



MEMEWEAVER: Inter-Meme Graph Reasoning for Sexism and Misogyny Detection

Paolo Italiani^{*1,2} David Gimeno-Gomez^{*2} Luca Ragazzi^{*1}
Gianluca Moro^{*1} Paolo Rosso^{*2,3}

¹Department of Computer Science and Engineering, University of Bologna, Italy

²PRHLT Research Center, Universitat Politècnica de València Spain,

³ValgrAI - Valencian Graduate School and Research Network of Artificial Intelligence, Spain
{paolo.italiani, l.ragazzi, gianluca.moro}@unibo.it, {dagigo1, proso}@dsic.upv.es

Abstract

Women are twice as likely as men to face online harassment due to their gender. Despite recent advances in multimodal content moderation, most approaches still overlook the social dynamics behind this phenomenon, where perpetrators reinforce prejudices and group identity within like-minded communities. Graph-based methods offer a promising way to capture such interactions, yet existing solutions remain limited by heuristic graph construction, shallow modality fusion, and instance-level reasoning. In this work, we present MEMEWEAVER, an end-to-end trainable multimodal framework for detecting sexism and misogyny through a novel inter-meme graph reasoning mechanism. We systematically evaluate multiple visual-textual fusion strategies and show that our approach consistently outperforms state-of-the-art baselines on the MAMI and EXIST benchmarks, while achieving faster training convergence. Further analyses reveal that the learned graph structure captures semantically meaningful patterns, offering valuable insights into the relational nature of online hate.¹

1 Introduction

Hate speech on social media is rarely an isolated act: it spreads through interactions among users. Psychological research (Walther, 2022) shows that perpetrators often gain social approval by sharing content with like-minded individuals, reinforcing prejudices and group identity. These dynamics, compounded with user anonymity (Kowalski and Whittaker, 2015), amplify harmful narratives and foster in-group codes and jokes, making content subtler, context-dependent, and harder to detect.

Within the spectrum of hate speech (Herz and Molnár, 2012), women remain one of the most disproportionately targeted groups. They are twice as

^{*}Equal contribution (co-first authors).

¹The code is publicly available at <https://github.com/disi-unibo-nlp/meme-weaver>

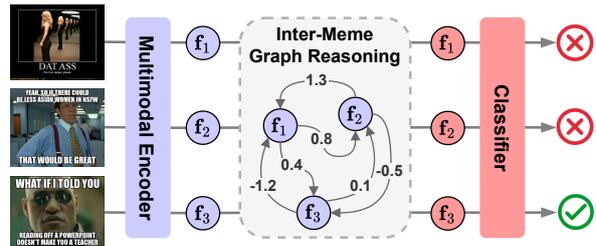


Figure 1: Overview of MEMEWEAVER, our graph-based framework for end-to-end sexism/misogyny detection in memes. Motivated by the social dynamics of online hate, it models batches of memes (3 shown) as inter-connected nodes for more effective learning.

likely as men to be harassed online due to their gender (Duggan, 2017), often within communities where such hateful discourses are normalized and perpetuated (Fontanella et al., 2024). Such abuse manifests primarily as *sexism*—prejudice or discrimination often rooted in stereotypes and socially accepted forms of bias (Glick and Fiske, 2001)—and *misogyny*, marked by deeper hostility, contempt, and explicit aggression (Richardson-Self, 2018). Persistent exposure to both has been shown to harm women’s self-esteem, restrict career ambitions, and reinforce traditional gender roles (Bradley-Geist et al., 2015). The rapid dissemination of such content in online communities makes effective regulation challenging (Herz and Molnár, 2012), motivating the development of automated detection methods (Luo et al., 2025).

Despite recent progress, most computational approaches overlook the social dynamics underlying online hate. Graph-based models (Zhang et al., 2019; Rehman et al., 2025; Xu et al., 2025) provide a promising avenue by capturing inter-sample relationships. However, existing solutions still rely on heuristic preprocessing, shallow modality fusion, and instance-level reasoning, which limit generalization and adaptability across domains.

In this paper, we introduce MEMEWEAVER (Fig-

ure 1), an end-to-end trainable framework for detecting sexism and misogyny in memes. Our approach integrates three key ideas. *First*, memes are inherently multimodal and context-dependent: they combine text embedded in images with visual scenes, and their meaning often hinges on humor, irony, or cultural references that may obscure harmful intent (Hodson et al., 2010; Drucker et al., 2014; Plaza et al., 2024; Cocchieri et al., 2025). To address this, we verbalize the text content of each meme and optionally enrich it with captions generated by a multimodal LLM, capturing both surface descriptions and higher-level interpretations. *Second*, we encode text and images with CLIP-based encoders and explore alternative fusion strategies that move beyond simple concatenation, enabling richer multimodal representations. *Third*, and most importantly, we introduce a novel Inter-Meme Graph Reasoning (IMGR) mechanism that models relationships across memes within each training batch. This design allows the system to automatically “weave” latent affinities between memes, capturing relational patterns without relying on heuristic graph construction.

We conduct extensive experiments on two complementary meme benchmarks: MAMI, a dataset targeting misogyny in English memes, and EXIST, a multilingual corpus focused on sexism with both English and Spanish memes. We compare MEMEWEAVER against established baselines as well as recent large multimodal language models (LLMs). Our main contributions are threefold:

- **State-of-the-art results.** MEMEWEAVER consistently outperforms strong baselines on both datasets, improving over CLIP fine-tuning by ≈ 5 points on average while converging faster during training.
- **Fusion analysis.** We evaluate alternative text-image fusion mechanisms, showing that their effectiveness depends on dataset characteristics and that higher model complexity does not necessarily yield better performance.
- **Graph insights.** We analyze the learned embeddings and inter-meme graph structure, revealing semantically meaningful affinity patterns that offer insights into the relational dynamics of online hate.

2 Related Work

Long-standing research has addressed the detection of sexism and misogyny in online content. This

section outlines the shift from early text-only methods to recent multimodal approaches, highlighting open challenges in graph-based reasoning.

Sexism and Misogyny Detection Early work relied on statistical text features (e.g., n-grams, TF-IDF) with classical machine learning classifiers (Jha and Mamidi, 2017; Anzovino et al., 2018). Deep learning, and in particular pre-trained transformers such as BERT, improved contextual understanding (Parikh et al., 2021; de Paula et al., 2023). More recently, prompt-based techniques with LLMs have been explored for nuanced comprehension (Tian et al., 2023; Samani et al., 2025). Given the multimodal nature of social media posts (Luo et al., 2025), research has shifted to text-vision integration (Gomez et al., 2020), especially in memes (Fersini et al., 2022; Plaza et al., 2025). Captioning-based methods (Rizzi et al., 2023; Plaza et al., 2024), cross-modal contrastive learning (Rizzi et al., 2024), multimodal LLMs (Cao et al., 2022; Xu et al., 2025), and prompt-based visual QA approaches such as ProCap (Cao et al., 2023) have been widely adopted. More recent studies extend to in-the-wild videos, incorporating voice and visual cues (Arcos and Rosso, 2024; Plaza et al., 2025). Despite these advances, systems still struggle with implicit, context-dependent sexism and with biases in data and models, pointing to the need for richer reasoning mechanisms such as graph-based approaches.

