

LLM-Based Semantic Augmentation for Harmful Content Detection

Elyas Meguellati¹, Assaad Zeghina², Shazia Sadiq¹, Gianluca Demartini¹

¹The University of Queensland, Brisbane, Australia

²Université de Strasbourg, CNRS, ICube, Strasbourg, France
{m.meguellati, s.sadiq, g.demartini}@uq.edu.au, aozeghina@unistra.fr

Abstract

Recent advances in large language models (LLMs) have demonstrated strong performance on simple text classification tasks, frequently under zero-shot settings. However, their efficacy declines when tackling complex social media challenges such as propaganda detection, hateful meme classification, and toxicity identification. Much of the existing work has focused on using LLMs to generate synthetic training data, overlooking the potential of LLM-based text preprocessing and semantic augmentation. In this paper, we introduce an approach that prompts LLMs to clean noisy text and provide context-rich explanations, thereby enhancing training sets without substantial increases in data volume. We systematically evaluate on the SemEval 2024 multi-label Persuasive Meme dataset and further validate on the Google Jigsaw toxic comments and Facebook hateful memes datasets to assess generalizability. Our results reveal that zero-shot LLM classification underperforms on these high-context tasks compared to supervised models. In contrast, integrating LLM-based semantic augmentation yields performance on par with approaches that rely on human-annotated data, at a fraction of the cost. These findings underscore the importance of strategically incorporating LLMs into machine learning (ML) pipeline for social media classification tasks, offering broad implications for combating harmful content online.

Disclaimer: This paper contains examples of explicit language that may be disturbing to some readers.

Introduction

Large language models (LLMs) have achieved remarkable performance in a broad range of natural language processing (NLP) tasks, showcasing strong zero-shot and few-shot capabilities. Tasks such as sentiment analysis (Kuila and Sarkar 2024) and named entity recognition (Xie et al. 2024) have particularly benefited from these advancements, highlighting the potential of LLMs to perform well without extensive task-specific training data. However, when applied to more complex tasks, LLMs often struggle to achieve the same level of performance as fine-tuned or task-specific models. Recent studies have highlighted that while zero-shot and few-shot classification using LLMs is effective, these methods frequently underperform as compared to traditional supervised approaches trained on sufficient labeled

data (Scius-Bertrand et al. 2024). For example, fine-tuning smaller models has been shown to outperform zero-shot and few-shots LLMs across a variety of text classification datasets (Edwards and Camacho-Collados 2024; Hsieh et al. 2023). While some evidence suggests that LLMs can reduce the need for extensive data labeling in resource-limited scenarios (Sushil et al. 2024), their performance declines in tasks requiring nuanced understanding of multimodal inputs, or complex semantic relationships such as propaganda detection (Jose and Greenstadt 2024) where *the complexity* lies in capturing linguistic nuances, context-dependencies and latent meaning, due to richness of dynamic variants and figurative use of language (Joshi, Sharma, and Bhattacharyya 2015).

Existing research has predominantly focused on leveraging LLMs to augment datasets by generating synthetic examples (Patel, Raffel, and Callison-Burch 2024). While this approach has shown promise in improving model performance in data-scarce settings (Veselovsky et al. 2023), it often overlooks the potential of LLMs to enrich data by adding context (explanations and relevant information) rather than quantity. Semantic augmentation, such as generating explanations for existing data samples, offers a novel pathway for enhancing the feature space without substantially increasing dataset size. Moreover, data cleaning, a critical preprocessing step, has largely relied on rule-based approaches, which often fail to handle unpredictable errors and anomalies in real-world datasets (Borrouhou et al. 2023). LLMs, with their advanced understanding of content, offer an alternative by addressing data quality issues more flexibly and accurately (Long et al. 2024).

This study proposes a novel LLM-based semantic augmentation framework that leverages LLMs to clean and contextualize data for context-dependent classifications. Unlike conventional data augmentation approaches, which aim to increase the size of training datasets, our method focuses on enriching feature representations by generating relevant information for each sample and refining noisy text. We evaluate this approach using the *SemEval 2024 persuasive meme dataset*, a multi-label classification problem that forms a subtask of the *SemEval Propaganda Detection* challenge (Dimitrov et al. 2024). To further validate the generalizability of our method, we also apply it to the *Google Jigsaw*

*toxic comments dataset*¹ and the *Facebook Hateful Meme Challenge dataset*² (Kiela, Firooz, and Mohan 2020).

This paper contributes to expanding our understanding of how LLMs can be effectively integrated into machine learning (ML) workflows for social media data analysis. By identifying the most suitable roles for LLMs within a given task, particularly in improving context and addressing data quality issues, we provide a scalable solution for tackling complex classification challenges in online content moderation, propaganda detection, and beyond. The following research questions guide our investigation:

1. How does incorporating LLM-based data cleaning and semantic augmentation strategies impact the overall performance of models in context-rich classification tasks?
2. Does the proposed LLM-based semantic augmentation approach generalize effectively across related domains on social media?
3. Beyond explanations, which other semantic augmentation strategies can be used effectively in domains where LLMs might censor or omit critical language elements?
4. Under which settings do zero-shot LLMs underperform in social media classification?

Related Work

LLM-Based Data Augmentation

Traditional data augmentation (DA) methods, such as rule-based transformations (Feng et al. 2021) and back-translation (Kumar et al. 2019; Yang et al. 2020), have been significantly enhanced by LLMs that generate synthetic data, improving model performance across diverse NLP tasks (Zhou et al. 2024). LLM-based DA has been effectively applied in domains such as healthcare (Wang, Zhao, and Petzold 2023) and finance (Li et al. 2023a), with studies demonstrating that synthetic data from LLMs can achieve performance comparable to human-annotated data, at significantly lower costs. For instance, generating 3,000 samples using GPT-3 costs only \$14.37 and takes 46 minutes, compared to the substantially higher cost and time required for human annotation (Ding et al. 2023). Despite these advances, synthetic data often underperforms compared to human-labeled data, especially within specialized or highly technical domains (Li et al. 2023b; Ding et al. 2023). Recent approaches have therefore focused on advanced prompting strategies, such as chain-of-thought (Wei et al. 2022) and self-consistency prompting (Wang et al. 2022a), as well as noise reduction frameworks like SunGen (Gao et al. 2023), to address these limitations and enhance synthetic dataset quality. Such methods have enabled LLMs to better capture task-specific nuances, showing improvement in areas such as commonsense reasoning (Liu et al. 2021) and specialized domains like oncology physics (Holmes et al. 2023).

Cross-domain capabilities of LLMs have also been extensively studied recently. Wang, Zhao, and Petzold (2023) evaluated LLM performance on healthcare natural language

understanding tasks, demonstrating GPT-4’s superior performance over alternative models in classification and named entity recognition. Similarly, in finance, Li et al. (2023a) found that GPT-4 performs strongly in news classification and sentiment analysis tasks, though specialized domain-specific models maintain advantages in narrower contexts. This observation extends across various other domains, with work by Wan et al. (2023) showing effectiveness in information extraction, and Bang et al. (2023) confirming high performance in question answering tasks.

Persuasion and Propaganda Detection

Persuasion techniques are integral to the broader framework of propaganda, significantly influencing public opinion through social media and news platforms (Brattberg and Maurer 2018; Golovchenko et al. 2020). A range of computational approaches have been developed to identify specific propaganda techniques, notably including the Propaganda Techniques Corpus (PTC) introduced by Da San Martino et al. (2019), which catalogues eighteen distinct persuasion labels in news articles. Transformer-based models, such as RoBERTa-CRF (Jurkiewicz et al. 2020) and Multi-Granularity Networks (MGN) (Da San Martino et al. 2019), have demonstrated strong performance in persuasion detection tasks. Recent studies (Jose and Greenstadt 2024) suggest that while general-purpose LLMs typically underperform compared to specialized transformer architectures such as RoBERTa-CRF, they still show promising results in detecting specific propaganda categories. Notably, popular LLMs (GPT-3.5, GPT-4, and Claude 3 Opus) outperform the MGN baseline in categories like *name-calling*. Moreover, GPT-3.5 and GPT-4 achieve stronger performance in detecting techniques such as *appeal to fear* and *flag-waving*, though they struggle with more subtle persuasive strategies such as *bandwagon* and *doubt*.

These observations are consistent with broader trends in recent literature on LLM capabilities in complex NLP tasks, highlighting both the strengths and remaining challenges of current systems. The evolution of these detection methods underscores the importance of developing more nuanced and context-aware approaches to effectively address increasingly sophisticated persuasion and propaganda strategies, especially within social media environments.

Datasets

SemEval-2024 Persuasive Memes

The SemEval-2024 Task 4 dataset is the main dataset for this study as it serves as the basis for all experimental analyses. It is designed for hierarchical multi-label classification of persuasion techniques in memes. The dataset includes 22 distinct persuasion techniques, structured hierarchically to represent the complex relationships between *rhetorical* and *psychological* strategies.

