

M-QUEST - Meme Question-Understanding Evaluation on Semantics and Toxicity

Stefano De Giorgis Ting-Chih Chen Filip Ilievski

Vrije Universiteit Amsterdam

De Boelelaan 1105

{s.degiorgis, t.c.chen, f.ilievski}@vu.nl

Abstract

Internet memes are a powerful form of on-line communication, yet their multimodal nature and reliance on commonsense and cultural knowledge make automated toxicity detection challenging. Prior work has studied individual dimensions of memes: textual content via optical character recognition, visual content via object recognition, and toxicity via hate speech detection, but lacks a unified architecture that formally identifies the semantic dimensions contributing to meme toxicity. We present three contributions: (1) a semantic framework consolidating 10 dimensions necessary for meme interpretation; (2) a corresponding benchmark, M-QUEST, of 1, 109 question-answer pairs for 307 memes, divided in two subsets: M-QUEST-qwen and M-QUEST-gemini generated via a semi-automatic process grounded in the framework, respectively with Qwen3-VL-8B-Instruct, and Gemini3.6-Flash; and (3) an empirical evaluation of 12 VLMs, 10 open-weight and 2 closed proprietary, on M-QUEST. The results show that the instruction-tuned models with reasoning capabilities significantly outperform others, whereas pragmatic inference remains challenging across models.

Disclaimer: This paper contains discriminatory and explicit content that may be disturbing to some readers.

1 Introduction

Internet memes are “linguistic, image, audio, and video texts created, circulated, and transformed by countless cultural participants” (Milner, 2018). Memes are currently a key form of communication in public and private online conversations. From funny content, they become actual “modes of hypersignification, wherein the code itself becomes the focus of attention” (Shifman, 2014). The communicative power of memes resides in their ability to condense complex ideas into shareable formats (Onielfa et al., 2022). This power is sustained

by two factors: (i) resilience over time due to their multimodal nature, and (ii) traction gained from cultural and social overproduction, amplified by users consuming multimodal content in binge-watching patterns on social media (Shifman, 2019). While memes usually express humor, many “harmful memes” (Pramanick et al., 2021a) or “hateful memes” (Qu et al., 2023), jointly termed “toxic” memes (Pandiani et al., 2025), are designed or used for malicious purposes. Their scale, velocity, and variety make identifying toxicity in memes challenging for content moderators, and inspire work on AI-assisted toxicity detection (Ilievski, 2024).

Early AI research treated meme toxicity as a black-box classification task (Kiela et al., 2020; Fersini et al., 2022). Yet, to paraphrase Tolstoy (1878): “all happy memes are alike, each toxic meme is toxic in its own way.” Identifying toxicity features requires a conceptual decomposition that goes beyond the meme as a material object (image + text) to the meme as a cultural object: a relational entity connecting observers, their background knowledge, and the broader cultural context implicated in its interpretation (Bollici et al., 2026).

This recognition of the multi-faceted nature of meme toxicity has led to recent research at the intersection of toxicity and reasoning (Pandiani et al., 2025). Semantic frameworks that characterize memes have been proposed in the form of theoretical frameworks (Bollici et al., 2026) or as knowledge graphs (Tommasini et al., 2023; Martinez Pandiani et al., 2025). Meanwhile, benchmarks evaluate the ability of vision-language models (VLMs) to assess and explain toxicity, in relation to their targets (Pramanick et al., 2021b), entity roles (Sharma et al., 2022b), background knowledge (Bhowmick et al., 2021), and emotions (Alzu’bi et al., 2023). Other benchmarks aim at a more comprehensive set of features (Nguyen et al., 2025), or alignment with human values (Lin et al., 2024).

Two limitations of prior work emerge. First, as

Table 1: **Meme dimensions have been sparsely addressed in prior work, lacking a comprehensive view.** We summarise the most relevant works overlapping with our evaluation framework and M-QUEST benchmark. See Appendix for the extended Table 4 with 40 analysed works.

Paper Reference	Toxicity	Textual Material	Visual Material	Scene	Background Knowledge	Overall Intent	Emotion	Target Community	Semiotic Projection	Analogical Mapping
(Hee et al., 2023)	✓	✓	✓	✗	✓	✓	✓	✓	✗	✗
(Qu et al., 2022)	✓	✓	✓	✓	✓	✓	✓	✗	✗	✓
(Rajput et al., 2022)	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
(Gasparini et al., 2022)	✓	✓	✗	✗	✗	✓	✗	✓	✗	✗
(Ramamoorthy et al., 2022)	✗	✓	✓	✓	✗	✓	✓	✗	✗	✗
(Kiela et al., 2020)	✓	✓	✓	✗	✗	✓	✗	✓	✗	✗
(Hee et al., 2023)	✓	✓	✓	✗	✓	✗	✓	✓	✗	✗
(Fersini et al., 2022)	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
(Sharma et al., 2022a)	✓	✓	✓	✓	✓	✗	✗	✓	✗	✗
(Nguyen et al., 2025)	✗	✓	✓	✗	✓	✓	✓	✓	✗	✗
M-QUEST (ours)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

shown in Table 1, *evaluation dimensions across meme toxicity benchmarks remain discrepant and seemingly ad-hoc*, highlighting the need for comprehensive frameworks grounded in theoretical conceptualizations (Bollici et al., 2026). Second, there is a *lack of separation between toxicity identification and reasoning*: analogously to work that separates perception from reasoning (Jiang et al., 2024), principled evaluation must isolate individual toxicity causes to expose the granular capabilities and limitations of state-of-the-art VLMs.

