

HateMirage: An Explainable Multi-Dimensional Dataset for Decoding Faux Hate and Subtle Online Abuse

Sai Kartheek Reddy Kasu^{1*}, Shankar Biradar², Sunil Saumya¹, Md. Shad Akhtar³

¹Indian Institute of Information Technology Dharwad

²Manipal Institute of Technology, Manipal Academy of Higher Education Manipal, Karnataka, India

³Indraprastha Institute of Information Technology Delhi

saikartheekreddykasu@gmail.com, shankar.biradar@manipal.edu,

sunil.saumya@iiitdwd.ac.in, shad.akhtar@iiitd.ac.in

Abstract

Subtle and indirect hate speech remains an underexplored challenge in online safety research, particularly when harmful intent is embedded within misleading or manipulative narratives. Existing hate speech datasets primarily capture overt toxicity, underrepresenting the nuanced ways misinformation can incite or normalize hate. To address this gap, we present *HateMirage*, a novel dataset of Faux Hate comments designed to advance reasoning and explainability research on hate emerging from fake or distorted narratives. The dataset was constructed by identifying widely debunked misinformation claims from fact-checking sources and tracing related YouTube discussions, resulting in 4,530 user comments. Each comment is annotated along three interpretable dimensions: Target (who is affected), Intent (the underlying motivation or goal behind the comment), and Implication (its potential social impact). Unlike prior explainability datasets such as HateXplain and HARE, which offer token-level or single-dimensional reasoning, HateMirage introduces a multi-dimensional explanation framework that captures the interplay between misinformation, harm, and social consequence. We benchmark multiple open-source language models on HateMirage using ROUGE-L F1 and Sentence-BERT similarity to assess explanation coherence. Results suggest that explanation quality may depend more on pretraining diversity and reasoning-oriented data rather than on model scale alone. By coupling misinformation reasoning with harm attribution, HateMirage establishes a new benchmark for interpretable hate detection and responsible AI research.

Keywords: Faux Hate; Hate Speech; Misinformation; Explainable NLP; Social Media Safety; Benchmark Dataset

1. Introduction

This paper contains examples of faux hate comments, included solely for research purposes. Some readers may find the content offensive or distressing. Reader discretion is advised.

Online platforms such as YouTube, Facebook, and X (formerly Twitter) have become central spaces for communication and opinion exchange but also amplify harmful content, including hate speech and misinformation (Schmidt and Wiegand, 2017; Fortuna and Nunes, 2018). Despite advances in moderation, subtle and indirect hate, where hostility is masked through humor, irony, or fabricated narratives, remains difficult to detect. Such content often avoids explicit slurs or profanity, making it ambiguous even for human annotators and challenging for automated systems.

Existing hate speech datasets, including HateXplain (Mathew et al., 2021), Hostility Detection (Bhardwaj et al., 2020), studies on geographic and cultural bias (Tonneau et al., 2024), and SemEval series corpora (Zampieri et al., 2020; Premjith et al., 2024), have enabled progress in classification and target identification. However, they largely focus on overt hate and surface-level toxicity cues. This lim-

its model interpretability and fails to capture forms of hate that are manufactured, fueled, or triggered by misinformation.

Building on the challenges of detecting subtle, context-dependent harm, recent work has identified *Faux Hate* (Biradar et al., 2025): hate speech expressed indirectly through misinformation or deceptive narratives. Unlike traditional hate, which uses overt hostility, Faux Hate embeds prejudice within misleading claims, making its harmful intent implicit. For example, a comment may falsely accuse a community of spreading disease, conveying resentment while hiding the premise's falsity. Detecting such content requires contextual reasoning and background knowledge, which makes automated detection difficult. For instance, Table 1 contrasts traditional hatespeeches, which relies on direct attacks, with Faux Hate, which embeds prejudice in misleading or fabricated narratives, highlighting the need for contextual reasoning to interpret it effectively.

Prior work (Biradar et al., 2024, 2025) mainly focused on the classification of Faux Hate, leaving the task of explaining its reasoning and potential harm largely unexplored. This limitation hinders a deeper understanding of how misinformation shapes and triggers hate. To address this gap,

*Work done at FLaMe.nlp Lab, IIIT Delhi.

Typical Hateful Comment

- I dislike the skin colour of people from <country>.
- <country> is dishonest toward everyone; we should boycott it.
- <group/religion> people are ruining our society; they should leave.

Faux Hate Comments

- <country> is spreading COVID on purpose to harm the world.
 - It's not love, it's a plan; <group/religion> use relationships to change people.
 - The virus came from <country> labs; they engineered it to destroy us.
-

Table 1: Illustrative examples comparing Traditional Hate comments with Faux Hate comments. Placeholders such as COUNTRY and GROUP/RELIGION denote target entities abstractly.

we introduce *HateMirage*¹, a dataset for reasoning and explainability in misinformation-driven hate. It contains 4,530 comments collected using a two-step process: we first gathered widely debunked claims from fact-checking websites (e.g., *AltNews*², *FactChecker*³), then identified YouTube videos related to these claims and extracted associated comments. Each comment is annotated with Target, Intent, and Implication, capturing the reasoning that links deceptive narratives to harm.