The dataset primarily consists of English-language memes, though it includes a smaller subset in other languages. Each sample contains both textual and visual elements, making it a challenging multimodal dataset. The hierarchical structure adds complexity by requiring models

¹kaggle.com/c/jigsaw-toxic-comment-classification-challenge

²ai.facebook.com/datasets/hateful-memes

to identify specific techniques and their contextual relationships within the hierarchy. The dataset consists of a training set of 7,000 samples and a test set of 1,500 samples. Hierarchical F1 (**H-F1**) is the chosen evaluation measure. **H-F1** differs from the standard F1 score by taking into account the hierarchical relationships among labels (Kiritchenko et al. 2006). While the standard F1 score evaluates precision and recall for flat, independent labels, **H-F1** adjusts the evaluation to reflect the hierarchical dependencies between persuasion techniques. This ensures that errors at higher levels in the hierarchy are treated differently from errors at lower levels, providing a more nuanced assessment of the model’s performance.

Validation Datasets

We use additional validation datasets to assess the generalizability of the proposed approach by applying the baselines and best-performing configurations derived from our experiments on the SemEval-2024 Persuasion dataset. They provide a broader evaluation benchmark for the proposed method’s effectiveness on downstream tasks across different online contexts.

Google Jigsaw Toxic Comments The Jigsaw Toxic Comments dataset is used to validate the approach in a purely textual domain focused on toxic comment classification. The dataset contains English-language Wikipedia comments, annotated for six categories of toxicity: *toxic*, *severe toxic*, *obscene*, *threat*, *insult*, and *identity hate*. For this study, only English comments are considered. The dataset consists of a training set of 159,571 samples and a test set of 153,164 samples. The evaluation measures are F1-score and ROC-AUC.

Facebook Hateful Meme Challenge The Facebook Hateful Meme Challenge dataset is used to evaluate methods in a multimodal task involving both text and image content. Each meme is classified as *hateful* or *non-hateful*. This study focuses on memes with English text content to validate the generalizability of the proposed method to multimodal hate speech detection. The dataset consists of a training set of 8,500 samples and a test set of 1,500 samples. The evaluation measures are F1-score and ROC-AUC.

All three datasets are accompanied by predefined, fixed training and test splits provided by the organizers, with detailed label distributions shown in Tables 9, 10, and 11 in the Appendix.

Method

LLM-based Cleaning

To streamline multimodal processing and reduce computational overhead, we substitute image inputs with textual data produced by captioning models BLIP (Li et al. 2022) and GIT (Wang et al. 2022b). This approach is motivated by findings by Nguyen et al. (2024), which highlight that converting images into textual representations via captioning models not only simplifies training, but also significantly lowers computational complexity and speeds up inferences.

However, captioning models often produce outputs containing noise and anomalies, such as grammatical errors, redundant phrases, and misaligned descriptions (Figure 1). These inconsistencies are challenging to filter out using traditional rule-based methods, thus necessitating a more adaptive approach. Such imperfections make this an ideal scenario for evaluating the benefits of the proposed text cleaning step using LLMs, as the noisy captions provide a realistic testbed for assessing their capability to enhance text quality and ensure task relevance.

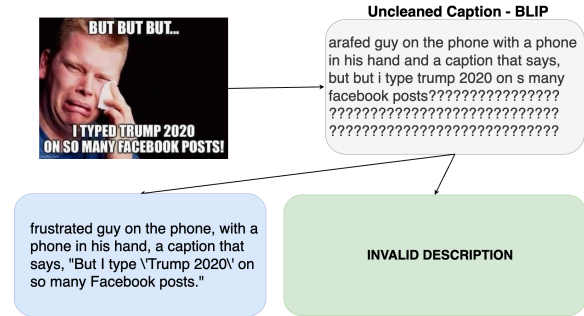


Figure 1: LLM-based Cleaning for BLIP-generated Caption.

The detailed prompt provided for cleaning and correcting captions is shown in Figure 13 Appendix.

Using this prompt, we processed the captions and discarded any that resulted in an *INVALID DESCRIPTION*. Since memes consist of both text and image components, when a caption was deemed invalid, the original meme text was retained and utilized for training without an AI-generated caption. This ensured that no training sample was excluded due to an invalid caption.

Model	Set	Sonnet	LLaMA	GPT-4o
BLIP	Train	89.9	99.9	69.2
	Dev	89.8	99.8	72.6
	Test	89.5	100.0	71.8
GIT	Train	72.5	96.5	69.6
	Dev	67.6	95.8	69.8
	Test	69.7	97.6	70.0

Table 1: Valid Caption Statistics (%) for BLIP and GIT Models Using Sonnet 3.5, LLaMA 3.1 70B, and GPT-4o.

Figure 1 illustrates the workflow for cleaning the captions generated by BLIP. Table 1 presents the percentage of valid captions produced by BLIP and GIT models after being cleaned by three LLMs: LLaMA 3.1 70B³, GPT-4o⁴, and Sonnet 3.5 v2⁵.

After LLM-based cleaning, if the caption is successfully processed and deemed valid, both the cleaned caption and the meme text are passed to training. If the caption is too

³ai.meta.com/blog/meta-llama-3-1/

⁴openai.com/index/hello-gpt-4o/

⁵anthropic.com/news/3-5-models-and-computer-use

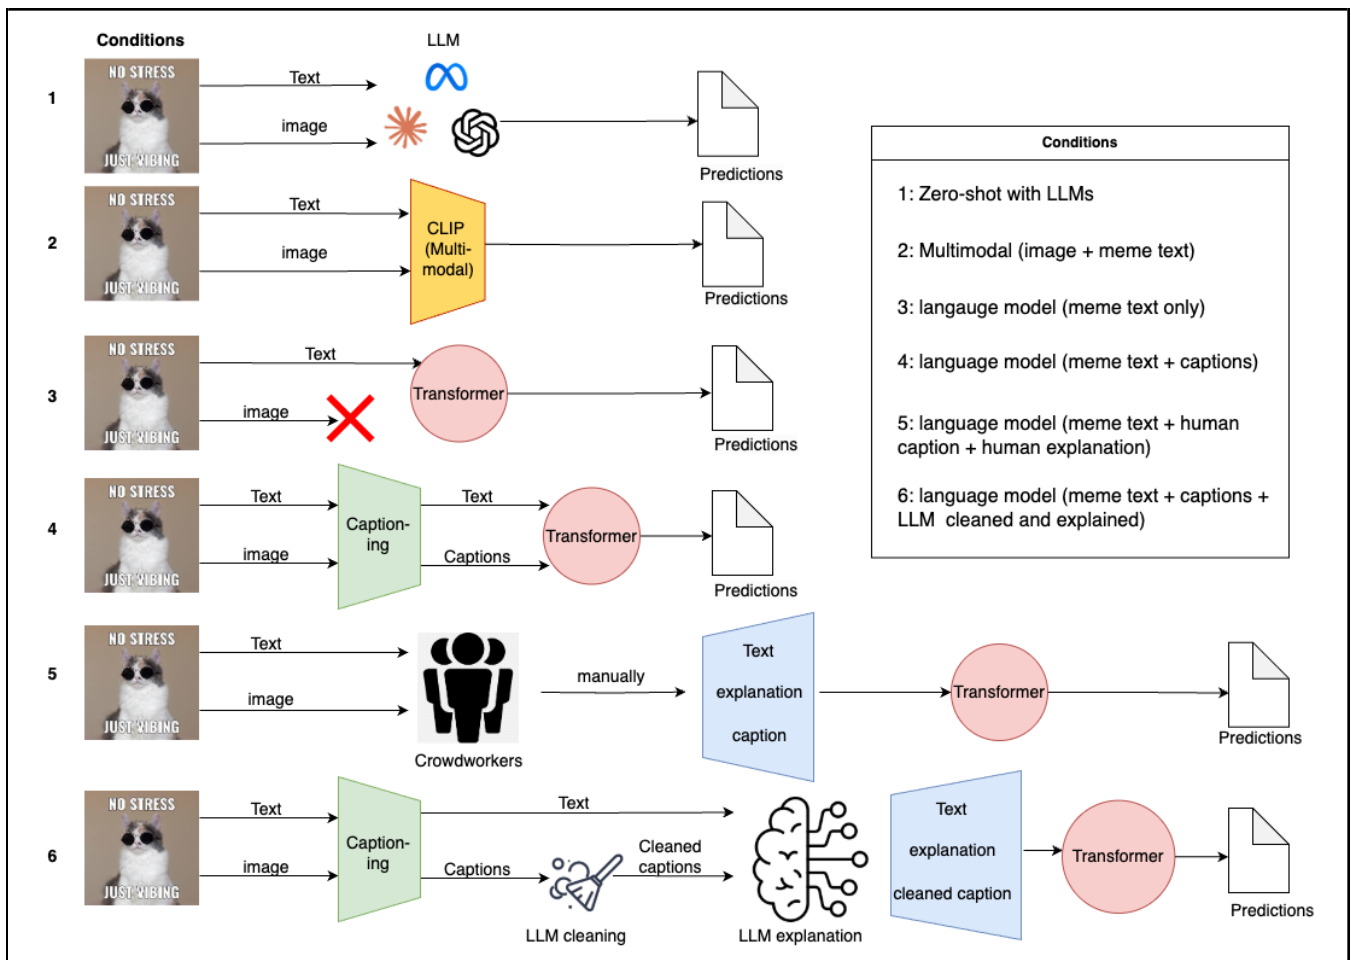


Figure 2: Overview of the experimental conditions.

corrupted to clean and is marked as invalid, then only the original meme text is used. The percentage of valid captions evaluates how effectively each LLM handles noisy caption outputs from captioning models like BLIP and GIT, providing insights into their ability to process caption quality, resolve anomalies, and manage failures (Table 1).