Multimodal Graph Modeling Graph-based models, particularly GCNs (Zhang et al., 2019), are increasingly used to capture structural and semantic relations across modalities, with applications in tasks such as summarization and QA (Moro et al., 2023c, 2024; Ragazzi et al., 2025), emotion recognition (Zhu et al., 2022) and scene-text retrieval (Mafla et al., 2021; Li et al., 2024). Similarly, graph-based approaches have been applied to network intrusion detection (Cerroni et al., 2013, 2015). In online hate speech detection, they have also shown promise. Xu et al. (2025) combined LLMs with hypergraph learning to integrate text and visual features and capture implicit semantics, while Rehman et al. (2025) enriched unimodal representations through gated cross-attention rather than simple concatenation. However, both approaches relied on similarity-based heuristics to define the graph. Only Hebert et al. (2024) introduced learnable components, though their attention-based method still depends on heuristics such as

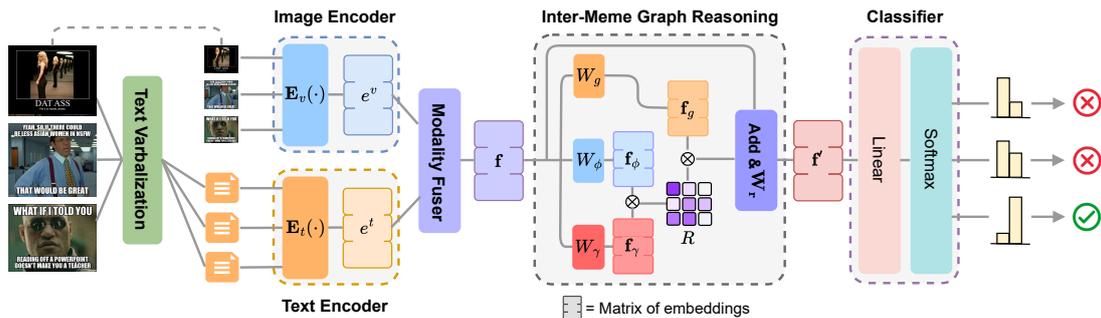


Figure 2: **Architecture of MEMEWEAVER.** Each meme undergoes text extraction (OCR, optionally enriched with LLM-generated captions) and is encoded by separate text and image encoders. The embeddings are fused, refined via inter-meme graph reasoning within each batch, and classified. Example shown with batch size $m = 3$.

the number of user mentions in social media posts.

Research Gaps Despite recent progress, two challenges remain central. First, most multimodal hate speech detection methods—including those targeting sexism and misogyny—rely on shallow fusion strategies, often limited to feature concatenation. Second, existing graph-based approaches typically construct structures through heuristic rules, restricting adaptability and limiting their ability to capture inter-sample relations. These limitations call for fully learnable frameworks that can jointly model multimodal content and social interactions in an end-to-end manner. Our work explores whether such a design can uncover generalizable relational patterns or whether optimal configurations are inherently domain-specific.

3 MEMEWEAVER

We introduce MEMEWEAVER, a framework for detecting sexist and misogynistic memes, which “weaves” together textual and visual features while modeling relations among memes at the batch level.

3.1 Problem Statement

Let $\mathcal{X} = \{x_1, \dots, x_n\}$ be a multimodal dataset of memes, processed in mini-batches of size m . Each instance $x_j = (t_j, v_j)$ consists of a textual component t_j (OCR-extracted text, optionally enriched with an LLM-generated caption) and an image v_j , annotated with a binary label $y_j \in \{0, 1\}$, where 1 denotes a sexist/misogynistic meme and 0 otherwise. The goal is to train a multimodal classifier f that, given a novel meme (t, v) , predicts its label y .

3.2 Model Architecture

Figure 2 shows the architecture of MEMEWEAVER. Each meme is an image that includes both a vi-

sual scene and overlaid text. Our pipeline begins with a verbalization step, which extracts the text embedded in the image.² This content may be optionally enriched with an LLM-generated caption. The extracted text is then encoded by a text encoder, while the complete meme image is processed by a separate image encoder. The two encoders yield modality-specific representations, which are subsequently fused into a multimodal embedding. To capture inter-sample relationships, we weave connections among memes within each training batch, modeling the batch as a fully connected graph. Finally, the context-enriched embeddings are fed into a classifier, with all modules trained end-to-end.

LLM-based Captioning Inspired by Xu et al. (2025), we leverage a multimodal LLM to generate not only surface-level descriptions but also high-level interpretations of memes. Specifically, we used Qwen2.5-VL-7B-Instruct³ and designed two prompt strategies (see Figure 3). The generated captions are concatenated with the OCR-extracted meme text, separated by the special token [CPT]; we evaluate both OCR-only and OCR+caption variants. Examples are provided in Appendix C.

Text Encoder Given a batch of memes, we process their textual components $\{t_1, \dots, t_m\}$ with a transformer-based text encoder $\mathbf{E}_t(\cdot)$, yielding hidden representations:

$$e_j^t = \mathbf{E}_t(t_j), \quad j = 1, \dots, m$$

$$e^t = [e_1^t, \dots, e_m^t] \in \mathbb{R}^{m \times \frac{d}{2}} \quad (1)$$

where $d = 1536$ denotes the overall multimodal hidden dimensionality, so that each encoder outputs

²In the datasets used for our experiments, this text annotation is already provided by the dataset creators.

³Run with default precision and decoding settings from the official configuration file.

Prompt A (Surface-level description)

Describe this image without including what text reads and credit sources.

Prompt B (Multimodal semantic inference)

You are a helpful assistant designed to detect {HATE_TYPE} expressions or behaviours in a meme, i.e., it is {HATE_TYPE} itself, describes a {HATE_TYPE} situation or criticizes a {HATE_TYPE} behaviour. Infer the implicit semantic information of the meme, considering that it may or may not contain {HATE_TYPE} content. Please be concise (no more than three sentences) while including all relevant information.

Figure 3: **Prompts used for LLM-based meme captioning.** The HATE_TYPE placeholder was set to *misogynistic* or *sexist* depending on the downstream dataset.

vectors of size $d/2$.

Image Encoder Similarly, given the same batch of memes, we process the meme images $\{v_1, \dots, v_m\}$ with an image encoder $\mathbf{E}_v(\cdot)$, yielding hidden representations:

$$\begin{aligned} e_j^v &= \mathbf{E}_v(v_j), \quad j = 1, \dots, m \\ e^v &= [e_1^v, \dots, e_m^v] \in \mathbb{R}^{m \times \frac{d}{2}} \end{aligned} \quad (2)$$

with the same dimensionality convention as in the text encoder. Both the text and image encoders are initialized from CLIP and fine-tuned jointly with the rest of the architecture.

Modality Fuser In addition to simple concatenation, we consider two fusion mechanisms to combine the text and image embeddings into unified multimodal representations $\mathbf{f} = [f_1, \dots, f_m]$.

(1) *Multi-modal Factorized Bilinear (MFB) pooling* (Kim et al., 2017) projects each modality into a shared space and applies element-wise interaction:

$$f_j = (e_j^t U^\top) \circ (e_j^v V^\top), \quad j = 1, \dots, m \quad (3)$$

where $U, V \in \mathbb{R}^{\frac{d}{2} \times d}$ are trainable projection matrices and \circ denotes the Hadamard product.

(2) *Gated Multimodal Unit (GMU)* (Ovalle et al., 2017) learns a gating vector to adaptively control each modality’s contribution:

$$\begin{aligned} z_j &= \sigma([e_j^t, e_j^v] U_z), \\ f_j &= z_j \odot \tanh(e_j^v U_v^\top) \\ &\quad + (1 - z_j) \odot \tanh(e_j^t U_t^\top), \quad j = 1, \dots, m \end{aligned} \quad (4)$$

where $U_t, U_v \in \mathbb{R}^{\frac{d}{2} \times d}$ and $U_z \in \mathbb{R}^{d \times d}$ are trainable matrices, σ is the sigmoid, $[\cdot, \cdot]$ denotes concatenation, and \odot denotes multiplication.

Inter-Meme Graph Reasoning (IMGR) Inspired by prior work on relational reasoning (Li et al., 2019; Mafra et al., 2021), we enhance multimodal representations by modeling relationships *across* memes in a batch, rather than within a single instance (e.g., linking image regions or textual tokens). We compute an affinity matrix $R \in \mathbb{R}^{m \times m}$ capturing pairwise meme similarities:

$$R = (\mathbf{f} W_\phi)(\mathbf{f} W_\gamma)^\top \quad (5)$$

where $\mathbf{f} = [f_1, \dots, f_m] \in \mathbb{R}^{m \times d}$ are fused embeddings and $W_\phi, W_\gamma \in \mathbb{R}^{d \times d}$ are trainable projections. This yields a fully connected graph $G = (\mathbf{f}, R)$, with nodes f_i (memes) and edge weights R_{ij} (their relations). Following Li et al. (2019), we then apply a graph-based message passing layer with residual connections:

$$\mathbf{f}' = (\hat{R} \mathbf{f} W_g) W_r + \mathbf{f} \quad (6)$$

where $W_r, W_g \in \mathbb{R}^{d \times d}$ are trainable projections and $\hat{R} = \frac{R}{m}$ is the uniform rescaled affinity matrix. The final representations $\mathbf{f}' = [f'_1, \dots, f'_m]$ are enriched with batch-level context, enabling each meme to incorporate information from its peers.