We benchmarked several open-source language models ranging from 1B to 8B parameters for the structured explanation generation task. The *HateMirage* dataset serves as a unified testbed, providing a consistent setup to evaluate model reasoning across the Target, Intent, and Implication dimensions. Each model was evaluated in two settings: zero-shot and RAG-based. Performance was measured using Sentence-BERT similarity and ROUGE-L F1 to assess semantic and lexical coherence. The analysis highlights variations across model architectures and settings, reflecting the complexity of reasoning required for multi-dimensional explanation generation.

HateMirage thus makes implicit harm explicit, supporting research on interpretable NLP, explanation generation, and context-aware moderation (Mathew et al., 2021; Yang et al., 2023). By exposing the logic connecting misinformation to prejudice, it enables models and analysts to justify predictions, which is critical for trustworthy content moderation.

¹The term reflects the illusory nature of Faux Hate, where hateful expression appears genuine but is in fact shaped or amplified by fake narratives.

²<https://www.altnews.in/>

³<https://www.factchecker.in/>

Contributions: The key contributions of this work are summarized as follows:

- **HateMirage Dataset:** We present *HateMirage*, a dataset of 4,530 Faux Hate comments, annotated with multi-dimensional explanations across Target, Intent, and Implication. These annotations capture the implicit narratives and reasoning underlying misinformation-driven hate speech, enabling interpretable analysis beyond surface-level classification.
- **Model Benchmarking:** We evaluate a range of small and large language models on the task of explanation generation for Faux Hate, establishing initial baselines and highlighting the challenges of structured reasoning in this domain.
- **Research Testbed for Explainable NLP:** By explicitly linking misinformation cues to hateful discourse, *HateMirage* provides a valuable resource for studying reasoning, interpretability, and responsible AI in online communities.

Reproducibility: Prompt templates, annotation guidelines, RAG configuration, source code, and supporting fabricated claims are publicly available at <https://github.com/Sai-Kartheek-Reddy/HateMirage>. The *HateMirage* dataset is available for academic and research purposes only upon acceptance of the data usage agreement ([link](#)).

2. Literature Review

Current literature on online harmful content has largely treated hate speech and misinformation as separate domains, focusing on developing models and resources to detect either hate or fake content individually. However, the intersection of these two phenomena, where users propagate hateful comments in the context of widely debunked claims, remains largely underexplored (Biradar et al., 2025; Mosleh et al., 2024). Understanding this intersection is crucial, as it captures nuanced forms of social media manipulation and the complex interplay between intent, target, and potential consequences. The *HateMirage* dataset is designed to address this gap by providing multi-dimensional annotations that reflect both hate and misinformation in a unified framework.

Hate Speech: While several prior datasets have contributed significantly to hate speech and misinformation research, they remain limited in scope with respect to structured explanatory annotations as shown in Table 2. *HateXplain* (Mathew et al., 2021) focuses purely on hate speech and includes token-level rationales highlighting which parts of a comment are hateful, yet it does not explain why

Dataset	Hate	Fake	Explanation	Interdisciplinary
HateXplain (Mathew et al., 2021)	✓	✗	✗	✗
HatEval (Basile et al., 2019)	✓	✗	✗	✗
Implicit Hate (ElSherief et al., 2021)	✓	✗	✓	✗
HARE (Yang et al., 2023)	✓	✗	✓	✗
CoAID (Cui and Lee, 2020)	✗	✓	✗	✗
FNVE (Chen et al., 2025)	✗	✓	✓	✗
Hostile (Bhardwaj et al., 2020)	✓	✓	✗	✗
Deceptive Humor (Kasu et al., 2025)	✓	✓	✗	✓
HateMirage (Proposed)	✓	✓	✓	✓

Table 2: Comparison of existing datasets across four key dimensions: hate speech labeling, fake/misinformation labeling, structured/multi-dimensional explanatory annotations, and interdisciplinary scope. The interdisciplinary attribute denotes the integration of two distinct domains within a single dataset, allowing their interaction to be studied rather than treated separately. ✓ indicates the presence of a feature, while ✗ indicates its absence. Unlike prior resources, HateMirage incorporates all four dimensions, supporting richer analysis and interpretability.

such comments are made or what intent and implication they convey. HatEval (Basile et al., 2019), another benchmark dataset, provides hate labels for multilingual data but lacks any explanatory or interpretive layer, restricting analysis to surface-level classification. Implicit Hate (ElSherief et al., 2021) advances the field by capturing latent and indirect hate through contextual reasoning, but it still does not link hate expression to misinformation cues or structured components like target or intent. Similarly, HARE (Yang et al., 2023) offers explanatory information for hate categories but remains confined to hateful content without incorporating misinformation or interdisciplinary context.

Misinformation: On the misinformation side, CoAID (Cui and Lee, 2020) and FNVE (Chen et al., 2025) provide extensive fake news data, with the latter including multimodal evidence and explanatory text for fact-checking. However, both lack any representation of hate or harmful intent, making them unsuitable for studying hate arising from false narratives. Hostile (Bhardwaj et al., 2020) combines hate and fake content but treats them as independent categories; the hate content is not causally derived from misinformation and vice versa. Deceptive Humor (Kasu et al., 2025), though conceptually closer to Faux Hate, presents misinformation veiled as humor, which complicates detection and amplifies misinformation propagation along with harmful reactions, yet it too omits structured explanatory dimensions.