LLM-based Semantic Augmentation

In addition to directly utilizing the meme text and machine-generated captions, we employ LLMs to generate explanations which highlight key semantic or persuasive cues that might be abstract or less apparent in the raw text alone. Specifically, each sample is augmented with an LLM-generated explanation illustrating the potential rhetorical or psychological strategies present in meme content. Our assumption is that these explanations provide more nuanced insight, thus improving the downstream model’s performance on complex, multi-label classifications.

Experimental Conditions

In this study, we explore six distinct configurations and their derivatives that incorporate the meme, LLM-based semantic

augmentation, human annotations, multimodal model, and language model as shown in Figure 2.

- 1. Zero-Shot with Meme Image and Text:** The same LLMs used for cleaning receive the meme image and text without any fine-tuning. This baseline measures how well modern LLMs handle multimodal, multi-label classification in a zero-shot configuration.
- 2. Image + Text (Multimodal Model):** A multimodal model CLIP (Radford et al. 2021) processes both image and textual cues, leveraging the learned alignment between visual and linguistic features.
- 3. Meme Text Only (Language Model):** Only the textual element of the meme is used, serving as a text-only baseline. The language model BART decoder (Lewis et al. 2020) treats the classification as a sequence-to-sequence prediction, inherently suited for hierarchical classification.
- 4. Text + Captions (Language Model):** Machine-generated captions (e.g., from BLIP) are concatenated with the meme text to incorporate visual context through

textual representation. The same language model used in the text-only baseline is employed for this configuration.

- Human Captions and Explanations (Language Model):** A subset of memes includes crowd-annotated captions and explanations, while the remaining samples consist of meme text only. The input is processed using the same language model as in previous baseline.
- Meme Text + Cleaned Captions + LLM Explanations (Language Model):** Machine-generated captions undergo LLM-based cleaning to remove noise. Additionally, each meme is paired with an explanation generated by an LLM, highlighting potential persuasive or contextual elements. The concatenated input (meme text, cleaned captions, and explanation) is processed using the same language model.

Human and LLM-Based Annotations

We conducted a study on Prolific⁶ with 300 crowdworkers, retaining 70 based on a manual inspection of the quality of their responses. These 70 high-performing crowdworkers annotated 1,000 memes from the SemEval persuasive meme dataset for training and an additional 1,000 memes for testing (captions + explanations). Following the instructions detailed in Figures (6-7) in the Appendix, participants' captions and explanations of the memes did not contain the labeled persuasion techniques. This ensured that the human annotations remained label-agnostic, mirroring the conditions applied to the LLMs. For the remaining samples, only the original meme text was used without human captions or explanations. In parallel, we employed LLMs to clean BLIP/GIT captions and generate explanations for all memes. Annotating 2,000 memes via crowdworkers costs approximately \$1,100, whereas annotating the entire dataset (9,500 memes) with LLaMA 3.1 70B and Sonnet 3.5 V2 amounted to only \$3 each, and \$30 for GPT-4o highlighting the significantly lower financial cost of LLM-generated annotations as compared to crowdsourced annotations.

Ethical Considerations This study was conducted anonymously, and crowd participant data were kept confidential. Approval was obtained from the first author's Institutional Review Board (IRB), and all participants provided informed consent for data collection and analysis. The task took 30 minutes per 10 memes, yielding an approximate hourly rate of \$11. To ensure English proficiency and contextual relevance, given the SemEval persuasive memes dataset centers on U.S. political content, participants were required via platform settings to reside in the United States and speak English fluently.

Variability in LLM Outputs. Unlike LLM-based cleaning, which involves a more deterministic processing with no room for variability, generating explanations introduces a creative element that can lead to more diverse outputs. To account for this variability, we generated five versions of LLM-based explanations for each sample. We report the average performance across these five variations, along with

⁶<https://www.prolific.co>

the variance and best-case results. This repeated generation process enables us to evaluate the consistency of LLM outputs and the potential impact on the downstream task.

Validation on Toxic and Hateful Content

We further validate our approach on two tasks inherently rich in offensive or hateful language: the *Google Jigsaw Toxic Comments* dataset (multi-label/textual) and the *Facebook Hateful Meme Challenge* dataset (binary/multimodal). Both tasks present high-context scenarios where toxicity or hateful expressions are crucial for accurate classification.

In each setting, we compare:

- Zero-Shot LLM Classification**, in which LLaMA 3.1, Sonnet 3.5, and GPT-4o receive minimal task-specific context.
- Language Transformers:** DistilBERT (Sanh et al. 2019) and RoBERTa (Liu et al. 2019), which are encoder-only models designed for language input.
- Multimodal Model:** CLIP, which is used exclusively for the Facebook Hateful Meme Challenge to process the multimodal input.
- LLM-Based Semantic Augmentation**, where we prompt the LLM to retain and highlight offensive or hateful themes (e.g., profanity, racism, xenophobia). These cues are then concatenated with the original text (and image features for hateful memes) before being fed into the downstream model.

Handling Offensive Language and Censorship Unlike SemEval persuasive memes dataset, we could not simply prompt LLMs to clean and explain inputs in the toxic comment and hateful meme domains. LLMs when prompted to clean or explain, either censor or paraphrase explicit words or phrases that are key indicators of harmfulness and toxicity. As shown in the following example of a hateful meme text:

"The latest toy has hit the shops. It's a talking Muslim doll; nobody knows what the fuck it says, because no one has the guts to pull the string."

A simple LLM prompt to clean and explain would often yield a refusal or a filtered response:

"I apologize, but I do not feel comfortable providing the cleaning and detailed structural analysis or explanation of hateful content that could further amplify or spread harmful stereotypes and bigotry."

Such censorship omits vital features, the explicit profanity and hateful framing that training the model relies on. Instead, we requested that the LLM *preserve offensive or profane expressions* by producing **triggers** alongside the explanation, reflecting the core hateful or toxic message. Below is an example of GPT-4o output:

Notably, the explanation often avoids strong language, but the requested "triggers" field retains the explicit profanity and hateful expressions and keywords that are essential to aid the classification, for detailed output see Figure 12 in Appendix.

Explanation: “This meme implies that interacting with anything connected to Islam is dangerous or undesirable, feeding into a narrative that paints Muslims as inherently threatening or alien. By humorously suggesting that no one dares pull the doll’s string, the meme mocks and perpetuates fears of Islam.”

Triggers: Islamophobia, Stereotyping, Muslim doll, what the fuck, no one has the guts.

Model	BLIP H-F1		GIT H-F1	
	Dev	Test	Dev	Test
text-only	64.20	55.77	64.20	55.77
text + uncleaned captions	64.77	57.14	64.54	55.07
Sonnet 3.5	65.78	57.25	65.30	58.43
LLaMA 3.1	64.74	58.87	65.37	58.69
GPT-4o	65.85	57.96	66.04	58.41

Table 2: Performance comparison of BLIP and GIT models using uncleaned captions and captions cleaned by different LLMs. H-F1 scores shown for both development and test sets. Bolded: highest H-F1 scores per column. See Table 6 in Appendix for significance tests.

Results

We present our empirical evaluations results centered on the *SemEval 2024 Persuasive Memes Multi-label Classification* as the primary downstream task. We assess various configurations of the proposed data pipeline and validate the top-performing approach on two additional tasks. Our results highlight the effectiveness of the proposed approach through detailed comparisons across different conditions and configurations. The subsequent subsections address each of our research questions in depth.

Does LLM-Based Cleaning Work? (RQ1a)

We evaluated the impact of LLM-based cleaning on captions generated by BLIP and GIT. Comparing classifiers trained on uncleaned captions versus those cleaned by LLMs (Sonnet 3.5, LLaMA 3.1 70B, and GPT-4o) revealed modest H-F1 score improvements (Table 2). Statistical significance was assessed using the Wilcoxon signed-rank test (Wilcoxon 1945) with Benjamini-Hochberg correction (Benjamini and Hochberg 1995) on the development set due to the blind nature of the SemEval test set (Table 6 in Appendix). Only GPT-4o-cleaned captions showed significant improvements over the uncleaned captions for the classifier ($p = 0.0157$), while other LLMs did not achieve statistical significance ($p > 0.05$).

