Craft a precise and engaging news summary in {language} that primarily summarizes the main point of the article and image. Use insights from user comments only as secondary, supplementary information. Follow these guidelines:</i></p> <ol style="list-style-type: none"> <li><i>1. Focus predominantly (about 80-90%) on the key information and main message from the article and image.</i></li> <li><i>2. If relevant, subtly incorporate a minor element (about 10-20%) from the user comments to reflect public sentiment or add context.</i></li> <li><i>3. Use active voice and clear, impactful language.</i></li> <li><i>4. Ensure the summary is factual and directly related to the article's main content and the key elements of the image.</i></li> <li><i>5. Avoid overemphasizing or sensationalizing elements from user comments.</i></li> <li><i>6. If an image is provided, balance the summary to reflect both textual and visual information when applicable.</i></li> </ol>

Figure 3: Summarization prompts: (a) only article text (b) using image and comments with the article text.

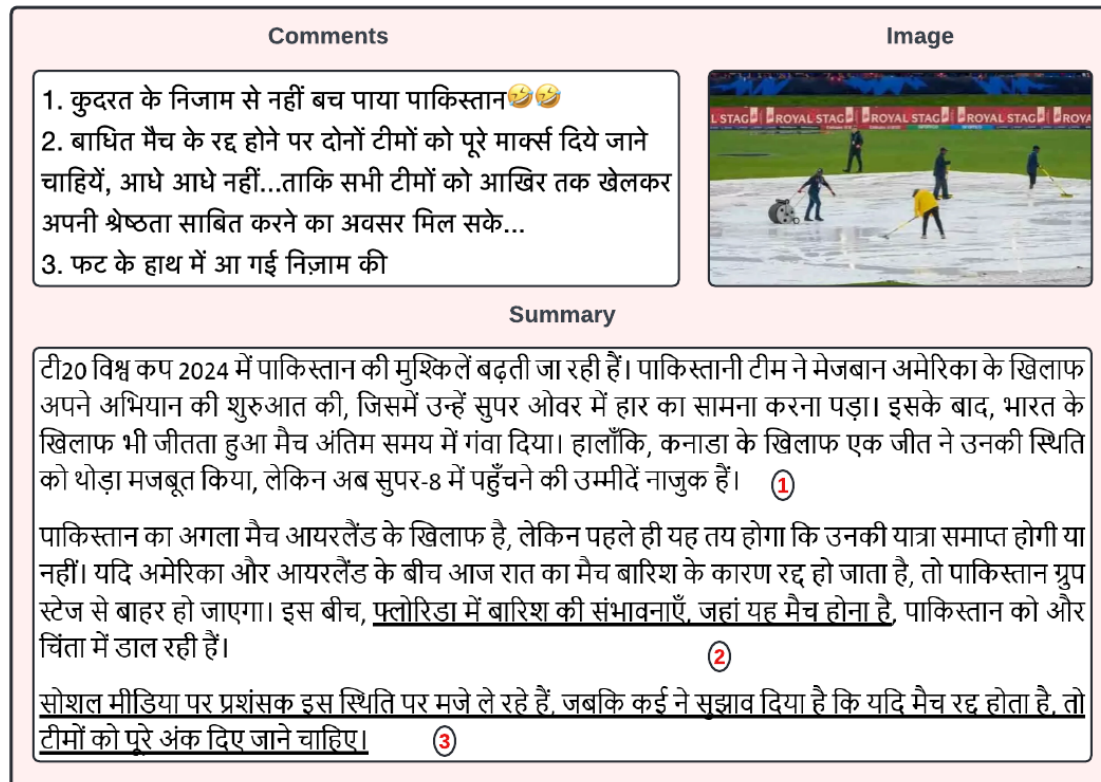


Figure 4: Demonstrates the significance of comments and images in the summary generated by GPT-4-turbo

#### A.4 Challenges faced during Summarization:

1. During our summarization process, we observed minor inconsistencies in the language used in the generated summaries. Despite instructing the model to produce summaries in the original content language, 1-2% of summaries were still generated in English. This issue was comparatively pronounced with the LLAMA model, where not only did the number of English summaries slightly increase, but very few summaries were also generated in incorrect languages. To address this issue, we focused on iteratively rerunning the summarization process for only those articles with erroneous summaries through which we could correct the errors.
2. Another challenge we encountered with the LLAMA model was the generation of incomplete summaries. About 0.5-1% of the summaries were reduced to a single line, failing to capture the full essence of the original content. To mitigate this, we resorted to rerunning the summarization code for those rows with incomplete summaries.