Classifier and Loss Function The classifier consists of a linear projection $W_c \in \mathbb{R}^{d \times 2}$ producing logits, followed by a softmax that yields probabilities $P = [p_1, \dots, p_m] \in \mathbb{R}^{m \times 2}$, where each $p_j \in \mathbb{R}^2$ is the predicted distribution for meme j . All modules of our architecture are trained end-to-end using cross-entropy loss:

$$\mathcal{L}_{ce} = -\frac{1}{m} \sum_{j=1}^m \sum_{c \in \{0,1\}} \mathbf{1}[y_j = c] \log p_{j,c} \quad (7)$$

where $y_j \in \{0, 1\}$ is the gold label for instance j .

4 Experimental Setup

4.1 Datasets

To capture domain and cultural diversity, we select two distinct multimodal meme datasets for misogyny and sexism detection, framing the task as binary classification following prior approaches (Cao et al., 2023; Xu et al., 2025). This setting implicitly introduces a cross-domain learning scenario, where differences in platform, language, annotation

Dataset	No. of Samples (% Misogyny/Sexism)		
	Training	Validation	Test
MAMI	9,000 (50.0%)	1,000 (50.0%)	1,000 (50.0%)
EXIST	3,235 (65.7%)	404 (67.3%)	405 (65.7%)

Table 1: **Dataset statistics.** We highlight the percentage of sexist and misogynistic memes in each split.

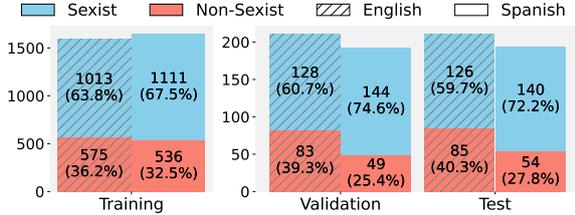


Figure 4: **Language distribution in EXIST.** Ratio of sexist vs. non-sexist memes across splits.

schemes, and cultural context induce distribution shifts similar to those studied in Domeniconi et al. (2014a,b, 2016, 2017); Moro et al. (2018).

(1) **MAMI** (Fersini et al., 2022) is a curated SemEval dataset focused specifically on *misogyny* in memes, collected from Twitter, Reddit, and meme-focused websites. It provides fine-grained subtype annotations such as shaming, stereotyping, objectification, and incitement to violence; we assign the misogyny label if a meme belongs to any subtype. This widely used benchmark serves as our primary dataset for misogyny detection in English memes.

(2) **EXIST** (Plaza et al., 2025) is a CLEF shared task on *sexism* detection across social media, including tweets and videos. Here we focus on Task 2, targeting memes. The dataset was built from 250 sexism-related queries on Google Images, and includes subtype labels such as ideological inequality, stereotyping and dominance, objectification, and sexual violence. Since each meme was annotated by three annotators, we derived the final sexism label via majority voting. Unlike MAMI, EXIST contains memes in both English and Spanish, allowing cross-lingual and cross-cultural evaluation.

Table 1 shows the dataset distributions. MAMI is balanced (50% misogynistic memes), while EXIST is skewed toward sexist content ($\approx 66\%$). Figure 4 further breaks down EXIST by language and class across splits: the training set is relatively balanced, but the skew becomes more pronounced in the validation and test sets, especially for Spanish. For MAMI we rely on the official splits, whereas for EXIST we define a custom 80/10/10 train/val/test

Setting	Range	MAMI	EXIST
Training batch	[10, 180]	20	64
Inference batch	[10, 180]	41	27
Class. threshold	[0.010, 0.999]	0.657	0.434
Modality fusion	N/A	MFB	Concat
Image captioning	N/A	–	Prompt A

Table 2: **Optimal settings.** Ranges denote search intervals; final values were selected on the validation sets.

split, since no validation set is provided and the official test set is not publicly available.⁴ Both datasets already come with OCR-extracted text.

4.2 Baselines

We use the widely adopted multimodal model CLIP (Radford et al., 2021)⁵ as the main backbone of MEMEWEAVER. We choose CLIP over larger multimodal LLMs because its lightweight design allows for larger effective batch sizes, thereby facilitating richer inter-meme relationships within IMGR. Unlike for EXIST, we use the original test set for MAMI, and thus report results against established baselines: EF-CaTrBERT (Khan and Fu, 2021), PromptHate (Cao et al., 2022), ProCap (Cao et al., 2023), and HyperHatePrompt (Xu et al., 2025). In addition, we compare against recent multimodal LLMs, such as Qwen2.5-VL (Bai et al., 2025), Phi-4 (Abdin et al., 2024), and GPT-4o-mini (Hurst et al., 2024), evaluated in a zero-shot setting using instruction-following prompts (details are provided in Appendix B).

4.3 Implementation Details

Environment All experiments run on a workstation equipped with an NVIDIA GeForce RTX 3090 GPU (24 GB VRAM), 64 GB system RAM, and an Intel® Core™ i9-10900X CPU @ 3.70GHz. Our implementation builds on Python 3.10.12, PyTorch 2.2.1, and HuggingFace Transformers 4.44.0.

Multilingual Setup For EXIST, we translated the Spanish split into English using Google Translate, since CLIP is English-only. Given the short and simple nature of the texts, automatic translation is highly reliable. We manually inspected the entire translated set and found no systematic errors or meaning distortions, ensuring negligible translation bias. Additionally, we conducted an automatic evaluation of the translation quality (see details in Ap-

⁴This ensures consistent internal evaluation, though results are not directly comparable to leaderboard scores.

⁵HuggingFace checkpoint: openai/clip-vit-large-patch14

Method	MAMI <i>English</i>			<i>All</i>			EXIST <i>English</i>			<i>Spanish</i>		
	Acc	F1	AUC	Acc	F1	AUC	Acc	F1	AUC	Acc	F1	AUC
Phi-4	69.8	70.0	-	61.4	62.5	-	66.4	65.8	-	58.3	56.3	-
GPT-4o-mini	71.7	73.1	-	64.4	64.4	-	70.1	70.1	-	57.8	58.2	-
Qwen2.5-VL	74.4	74.8	-	71.4	<u>70.0</u>	-	72.5	<u>72.1</u>	-	70.1	66.7	-
MEMEWEAVER												
w/ IMGR	77.6*	77.4*	<u>83.4*</u>	73.3	67.9*	<u>76.9*</u>	71.6	68.4*	<u>77.0*</u>	75.3	66.3*	76.5*
w/o IMGR	71.9	71.8	79.4	71.1	60.5	71.2	69.2	61.6	74.8	73.2	58.3	65.9
+ Prompt A												
w/ IMGR	74.8	74.4	80.1	76.3	71.6*	75.7*	76.3*	73.6*	72.9	<u>75.3</u>	<u>67.2</u>	66.4
w/o IMGR	69.2	68.9	77.5	72.1	61.8	73.5	68.2	59.5	75.1	76.3	64.4	72.1
+ Prompt B												
w/ IMGR	<u>76.3*</u>	<u>76.6*</u>	83.5*	<u>73.6</u>	67.3	78.9*	<u>73.9</u>	71.3	80.4*	73.2	59.1	<u>76.4</u>
w/o IMGR	<u>72.7</u>	<u>71.9</u>	78.8	<u>73.3</u>	69.7	74.3	<u>73.5</u>	71.0	76.3	73.2	67.7	<u>72.3</u>

Table 3: **Comparison to state of the art on test sets.** Best results in bold, second best underlined. * denotes statistically significant improvements ($\alpha = 0.05$, Koehn, 2004). Ablation studies confirm the importance of IMGR.