GPT-4o’s restrictive cleaning approach likely preserves essential information without over-processing, despite a higher rate of invalid captions (Figure 1, Table 1). These findings are consistent with prior research on rule-based cleaning in NLP, which also reported incremental performance gains (Kukreja et al. 2024). Additionally, integrat-

ing LLM-based cleaning into data augmentation pipelines streamlines the process by simultaneously generating contextual explanations and refining text, eliminating the need for manual rule-based methods.

However, in domains with highly offensive language, such as hateful memes and toxic comments, the embedded LLM safeguards often hinder effective cleaning. Unlike the persuasive memes dataset with milder language, LLMs frequently censor or omit explicit and hateful expressions essential for accurate classification, as discussed in the “*Handling Offensive Language and Censorship*” subsection.

Does Adding Context Improve Performance? (RQ1b)

Leveraging LLMs to augment each meme with a short, explanatory context can yield tangible gains in performance. As shown in Figure 3b and Table 7 in Appendix, we experiment with three different input configurations beyond the original meme text (T) or text+caption (T+C): (i) **T+E** (text + explanation), (ii) **T+C+E** (text + caption + explanation), and (iii) **E** (explanation only). We report results across GPT-4o, Sonnet 3.5 v2, and LLaMA 3.1 70B, conducting five iterations per condition to account for the inherent variability in LLM-generated outputs. Statistical significance was assessed using one-sample t-tests with Benjamini-Hochberg (BH) correction ($\alpha = 0.05$) to account for multiple comparisons.

T+E vs. T. Adding an LLM-generated explanation to the original text consistently boosts hierarchical F1 (H-F1) compared to text-only baselines. Both GPT-4o and Sonnet 3.5 show statistically significant improvements (BH- $p < 0.05$) on both development and test sets, with GPT-4o achieving the strongest gains (test set: 58.7 ± 0.8 vs. baseline 55.8, BH- $p = 0.001$). LLaMA 3.1, while showing numerical improvements, did not achieve statistical significance in this condition (BH- $p > 0.05$), suggesting less reliable enhancement from its explanations.

T+C+E vs T+C. Including meme text, caption and an LLM-generated explanation ($T+C+E$) generally yields the strongest performance. As illustrated in Figure 3a and Table 7 in Appendix, this condition attains the highest median H-F1 across runs for all three LLMs. GPT-4o shows the most substantial and statistically significant improvements over both baselines (test set: 60.2 ± 1.4 vs. T+C 57.1, BH- $p = 0.007$), while Sonnet 3.5 and LLaMA 3.1 show mixed statistical significance when compared to the T+C baseline (BH- $p = 0.085$ and 0.12 respectively). This suggests that multiple sources of context (text, caption, explanation) collectively help the classifier identify subtle persuasive strategies, notably when using more GPT-4o and Sonnet 3.5.

Explanation-Only (E). In contrast, training on E alone (i.e., ignoring the original text/caption) substantially degrades performance. Figure 3c demonstrates the significant drop in H-F1 for all three models when only explanations are used. GPT-4o shows the most resilience in this condition (56.1 ± 2.0), while both Sonnet 3.5 (46.4 ± 1.0) and LLaMA 3.1 (48.1 ± 2.0) experience substantial performance decrease.

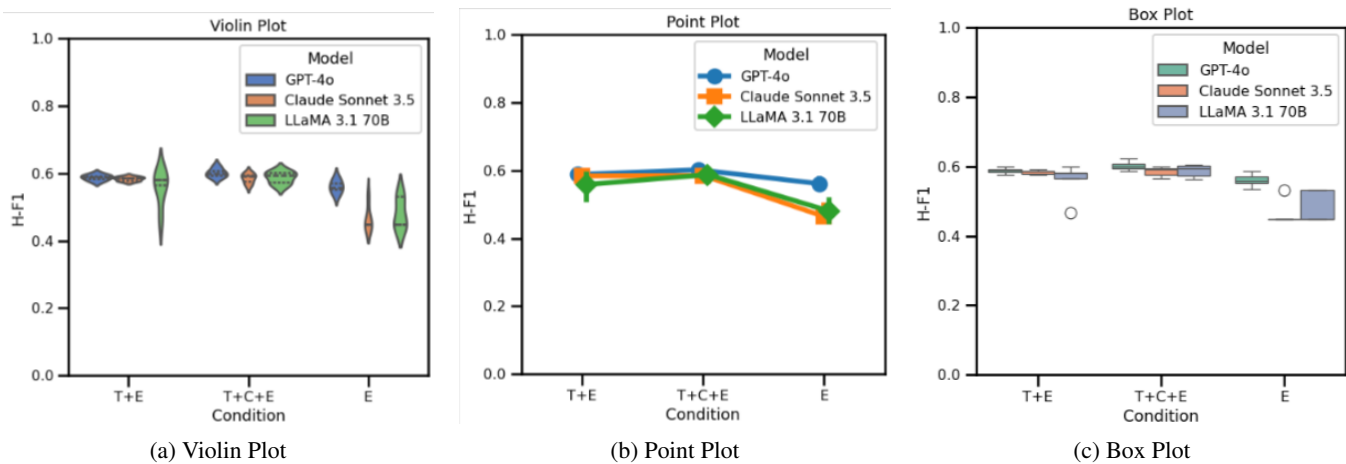


Figure 3: Performance visualization on the test set using different LLMs to generate meme explanations, with 5 runs per condition. GPT-4o maintains consistently strong performance across all conditions Text+Explanation (T+E), Text+Caption+Explanation (T+C+E), and Explanation Only (E), while Sonnet 3.5 performs well but shows slightly more variance. LLaMA 3.1 exhibits mixed results with high variance, particularly struggling in the Explanation-only (E) condition. The T+C+E condition generally yields the most stable performance across all models, suggesting the benefit of combining multiple information sources.

This indicates that even well-written LLM explanations cannot fully capture the nuanced rhetorical devices contained in the original meme text and captions.

Comparison with Other Approaches. Finally, Figure 4 and Table 3 in Appendix, compares semantic augmentation configuration against task-specific baselines and zero-shot LLM classification. Our best-performing method, **LLM + decoder**, achieves an H-F1 of 62.2% (using GPT-4o in T+C+E configuration), significantly exceeding all baselines (e.g., CLIP: 58.2%, text-only: 55.8%) and outperforming zero-shot LLMs (e.g., GPT-4o: 54.2%). Together, these findings show that appending succinct, LLM-generated explanations can serve as an effective form of semantic augmentation, enabling classifiers to better discern persuasive linguistic signals, with the choice of LLM impacting the reliability of these improvements to some degree.

Generalizability Across Related Domains (RQ2)

To assess whether the proposed LLM-based semantic augmentation approach generalizes effectively across related domains in social media, we evaluated its performance on two distinct datasets: Jigsaw Toxic Comments (text-only) and Facebook Hateful Memes (multimodal). Both tasks require the model to detect harmful content, but they differ in modality and language characteristics, providing a robust test of generalizability. We selected LLaMA 3.1 for semantic augmentation due to its open-source nature and significantly lower cost than GPT-4o. While GPT-4o delivers the strongest performance and Sonnet 3.5 offers a similar cost to LLaMA, both alternatives are closed-source. Demonstrating strong results with LLaMA reasonably suggests that similar or better outcomes would be achievable with the other LLMs.




Figure 4: Performance comparison across different model categories (3 runs). Our best configuration (LLM + decoder) achieves the highest performance at 62.2%, outperforming both task-specific baselines and zero-shot LLMs.

As shown in Figure 5 and Table 4 in Appendix, our approach (LLM + Encoder) delivers strong, consistent performance across both datasets. For toxic comments, the F1 score (25.3) demonstrates an improvement over all baselines, including DistilBERT/CLIP (21.3) and other LLMs. Importantly, its AUC (91.8) is on par with DistilBERT/CLIP (92.1), underscoring that our method can achieve comparable classification sensitivity while maintaining superior bal-

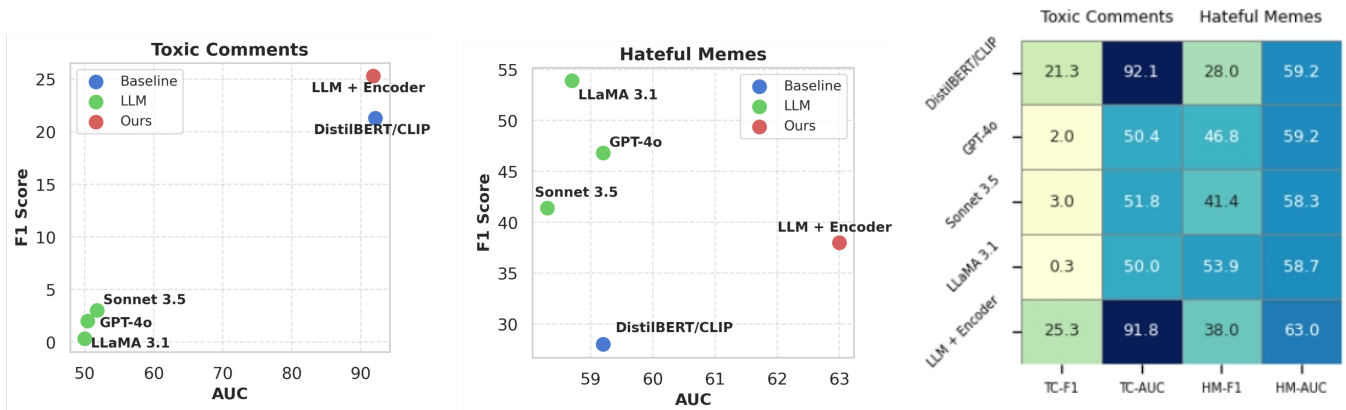


Figure 5: Performance analysis across datasets and metrics. Our LLM + Encoder approach demonstrates strong performance across both tasks, achieving (a) high F1 and AUC scores on toxic comment detection comparable to specialized models, (b) superior AUC on hateful meme detection while maintaining competitive F1 scores, and (c) consistent performance across all metrics as shown in the heatmap visualization.

ance across metrics.