3. Some LLaMA-generated summaries included words combining English and the original language, rather than using only the original language. This issue appears to stem from the lightweight nature of the LLaMA model, as no such problems were encountered when using GPT. Consequently, these mixed-language words in summaries led to noticeable errors, as highlighted in Figure 5.

The accurate summaries generated by the LLAMA model (Figure 5) can be compared with the GPT-4-turbo outputs (Figure 4) to recognize that the GPT-4-turbo summaries were more consistent and coherent and encountered fewer challenges while developing the summaries. This was expected as the LLAMA model was relatively lightweight compared to GPT-4-turbo.

#### A.5 Headline Generation using GPT

The hyperparameters used while generating the headline were:

- **Max output token:** 100
- **Model used:** gpt-4-turbo
- **Temperature:** 0.7

Llama3 model	
Comments	Comments
<p>1. कुदरत के निजाम से नहीं बच पाया पाकिस्तान 🤔</p> <p>2. बाधित मैच के रद्द होने पर दोनों टीमों को पूरे मार्क्स दिये जाने चाहियें, आधे आधे नहीं...ताकि सभी टीमों को आखिर तक खेलकर अपनी श्रेष्ठता साबित करने का अवसर मिल सके...</p> <p>2. फट के हाथ में आ गई निज़ाम की</p>	<p>1. जब सरस्वती का जगत में ऐसे अपमान होगा तो ऐसे ही करेंगे धर्म और हम टूटेंगे हाथ पैर।</p> <p>2. विद्या को एक साधारण मनुष्य से तौलने का अहंकार मन में आया कैसे।</p> <p>3. ऐसे ढोंग कृत्य को भगवान भी नाराज हो जाते हैं</p> <p>4. बहुत ही बे हूदा बातें सुनने में आ रही हैं हमारे जमाने में काफी किताबें मां सरस्वती का रूप माना जाता है हम लोग पुजते हैं , और ये नया जमाना सरस्वती जी के साथ अपना वज़न तौला जा रहा है कैसे लोग सठिया गये हैं।</p>
Good Summary	Erraneous Summary
<p>पाकिस्तानी टीम टी20 वर्ल्ड कप 2024 के ग्रुप स्टेज से ही बाहर होने की कगार पर है। टीम ने पहले मैच में USA के हार्थी सुपर ओवर में हारा था। इसके बाद उसने भारत के खिलाफ जीता हुआ मैच अंतिम समय में गंवा दिया। टीम ने अपना अगला मैच कनाडा के खिलाफ जीता है, लेकिन उसके लिए सुपर-8 में पहुंचने की उम्मीदों पर संकट के बादल मंडरा रहे हैं। पाकिस्तानी टीम का अगला मैच आयरलैंड से होना है, लेकिन अगर आज का मैच रद्द हो जाता है तो टीम ग्रुप स्टेज से ही टूर्नामेंट से बाहर हो जाएगी। सुपर-8 में पहुंचने की संभावना अब मिट चुकी है, लेकिन अगर आज का मैच बारिश से रद्द हो जाता है तो USA और आयरलैंड दोनों को 1-1 अंक मिल जाएंगे और पाकिस्तान की टीम टूर्नामेंट से बाहर हो जाएगी।</p>	<p>महाराष्ट्र के बुलढाना जिले के सांसद और पहली बार मंत्री बने शिवसेना शिंदे गुट के प्रतानराव जाधव ने अपना वजन तौलने का एक नाटकीय स्वरूप प्रदान किया. मंत्री ने तराजू से अपना वजन तौलकर गरीब बच्चों में किताबें बांटी जानी थीं. इस दौरान तराजू के एक पलड़े में स्थान लेने की कोशिश करने पर मंत्रीजी ने पलड़े का एक हिस्सा टूट होने के बावजूद कोई चोट नहीं झेली. कुछ लोगों ने इस घटना को संस्कृति के अपमान के रूप में देखा और भगवान से श-lined.</p>

Figure 5: Summaries generated by the LLAMA model. Additionally, the good and erroneous summaries are compared

### A.5.1 Challenges Faced During Headline Generation

- Several iterations of prompt refinement were done in order to come up with the best prompt.