	Acc	F1	AUC
Prior Work			
EF-CaTrBERT (2021)	67.8	67.2	74.9
PromptHate (2022)	71.1	70.8	80.8
Pro-Cap (2023)	73.3	72.4	83.2
HyperHatePrompt (2025)	<u>75.3</u>	<u>75.1</u>	84.3
MEMEWEAVER (Ours)	77.6	77.4	<u>83.4</u>

Table 4: **Our best MEMEWEAVER vs. state of the art on MAMI.** Best scores in bold, second-best underlined.

pendix A). This choice enables a unified backbone across datasets, avoiding confounding effects from heterogeneous multilingual encoders and allowing a fair assessment of IMGR. It also represents a practical strategy for handling multilingual memes when only monolingual encoders are available.

Method Settings For the EXIST dataset, which is relatively low-resource in terms of available training data—thus also simulating a more realistic scenario with limited supervision (Moro and Ragazzi, 2022, 2023)—we initialize from the checkpoint pretrained on MAMI before fine-tuning. All models are trained for 3 epochs with the AdamW optimizer (learning rate $5e-6$) and a linear scheduler, using a fixed random seed of 42 for reproducibility. Training and inference batch sizes are tuned on the validation set, as they directly determine how many instances can interact within the IMGR module and thus critically impact MEMEWEAVER’s performance. Finally, rather than adopting the default 0.5 classification threshold, we follow Xu et al. (2025) and select the cutoff that maximizes F1 on the validation set. Table 2 summarizes all final

hyperparameters, including not only optimization settings but also the selected fusion strategy and the use of LLM-based captioning along with the best-performing prompt.

Evaluation Metrics For each model, we retain the checkpoint with the highest macro-F1 score on the validation set. We report three complementary metrics: Accuracy (Acc), macro-F1, and Area Under the Curve (AUC). Among these, macro-F1 serves as our primary metric, as it equally weights both classes and is therefore more reliable under the label imbalance in the EXIST dataset.⁶

5 Results & Discussion

Performance Comparison Table 3 shows that MEMEWEAVER with IMGR consistently outperforms ablated variants without graph reasoning, particularly on MAMI and EXIST-English, with gains of 7-14 F1 points depending on the configuration, yielding more discriminative feature representations across memes. This effectiveness is further supported by Table 4, where we establish a new state of the art on MAMI over prior work.

Performance differences across datasets reflect the subtler, more complex nature of sexist discourse in EXIST compared to the more explicit, hate-driven memes in MAMI. This may also explain why incorporating complementary image captioning is more beneficial in the sexism case. Yet, the limited impact of semantic captioning (Prompt B) aligns with the weaker performance of our zero-shot LLM experiments, suggesting that even these

⁶For generative LLMs, which output hard labels rather than probability distributions, AUC is not reported.

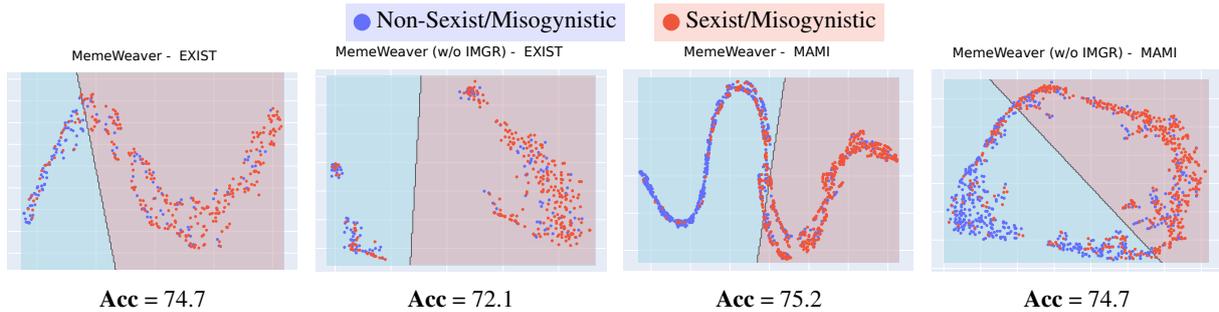


Figure 5: **Embedding space analysis.** 2D visualizations of MAMI and EXIST test set embeddings f' and f , generated with and without MEMEWEAVER’s IMGR, respectively. Decision boundaries from a linear classifier show higher accuracy on f' (see scores below plots), underscoring the benefit of IMGR’s representations.

powerful systems struggle with context-dependent social phenomena. This contrasts with prior studies that emphasized captioning as crucial for multimodal hate detection (Cao et al., 2022, 2023). Furthermore, beyond improved overall performance, our method exhibits greater robustness across languages compared to LLMs. Specifically, LLMs show a substantial cross-lingual discrepancy, with an average F1 gap of 8.9 ± 2.7 between the English and Spanish EXIST partitions.

Overall, the results highlight the central role of learnable graph reasoning and robust multimodal fusion, which proved to adapt more effectively than both LLMs and hand-crafted heuristics.

Training Convergence Beyond boosting performance, the IMGR module substantially improves CLIP’s learning dynamics. As shown in Figure 6, the standard CLIP (dash lines) exhibits a roughly linear, gradual decrease in training loss over the full schedule. In contrast, MEMEWEAVER plunges almost immediately—especially on MAMI—halving its loss in under 20% of the training steps and reaching a stable plateau by 40%. Although the effect is less pronounced on the EXIST corpus, it still yields a faster descent and a lower steady-state loss than the vanilla model. This faster convergence also translates into reduced computational cost and energy consumption, making MEMEWEAVER a more environmentally sustainable solution—an increasingly important criterion that should be weighed alongside accuracy in model evaluation (Moro et al., 2023b; Ragazzi et al., 2024).

Modality Fusion Strategy Table 5 shows the impact of different fusion strategies on our best configurations across both datasets. On EXIST, naive concatenation performs surprisingly well, outperforming both MFB and GMU. In particular, GMU

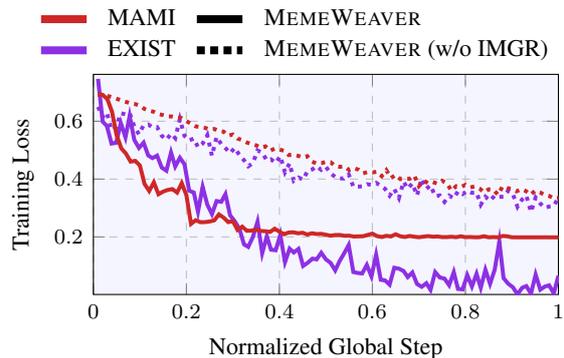


Figure 6: **Training loss evolution of MEMEWEAVER with and without IMGR across datasets.** Loss is plotted over normalized global training steps.

Fusion	MAMI			EXIST		
	Acc	F1	AUC	Acc	F1	AUC
Concat	73.6	73.5	82.0	76.3	71.6	75.7
MFB	77.6	77.4	83.4	<u>74.6</u>	<u>70.5</u>	<u>75.7</u>
GMU	<u>74.4</u>	<u>74.3</u>	<u>82.1</u>	71.6	60.3	73.6

Table 5: **Comparison of modality fusion strategies.** Best results are in bold and second-best are underlined.

performs poorly, suggesting that gated mechanisms may be unstable in low-resource settings, where dynamic reweighting can amplify noise. In contrast, both MFB and GMU yield substantial gains on MAMI, with MFB achieving the best overall results. These findings highlight the data-dependent nature of fusion effectiveness and stress the need for systematic evaluation, rather than assuming that higher model complexity ensures better integration.

Inference Batch Size Batch size is a key hyperparameter, as it determines graph depth and the number of inter-instance relations. Figure 7 shows that F1 improves with larger inference-time batches, plateauing around 20. On MAMI, perfor-

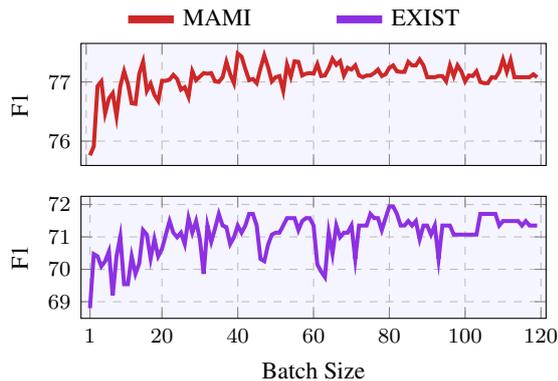


Figure 7: F1 scores evolution as inference batch size increases for MAMI (top) and EXIST (bottom).

mance remains stable beyond this point, while on EXIST it fluctuates more, likely due to the smaller training set and greater sensitivity to this hyperparameter. However, the model stays competitive even with small batches: with size one, performance drops only slightly, especially on MAMI, despite the absence of inter-instance connections. We attribute this robustness to training, where MEMEWEAVER reaches a much lower loss minimum than when trained without IMGR (see Figure 6). Thus, even when reasoning cannot be fully exploited at inference time, the model continues to benefit from its stronger training dynamics.