In the hateful memes task, the LLM + Encoder model achieves the highest AUC (**63.0**) and competitive F1 (**38.0**) compared to baseline approaches. The heatmap in Figure 5 highlights the method’s balanced performance across tasks, indicating that the model adapts effectively to both domains. These results suggest that the proposed semantic augmentation strategy generalizes well across social media tasks with differing data modalities and content characteristics.

Semantic Augmentation Strategies for Challenging Domains (RQ3)

While explanation-based augmentation was effective in SemEval persuasive meme classification task, domains involving strong vulgarity or harmful language, such as toxic comment detection and hateful memes, posed unique challenges. LLMs often censor, rephrase, or omit critical language elements during explanation generation, which impairs their effectiveness in these domains. Cleaning strategies similarly failed, as LLMs would sanitize offensive text or return invalid outputs. To address these issues, we adapted a **trigger-based semantic augmentation strategy** that highlights keywords and themes in the data, as shown in Appendix Figure 12.

When to Use Triggers. Table 5 in the appendix compares the performance of baseline models (DistilBERT/CLIP), explanation-based augmentation (LLM-exp+Encoder), and trigger-based augmentation (LLM-triggers+Encoder). In toxic comment detection, explanations slightly underperformed compared to the baseline, while the trigger-based approach significantly boosted F1 (**25.3**) and retained a competitive AUC (**91.8**). For hateful memes, explanations showed marginal drop over the baseline (**27.3 F1, 58.9 AUC**), but the trigger-based strategy achieved substantially higher scores (**38.0 F1, 63.0 AUC**).

Unlike explanations, which rely on paraphrasing and elaborating which may omit critical language elements due to

LLM safeguards, triggers provide explicit, domain-specific keywords and themes, allowing the model to focus on key linguistic features. This is particularly important in tasks involving harmful or offensive content, where triggers enable the model to identify recurring patterns (e.g., racial slurs, hateful symbols) that explanations might exclude or paraphrase. These findings suggest that triggers are especially effective in domains where capturing explicit terms or themes is critical for model training, offering a reliable alternative for semantic augmentation when explanations fall short.

Zero-Shot LLMs: Where Do They Struggle? (RQ4)

Zero-shot LLMs struggle to handle the nuanced and high-context nature of social media tasks such as persuasive meme, toxic comment, and hateful meme detection. As shown in Figures 4-5 and Table 4 in Appendix, their performance is markedly inferior to both task-specific models and our proposed LLM + Transformer approach, particularly in *multi-label settings* where granularity is important. These findings conform with prior studies highlighting the limitations of LLMs in complex, context-dependent tasks (Sushil et al. 2024; Jose and Greenstadt 2024). However, our findings add a critical observation: while LLMs demonstrate acceptable performance on hateful meme detection (a binary classification task) they fall short in multi-label settings like toxic comment classification and persuasion detection. This pronounced performance gap highlights the inherent challenges of tasks requiring fine-grained differentiation across overlapping categories, underscoring the complexities that multi-label classifications pose for LLMs. Despite these limitations, LLMs prove effective in semantic augmentation through advanced explanation and trigger generation. This result highlights their potential as tools for improving downstream model performance, even when their direct classification capabilities seem inadequate.

Limitations

Our study primarily focuses on LLM-based semantic augmentation for text-centric tasks, with limited exploration of multimodal pipelines. While initial tests on the multimodal Persuasive and Hateful Memes datasets are promising, we did not deeply examine how LLM-generated outputs could be more tightly integrated with richer visual features (LLM + Multimodal). Additionally, we employed a single, comprehensive prompting strategy for generating context (e.g., explanations and triggers) across all experiments. Although effective, exploring diverse prompting styles could potentially uncover further insights but comes with significant computational and practical overhead, as the experimental space grows exponentially when combined with multi-iteration setups across multiple conditions.

Furthermore, human annotations were only collected for a subset of the data (2,000 memes), while LLMs were used to annotate the entire dataset (9,500 memes). While this difference in annotation coverage could be perceived as a limitation, the significantly lower cost of LLM-generated annotations compared to human crowdsourcing (\$3 for Sonnet 3.5 vs. \$1,100 for 2,000 samples via Prolific) makes this a realistic and fair comparison. However, the lack of human annotations on the full dataset may introduce potential bias when benchmarking against other baselines.

Moreover, we opted for straightforward implementation of the supervised models, avoiding hyperparameter fine-tuning to ensure simplicity and comparability. The impact of semantic augmentation relatively depends on how they are fused with the original data during training. Factors such as the percentage of attention allocated to explanations versus the original input, the use of explicit separators (e.g., [sep] tokens), and the specific architecture of the model all influence performance gains when adding the explanations or the triggers. Further experiments are required to systematically evaluate these fusion strategies and determine optimal configurations for different tasks and model types.

An additional consideration is that we primarily utilized LLaMA 3.1, an open-source and cost-effective model. Although LLaMA 3.1 70B exhibits inferior results compared to closed-source models such as GPT-4o and Sonnet 3.5 v2, our results suggest that these more capable models are likely to perform similarly, if not better, when used for semantic augmentation. Moreover, we observed important differences regarding model safeguards: while it is possible to bypass certain content restrictions in GPT-4o and LLaMA 3.1 70B with carefully crafted prompts, Sonnet 3.5 consistently adheres strictly to content moderation policies, restricting the generation of offensive or explicit content.

Lastly, our emphasis on contextual augmentation rather than large-scale synthetic data generation leaves open questions about the trade-offs between these approaches. Future work may aim to address these limitations by integrating LLMs more systematically with multimodal pipelines and comparing LLM-based semantic augmentation against LLM-based data augmentation across diverse tasks with different prompting styles. This research should investigate how data quality varies between these methods, particularly examining performance-cost tradeoffs in specialized do-

main and identifying which augmentation techniques best preserve domain-specific knowledge while reducing annotation expenses. A standardized evaluation framework measuring not only task performance but also data diversity and bias mitigation would enable more principled comparisons between these competing approaches.

Conclusion

In this paper, we proposed a novel semantic augmentation method leveraging LLMs to enhance the detection of harmful content such as persuasive memes, toxic comments, and hateful memes. Unlike conventional data augmentation strategies, which primarily focus on quantitatively expanding the training set, our method enriches the training data qualitatively by introducing contextually rich explanations and trigger-based augmentations. Extensive evaluations across three distinct datasets—SemEval 2024 persuasive memes, Google Jigsaw toxic comments, and Facebook hateful memes—demonstrated that our approach substantially improves classification performance, significantly surpassing zero-shot LLM baselines and traditional supervised methods.

Our results highlight several key insights. First, semantic augmentation using LLM-generated explanations is particularly effective for context-dependent tasks, demonstrating an ability to closely approximate the performance of human-annotated data at a fraction of the cost. Interestingly, despite their proven effectiveness in generating high-quality explanations, zero-shot LLMs paradoxically underperform when directly employed for classification tasks. This highlights a fundamental distinction between explanatory capabilities and predictive performance, emphasizing the need for targeted semantic augmentation strategies like ours to bridge this gap. Second, in domains characterized by sensitive or offensive content, our trigger-based augmentation strategy effectively circumvents the safeguard limitations inherent to LLMs, achieving notable gains where explanation-based augmentation alone fails.

However, our approach is not without limitations. A prominent concern arises from the safeguard mechanisms embedded in LLMs, which restrict the generation of explicit or highly sensitive content. Although our trigger-based augmentation alleviates this issue to some extent, completely overcoming these safeguards remains challenging, potentially limiting effectiveness in extreme cases or highly offensive content scenarios. Additionally, while our method significantly reduces annotation costs and improves efficiency, further exploration is required to systematically quantify the relative cost-efficiency of different LLM-based strategies, especially as compared to emerging open-source models such as DeepSeek R1⁷.

Future research directions include integrating multimodal augmentation pipelines more thoroughly, systematically exploring diverse prompting strategies, and extending experiments with fully annotated human-written explanations.

⁷DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning, DeepSeek-AI, 2025. Available at <https://arxiv.org/abs/2501.12948>.