### A.5.2 Pre-processing of the Data Included:

1. **Removal of redundant data:** During the data collection process, redundant information such as the number of likes, shares, and other extraneous text displayed on the website was inadvertently scraped alongside the actual news articles. To ensure data quality, we meticulously processed the scraped content, filtering out these unnecessary elements to retain only the relevant news articles in the final dataset.
2. **Removal of noise characters:** During the scraping process, extraneous characters—likely resulting from the dynamic nature of the website—were occasionally captured alongside the actual news article content. To maintain the integrity of the dataset, we carefully removed these irrelevant characters, ensuring that only the authentic article content was retained.
3. **Incomplete data:** At times, only partial news articles were scraped—such as individual paragraphs or sentences—rather than the complete article. This issue arose due to irregular spacing and indentation on the website, combined with the limitations of Selenium in capturing full content under these condi-

tions. To address this, we manually scraped the complete articles in such cases to ensure the dataset was comprehensive and accurate.

### A.6 Headline Comparison

#### A.6.1 Headline (w/o comments,images) as shown in the Figure 6.

Original Headline: USA vs IRE:पाकिस्तान के खिलाफ कुदरत का निजाम, एक झटके में 3 टीमों का काम तमाम  
GPT Headline: टी20 वर्ल्ड कप 2024: पाकिस्तान ग्रुप स्टेज से बाहर होने के कगार पर, बारिश और उलटफेर से बड़ी मुश्किलें

Figure 6: Headline (w/o comments,images).

- **Good Aspects:** The GPT headline provides a relevant and clear update on Pakistan's situation in the T20 World Cup, highlighting the impact of rain and upsets. It accurately reflects the severity of the situation.
- **Bad Aspects:** The GPT headline does not capture the specific match between the USA and Ireland or the broader dramatic impact on multiple teams as described in the original headline. It focuses more on the overall tournament context rather than the immediate consequences of the rain.

#### A.6.2 Headline with relevant comments as shown in the Figure 7.

- **Good Aspects:** The GPT headline provides a clear and relevant summary of the consequences of rain on a cricket match and its impact on Pakistan's chances. It is accurate



Original Headline: USA vs IRE:पाकिस्तान के खिलाफ कुदरत का निजाम, एक झटके में 3 टीमों का काम तमाम  
 GPT Headline: बारिश के कारण भारत और आयरलैंड का मैच प्रभावित हुआ, जिससे पाकिस्तान के सुपर-8 में पहुंचने की संभावना घट गई।

Figure 7: Headline with relevant comments.

in terms of the consequences but misses the original context.

- **Bad Aspects:** The GPT headline incorrectly names the teams involved, which is a significant deviation from the original headline. Additionally, it does not capture the broader, dramatic impact on multiple teams as described in the original.

### A.6.3 Prompting Detail for Article and all comments inclusive:

After experimenting with several prompts to ensure accurate, robust summaries, the prompt given in Fig 8 below was finalized. The prompt has been crafted so that the model captures all the essential details of the article and the relevant user comments in the summary generated.

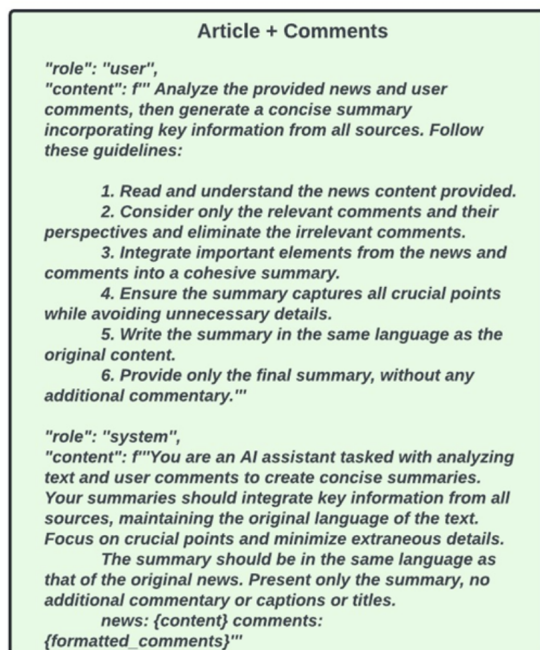


Figure 8: Prompts for Article + Comments

### A.7 Human Evaluation

To complement automated evaluation metrics like ROUGE and BERTScore, we conducted human assessments with 10 participants (postgraduate students) to evaluate the quality of summaries produced using the COSMMIC dataset across different

configurations for 10 articles in the Hindi language and their corresponding summaries. The evaluators rated the outputs based on key criteria, including grammatical accuracy (Gr-Co), relevancy, and overall summary quality (SQ), using a 1-to-5 scale. These evaluations encompassed both text-only and text-with-images scenarios, assessing the quality of summaries in two contexts: how effectively they conveyed information when paired with the corresponding article image, and how they performed when presented without any visual support. To minimize evaluator bias and enhance reliability, the scores were averaged. This thorough evaluation underscores the advantages of configurations that leverage richer input modalities, such as user comments and images. The evaluations were involved in summarization and headline generation for both the models as shown in Tables 8, 9, 10, 11.