In contrast, HateMirage deals with faux hate comments (i.e, fake = 1 and hate = 1) with multi-dimensional structured explanations spanning target, intent, and implication. This comprehensive design enables deeper analysis of how misinforma-

tion shapes hateful discourse and supports interpretability across social, psychological, and linguistic perspectives.

3. Dataset Overview

In this section, we present a comprehensive overview of the HateMirage dataset. We describe the process of collecting widely debunked fake claims from reputable fact-checking websites, followed by scraping user comments from international English news channels. We explain the automated labeling procedure for identifying fake and hate comments, and detail the generation of structured explanations capturing the multi-dimensional aspects of each comment, including target, intent, and implication. Finally, we provide quantitative statistics and representative examples to illustrate the linguistic complexity and diversity of the dataset.

3.1. Collection of Debunked Fake Claims

The foundation of the HateMirage dataset is a curated set of widely debunked fake claims. These claims were sourced from reputable fact-checking websites such as AltNews⁴ and FactCheckers⁵, which maintain comprehensive repositories of misinformation that has been verified and refuted. We focused on claims that were widely circulated on social media to ensure the dataset captures content that is both impactful and relevant. Each claim was carefully reviewed to confirm its status as false or misleading, ensuring that subsequent data collection and analysis are grounded in verified information. This curated set of claims serves as the starting point for gathering user comments and generating structured explanations, providing a reliable basis for studying the intersection of misinformation and hate speech.

3.2. Data Collection and Labeling

We collect user comments from the YouTube sections of reputable international English news channels, capturing both purely English comments and code-mixed Hindi-English comments to reflect the linguistic diversity of social media discourse. International channels were chosen to reduce potential bias that may arise from local news coverage favoring one side. Each comment is labeled with two binary annotations: Hate (0/1) and Fake (0/1). The labeling process is automated using a GPT-4 model (Achiam et al., 2023), which is augmented with Retrieval-Augmented Generation (RAG) (Lewis et al., 2020) to provide additional grounded context for more accurate predictions. To

⁴<https://www.altnews.in/>

⁵<https://www.factchecker.in/>

maintain high data quality, we incorporate a human-in-the-loop validation stage even in the presence of RAG. While retrieval provides supporting context, model outputs can still suffer from hallucinations, cultural misinterpretations, or sensitivity to code-mixed language. Therefore, randomly sampled annotations are manually reviewed by human evaluators to verify label correctness, resolve ambiguous cases, and detect systematic biases. This process acts as a quality control mechanism, increasing confidence in the reliability of the automatically generated labels and ensuring that errors do not propagate into the final dataset.

3.3. Structured Explanation Generation

To better understand the underlying dynamics of fake and hateful comments, we first curated only the *Faux Hate* instances, i.e., comments labeled as both *fake* = 1 and *hate* = 1. These instances were drawn from both the existing FEUD dataset (Biradar et al., 2025), previously collected from Twitter and YouTube, as well as additional comments newly collected from YouTube channels. The collected Hindi-English code-mixed comments were then translated into English using the GPT-4 model, leveraging its extensive pre-trained multilingual understanding. We intentionally avoided text preprocessing operations such as punctuation or special character removal to preserve the raw linguistic style and contextual essence typical of social media discourse. English translation was preferred since most LLMs demonstrate stronger performance and broader support for English text, ensuring consistency in downstream processing and evaluation.

Building on this curated dataset, we aim to move beyond binary classification and delve into the underlying mechanisms of hate expression. To achieve this, we generate structured explanations that capture three key dimensions: the *Target* of the comment, the *Intent* behind it, and its potential *Implications* on social discourse. These explanations offer a multi-dimensional understanding of faux hate behavior, enabling a richer analysis of user motivations, the nature of attacks, and their societal impact.

We employ GPT-4 to generate structured explanations due to its extensive pretrained world knowledge, strong reasoning abilities, and robust language understanding. Faux hate comments often embed misinformation implicitly and rely on nuanced socio-political or cultural references, which makes simple pattern matching or smaller models insufficient. GPT-4 can effectively infer unstated targets, identify underlying intent, and articulate broader implications without task-specific fine-tuning. To ensure factual grounding, we adopt a RAG approach that supplements GPT-4 with relevant fact-checked evidence during generation.

Quality Validation (Human-in-the-Loop): The explanation generation process includes a human-in-the-loop component, where sampled outputs are continuously monitored and verified during generation to maintain quality, reliability, and consistency. Quality is evaluated across three objective dimensions: factual consistency, ensuring that each explanation aligns with verified evidence without introducing unsupported claims; semantic coherence, requiring logical alignment between the comment and its *Target*, *Intent*, and *Implication* fields; and linguistic clarity, emphasizing grammatical correctness, conciseness, and interpretability. We discard the samples that do not meet the above quality criteria.

To operationalize this workflow, two reviewers (one PhD scholar specializing in hate speech and one graduate student) participated in validation. As the dataset was generated in batches, we randomly sampled 10–15% of instances from each batch for manual inspection. Reviewers assessed explanations for factual consistency, semantic coherence, and linguistic clarity against the retrieved evidence. When inconsistencies were identified, we refined the retrieval pool and regenerated the outputs. Disagreements were resolved through discussion, and explanations were manually consolidated when necessary. All reviewers provided informed consent. This iterative auditing process helped prevent systematic errors and improved retrieval coverage and explanation reliability.