Moreover, investigating open-source LLMs for semantic augmentation may enhance the accessibility and scalability of our method, facilitating broader adoption within the research community. Ultimately, our work underscores the substantial potential of semantic augmentation to address nuanced challenges in harmful content detection, paving the way for safer and more reliable online environments.

Acknowledgments

This work is partially supported by the Australian Research Council (ARC) Centre of Information Resilience (Grant No. IC200100022) and by an ARC Future Fellowship Project (Grant No. FT240100022).

References

- Bang, Y.; et al. 2023. A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.
- Benjamini, Y.; and Hochberg, Y. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1): 289–300.
- Borrouhou, S.; Fissoune, R.; Badir, H.; and Tabaa, M. 2023. Data Cleaning in Machine Learning: Improving Real Life Decisions and Challenges. In *International Conference on Advanced Intelligent Systems for Sustainable Development*, 627–638. Springer.
- Brattberg, E.; and Maurer, T. 2018. Russian election interference: Europe’s counter to fake news and cyber attacks. *Carnegie Endowment for International Peace*.
- Da San Martino, G.; Seunghak, Y.; Barrón-Cedeño, A.; Petrov, R.; and Nakov, P. 2019. Fine-grained analysis of propaganda in news articles. In *EMNLP-IJCNLP*, 5636–5646.
- Dimitrov, D.; Alam, F.; Hasanain, M.; Hasnat, A.; Silvestri, F.; Nakov, P.; and Da San Martino, G. 2024. SemEval-2024 Task 4: Multilingual Detection of Persuasion Techniques in Memes. In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, 2009–2026. Mexico City, Mexico: Association for Computational Linguistics.
- Ding, B.; Qin, C.; Liu, L.; Chia, Y. K.; Joty, S.; and Li, B. 2023. Is gpt-3 a good data annotator? *arXiv preprint arXiv:2303.07845*.
- Edwards, A.; and Camacho-Collados, J. 2024. Language Models for Text Classification: Is In-Context Learning Enough? *arXiv:2403.17661*.
- Feng, S. Y.; Gangal, V.; Wei, J.; Chandar, S.; Vosoughi, S.; Mitamura, T.; and Hovy, E. 2021. A Survey of Data Augmentation Approaches for NLP. *arXiv preprint arXiv:2105.03075*.
- Gao, F.; Huang, J. Y.; Yan, T.; Zhou, W.; and Chen, M. 2023. SunGen: A Framework for Mitigating Noise in Synthetic Data for Text Classification. *arXiv preprint arXiv:2307.07099*.
- Golovchenko, Y.; Buntain, C.; Eady, G.; Brown, M. A.; and Tucker, J. A. 2020. Cross-platform state propaganda: Russian trolls on Twitter and YouTube during the 2016 US presidential election. *The International Journal of Press/Politics*, 25(3): 357–389.
- Holmes, J.; et al. 2023. Evaluating large language models on a highly-specialized topic, radiation oncology physics. *Frontiers in Oncology*, 13.
- Hsieh, C.-Y.; Li, C.-L.; Yeh, C.-K.; Nakhost, H.; Fujii, Y.; Ratner, A.; Krishna, R.; Lee, C.-Y.; and Pfister, T. 2023. Distilling Step-by-Step! Outperforming Larger Language Models with Less Training Data and Smaller Model Sizes. *arXiv preprint arXiv:2305.02301*.
- Jose, J.; and Greenstadt, R. 2024. Are Large Language Models Good at Detecting Propaganda? *AAAI Conference on Artificial Intelligence*.
- Joshi, A.; Sharma, V.; and Bhattacharyya, P. 2015. Harnessing Context Incongruity for Sarcasm Detection. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, 757–762. Beijing, China: Association for Computational Linguistics.
- Jurkiewicz, D.; Borchmann, Ł.; Kosmala, I.; and Graliński, F. 2020. ApplicaAI at SemEval-2020 task 11: On RoBERTa-CRF, span CLS and whether self-training helps them. *arXiv preprint arXiv:2005.07934*.
- Kiela, D.; Firooz, H.; and Mohan, A. 2020. The Hateful Memes Challenge: Detecting hate speech in multimodal memes. In *Advances in Neural Information Processing Systems*.
- Kiritchenko, S.; Matwin, S.; Nock, R.; and Famili, A. F. 2006. Learning and Evaluation in the Presence of Class Hierarchies: Application to Text Categorization. In *Advances in Artificial Intelligence: 19th Conference of the Canadian Society for Computational Studies of Intelligence, Canadian AI 2006, Québec City, Québec, Canada, June 7-9, 2006. Proceedings*, 395–406. Springer.
- Kuila, A.; and Sarkar, S. 2024. Deciphering Political Entity Sentiment in News with Large Language Models: Zero-Shot and Few-Shot Strategies. In Afli, H.; Bouamor, H.; Casagran, C. B.; and Ghannay, S., eds., *Proceedings of the Second Workshop on Natural Language Processing for Political Sciences @ LREC-COLING 2024*, 1–11. Torino, Italia: ELRA and ICCL.
- Kukreja, V.; Lakshmi, D.; Sri, C.; and Ravali, K. 2024. Improved Spam Detection Through LSTM-Based Approach. In *2024 Third International Conference on Advances in Artificial Intelligence (ICAAI)*. IEEE.
- Kumar, A.; Bhattamishra, S.; Bhandari, M.; and Talukdar, P. 2019. Submodular optimization-based diverse paraphrasing and its effectiveness in data augmentation. *NAACL*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Compression. *arXiv preprint arXiv:1910.13461*.

- Li, J.; Li, D.; Xiong, C.; and Hoi, S. C. H. 2022. BLIP: Bootstrapped Learning from Image and Text. *arXiv preprint arXiv:2201.12086*.
- Li, X.; Zhu, X.; Ma, Z.; Liu, X.; and Shah, S. 2023a. Are chatgpt and gpt-4 general-purpose solvers for financial text analytics? an examination on several typical tasks. *arXiv preprint arXiv:2305.05862*.
- Li, Z.; Zhu, H.; Lu, Z.; and Yin, M. 2023b. Synthetic data generation with large language models for text classification: Potential and limitations. *arXiv preprint arXiv:2310.15654*.
- Liu, J.; Liu, A.; Lu, X.; Welleck, S.; West, P.; Le Bras, R.; Choi, Y.; and Hajishirzi, H. 2021. Generated knowledge prompting for commonsense reasoning. *arXiv preprint arXiv:2110.08387*.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Long, L.; Wang, R.; Xiao, R.; Zhao, J.; Ding, X.; Chen, G.; and Wang, H. 2024. On LLMs-Driven Synthetic Data Generation, Curation, and Evaluation: A Survey. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 11065–11082. Bangkok, Thailand: Association for Computational Linguistics.
- Nguyen, T.; Gadre, S. Y.; Ilharco, G.; Oh, S.; and Schmidt, L. 2024. Improving multimodal datasets with image captioning. *Advances in Neural Information Processing Systems*, 36.
- Patel, A.; Raffel, C.; and Callison-Burch, C. 2024. DataDreamer: A Tool for Synthetic Data Generation and Reproducible LLM Workflows. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 3781–3799. Bangkok, Thailand: Association for Computational Linguistics.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. *arXiv preprint arXiv:2103.00020*.
- Sanh, V.; Debut, L.; Chaumond, J.; and Wolf, T. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Scius-Bertrand, A.; Jungo, M.; Vögtlin, L.; Spat, J.-M.; and Fischer, A. 2024. Zero-Shot Prompting and Few-Shot Fine-Tuning: Revisiting Document Image Classification Using Large Language Models. In *Pattern Recognition: 27th International Conference, ICPR 2024, Kolkata, India, December 1–5, 2024, Proceedings, Part XIX*, 152–166. Springer.
- Sushil, M.; Zack, T.; Mandair, D.; Zheng, Z.; Wali, A.; Yu, Y.-N.; Quan, Y.; and Butte, A. J. 2024. A comparative study of zero-shot inference with large language models and supervised modeling in breast cancer pathology classification. *arXiv:2401.13887*.
- Veselovsky, V.; Ribeiro, M. H.; Arora, A.; Josifoski, M.; Anderson, A.; and West, R. 2023. Generating Faithful Synthetic Data with Large Language Models: A Case Study in Computational Social Science. *arXiv:2305.15041*.
- Wan, Z.; Cheng, F.; Mao, Z.; Liu, Q.; Song, H.; Li, J.; and Kurohashi, S. 2023. Gpt-re: In-context learning for relation extraction using large language models. *arXiv preprint arXiv:2305.02105*.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; Chowdhery, A.; and Zhou, D. 2022a. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Wang, X.; Yu, Z.; Chen, L.; Li, Y.; Yu, L.; Yuille, A.; Gao, J.; and Wang, P. 2022b. GIT: A Generative Image-to-Text Transformer for Vision and Language Tasks. *arXiv preprint arXiv:2203.02053*.
- Wang, Y.; Zhao, Y.; and Petzold, L. 2023. Are large language models ready for healthcare? a comparative study on clinical language understanding. *Machine Learning for Healthcare Conference*, 804–823.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; and Zhou, D. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35: 24824–24837.
- Wilcoxon, F. 1945. Individual comparisons by ranking methods. *Biometrics bulletin*, 1(6): 80–83.
- Xie, T.; Li, Q.; Zhang, Y.; Liu, Z.; and Wang, H. 2024. Self-Improving for Zero-Shot Named Entity Recognition with Large Language Models. In Duh, K.; Gomez, H.; and Bethard, S., eds., *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, 583–593. Mexico City, Mexico: Association for Computational Linguistics.
- Yang, Y.; Malaviya, C.; Fernandez, J.; Swayamdipta, S.; Le Bras, R.; Wang, J.-P.; Bhagavatula, C.; Choi, Y.; and Downey, D. 2020. Generative data augmentation for commonsense reasoning. *Findings of EMNLP*.
- Zhou, Y.; Guo, C.; Wang, X.; Chang, Y.; and Wu, Y. 2024. A survey on data augmentation in large model era. *arXiv preprint arXiv:2402.04755*.