### A.8 Filtering and Considering Relevant Comments

**Manual Filtering Process:** To ensure that user comments included in the summaries enhanced content quality and relevance, a *manual filtering process* was employed, carried out by the same team of expert annotators responsible for summary creation. This process required an additional layer of effort and compensation, ensuring that only relevant comments were included while noisy or extraneous comments were excluded.

**Defining Noisy Comments:** Noisy comments were identified as those that detracted from the article's central themes or introduced irrelevant, redundant, or biased information. Examples included:

- **Spam-like entries:** Comments promoting unrelated products, services, or containing external links.
- **Generic phrases:** Statements such as "Nice article" or "Interesting post," which lacked substantive value.
- **Hostile or inappropriate remarks:** Overly aggressive or offensive comments not contributing meaningfully to the discussion.
- **Unrelated tangents:** Comments entirely disconnected from the topic being discussed.

**Steps in Manual Filtering:** The team of annotators undertook the following steps to manually filter comments:

Configuration	Gr-Co	Relevance	SQ	
			WO-IM	W-IM
Article Text Only	3.5	3.6	3.4	3.7
Article Text + Comments	3.7	3.9	3.6	3.9
Text + Relevant Comments	3.8	4.0	3.8	4.1
Text + Images	3.9	4.1	4.0	4.3
Text + Images + Comments	<b>4.0</b>	<b>4.2</b>	<b>4.1</b>	<b>4.4</b>

Table 8: Human evaluation scores for various metrics of GPT-based model using COSMMIC configurations. Gr-Co means grammatical correctness, SQ means summary quality (With Images (W-IM) and Without Images (WO-IM)).

Configuration	Gr-Co	Relevance	HQ	
			WO-IM	W-IM
Article Text Only	3.3	3.2	3.4	3.7
Article Text + Comments	3.2	3.4	3.5	3.8
Text + Relevant Comments	3.4	3.5	3.6	4.0
Text + Images	3.3	3.4	3.7	4.1
Text + Images + Comments	<b>3.5</b>	<b>3.6</b>	<b>3.9</b>	<b>4.2</b>

Table 9: Human evaluation scores for various metrics of GPT-based model using COSMMIC configurations for headline generation. Gr-Co means grammatical correctness, HQ means headline quality (With Images (W-IM) and Without Images (WO-IM)).

Configuration	Gr-Co	Relevance	SQ	
			WO-IM	W-IM
Article Text Only	2.3	2.4	2.2	2.5
Article Text + Comments	2.5	2.6	2.4	2.7
Text + Relevant Comments	<b>2.6</b>	<b>2.8</b>	<b>2.9</b>	<b>3.1</b>

Table 10: Human evaluation scores for various metrics of LLaMA 3-based model using COSMMIC configurations. Gr-Co means grammatical correctness, SQ means summary quality.

Configuration	Gr-Co	Relevance	HQ	
			WO-IM	W-IM
Article Text Only	2.1	2.2	2.1	2.3
Article Text + Comments	2.3	2.4	2.2	2.5
Text + Relevant Comments	<b>2.4</b>	<b>2.6</b>	<b>2.5</b>	<b>2.6</b>

Table 11: Human evaluation scores for various metrics of LLaMA 3-based model using COSMMIC configurations for headline generation. Gr-Co means grammatical correctness, HQ means headline quality.

1. **Initial Screening:** Annotators carefully read through the entire set of comments associated with each article to assess their relevance.
2. **Relevance Assessment:** Each comment was evaluated on its alignment with the article’s content and its potential to provide contextual insights. Annotators considered whether the comment added value to the discussion or reinforced key ideas presented in the article.
3. **Exclusion of Noisy Comments:** Spam, generic, or off-topic comments were removed to ensure that the retained comments maintained the summaries’ coherence and rele-

vance.