3.4. Representative Examples

The representative examples in Table 3 illustrate the multi-dimensional reasoning framework underlying HateMirage explanations. Each comment is deconstructed along three critical axes: *Target*, *Intent*, and *Implication*, which together delineate not just who is attacked, but why and with what social consequence. Identifying the *Target* demands recognizing both explicit mentions and implicit associations embedded within misinformation. Inferring the *Intent* requires discerning the author's underlying motive, whether to vilify, delegitimize, or polarize, often veiled beneath seemingly informational or conspiratorial tones. Finally, articulating the *Implication* captures the downstream effects on collective perception, social cohesion, or intergroup trust, bridging linguistic form and societal impact. Capturing these dimensions is inherently challenging, as they often co-occur in subtle, context-dependent ways, requiring interpretive reasoning beyond surface-level toxicity cues. Together, these structured explanations provide a granular view of how deceptive narratives operationalize hate, advancing both interpretability and accountability in computational social science.

Comment	Fake Claim Base	Target	Intent	Implication
<i>The coronavirus is just a scam planned by [Political Party] and [Country] to fool us all.</i>	"[Country] is Using Covid as a Biological Weapon"	[Political Party], [Country]	To link [Political Party] and [Country] to a conspiracy theory about COVID-19.	Could fuel public anger against [Political Party] and [Country].
<i>[Religious Group] and [Political Party], with [Person] in the mix, spread COVID to ruin [City] and harm [Country].</i>	"[Religious Group] and [Political Party] collaborated to spread COVID-19."	[Political Party], [Person], [Religious Group]	To accuse [Political Party], [Person], and [Religious Group] of intentionally worsening the pandemic.	Could stir up communal tensions and increase distrust in political and religious groups.
<i>[Religious Group] men are trapping innocent [Community] women in fake love to convert them. This is how [Country] is being taken over slowly.</i>	"Love Jihad: [Religious Group] men are systematically luring women from other communities to convert them."	[Religious Group]	To portray [Religious Group] as deceptive and predatory, reinforcing the fabricated 'Love Jihad' narrative.	Could deepen mistrust between communities, normalize religious profiling, and justify hostility toward [Religious Group].

Table 3: Representative examples of Faux Hate comments illustrating how explanations are derived from comments rooted in fake or manipulative narratives. Each example decomposes the comment into its underlying fake claim base, the entities targeted, inferred author intent, and the broader social or emotional implications.

3.5. Quantitative Analysis

Table 4 presents a quantitative summary of the HateMirage dataset across all fields: Comments, Target, Intent, and Implication. The dataset contains linguistically diverse user comments, ranging from short statements to lengthy posts (up to 800 tokens). The table reports key metrics for each field: Avg Tokens (average number of tokens per text), Median Tokens (median number of tokens per text), Min/Max Tokens (minimum and maximum tokens per text), Avg Sentence Length (average number of tokens per sentence), Max Sentence Length (longest sentence in tokens), Multi-Sentence % (percentage of texts containing more than one sentence), and Total Tokens (overall token count across all texts). Structured explanations are concise, with most Target annotations consisting of one or two tokens, while Intent and Implication texts typically span 14-15 tokens on average. Over 75% of Comments contain multiple sentences, highlighting the linguistic complexity present in user-generated content. These statistics illustrate the richness and diversity of the dataset, supporting both hate/fake classification and multi-dimensional explanatory analysis.

4. Benchmarking, Experimental Results and Analysis

We evaluate a diverse set of language models, ranging from small to large, to assess their ability to generate structured explanations for faux hate comments. For small language models (SLMs), we include LLaMA variants (Grattafiori et al., 2024) such as LLaMA-3.2-1B, LLaMA-3.2-3B, Qwen variants (Yang et al., 2025) such as Qwen-2.5-1.5B, Qwen-2.5-3B, and Microsoft Phi model (Abdin et al.,

2024) such as Phi-3-128k-3B. For large language models (LLMs), we consider Mistral-v0.3-7B (Jiang et al., 2023) and LLaMA-3.1-8B-Instruct. These models were chosen to represent diverse architectures, parameter scales, and training paradigms, allowing a systematic comparison of their reasoning and explanation generation capabilities.

The task is formulated as a structured explanation generation problem, where the model is prompted to extract and articulate the Target, Intent, and Implication fields for a given comment. In the zero-shot setting, models are provided only with the comment and a clear instruction describing each field’s meaning and expected format. No fine-tuning or supervised adaptation is performed. We use greedy decoding for all generations to ensure deterministic and reproducible outputs, avoiding variability introduced by sampling-based methods.

We evaluate models in two complementary settings. In the zero-shot setting, we examine the models’ internal reasoning abilities by relying solely on their pretrained knowledge and instructions, without any external context. In the RAG setting, we incorporate additional grounded knowledge, as faux hate often involves specific misinformation narratives that require external verification. We constructed a vector database using FAISS, populated with documents sourced from fact-checking platforms. All documents were preprocessed to remove noise and split into passage-level chunks suitable for retrieval.