Embedding Space Analysis To better understand the contribution of IMGR, we compare embeddings \mathbf{f} (without IMGR) and \mathbf{f}' (with IMGR). We first apply PCA (Abdi and Williams, 2010) to reduce embeddings to 50 dimensions, followed by t-SNE (Maaten and Hinton, 2008) for 2D projection. As shown in Figure 5, the \mathbf{f}' embeddings exhibit clearer separation w.r.t. ground-truth labels. To quantify this effect, we train a linear logistic classifier on the 2D PCA \rightarrow t-SNE embeddings. The 5-fold cross-validation accuracies, reported below each plot in Figure 5, confirm the superior discriminative power of \mathbf{f}' , supporting our hypothesis and explaining the performance gains of our variant.

Embedding Affinity Analysis To better understand the connections learned by IMGR, we examined the relationship between meme similarity and affinity scores R . Using fused multimodal embeddings \mathbf{f} , we computed pairwise cosine similarities across memes and grouped them by class alignment (both misogynistic/sexist, both non-misogynistic/non-sexist, or mixed). We deliberately avoided the refined embeddings \mathbf{f}' , as each

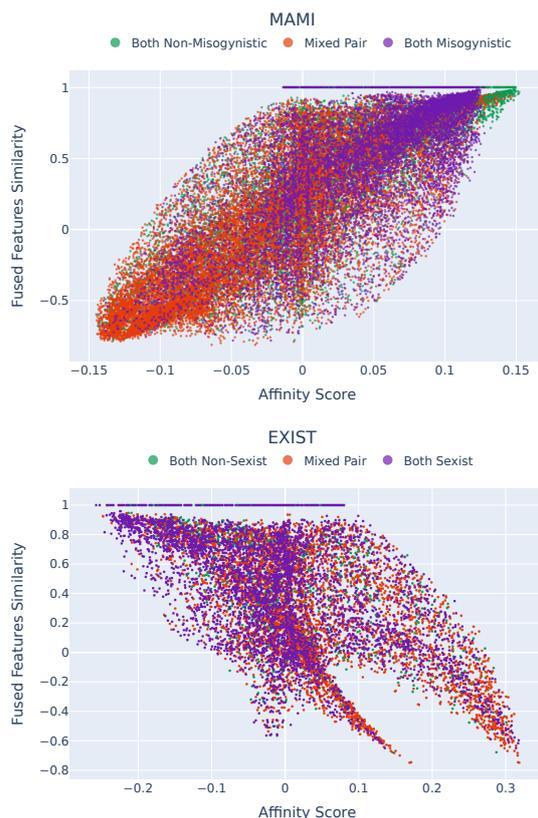


Figure 8: **Affinity patterns in IMGR.** Correlation between learned affinity scores and fused embedding similarity, with dataset-specific behaviors.

incorporates information from all others, which would distort pairwise similarity estimates.

As shown in Figure 8, MAMI exhibits a strong positive Pearson correlation ($r = 0.819$): similar memes tend to receive stronger affinity scores, reinforcing their connections. This contrasts with the heuristic of Xu et al. (2025), which links dissimilar instances to maximize contrastive learning. EXIST, instead, aligns more closely with that heuristic, though with a weaker negative correlation ($r = -0.5937$). A post hoc test further confirmed the importance of these learned patterns: simply inverting affinity signs at inference time caused a severe performance drop (≈ 30 points), suggesting that graph structures must adapt to dataset characteristics rather than fixed heuristics.

We also find that positive pairs cluster in the upper half of the figure, indicating thematic coherence within classes and reflecting the social interplay. Finally, both datasets contain many meme pairs with affinity scores near zero, producing a sparse graph that avoids over-connecting unrelated memes and preserves meaningful distinctions.

Backbone	MEME WEAVER	MAMI			EXIST		
		Acc	F1	AUC	Acc	F1	AUC
CLIP	✗	71.9	71.8	79.4	68.2	59.5	75.1
	✓	77.6	77.4	83.4	76.3	73.6	72.9
XLM-R/ViT	✗	66.8	66.8	74.7	71.1	66.1	67.2
	✓	72.5	72.4	78.6	72.5	63.5	71.4

Table 6: **Performance of MEMEWEAVER with different backbone architectures.** Best results are in bold.

Backbone Variant Beyond CLIP, we explore a custom backbone for MEMEWEAVER using XLM-RoBERTa (Conneau et al., 2020) as the text encoder and Vision Transformer (Dosovitskiy et al., 2021) as the image encoder. Table 6 compares the XLM-R/ViT variant to our best-performing CLIP-based settings (no prompts for MAMI and Prompt A for EXIST; see Table 3). Although the custom backbone does not surpass the CLIP-based model, MEMEWEAVER continues to outperform standard baselines on most metrics (with the exception of F1 on EXIST), highlighting the framework’s robustness and adaptability across architectures.

6 Conclusion

We introduced MEMEWEAVER, a fully learnable graph-based framework for multimodal detection of sexism and misogyny in memes. Our method consistently outperforms the state of the art on challenging benchmarks, highlighting the benefits of batch-level graph reasoning and advanced multimodal fusion. Analyses of embedding spaces and affinity patterns revealed meaningful structures and dataset-dependent behaviors, stressing the need to adapt modeling assumptions to context. These results open avenues for more nuanced relation modeling and fine-grained archetype categorization in multimodal hate speech detection (Rizzi et al., 2023), as well as extending MEMEWEAVER to broader hate detection tasks.

Limitations

Despite its contributions, our work has several limitations that warrant further exploration.

First, while we experimented with LLMs for image captioning, we did not integrate them as backbones for MEMEWEAVER, which could enhance semantic and contextual understanding; yet this was infeasible due to computational limits, particularly constraints on batch size, which is a critical hyperparameter of our graph-based framework.

Second, we evaluated multimodal LLMs in a

zero-shot setting, which is also their most common usage scenario; however, fine-tuned comparisons remain an interesting avenue for future research.

Third, we only tested two prompt types for captioning (surface-level vs. semantic inference) and did not explore other prompts or reasoning chains, which may improve visual–textual grounding.

Fourth, for EXIST-Spanish, we relied on machine translation, which we manually verified and found reliable; however, translation may still lose subtle cultural or linguistic cues, and future research should employ native multilingual encoders.

Finally, the inter-meme graph is built within fixed training batches, which ensures efficiency but may restrict the capture of global structures or cross-batch relations; future work should investigate scalable or memory-augmented graph reasoning to improve global coherence.

Ethical Considerations

While MEMEWEAVER improves automated detection of harmful content targeting women, it risks bias, including misflagging marginalized communities and missing culturally specific sexism. Automated moderation should therefore be paired with human oversight to prevent unjust removals or oversight of harmful content.

Acknowledgments

The work of Gómez was partially supported through Grant CIACIF/2021/295 funded by GVA. The work of Rosso was in the framework of the ANNOTATE-MULTI2 project (PID2024-156022OB-C32) funded by MICIU/AEI/10.13039/501100011033 and by ERDF/EU. The work of Italiani, Ragazzi, and Moro was partially supported by AI-PACT project (CUP B47H22004450008, B47H22004460001); National Plan PNC-I.1 DARE initiative (PNC0000002, CUP B53C22006450001); PNRR Extended Partnership FAIR (PE00000013, Spoke 8); 2024 Scientific Research and High Technology Program, project “AI analysis for risk assessment of empty lymph nodes in endometrial cancer surgery”, the Fondazione Cassa di Risparmio in Bologna; Chips JU TRISTAN project (G.A. 101095947). We extend our thanks to the Maggioli Group⁷ for partially supporting the Ph.D. scholarship granted to Italiani.