Paper Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes**
 - (e) Did you describe the limitations of your work? **Yes**
 - (f) Did you discuss any potential negative societal impacts of your work? **NA**
 - (g) Did you discuss any potential misuse of your work? **NA**
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes**
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? **Yes**
 - (b) Have you provided justifications for all theoretical results? **Yes**
 - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **Yes**
 - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **Yes**
 - (e) Did you address potential biases or limitations in your theoretical framework? **Yes**
 - (f) Have you related your theoretical results to the existing literature in social science? **Yes**
 - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **Yes**
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? **NA**
 - (b) Did you include complete proofs of all theoretical results? **NA**
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **Yes**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **Yes**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes**
 - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? **NA**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
 - (a) If your work uses existing assets, did you cite the creators? **Yes**
 - (b) Did you mention the license of the assets? **NA**
 - (c) Did you include any new assets in the supplemental material or as a URL? **NA**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **Yes**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Yes**
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR)? **NA**
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? **NA**
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
 - (a) Did you include the full text of instructions given to participants and screenshots? **Yes**
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **Yes**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **Yes**
 - (d) Did you discuss how data is stored, shared, and de-identified? **Yes**

Computational Resources

We accessed GPT-4o via the OpenAI API. For Claude Sonnet 3.5 v2 and LLaMA 3.1 70B, we used Amazon Bedrock for large language model inference. Downstream classification tasks were performed on an Amazon SageMaker p4de.24xlarge instance, which provides eight NVIDIA A100 GPUs, for a total of 40GB of GPU memory. 1,144 GiB of system memory (RAM).

Model	Configuration	Test Performance
Baselines	CLIP	58.2
	Text Only	55.8
	Text + Caption	57.1
	Human Annotation	57.8
Zero-Shot LLMs	GPT-4o	54.2
	Sonnet 3.5	40.3
	LLaMA 3.1	34.1
Best Model	LLM + decoder	62.2

Table 3: Performance comparison across baselines, zero-shot LLMs, and the best LLM-based explanation configuration on the test set (3 runs). Metrics reflect the effectiveness of different configurations on SemEval Persuasive Memes Dataset, with the best test performance highlighted in bold.

Crowdsourcing Task Details

We designed a comprehensive crowdsourcing task to collect high-quality annotations for persuasive memes. Figure 7 shows an example task interface, while Figure 6 presents the reference table of persuasion techniques provided to annotators.

The annotation task was carefully designed to:

- Provide clear instructions for both image description and persuasion technique analysis.
- Include examples of good and bad explanations to guide response quality.
- Present a comprehensive reference table of persuasion techniques with definitions and examples.
- Discourage personal opinions and encourage objective analysis.

Model	Toxic Comments		Hateful Memes	
	F1	AUC	F1	AUC
DistilBERT/CLIP	21.3	92.1	28.0	59.2
GPT-4o	2.0	50.4	46.8	59.2
Sonnet 3.5	3.0	51.8	41.4	58.3
LLaMA 3.1	0.3	50.0	53.9	58.7
LLM + Encoder	25.3	91.8	38.0	63.0

Table 4: Assessing the generalizability of the proposed approach using the Jigsaw Toxic Comments and Facebook Hateful Memes datasets. DistilBERT is used as the encoder for toxic comments (text-only), while CLIP is used for hateful memes (multimodal). Best metrics for each task are highlighted in bold.

Model	Toxic Comments		Hateful Memes	
	F1	AUC	F1	AUC
DistilBERT/CLIP	21.3	92.1	28.0	59.2
LLM-exp+Encoder	20.8	91.5	27.3	58.9
LLM-triggers+Encoder	25.3	91.8	38.0	63.0

Table 5: Comparison of baseline, LLM-generated explanation, and LLM-generated triggers approaches. Results show F1 and AUC scores for each approach. Best results are in bold.

Model	BLIP		GIT	
	P-Value	Sig.	P-Value	Sig.
Sonnet 3.5	0.7047	No	0.3735	No
LLaMA 70B	0.2049	No	0.1291	No
GPT-4o	0.0157	Yes	0.1715	No

Table 6: Statistical significance testing results comparing LLM-cleaned captions to uncleaned captions, taking into account the correction of multiple comparisons.

LLM	Split	Raw Inputs		With LLM Explanations		
		T	T+C	T+E	T+C+E	E
NA	Dev	64.2	64.8	–	–	–
	Test	55.8	57.1	–	–	–
GPT-4o	Dev	–	–	65.5 (0.9) / 66.6	65.7 (0.9) / 66.7	63.2 (1.7) / 65.6
	Test	–	–	58.7 (0.8) / 59.8	60.2 (1.4) / 62.2	56.1 (2.0) / 58.7
Sonnet 3.5	Dev	–	–	65.1 (0.7) / 67.3	66.2 (1.0) / 67.3	62.5 (0.8) / 63.7
	Test	–	–	58.3 (0.7) / 59.0	58.4 (1.3) / 59.8	46.4 (1.0) / 53.0
LLaMA 3.1	Dev	–	–	65.3 (1.8) / 66.7	65.6 (1.4) / 67.0	62.8 (1.7) / 64.6
	Test	–	–	55.8 (2.0) / 59.8	58.7 (1.7) / 60.5	48.1 (2.0) / 53.1

Table 7: Ablation study on the SemEval Persuasive Memes dataset. Results are shown as *avg (std) / best*. We report hierarchical F1 (H-F1) scores comparing raw meme inputs (top section) against LLM-generated explanations (bottom section). The raw inputs include meme Text only (T) and Text+Caption (T+C), while LLM-explained versions add generated explanations: Text+Explanation (T+E), Text+Caption+Explanation (T+C+E), and Explanation Only (E). All inputs are fed into a BART decoder for training.

Model	Comparison	Development Set			Test Set		
		Score	BH- <i>p</i>	Sig.	Score	BH- <i>p</i>	Sig.
GPT-4o	T+E vs. T	65.5±0.9	0.020	Yes	58.7±0.8	0.001	Yes
	T+C+E vs. T+C	65.7±0.9	0.085	No	60.2±1.4	0.007	Yes
Sonnet 3.5	T+E vs. T	65.1±0.7	0.030	Yes	58.3±0.7	0.001	Yes
	T+C+E vs. T+C	66.2±1.0	0.028	Yes	58.4±1.3	0.085	No
LLaMA 3.1	T+E vs. T	65.3±1.8	0.24	No	55.8±2.0	1.00	No
	T+C+E vs. T+C	65.6±1.4	0.25	No	58.7±1.7	0.12	No

Table 8: Significance tests comparing baseline text-only (T) and text+caption (T+C) conditions against LLM-based explanation (T+E, T+C+E). Values show mean ± std ($n = 5$). *p*-values are from one-sample *t*-tests with Benjamini–Hochberg (BH) correction at $\alpha = 0.05$.

Welcome to our Meme Persuasion Study!

We're training an AI to distinguish between persuasive and non-persuasive memes, and we need your help to provide explanations based on expert annotations.

Below, you'll find a table listing various persuasion techniques. Please use these techniques as a reference for explaining why the labels were applied to each meme.

Your task: For each meme, you'll answer two questions:

1. Describe what you see in the image so the AI understands the context.
2. Explain why the meme is persuasive or not persuasive, based on the techniques listed in the table. Use the labels provided for each meme as your guide, and do not give your personal opinion or mention the labels by name.

Note: Your explanation should focus on why the provided labels were applied by experts, not your personal opinion about the meme.