To that end, we provide *a comprehensive framework for connecting the toxicity of internet memes to underlying semantic dimensions and assessing the ability of VLMs to analyze and explain toxicity through these dimensions*. Our contribution is threefold: 1. A **framework** that formalizes and operationalizes a comprehensive set of 10 dimensions, providing a semantic map of the features in existing literature. 2. A compositional **benchmark** called M-QUEST with 1,109 paired multiple-choice questions about the toxicity and its underlying reasons of 307 memes.¹ M-QUEST is divided in two subsets: M-QUEST-qwen and M-QUEST-gemini, respectively built using Qwen3-VL-8B-Instruct, and Gemini3.6-Flash, and is built semi-automatically by combining our framework dimensions, the divergent generative abilities of VLMs, and the common sense of human judges. 3. An **empirical analysis** of the granular performance of 10 diverse

open-weight and 2 closed proprietary VLMs on the 10 dimensions of M-QUEST.

2 Related Works

As apparent in Table 1, existing toxicity work has focused on only a few dimensions beyond Textual and Visual Material. Our framework subsumes all of these while introducing *Semiotic Projection* - identifying which meme element the observer or external entities are projected onto - as a novel contribution absent from prior research on memes. We present benchmarks and approaches that ground our framework. More specific references for each dimension are given in Section 3.

Meme Benchmarks Several benchmarks address meme toxicity subtasks (Table 1).² The Hateful Memes challenge (Kiela et al., 2020) is a foundational and widely used benchmark to date. Gasparini et al. (2022) target misogyny specifically (800 memes), as does MAMI SemEval-2022 (Fersini et al., 2022), which adds misogyny subclassification. MOMENTA (Pramanick et al., 2021b) classifies meme harmfulness and its target. The GOAT-Benchmark (Lin et al., 2024) covers 6,000 memes across five hate speech types (hatefulness, misogyny, offensiveness, sarcasm, harmful content) via binary classification. The closest to our work is MemeQA (Nguyen et al., 2025), comprising 9,000 human-curated multiple-choice Q&A across seven dimensions.

¹<https://huggingface.co/datasets/vulr/M-QUEST>

²See (Pandiani et al., 2025) for a comprehensive survey.

Our benchmark M-QUEST subsumes a broader set of dimensions than all prior work, including MemeQA. It provides a more granular analysis by explicitly representing relations among elements within each dimension and across dimensions with respect to toxicity. Finally, it is compositional: toxicity labels are systematically linked to explanation questions drawn from the other nine dimensions.

3 Framework of Meme Dimensions

Table 1 consolidates the key semantic dimensions considered in prior work on meme toxicity and related tasks, highlighting that they have been considered in seemingly arbitrary subsets, lacking a comprehensive view. Based on this body of work, we design a framework including 10 dimensions. These dimensions represent cognitive elements employed in the sense-making process of understanding a meme and its toxicity. We now describe and motivate each dimension.

Toxicity Harmful, offensive, or problematic content in the meme. Following (Pandiani et al., 2025), we distinguish six subtypes: *Abuse* (threats or abuse toward individuals or communities), *Cyberbullying* (mocking based on protected characteristics), *Harmfulness* (potential to harm individuals, organizations, or society), *Hatefulness* (direct or indirect attacks based on protected characteristics), *Misogyny* (hate targeting women), and *Propaganda* (content designed to influence opinions or actions toward a goal). Our notion of toxicity is linked to - and grounded in - all other 9 dimensions.

Textual Material Any visible text in the meme. Virtually all meme analysis works incorporate this component, focusing on captions, overlay text, and speech bubbles, particularly those focused on hate speech detection (Rajput et al., 2022; Bhandari et al., 2023; Kiela et al., 2020). Our notion includes also diegetic text (any text that exists within the meme image).

Visual Material Any visual entity in the meme: people, objects, characters, graphics, and salient visual elements. Alongside Textual Material, it is the most prevalent dimension; several works analyze both jointly (Zhang et al., 2023; Sabat et al., 2019; Badour and Brown, 2021). Rather than treating memes as images with embedded text (Bhandari et al., 2023), our approach integrates both modalities in a unified Knowledge Base (KB).

Scene The meaningful interpretation of how visual entities relate to one another, encompassing spatial arrangements, positioning, and layout. It has been investigated to detect victims of hateful memes (Sharma et al., 2022a) and for general multimodal meme question answering (Agarwal et al., 2024). Our notion subsumes relational knowledge about interactions of visual entities, with the scope of making sense of visual clues (e.g. starting from identifying person, spoon and kitchen, we build in the KB the scene “person cooking” taking as roles the above mentioned visual entities).