4. **Final Review:** A second pass was conducted by another annotator to confirm that no valuable comments were mistakenly filtered out.

#### **Guidelines for Including Relevant Comments:**

The manual filtering process was guided by clear criteria for retaining comments that contributed to the quality of the summaries:

- *Content Relevance:* Comments directly related to the article’s subject matter were prioritized.
- *Contextual Depth:* Comments that elaborated

Setting	R1	R2	RL	BS
<b>Original (Only Article)</b>	44.18	18.94	33.47	76.61
<b>Translated (Hindi → English)</b>	35.52	12.03	25.91	69.82
<b>Back-Translated (Hindi)</b>	21.79	4.51	15.62	60.23
<b>Original (Article + Comments)</b>	50.98	24.82	38.34	78.22
<b>Translated (Hindi → English)</b>	40.21	15.49	29.12	71.31
<b>Back-Translated (Hindi)</b>	26.29	6.81	18.39	62.48
<b>Original (Article + Relevant Comments)</b>	52.13	25.52	39.43	78.71
<b>Translated (Hindi → English)</b>	41.49	16.01	30.18	72.10
<b>Back-Translated (Hindi)</b>	27.01	7.22	19.51	63.02

Table 12: Summarization Performance Across Different Translation Settings

on or clarified points in the article were included.

- *Collective Sentiment:* Comments reflecting collective user opinions or emphasizing key elements of the article were considered valuable for summarization.

**Quality Assurance and Compensation:** The annotators’ thorough approach to manual filtering ensured that the summaries accurately represented the articles while benefiting from the inclusion of relevant user comments. This additional task required substantial effort, and annotators were compensated at an additional rate of **USD 1.20 per 100 comments** for their work, recognizing the meticulous nature of the filtering process.

**Outcome:** This manual filtering process ensured that the final summaries captured the essence of the original articles while incorporating meaningful user contributions. By excluding noisy comments, the integrity and relevance of the dataset were preserved, resulting in summaries that were both concise and enriched by user perspectives.

### A.9 Back-Translation Study on Hindi Dataset for Summarization

To evaluate the impact of machine translation on linguistic fidelity in Reader-Aware Comment-Based Summarization (RACBS), we conducted a back-translation study on the Hindi dataset. This experiment aimed to determine whether translating Indian language content into English for processing and then back into the original language affects the quality of summarization and headline generation.

For this study, we first translated the Hindi articles, their corresponding comments, and the ground-truth summaries and headlines into English using GPT-4o. Once translated, we ran all

our summarization and headline-generation experiments on the English dataset. Following this, the outputs were back-translated into Hindi, and the same experimental pipeline was applied to the back-translated dataset.

The results on the back-translated Hindi dataset showed a noticeable drop in performance compared to the original Hindi dataset. This degradation highlights the loss of linguistic nuances and contextual accuracy introduced by machine translation. Such findings reinforce our claim that preserving the original language is crucial for maintaining the quality of Indian language summarization and headline generation.

Moreover, incorporating machine translation workflows would require extensive manual verification across multiple languages, which is resource-intensive and currently beyond our scope. This experiment further justifies our decision to exclude machine translation-based approaches and focus on direct processing in Indian languages for this resource paper.

We evaluated three settings: (i) **Original Hindi data**, (ii) **Translated Hindi-to-English data**, and (iii) **Back-Translated data** where the English outputs were retranslated into Hindi. The summarization performance was measured using ROUGE-1 (R1), ROUGE-2 (R2), ROUGE-L (RL), and BERTScore (BS). The results are summarized in Table 12.

#### A.9.1 Analysis

From Table 12, it is evident that the **original Hindi dataset consistently achieves the highest performance** across all metrics and settings. The **ROUGE and BERTScore values decline significantly when content is translated into English and further degrade when back-translated into**

Setting	R1	R2	RL	BS
<b>Original (Only Article)</b>	32.23	9.58	24.54	72.97
<b>Translated (Hindi → English)</b>	22.01	4.02	15.49	65.31
<b>Back-Translated (Hindi)</b>	10.39	1.21	7.32	55.08
<b>Original (Article + Comments)</b>	30.27	8.99	23.61	72.89
<b>Translated (Hindi → English)</b>	21.02	3.51	14.69	64.52
<b>Back-Translated (Hindi)</b>	9.79	0.91	6.81	54.02
<b>Original (Article + Relevant Comments)</b>	30.49	9.12	23.94	72.87
<b>Translated (Hindi → English)</b>	21.50	3.81	14.99	64.81
<b>Back-Translated (Hindi)</b>	10.11	1.11	7.12	54.59

Table 13: Headline Generation Performance Across Different Translation Settings

**Hindi**, reinforcing the importance of preserving the original language.

For summarization based only on the article, the original Hindi dataset achieves an **R1 score of 44.18**, whereas the translated version drops to **35.52**, and the back-translated version further declines to **21.79**. The degradation in **R2 and RL scores follows a similar pattern**, demonstrating that information loss occurs due to translation.