⁷<https://www.maggioli.com/who-we-are/company-profile>

References

- Hervé Abdi and Lynne J Williams. 2010. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459.
- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, and 1 others. 2024. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*.
- Maria Anzovino, Elisabetta Fersini, and Paolo Rosso. 2018. Automatic Identification and Classification of Misogynistic Language on Twitter. In *Proc. of NLDB*, pages 57–64. Springer.
- Iván Arcos and Paolo Rosso. 2024. Sexism identification on tiktok: A multimodal ai approach with text, audio, and video. In *Proc. of CLEF*, pages 61–73. Springer.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-v1 technical report. *arXiv preprint arXiv:2502.13923*.
- Jill C Bradley-Geist, Ivy Rivera, and Susan D Geringer. 2015. The Collateral Damage of Ambient Sexism: Observing Sexism Impacts Bystander Self-Esteem and Career Aspirations. *Sex Roles*, 73:29–42.
- Rui Cao, Ming Shan Hee, Adriel Kuek, Wen-Haw Chong, Roy Ka-Wei Lee, and Jing Jiang. 2023. ProCap: Leveraging a Frozen Vision-Language Model for Hateful Meme Detection. In *Proc. of the 31st ACM ICM*, pages 5244–5252.
- Rui Cao, Roy Ka-Wei Lee, Wen-Haw Chong, and Jing Jiang. 2022. Prompting for Multimodal Hateful Meme Classification. In *Proc. of EMNLP*, pages 321–332.
- Walter Cerroni, Gianluca Moro, Roberto Pasolini, and Marco Ramilli. 2015. Decentralized detection of network attacks through p2p data clustering of snmp data. *Computers & Security*, 52:1–16.
- Walter Cerroni, Gianluca Moro, Tommaso Pirini, and Marco Ramilli. 2013. Peer-to-peer data mining classifiers for decentralized detection of network attacks. In *Proceedings of the Twenty-Fourth Australasian Database Conference-Volume 137*, pages 101–107.
- Alessio Cocchieri, Luca Ragazzi, Paolo Italiani, Giuseppe Tagliavini, and Gianluca Moro. 2025. “what do you call a dog that is incontrovertibly true? dogma”: Testing LLM generalization through humor. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22922–22937, Vienna, Austria. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. *Unsupervised cross-lingual representation learning at scale*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 8440–8451. Association for Computational Linguistics.
- A de Paula, G Rizzi, E Fersini, and D Spina. 2023. AI-UPV at EXIST 2023–Sexism Characterization Using Large Language Models Under The Learning with Disagreements Regime. In *CEUR WORKSHOP PROCEEDINGS*, volume 3497, pages 985–999. CEUR-WS.
- Giacomo Domeniconi, Gianluca Moro, Andrea Pagliarani, Roberto Pasolini, and 1 others. 2017. On deep learning in cross-domain sentiment classification. In *Proceedings of the 9th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management-(Volume 1)*, volume 1, pages 50–60. SciTePress.
- Giacomo Domeniconi, Gianluca Moro, Roberto Pasolini, and Claudio Sartori. 2014a. Cross-domain text classification through iterative refining of target categories representations. In *International Conference on Knowledge Discovery and Information Retrieval*, volume 2, pages 31–42. SciTePress.
- Giacomo Domeniconi, Gianluca Moro, Roberto Pasolini, and Claudio Sartori. 2014b. Iterative refining of category profiles for nearest centroid cross-domain text classification. In *International Joint Conference on Knowledge Discovery, Knowledge Engineering, and Knowledge Management*, pages 50–67. Springer.
- Giacomo Domeniconi, Konstantinos Semertzidis, V. López, Elizabeth M. Daly, Spyros Kotoulas, and Gianluca Moro. 2016. A novel method for unsupervised and supervised conversational message thread detection. In *International Conference on Data Technologies and Applications*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Ari Drucker, Ofer Fein, Dafna Bergerbest, and Rachel Giora. 2014. On Sarcasm, Social Awareness, and Gender. *Humor*, 27(4):551–573.
- Maeve Duggan. 2017. Online Harassment. https://www.pewresearch.org/internet/wp-content/uploads/sites/9/2017/07/PI_2017.07.11_Online-Harassment_FINAL.pdf.
- Xiachong Feng, Xiaocheng Feng, and Bing Qin. 2022. Msamsun: Towards benchmarking multi-lingual dialogue summarization. In *DialDoc@ACL 2022*, pages 1–12. ACL.

- Elisabetta Fersini, Francesca Gasparini, Giulia Rizzi, Aurora Saibene, Berta Chulvi, Paolo Rosso, Alyssa Lees, and Jeffrey Sorensen. 2022. [SemEval-2022 Task 5: Multimedia Automatic Misogyny Identification](#). In *Proc. of the 16th SemEval@NAACL*, pages 533–549. ACL.
- Lara Fontanella, Berta Chulvi, Elisa Ignazzi, Annalina Sarra, and Alice Tontodimamma. 2024. How do we study misogyny in the digital age? A systematic literature review using a computational linguistic approach. *Humanities and Social Sciences Communications*, 11(1):1–15.
- Peter Glick and Susan T Fiske. 2001. Hostile and Benevolent Sexism as Complementary Justifications for Gender Inequality. *American psychologist*, 56(2):109.
- Raul Gomez, Jaume Gibert, Lluís Gomez, and Dimosthenis Karatzas. 2020. [Exploring Hate Speech Detection in Multimodal Publications](#). In *Proc. of the IEEE/CVF WACV*, pages 1470–1478.
- Liam Hebert, Gaurav Sahu, Yuxuan Guo, Nanda Kishore Sreenivas, Lukasz Golab, and Robin Cohen. 2024. Multi-Modal Discussion Transformer: Integrating Text, Images and Graph Transformers to Detect Hate Speech on Social Media. In *Proc. of AAAI*, volume 38, pages 22096–22104.
- Michael Herz and Péter Molnár. 2012. *The Content and Context of Hate Speech: Rethinking Regulation and Responses*. Cambridge University Press.
- Gordon Hodson, Jonathan Rush, and Cara C MacInnis. 2010. A Joke is Just a Joke (Except When It Isn't): Cavalier Humor Beliefs Facilitate the Expression of Group Dominance Motives. *Journal of personality and social psychology*, 99(4):660.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Akshita Jha and Radhika Mamidi. 2017. When Does a Compliment Become Sexist? Analysis and Classification of Ambivalent Sexism using Twitter Data. In *Proc. of the Workshop on NLP and Computational Social Science*, pages 7–16.
- Zaid Khan and Yun Fu. 2021. Exploiting BERT for Multimodal Target Sentiment Classification through Input Space Translation. In *Proc. of ACM ICM*, pages 3034–3042.
- Elizabeth C. Khoong, Eric Steinbrook, Catherine Brown, and Alicia Fernandez. 2019. [Assessing the use of google translate for spanish and chinese translations of emergency department discharge instructions](#). *JAMA Internal Medicine*, 179(4):580–582.
- Jin-Hwa Kim, Kyoung Woon On, Woosang Lim, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. 2017. [Hadamard product for low-rank bilinear pooling](#). In *Proc. of ICLR*. OpenReview.net.
- Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pages 388–395.
- Robin M Kowalski and Elizabeth Whittaker. 2015. Cyberbullying: Prevalence, Causes, and Consequences. *The Wiley handbook of Psychology, Technology, and Society*, pages 142–157.
- Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. 2019. Visual Semantic Reasoning for Image-Text Matching. In *Proc. of the IEEE/CVF ICCV*, pages 4654–4662.
- Qian Li, Lixin Su, Jiashu Zhao, Long Xia, Hengyi Cai, Suqi Cheng, Hengzhu Tang, Junfeng Wang, and Dawei Yin. 2024. Text-Video Retrieval via Multimodal Hypergraph Networks. In *Proc. of the ACM ICWSDM*, pages 369–377.
- Xuan Luo, Bin Liang, Qianlong Wang, Jing Li, Erik Cambria, Xiaojun Zhang, Yulan He, Min Yang, and Ruifeng Xu. 2025. [A Literature Survey on Multimodal and Multilingual Sexism Detection](#). *IEEE Trans. on Computational Social Systems*, pages 1–19.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.
- Andres Mafra, Sounak Dey, Ali Furkan Biten, Lluís Gomez, and Dimosthenis Karatzas. 2021. Multi-Modal Reasoning Graph for Scene-Text based Fine-Grained Image Classification and Retrieval. In *Proc. of the IEEE/CVF WACV*, pages 4023–4033.
- Gianluca Moro, Andrea Pagliarani, Roberto Pasolini, Claudio Sartori, and 1 others. 2018. Cross-domain & in-domain sentiment analysis with memory-based deep neural networks. In *Proceedings of the 10th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management, IC3K2018, Volume 1: KDIR, Seville, Spain, September 18-20, 2018.*, pages 125–136. SciTePress.
- Gianluca Moro, Nicola Piscaglia, Luca Ragazzi, and 1 others. 2023a. [Multi-language transfer learning for low-resource legal case summarization](#). *Artificial Intelligence and Law*, pages 1–29.
- Gianluca Moro and Luca Ragazzi. 2022. [Semantic self-segmentation for abstractive summarization of long documents in low-resource regimes](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):11085–11093.
- Gianluca Moro and Luca Ragazzi. 2023. [Align-then-abstract representation learning for low-resource summarization](#). *Neurocomputing*, 548:126356.