For each comment, we performed a semantic search using cosine similarity to retrieve the top five most relevant evidence documents. These retrieved documents were prepended as context in the input prompt, allowing the models to generate explanations of Target, Intent, and Implication grounded in verified factual information. To ensure

	Field	Avg Tokens	Median Tokens	Min Tokens	Max Tokens	Avg Sent	Max Sent Len	Multi-Sent(%)	Total Tokens
Train (3,624 samples)	Comments	30.56	29	3	800	2.82	240	76.49	110741
	Target	1.48	1	1	27	1.00	12	0.30	5338
	Intent	14.66	14	5	39	1.02	34	1.41	53101
	Implication	15.16	15	6	48	1.01	38	0.86	54878
Test (906 samples)	Comments	30.28	27	4	316	2.84	78	73.18	27438
	Target	1.51	1	1	44	1.01	24	0.44	1371
	Intent	14.57	14	5	40	1.02	33	1.43	13204
	Implication	15.14	14	5	40	1.01	40	0.66	13706
Overall (4,530 samples)	Comments	30.50	28	3	800	2.82	240	75.83	138179
	Target	1.48	1	1	44	1.00	24	0.33	6709
	Intent	14.65	14	5	40	1.02	34	1.41	66305
	Implication	15.15	15	5	48	1.01	40	0.82	68584

Table 4: Quantitative summary of the HateMirage dataset, reporting key statistics for the complete dataset and its train and test splits. Metrics include text length, sentence structure, and multi-sentence prevalence, providing an overview of the linguistic complexity and scale of the dataset.

fully deterministic and reproducible outputs across all benchmarked models, we used greedy decoding with temperature = 0 and do_sample = False for all inference tasks. This setup guarantees that results can be replicated consistently, supporting reproducibility and benchmarking across different model configurations.

For evaluation, we compare the generated Target, Intent, and Implication explanations against the gold annotations using multiple complementary metrics: Sentence-BERT similarity (Reimers and Gurevych, 2019) and ROUGE-L F1 (Lin, 2004). Sentence-BERT similarity captures semantic alignment, while ROUGE-L F1 measures lexical and structural overlap. Together, this combination provides a balanced assessment of both meaning and surface form across all explanation fields.

From Table 5, it is evident that model performance varies significantly across model size, architecture, and grounding strategy. In the Zero-Shot setting, the Phi-3-128k-3B-Instruct model clearly dominates in Target identification, achieving the highest SBERT Similarity (65.55%) and ROUGE-L F1 (50.36%). For the Intent component, Phi-3 again performs competitively, securing the second-best SBERT Similarity (61.11%) and the best ROUGE-L F1 (29.52%), indicating strong capability in capturing the author’s underlying motive. The Implication dimension, being the most abstract and context-dependent, is best handled by Mistral-v0.3-7B-Instruct, which leads both metrics (SBERT = 55.64%, ROUGE-L F1 = 17.39%). These observations suggest that while larger models exhibit stronger inferential reasoning, smaller models like Phi-3 demonstrate exceptional generalization in identifying factual and role-focused elements of faux hate explanations.

Under the RAG-based setting, performance trends remain consistent, with notable improvements in grounding-sensitive components. The Phi-3-128k-3B-Instruct model retains strong performance, achieving the second-best Target results (SBERT = 63.65%, ROUGE-L F1 = 47.81%), the highest Intent similarity (SBERT = 62.03%), and the

second-best Intent ROUGE-L (26.90%). For Implication, Mistral-v0.3-7B-Instruct again performs robustly (SBERT = 53.68%, ROUGE-L F1 = 16.17%), reflecting its capacity to reason about broader societal effects when supported by retrieval context.

Overall, while other models such as Qwen-2.5 and LLaMA-3 variants exhibit moderate yet balanced results, the Phi-3 model consistently shows strong performance across multiple dimensions. One possible explanation is that its training mixture includes substantial synthetic and reasoning-focused data, which may encourage structured and instruction-aligned outputs that resemble the format of our references.

At the same time, similarity to GPT-4 may also play a role. Prior documentation for the Phi family indicates influence from GPT-4 class teacher systems, and because our reference explanations are themselves generated by GPT-4, automatic metrics such as ROUGE-L can partly reward overlap in phrasing or discourse style in addition to the underlying quality of reasoning. Therefore, the observed advantage should be interpreted carefully. The results may reflect a combination of genuine interpretive capability and stylistic alignment with the reference generator.

All experiments are conducted using open-source models, with inference run on a cluster of four NVIDIA GPUs, each with 15GB of memory, totaling 60GB. Both small and large language models are evaluated under identical experimental conditions to ensure a fair comparison.

4.1. Error Analysis

To better understand the limitations of the benchmarked models in the Faux Hate Explanation task, we conducted a qualitative error analysis on selected examples across the three explanation dimensions: Target, Intent, and Implication (see Table 6). For Target prediction, qualitative inspection suggests that in several instances the system relies on surface-level cues such as prominent nouns or early-mentioned entities, which can lead to miss-