When incorporating **comments** into summarization, performance improves across all settings. The original Hindi dataset (Article + Comments) achieves an **R1 of 50.98**, which is significantly higher than both the translated (**40.21**) and back-translated (**26.29**) versions. Similarly, the **Article + Relevant Comments** setting yields the best results, with an **R1 of 52.13** in the original Hindi dataset, compared to **41.49** (translated) and **27.01** (back-translated).

These findings emphasize that **machine translation introduces considerable degradation in summarization quality**, likely due to the loss of linguistic nuances and contextual cues specific to Hindi. The back-translated results, performing the worst in all cases, further confirm that relying on translation workflows can negatively impact summarization tasks.

This experiment strongly supports our decision to **focus on direct summarization in Indian languages rather than incorporating machine translation workflows**, as doing so would compromise linguistic accuracy and overall performance.

#### A.10 Back-Translation Study on Hindi Dataset for Headline Generation

To further assess the impact of machine translation, we conducted a similar back-translation study on the **headline generation** task. This experiment fol-

lowed the same methodology as summarization: (i) processing the **original Hindi dataset**, (ii) translating it into **English for experimentation**, and (iii) **back-translating** the English results into Hindi to examine performance degradation.

Table 13 presents the results across different settings, evaluated using ROUGE-1 (R1), ROUGE-2 (R2), ROUGE-L (RL), and BERTScore (BS).

##### A.10.1 Analysis

The results in Table 13 further reinforce the observations made for summarization. The **original Hindi dataset consistently achieves the best performance**, while both **translation to English and back-translation to Hindi lead to severe performance degradation**.

For headline generation using only the article, the original Hindi dataset achieves an **R1 score of 32.23**, but this drops to **22.01** after translation into English and further declines to **10.39** after back-translation. The same trend is observed in **R2 and RL scores**, confirming a significant loss of key information and linguistic quality during translation.

When comments are included, headline generation performance improves slightly, similar to the summarization task. However, even with additional context, **back-translated Hindi results remain substantially weaker**. For instance, in the **Article + Relevant Comments** setting, the **original Hindi dataset achieves an R1 of 30.49**, which is much higher than both the translated (**21.50**) and back-translated (**10.11**) versions.

These findings indicate that **headline generation is even more sensitive to translation degradation than summarization**. This is expected, as headlines require **high linguistic precision and contextual awareness**, which are often distorted



through translation workflows.

Overall, these results strongly validate our claim that **machine translation is not a viable preprocessing step for our task**. Instead, our dataset should be leveraged directly in its original language to ensure linguistic integrity and maintain the highest possible performance across summarization and headline generation tasks.

#### A.11 An Intriguingly Intuitive Question

**Question:** If the headline and summary are written prior to the publication of an article, the scenario of comment-sensitive headline/summary generation can appear non-intuitive. How can the comment-sensitive headline or summary of an article be generated without relying on user comments?

**Answer:** It is important to recognize the dynamic nature of modern digital media, where articles, summaries, and headlines are often updated in response to user feedback, emerging developments, or shifts in public sentiment. This dynamic environment introduces challenges for comment-sensitive generation, as summaries and headlines must adapt to an ever-evolving flow of information, including real-time user interactions. Therefore, a forward-looking strategy could involve creating standalone summaries of relevant user comments (if present) for each article and storing them as part of a dynamic feedback repository. These stored summaries could then be utilized when publishing related follow-up articles, enabling the generation of headlines and summaries that reflect updated contexts, user sentiment, or emerging trends. By leveraging previously collected user feedback, this method would facilitate the integration of comment sensitivity without requiring real-time comment data for every article. Such a strategy enhances scalability and allows the summarization process to remain adaptable over time, addressing the complexities of a constantly changing digital media landscape. This highlights the potential for innovative comment-sensitive headline and summary generation that balances real-time relevance with long-term adaptability.