- Gianluca Moro, Luca Ragazzi, and Lorenzo Valgimigli. 2023b. [Carburacy: Summarization models tuning and comparison in eco-sustainable regimes with a novel carbon-aware accuracy](#). In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pages 14417–14425. AAAI Press.
- Gianluca Moro, Luca Ragazzi, and Lorenzo Valgimigli. 2023c. [Graph-based abstractive summarization of extracted essential knowledge for low-resource scenarios](#). In *ECAL, September 30 - October 4, 2023*, volume 372 of *FAIA*, pages 1747–1754. IOS Press.
- Gianluca Moro, Luca Ragazzi, Lorenzo Valgimigli, Fabian Vincenzi, and Davide Freddi. 2024. [Revelio: Interpretable long-form question answering](#). In *The Second Tiny Papers Track at ICLR 2024, Tiny Papers @ ICLR 2024, Vienna, Austria, May 11, 2024*. OpenReview.net.
- John Edison Arevalo Ovalle, Tamar Solorio, Manuel Montes-y-Gómez, and Fabio A. González. 2017. [Gated multimodal units for information fusion](#). In *Proc. of ICLR*. OpenReview.net.
- Pulkit Parikh, Harika Abburi, Niyati Chhaya, Manish Gupta, and Vasudeva Varma. 2021. Categorizing Sexism and Misogyny through Neural Approaches. *ACM Trans. on Web*, 15(4):1–31.
- Laura Plaza, Jorge Carrillo-de-Albornoz, Iván Árcos, Paolo Rosso, Damiano Spina, Enrique Amigó, Julio Gonzalo, and Roser Morante. 2025. [EXIST 2025: Learning with Disagreement for Sexism Identification and Characterization in Tweets, Memes, and TikTok Videos](#). In *Proc. of the 47th ECIR*, volume 15576 of *Lecture Notes in Computer Science*, pages 442–449. Springer.
- Laura Plaza, Jorge Carrillo-de-Albornoz, Víctor Ruiz, Alba Maeso, Berta Chulvi, Paolo Rosso, Enrique Amigó, Julio Gonzalo, Roser Morante, and Damiano Spina. 2024. Overview of EXIST 2024 – Learning with Disagreement for Sexism Identification and Characterization in Social Networks and Memes (Extended Overview). In *Working Notes of CLEF*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR.
- Luca Ragazzi, Paolo Italiani, Gianluca Moro, and Mattia Panni. 2024. [What are you token about? differentiable perturbed top-k token selection for scientific document summarization](#). In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 9427–9440. Association for Computational Linguistics.
- Luca Ragazzi, Gianluca Moro, Lorenzo Valgimigli, and Riccardo Fiorani. 2025. [Cross-document distillation via graph-based summarization of extracted essential knowledge](#). *IEEE Transactions on Audio, Speech and Language Processing*, 33:518–527.
- Mohammad Zia Ur Rehman, Sufyaan Zahoor, Areeb Manzoor, Musharaf Maqbool, and Nagendra Kumar. 2025. A Context-Aware Attention and Graph Neural Network-based Multimodal Framework for Misogyny Detection. *Information Processing & Management*, 62(1):103895.
- Louise Richardson-Self. 2018. Woman-Hating: On Misogyny, Sexism, and Hate Speech. *Hypatia*, 33(2):256–272.
- Giulia Rizzi, Francesca Gasparini, Aurora Saibene, Paolo Rosso, and Elisabetta Fersini. 2023. Recognizing misogynous memes: Biased models and tricky archetypes. *Information Processing & Management*, 60(5):103474.
- Giulia Rizzi, David Gimeno-Gómez, Elisabetta Fersini, and Carlos-D Martínez-Hinarejos. 2024. PINK at EXIST2024: A Cross-Lingual and Multi-Modal Transformer Approach for Sexism Detection in Memes. *Working Notes of CLEF*.
- Ali Riahi Samani, Tianhao Wang, Kangshuo Li, and Feng Chen. 2025. Large Language Models with Reinforcement Learning from Human Feedback Approach for Enhancing Explainable Sexism Detection. In *Proc. of the 31st ICCL*, pages 6230–6243.
- Breena R. Taira, Vanessa Kreger, Andrea Orue, and Lisa C. Diamond. 2021. [A pragmatic assessment of google translate for emergency department instructions](#). *Journal of General Internal Medicine*, 36(11):3361–3365.
- Lin Tian, Nannan Huang, and Xiuzhen Zhang. 2023. Efficient multilingual sexism detection via large language model cascades. In *CLEF (Working Notes)*, pages 1083–1090.
- Joseph B. Walther. 2022. Social Media and Online Hate. *Current Opinion in Psychology*, 45:101298.
- Bo Xu, Erchen Yu, Jiahui Zhou, Hongfei Lin, and Linlin Zong. 2025. HyperHatePrompt: A Hypergraph-based Prompting Fusion Model for Multimodal Hate Detection. In *Proc. of the 31st ICCL*, pages 3825–3835. ACL.
- Si Zhang, Hanghang Tong, Jiejun Xu, and Ross Maciejewski. 2019. Graph Convolutional Networks: A Comprehensive Review. *Computational Social Networks*, 6(1):1–23.

Junnan Zhu, Qian Wang, Yining Wang, Yu Zhou, and 1 others. 2019. *NCLS: Neural cross-lingual summarization*. In *EMNLP-IJCNLP*, pages 3054–3064, Hong Kong, China. ACL.

Tong Zhu, Leida Li, Jufeng Yang, Sicheng Zhao, and Xiao Xiao. 2022. Multimodal Emotion Classification with Multi-Level Semantic Reasoning Network. *IEEE Trans. on Multimedia*, 25:6868–6880.

A Translation Quality

Assessing translation quality is important for evaluating the reliability of the Spanish-to-English preprocessing step. However, gold-standard English references are not available for the EXIST dataset, which prevents the direct evaluation for Spanish-to-English translation. To address this limitation, we estimated translation quality using a standard *back-translation* protocol (Zhu et al., 2019; Feng et al., 2022). Specifically, the original Spanish text (Spanish₁) was translated into English using Google Translate, and the resulting English text was then translated back into Spanish (Spanish₂). Translation quality was assessed by comparing Spanish₁ and Spanish₂ using automatic overlap-based metrics. Using this approach, we obtained a BLEU-1 score of 57.30 and a ROUGE-1 score of 76.44. These results indicate a high degree of lexical overlap between the original and back-translated Spanish text, suggesting that the English intermediate translation preserved most of the core semantic content. This outcome is consistent with the characteristics of EXIST, where meme texts are typically short, informal, and syntactically simple.

Moreover, the observed translation quality is consistent with prior findings in the literature, which have consistently shown that Google Translate achieves particularly high accuracy for Spanish-to-English translation, one of its strongest-performing language pairs (Khoong et al., 2019; Taira et al., 2021; Moro et al., 2023a). Overall, these results suggest that the translation step introduces limited distortion and does not compromise the reliability of the downstream analysis.