Model	Target		Intent		Implication	
	SBERT	R-L (F1)	SBERT	R-L (F1)	SBERT	R-L (F1)
LLaMA-3.2-1B-Instruct	50.19	30.91	57.99	26.25	47.04	12.79
LLaMA-3.2-3B-Instruct	53.43	34.62	57.67	26.81	47.76	16.36
LLaMA-3.1-8B-Instruct	48.29	21.39	56.68	27.95	51.24	15.62
Qwen-2.5-1.5B	53.57	35.93	57.86	25.56	43.41	14.45
Qwen-2.5-3B-Instruct	52.91	36.13	59.96	19.06	53.09	13.75
Phi-3-128k-3b-Instruct	65.55	50.36	61.11 [†]	29.52	50.39	17.27 [†]
Mistral-v0.3-7B-Instruct	59.81	40.38	60.75	27.74	55.64	17.39
LLaMA-3.2-1B-Instruct	52.43	30.92	57.77	25.14	47.38	13.20
LLaMA-3.2-3B-Instruct	49.55	28.50	55.98	25.80	48.74	16.42
LLaMA-3.1-8B-Instruct	40.37	14.80	59.00	26.70	52.17	14.67
Qwen-2.5-1.5B	62.18	47.33	55.58	21.82	44.31	13.71
Qwen-2.5-3B-Instruct	53.86	37.97	61.03	18.15	51.31	12.82
Phi-3-128k-3b-Instruct	63.65 [†]	47.81 [†]	62.03	26.90 [†]	51.20	16.25
Mistral-v0.3-7B-Instruct	63.06	44.27	60.03	26.38	53.68 [†]	16.17

Table 5: Comparison of LLM performance across Zero-Shot and RAG-based settings. Highest results are highlighted in **bold**, and second-best results are marked with [†]. (Metrics reported in %)

ing the true focus of the comment. For example, in sarcastic or indirect messages, the model may assign the label `<organization>` even when the criticism is directed toward a journalist or media actor.

For Intent, we observe multiple cases where the model produces generic or templated explanations such as “spreading hateful or fake narratives,” which may overlook more nuanced motives including sarcasm, racialized mockery, or political positioning. Implication appears to be the most challenging dimension. In several reviewed examples, models generate broad or underspecified outcomes, for instance “promoting general negativity or unrest,” while failing to capture more specific social consequences such as scapegoating, perceived institutional bias, or the reinforcement of double standards. This observation is consistent with our human evaluation results, where Implication also showed comparatively lower agreement than Target and Intent.

Overall, these patterns indicate a tendency toward overgeneralization and difficulty in capturing pragmatic and context-dependent meaning, underscoring the inherent complexity of multi-dimensional reasoning in Faux Hate Explanation.

4.2. Human Evaluation

The HateMirage dataset was generated using the GPT-4 model augmented with contextual information through Retrieval-Augmented Generation (RAG). To assess the quality of GPT-4-generated explanations in the HateMirage dataset, we first compared them against human-written explanations. For this, 500 randomly selected records were independently annotated by trained human annotators, who produced their own explanations for Target, Intent, and Implication. The GPT-4 outputs were then quantitatively evaluated using

SBERT cosine similarity and ROUGE-L F1 scores against the human-written references as shown in Table 7. SBERT primarily reflects semantic agreement, whereas ROUGE-L captures lexical and phrasal overlap. Because multiple valid explanations can differ in wording while preserving meaning, it is possible to observe relatively high SBERT values alongside more moderate ROUGE scores. Notably, 208 of 500 records surpassed a 90% similarity threshold for Target, only 7 for Intent, and none for Implication, highlighting that GPT-4 reliably identifies entities and general motives, while reasoning over broader causal consequences and nuanced implications remains challenging. This pattern is consistent with our independent human rating study, where Implication likewise received lower agreement relative to Target and Intent.

	SBERT	R-L (F1)
Target	0.7479	0.6293
Intent	0.7020	0.3458
Implication	0.5780	0.2013

Table 7: Human evaluation results comparing GPT-4 generated explanations with human-written explanations for 500 randomly sampled records.

In addition, we conducted a quantitative human assessment on 100 randomly selected records from the dataset, independently annotated by two human annotators. For each record, annotators rated the GPT-4-generated explanations for Target, Intent, and Implication on a scale from 1 (poor) to 5 (excellent). As shown in Figure 1, the majority of ratings for Target fall between 4 and 5, Intent is close to 4, and Implication ranges between 3 and 4. These results indicate that while GPT-4 effectively captures entities and general motives, capturing broader causal consequences and nu-

	Comment	Ground Truth	Model Prediction
Target	<i>During the time of Corona, you also actively spread information about the <organization>. Have some shame for you were metaphorically "hit with a shoe on your face." Please report the news with a bit of discretion.</i>	Media / journalist	<organization>
	<i>Hey Brother, <organization> is spreading the coronavirus here because you're not paying attention to it.</i>	<organization>	Brother
Intent	<i><@user> Were the people from the group in Markaz there to catch corona?</i>	To question the motives of the <organization> gathering and imply reckless behavior.	The commenter is questioning the intentions of people in Markaz, suggesting they might be there to spread misinformation about the COVID-19 pandemic.
	<i>I still advocate calling it the "kung-flu". And label anyone suspected of being infected as "Flu Kang".</i>	To use racially insensitive language to mock and stigmatize people, especially Asians, in relation to COVID-19.	To spread hateful and fake narratives about the flu.
Implication	<i><@user> Why.. Did they have corona too? Oh, so you all have been spreading this disease since then... We've been blaming the <organization> for no reason... This has been going on traditionally.</i>	It implies that the criticism directed toward <organization> was unwarranted and highlights bias or unfair treatment.	The comment could be interpreted as promoting general negativity or unrest, possibly encouraging harmful behavior.
	<i>If <organization> is responsible, then Namaste Trump is more responsible. <person> had given a warning on February 11 regarding Corona. He was mocked then. Who will you hold responsible for Namaste Trump? Trump or our PM? #shameOnRupani</i>	This criticism aims to highlight perceived double standards and questions leadership decisions in handling the COVID-19 situation.	The comment could be interpreted as drawing a general comparison between different parties, possibly implying blame or criticism.