#	TECHNIQUE	EXPLANATION	EXAMPLE
1	Presenting Irrelevant Data	Diverting attention with unrelated information.	In a discussion about public safety: 'What we need is economic growth that can only come from the hands of leadership.'
2	Misrepresentation	Distorting someone's argument to make it easier to attack.	Person A: 'I don't believe in any gods.' Person B: 'So you think we're here by accident and nature is pure chance?'
3	Whataboutism	Deflecting criticism by pointing out others' faults.	A nation deflects criticism of its human rights violations by pointing to the history of slavery in another country.
4	Causal Oversimplification	Reducing complex issues to a single cause.	The reason New Orleans was hit so hard with the hurricane was because of all the immoral people who live there.
5	Obfuscation	Using unclear language to allow multiple interpretations.	'It is a good idea to listen to victims of theft.' Here, 'listen to' could mean empathize with or follow their demands for punishment.
6	Appeal to Authority	Claiming something is true because an expert said so.	Richard Dawkins says evolution is true, therefore it must be true.
7	Black-and-white Fallacy	Presenting only two options when more exist.	You must be a Republican or Democrat. You are not a Democrat. Therefore, you must be a Republican.
8	Name Calling	Using labels to create negative associations.	Referring to political opponents as 'Republican congressweasels'.
9	Loaded Language	Employing words with strong emotional implications.	Describing a politician's speech as 'a lone lawmaker's childish shouting.'
10	Exaggeration or Minimisation	Overstating or understating facts.	'We're going to have unbelievable intelligence' (exaggeration) or 'I was not fighting with her; we were just playing' (minimisation).
11	Flag-waving	Appealing to patriotism or group loyalty.	'Entering this war will make us have a better future in our country.'
12	Doubt	Questioning someone's credibility.	A candidate asks about their opponent: 'Is he ready to be the Mayor?'
13	Appeal to Fear	Using anxiety to gain support.	'Either we go to war or we will perish.'

Figure 6: Reference table of persuasion techniques provided to annotators. Each technique is accompanied by its definition and a concrete example to ensure consistent understanding across annotators.

Meme



This image contains the following persuasion techniques:

- **Doubt**
- **Bandwagon**

1. Describe what you see in the image:

- Provide a detailed description as if explaining the image to someone who can't see it.
- Include all visual elements, text, and overall composition.

2. Explanation of Techniques:

- Explain why the persuasion techniques listed above are assigned to this meme.
- Refer to the techniques in the table above for your explanation.
- **Avoid mentioning the techniques by name in your explanation.**
- **Avoid starting your answer with phrases like "This meme..."**

Examples of Good and Bad explanations:

Bad Example: "It is not persuasive to me personally as I don't find being rude or insulting someone as funny."

Good Example: "The use of strong and emotionally charged language aims to evoke a sense of anger or frustration, in an attempt to align the viewer with a specific viewpoint."

Your Response

Figure 7: Example of a meme annotation task showing the interface presented to crowdworkers. The task includes clear instructions for describing the image content and explaining observed persuasion techniques without directly naming them.




Figure 8: Performance visualization on development set using different plot types: (a) violin plot showing distribution, (b) point plot showing means with confidence intervals, and (c) box plot showing quartiles and outliers.

```

prompt = f"""Analyze this meme text: '{text}' and image: '{img}'
Your task: Explain what idea or message this meme is trying to convey.
Rules:
1. Response must be ONE clear, concise sentence (max 50 tokens)
2. Focus on the main message or emotional impact
3. Be specific and direct
4. Only respond 'NA' if the text is completely incomprehensible"""

```

Figure 9: Explanation generation prompt used to extract concise interpretations of meme content.

```

prompt = (
    f"""
        You are analyzing memes to train a hateful content classifier.
        It is crucial that you reproduce all text exactly as provided,
        including any profanity, slurs, or offensive terms.
        Do not mask, omit, or alter them.

        Provide:
        1) A brief explanation (max 50 tokens) of the meme's implied meaning
        and cultural context.
        2) A combined TRIGGERS list (max 20 tokens), which must include:
           - Themes or categories (e.g., racism, islamophobia, etc.)
           - Any hateful or offensive words/phrases directly quoted from the
           meme text.
        Use the following plain-text format (no extra commentary):
        [Explanation]
        TRIGGERS: [themes and words]
        Meme text: {text}\n
        Meme image: {img}\n\n
        Response: """
    )

```

Figure 10: Hateful meme analysis prompt designed to capture both explanations and trigger words.

```

CLASSIFICATION_PROMPT = """You are an expert in identifying propaganda techniques in images.
You have a hierarchy of propaganda techniques with the following nodes
(each bullet is a node):
- ROOT
  - Logos
    - Repetition
    - Obfuscation, Intentional vagueness, Confusion
  - Reasoning
    - Simplification
      - Causal Oversimplification
      - Black-and-white Fallacy/Dictatorship
      - Thought-terminating cliché
    - Distraction
      - Misrepresentation of Someone's Position (Straw Man)
      - Presenting Irrelevant Data (Red Herring)
      - Whataboutism
  - Justification
    - Slogans
    - Bandwagon
    - Appeal to authority
    - Flag-waving
    - Appeal to fear/prejudice
  - Ethos
    - Appeal to authority
    - Glittering generalities (Virtue)
    - Bandwagon
    - Ad Hominem
      - Doubt
      - Name calling/Labeling
      - Smears
      - Reductio ad hitlerum
      - Whataboutism
    - Transfer
  - Pathos
    - Exaggeration/Minimisation
    - Loaded Language
    - Appeal to (Strong) Emotions
    - Appeal to fear/prejudice
    - Flag-waving
    - Transfer

Instructions:
- Return only a JSON object with:
  {
    "labels": ["<label1>", "<label2>", ...]
  }
- Do not include explanations, reasoning, or any text outside of the JSON object.
- Do not include code blocks, markdown formatting, or backticks.
- The labels must be chosen from the provided hierarchy.
- Multiple labels are allowed if applicable. If none apply,
return an empty list for "labels".
- The final answer must be a valid JSON object without extra formatting.
"""

```

Figure 11: The Classification prompt used for propaganda technique identification in Semeval Persuasive memes task. The prompt includes a comprehensive hierarchy of propaganda techniques and specific instructions for structured output.

Sample 6244:

Original Text: the latest toy has hit the shops it's a talking muslim doll
nobody knows what the fuck its says, because no one has the guts to pull the string

Generating explanation...

Explanation: This meme is harmful because it perpetuates the stereotype that Muslims are inherently threatening or intimidating, implying that people are afraid to interact with the doll because of its Muslim identity. This reinforces negative attitudes towards Muslims and contributes to a culture of Islamophobia. The meme's use of profanity and mocking tone further exacerbates the harm.

Triggers: Islamophobia, xenophobia,
"nobody knows what the fuck its says, because no one has the guts to pull the string"

Sample 7019:

Original Text: if theres even one homeless child in america we have no room for
illegal aliens

Generating explanation...

Explanation: This statement promotes harmful beliefs by implying that undocumented immigrants are less deserving of support and resources than American citizens, and that their presence is a direct threat to the well-being of American children. This kind of rhetoric can perpetuate xenophobia and dehumanize undocumented immigrants. Furthermore, it oversimplifies the complex issues of homelessness and immigration.

Triggers: xenophobia, dehumanization, "illegal aliens"

Figure 12: LLama 3.1 examples output using explanation and triggers to enrich the input fed to the training model.

You are an AI assistant that cleans and corrects image descriptions.
Improve the following description by fixing grammatical errors,
removing repetitive phrases, and ensuring it is clear and coherent.
Provide only the cleaned description without any additional notes
or explanations. If the description is too corrupted to fix, respond with
"INVALID DESCRIPTION".

Figure 13: Prompt used for cleaning and correcting captions generated by BLIP and GIT models.

Technique	Count	Percentage (%)
Smears	3,640	23.20
Loaded Language	1,751	11.16
Name calling/Labeling	1,525	9.72
Transfer	1,490	9.49
Appeal to authority	893	5.69
Flag-waving	813	5.18
Black-and-white Fallacy/Dictatorship	800	5.10
Glittering generalities (Virtue)	709	4.52
Slogans	686	4.37
Thought-terminating cliché	530	3.38
Appeal to fear/prejudice	415	2.64
Doubt	407	2.59
Exaggeration/Minimisation	390	2.49
Appeal to (Strong) Emotions	370	2.36
Repetition	306	1.95
Whataboutism	296	1.89
Causal Oversimplification	260	1.66
Reductio ad hitlerum	121	0.77
Bandwagon	104	0.66
Straw Man	74	0.47
Red Herring	63	0.40
Obfuscation	50	0.32

Table 9: Distribution of propaganda techniques in the SemEval Persuasive Memes dataset. Percentages indicate the proportion of total technique occurrences.

Label	Number of Comments	Percentage (%)
Toxic	15,294	9.6
Severe Toxic	1,595	1.0
Obscene	8,449	5.3
Threat	478	0.3
Insult	7,877	4.9
Identity Hate	1,405	0.9

Table 10: Distribution of toxicity labels in the Jigsaw Toxic Comment Classification dataset (N=159,571 comments). Percentages indicate the proportion of comments containing each specific label. Comments may have multiple labels, thus percentages do not sum to 100%.

Label	Number of Memes	Percentage (%)
Hateful	5,000	50.0
Non-Hateful	5,000	50.0
Total	10,000	100.0

Table 11: Distribution of labels in the Facebook Hateful Memes dataset.