B Prompts

Figure 9 reports the prompts used for zero-shot experiments using LLMs. For the sake of reproducibility, the definitions of “misogynistic” and “sexist” included in these prompts were taken directly from the original MAMI and EXIST task guideline descriptions.

MAMI

Consider the following definitions. 1. ‘MISOGYNISTIC’ - ‘The meme expresses hatred, hostility, or contempt toward women, including insulting, demeaning, or promoting harm or discrimination against them.’ 2. ‘NON-MISOGYNISTIC’ - ‘The meme does not express hatred, hostility, or contempt, and does not promote harm, discrimination, or demeaning attitudes against women.’ Classify the input meme as ‘MISOGYNISTIC’ or ‘NON-MISOGYNISTIC’ based on the above definitions, considering the image as well as the extracted text from the image. Provide the answer as either ‘MISOGYNISTIC’ or ‘NON-MISOGYNISTIC’ only. Extracted text from the image: {ocr_text}.

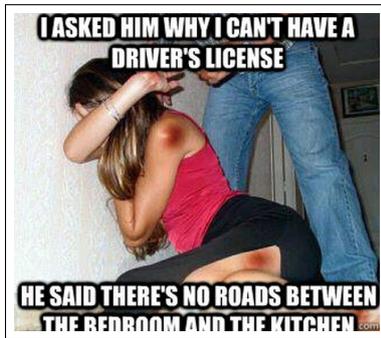
EXIST

Consider the following definitions: 1. ‘SEXIST’ - “The meme is sexist itself, describes a sexist situation or criticizes a sexist behavior.” 2. ‘NON-SEXIST’ - “The meme does not prejudice, underestimate, or discriminate women.” Also consider the extracted text and the prompt-based caption: Extracted text: {ocr_text}. Prompt-based caption: {caption_text}. Classify the input meme as ‘SEXIST’ or ‘NON-SEXIST’ based on the above definitions considering the image and text. Provide only ‘SEXIST’ or ‘NON-SEXIST.’

Figure 9: Zero-shot LLM prompts across datasets.

C LLM-based Image Captioning

Figures 10 and 11 provide examples of our LLM-based captioning strategies, contrasting the surface-level description (Prompt A) with the deeper semantic inference (Prompt B).



Label: Misogynistic

Prompt A Output (Surface-level description)

The image shows a woman in a red tank top and black pants, appearing **distressed** with her head down and hands covering her face. A person wearing blue jeans is standing behind her, holding her by the shoulders. The setting appears to be indoors, possibly a bedroom or a similar private space.

Prompt B Output (Multimodal semantic inference)

The meme contains a humorous and somewhat **crude joke** about a relationship dynamic, where one partner is being told they can't have a driver's license because there are no roads between the bedroom and the kitchen. This type of humor often relies on **stereotypes** and can be seen as **sexist**, as it implies that relationships should be confined to domestic spaces without broader societal implications. The use of such language can perpetuate gender stereotypes and limit perceptions of what relationships and personal freedom should encompass.



Label: Non-Misogynistic

Prompt A Output (Surface-level description)

The image shows an extremely large sandwich stacked high with various ingredients, including what appears to be pulled meat, bacon, cheese, and possibly some sauce. The sandwich is placed on a plate, and the background suggests an outdoor setting with a brick wall and a dark surface beneath the plate. The overall presentation emphasizes the size and abundance of the sandwich.

Prompt B Output (Multimodal semantic inference)

The meme humorously exaggerates the appeal of a sandwich by labeling it as "every man wants," which could be seen as a **stereotype**. However, the intent appears to be **playful rather than derogatory**, and the use of "man" might be intended **inclusively** to refer to any person. The focus is on the sandwich's excessive layers, suggesting it's a highly desirable food item.



Label: Misogynistic

Prompt A Output (Surface-level description)

The image shows a person lying on a lounge chair, wearing a bikini top and high-heeled shoes. The background appears to be an outdoor setting with people in the distance. The text at the bottom of the image reads "PLASTIC not just for mannequins anymore".

Prompt B Output (Multimodal semantic inference)

The meme uses a **provocative** image and text to make a statement about plastic surgery, implying that it is now common among women. The use of the word "plastic" in this context can be seen as a commentary on the perceived **superficiality and artificiality** of modern beauty standards. However, the meme does not explicitly criticize or promote any particular behavior but rather presents a **satirical view on societal norms** regarding appearance.

Figure 10: Illustrative MAMI examples showing the outputs of our prompt-based LLM captioning strategies. *Left:* meme image. *Right:* corresponding surface-level and deep-inference captions (Prompt A and Prompt B).



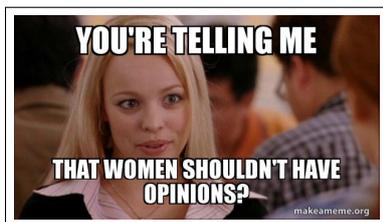
Label: Sexist

Prompt A Output (Surface-level description)

The image is a meme featuring characters from the animated TV show "The Simpsons." The top half of the image shows four adult female characters standing together, labeled "Expectativa" (Expectation), while the bottom half depicts a chaotic scene with various characters in costumes and masks, labeled "Realidad" (Reality). The contrast between the two halves **humorously** illustrates the difference between expectations and reality, particularly in the context of the Women's Day celebration.

Prompt B Output (Multimodal semantic inference)

The meme uses characters from "The Simpsons" to depict the "Expectativa" (expectation) and "Realidad" (reality) of the "Día de la Mujer" (International Women's Day). The top part shows four women in professional attire, representing the idealized expectation for International Women's Day. The bottom part shows a chaotic scene with various characters in costumes, suggesting the reality of the day. [...] This comparison **might be interpreted as a critique** of how International Women's Day is sometimes celebrated or perceived. However, this interpretation is subjective and depends on the viewer's perspective. In general, the meme does **not contain overtly sexist content** but rather uses humor to highlight a contrast. It's important to note that humor can be culturally specific and may not be universally understood or appreciated.



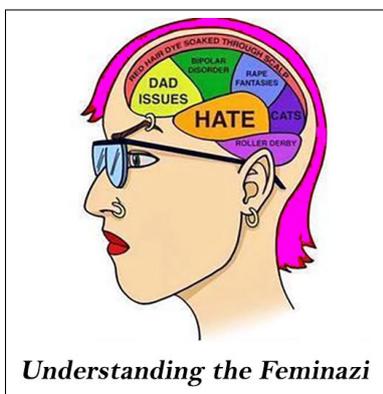
Label: Non-Sexist

Prompt A Output (Surface-level description)

The image shows a person with blonde hair looking surprised or confused, with another person partially visible in the background. The setting appears to be indoors, possibly a social gathering or meeting.

Prompt B Output (Multimodal semantic inference)

The meme could be interpreted as **promoting gender stereotypes**, specifically suggesting that women should not have opinions. **This is a form of sexism** as it implies that women's perspectives and thoughts are less valid or important than those of men. The meme expresses disbelief or frustration, which can be seen as a criticism of the idea that women should be limited in their ability to have opinions. Therefore, the meme **does contain sexist content** by implicitly promoting a gender-based limitation on opinion formation.



Understanding the Feminazi

Label: Sexist

Prompt A Output (Surface-level description)

The image shows a person lying on a lounge chair, wearing a bikini top and high-heeled shoes. The background appears to be an outdoor setting with people in the distance. The text at the bottom of the image reads "PLASTIC not just for mannequins anymore".

Prompt B Output (Multimodal semantic inference)

The meme uses a **provocative** image and text to make a statement about plastic surgery, implying that it is now common among women. The use of the word "plastic" in this context can be seen as a commentary on the perceived **superficiality and artificiality** of modern beauty standards. However, the meme does not explicitly criticize or promote any particular behavior but rather presents a **satirical view on societal norms** regarding appearance.

Figure 11: Illustrative EXIST examples showing the outputs of our prompt-based LLM captioning strategies. Left: meme image. Right: corresponding surface-level and deep-inference captions (Prompt A and Prompt B).