Table 6: Error analysis showing model predictions for Target, Intent, and Implication. Placeholders such as <organization> and <user> are used to anonymize entities and generalize examples.

anced societal implications remains more difficult, emphasizing the inherent challenge of reasoning over downstream effects in faux hate content.

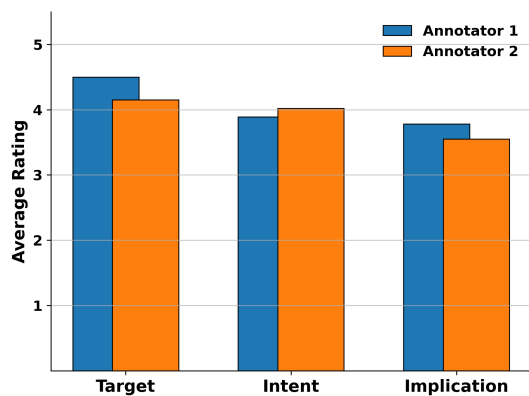


Figure 1: Comparison of average ratings from two annotators for GPT-4-generated explanations across Target, Intent, and Implication.

5. Conclusion and Future Work

In this work, we introduced *HateMirage*, a novel dataset designed to advance the study of subtle and deceptive forms of harmful online content. Unlike traditional hate speech or misinformation corpora, *HateMirage* focused exclusively on Faux

Hate, the content where hateful narratives emerge indirectly through misinformation or deceptive framing. Each comment was annotated with structured, multi-dimensional explanations capturing *Target*, *Intent*, and *Implication*, thereby uncovering the implicit layers of meaning that are often missed by surface-level analysis. By providing these rich explanatory annotations, *HateMirage* aimed to support both research and moderation efforts by making the underlying harmful reasoning more transparent and interpretable, even to non-expert users. We further benchmarked a range of language models on explanation generation tasks, establishing initial baselines and revealing key challenges in structured reasoning over misinformation-driven hate. These results highlighted significant gaps between current model capabilities and the complexity of reasoning required to interpret Faux Hate.

In future work, we plan to extend *HateMirage* to the multimodal domain by incorporating memes and visual narratives, enabling the study of how text-image interactions shape deceptive hateful discourse across different social media platforms. Additionally, we aim to explore automated methods for generating and evaluating multi-dimensional explanations, including explanation faithfulness metrics and model-based evaluators that move beyond lexical similarity to better assess reasoning consistency and grounding, paving the way for more interpretable and robust moderation systems.

6. Limitations

While the HateMirage dataset is a unique resource providing structured, multi-dimensional explanations for Faux Hate comments, certain limitations must be acknowledged.

The primary limitation lies in the synthetic nature of the explanations. The annotations were generated using GPT-4 with RAG. This approach was deliberately chosen to minimize the exposure of human annotators to highly distressing and toxic content, as Faux Hate can be psychologically disturbing; accordingly, human verification was conducted on a small, random subset of 500 records (approximately 11% of the total dataset). Despite human-in-the-loop validation, these model-generated explanations may lack the full subtlety and emotional variability of purely human interpretations. This is particularly notable in the abstract Implication field, which proved most difficult for the model to infer accurately against human-written gold standards. Furthermore, the corpus is derived specifically from YouTube comments on international English news channels, meaning it may not fully represent the vernacular and platform-dependent characteristics of Faux Hate spread across all social media environments. Despite these limitations, HateMirage provides a valuable foundation for studying Faux Hate and developing reasoning and moderation-oriented systems.

7. Ethics Statement

The HateMirage dataset is constructed from publicly available YouTube comments posted under international news videos. Data collection strictly followed YouTube's Terms of Service⁶ and Privacy Policy⁷. Only publicly accessible comments were included; no private or restricted data were used. Given the sensitive nature of the content, strict anonymization procedures were applied: all personally identifiable information (PII), including usernames, profile links, and timestamps, was not collected, and named entities were masked where appropriate to minimize risks of re-identification and to avoid amplifying harmful narratives. The dataset is intended solely for academic research on explainable NLP, misinformation, and hate speech. Public redistribution is strictly prohibited. Access will be provided under a controlled, research-only license requiring requesters to agree to responsible-use terms, and any form of commercial use or generative model training outside a research setting is prohibited.

⁶<https://www.youtube.com/t/terms>

⁷<https://policies.google.com/privacy>

Broader Impact: By focusing on explanatory reasoning rather than simple classification, this work aims to enable more transparent and accountable moderation systems and support research on the interplay between misinformation and hateful discourse. Structured explanations can help moderation teams, fact-checkers, and researchers better understand why content is harmful, thereby promoting safer online spaces and more interpretable AI systems. However, potential risks remain: malicious actors could misuse structured explanations to craft more sophisticated hateful narratives or target vulnerable groups. To mitigate this, we restrict dataset access, anonymize all entities, and prohibit redistribution or use in generative settings that could amplify harm. Overall, this work seeks to advance responsible scientific analysis of nuanced online harm while safeguarding user privacy, minimizing misuse, and reflecting critically on the societal implications of explainable NLP systems in sensitive domains.

Acknowledgement

M. S. Akhtar acknowledges the partial support of Infosys Center of AI (CAI) at IIT Delhi.

8. Bibliographical References

- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. 2024. Phi-4 technical report. [arXiv preprint arXiv:2412.08905](https://arxiv.org/abs/2412.08905).
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. [arXiv preprint arXiv:2303.08774](https://arxiv.org/abs/2303.08774).
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In [Proceedings of the 13th international workshop on semantic evaluation](https://www.aclweb.org/anthology/W19-07), pages 54–63.
- Shankar Biradar, Kasu Sai Kartheek Reddy, Sunil Saumya, and Md Shad Akhtar. 2024. Proceedings of the 21st international conference on natural language processing (icon): Shared task on decoding fake narratives in spreading hateful stories (faux-hate). In [Proceedings of the 21st](https://www.aclweb.org/anthology/I24-1)

- International Conference on Natural Language Processing (ICON): Shared Task on Decoding Fake Narratives in Spreading Hateful Stories (Faux-Hate), pages 1–5.
- Paula Fortuna and Sérgio Nunes. 2018. A survey on automatic detection of hate speech in text. *Acm Computing Surveys (Csur)*, 51(4):1–30.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. [arXiv preprint arXiv:2407.21783](#).
- Dongsheng Jiang, Yuchen Liu, Songlin Liu, Jin’e Zhao, Hao Zhang, Zhen Gao, Xiaopeng Zhang, Jin Li, and Hongkai Xiong. 2023. From clip to dino: Visual encoders shout in multimodal large language models. [arXiv preprint arXiv:2310.08825](#).
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Mohsen Mosleh, Rocky Cole, and David G Rand. 2024. Misinformation and harmful language are interconnected, rather than distinct, challenges. *PNAS nexus*, 3(3):pgae111.
- B Premjith, Bharathi Raja Chakravarthi, Prasanna Kumar Kumaresan, Saranya Rajiakodi, Sai Karnati, Sai Mangamuru, and Chandu Janakiram. 2024. Findings of the shared task on hate and offensive language detection in telugu codemixed text (hold-telugu)@ dravidianlangtech 2024. In *Proceedings of the Fourth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*, pages 49–55.
- Nils Reimers and Iryna Gurevych. 2019. Sentencebert: Sentence embeddings using siamese bert-networks. [arXiv preprint arXiv:1908.10084](#).
- Anna Schmidt and Michael Wiegand. 2017. A survey on hate speech detection using natural language processing. In *Proceedings of the fifth international workshop on natural language processing for social media*, pages 1–10.
- Manuel Tonneau, Diyi Liu, Samuel Fraiberger, Ralph Schroeder, Scott A Hale, and Paul Röttger. 2024. From languages to geographies: Towards evaluating cultural bias in hate speech datasets. In *Proceedings of the 8th Workshop on Online Abuse and Harms (WOAH 2024)*, pages 283–311.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. [arXiv preprint arXiv:2505.09388](#).
- Yongjin Yang, Joonkee Kim, Yujin Kim, Namgyu Ho, James Thorne, and Se-Young Yun. 2023. Hare: Explainable hate speech detection with step-by-step reasoning. [arXiv preprint arXiv:2311.00321](#).
- Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. Semeval-2020 task 12: Multilingual offensive language identification in social media (offenseval 2020). [arXiv preprint arXiv:2006.07235](#).

9. Language Resource References

- Mohit Bhardwaj, Md Shad Akhtar, Asif Ekbal, Amitava Das, and Tanmoy Chakraborty. 2020. Hostility detection dataset in hindi. [arXiv preprint arXiv:2011.03588](#).
- Shankar Biradar, Sunil Saumya, and Arun Chauhan. 2025. Faux hate: unravelling the web of fake narratives in spreading hateful stories: a multi-label and multi-class dataset in cross-lingual hindi-english code-mixed text. *Language Resources and Evaluation*, 59(1):477–508.
- Lizhi Chen, Zhong Qian, Peifeng Li, and Qiaoming Zhu. 2025. Multimodal fake news video explanation: Dataset, analysis and evaluation. [arXiv preprint arXiv:2501.08514](#).
- Limeng Cui and Dongwon Lee. 2020. Coaid: Covid-19 healthcare misinformation dataset. [arXiv preprint arXiv:2006.00885](#).
- Mai ElSherief, Caleb Ziems, David Muchlinski, Vaishnavi Anupindi, Jordyn Seybolt, Munmun De Choudhury, and Diyi Yang. 2021. [Latent hatred: A benchmark for understanding implicit hate speech](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 345–363, Online

and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Sai Kartheek Reddy Kasu, Shankar Biradar, and Sunil Saumya. 2025. Deceptive humor: A synthetic multilingual benchmark dataset for bridging fabricated claims with humorous content. arXiv preprint arXiv:2503.16031.

Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2021. Hatexplain: A benchmark dataset for explainable hate speech detection. In Proceedings of the AAAI conference on artificial intelligence, volume 35, pages 14867–14875.