Background Knowledge Any external knowledge required to interpret the meme that is not directly observable. Our notion subsumes public personas (Bhowmick et al., 2021), contextual information (Shang et al., 2021b), factual correctness (Jabiyev et al., 2021), real-world named entities (Sharma et al., 2022a) and social entities (Pramanick et al., 2021a). It expands over specific groups (Qu et al., 2023), while extending coverage to general stereotypes, events, concepts, beliefs, and implicit references (e.g., actor names, historical events, political parties, cultural phenomena, or fictional references).

Overall Intent The communicative purpose behind the meme. The *Overall Intent* determines *underlying reasons* (Hee et al., 2023) for which a meme is meant to be toxic, political, nonsensical, etc. Our notion thus incorporates DisInfoMeme’s focus on intentional misinformation (Qu et al., 2022) and Agarwal et al. (2024)’s notion of intended effects, expanding over Nguyen et al. (2025), further encompassing the meme’s creator intentions, message goals, and purpose indicators.

Emotion Emotion conveyed by the meme. *Sentiment analysis* and *emotion detection*, are observed in a considerable body of work analyzing memes, including the Memotion 1.0 and 2.0 datasets (Sharma et al., 2020; Ramamoorthy et al., 2022). They are often integrated as plain labels of the meme (Maity et al., 2022), or to capture the emotion conveyed by the meme’s visual or textual component. Notably, the Misogyny Memes benchmark (Gasparini et al., 2022) incorporates emotions among the features they encompass. Beyond using emotion analysis to detect offensive intent (Kumari et al., 2023; Alzu’bi et al., 2023), we include neuro-physiological expressors: facial expressions, emotional gestures, and affective language.




Figure 1: M-QUEST-qwen framework: a meme is given as input to a VLM, which generates a Knowledge Base (KB), used by a second VLM module to generate Q&A sets based on KB. Running example in Appendix A.3.

Target Community The intended audience or community targeted by the meme. While prior work on hateful memes focuses on specific individuals targeted by the attack of the meme (Mathias et al., 2021), being it towards individuals and organizations (Jha et al., 2024), or communities (Pranick et al., 2021a,b), our notion broadens the concept to include community markers, demographic cues, cultural references, and audience signals.

Semiotic Projection A novel dimension capturing the projection of the meme observer onto a meme element (typically visual, textual, or scene). While prior work addresses narrative role projection - Hero, Villain, and Victim - onto meme elements (Sharma et al., 2022b; Bollici et al., 2026), the semiotic triangle formed by (1) a visual entity, (2) its role in the meme, and (3) the observer’s projection onto it remains underexplored. For instance, in the Disaster Girl meme,³ the text Me juxtaposed with the girl projects the observer onto an apparently innocent character (the girl being a child) whose smirk hints at complicity in the house on fire in the background. The resulting inference is a visual manifestation of *Schadenfreude*: pleasure derived from another entity’s misfortune.

Analogical Mapping A mapping between two meme entities - typically visual and textual - follow-

ing the pattern X is Y . Inspired by analogical thinking (Holyoak and Thagard, 1996), the Structure-Mapping Engine (Forbus et al., 2017), and cognitive metaphor theory (Lakoff and Johnson, 2024), we unify analogical and metaphorical mappings under a single dimension. For instance, the text Me juxtaposed with a perplexed child, maps Me onto child being perplexed, yielding *I am perplexed*.⁴ Spatial proximity commonly signals such mappings. A notable work investigating this dimension is Met-meme (Xu et al., 2022), a multimodal dataset rich in metaphorical features. More principled theoretical proposals (Bollici et al., 2026) and ontologies (Tommasini et al., 2023; Martinez Pandiani et al., 2025) are emerging, although to date they have not been operationalized into practical benchmarks or methods, and the explicit focus on mapping structure and semantic gain is a novelty of this work.

4 M-QUEST Benchmark Construction

In this section, we describe our semi-automatic benchmark construction method, starting with automatic Q&A generation, followed by human validation. Finally, we describe the resulting benchmark statistics. We implemented this method with two models: Qwen3-VL-8B-Instruct, and Gemini3.6-Flash, obtaining respectively two benchmark subsets: M-QUEST-qwen and M-QUEST-gemini.

³<https://knowyourmeme.com/memes/disaster-girl>

⁴https://en.wikipedia.org/wiki/Side_Eyeing_Chloe

Table 2: Q&As distribution across dimensions for different benchmark subsets. M-QUEST-qwen is generated from an initial toxicity-focused set of 307 memes, later extended with additional dimension-specific Q&As; M-QUEST-gemini is generated from a subset of 232 memes out of the M-QUEST-qwen’s 307. For more details about M-QUESTcomposition see Appendix A.4.

M-QUEST subset	Toxicity	Textual Material	Visual Material	Scene	Background Knowledge	Overall Intent	Emotion	Analogical Mapping	Target Community	Semiotic Projection	Unique Questions	Total Memes	Avg. Q&As
-qwen	307	45	49	44	35	35	30	28	20	16	609	307	1.98
-gemini	50	50	50	50	50	50	50	50	50	50	500	232	2.16
M-QUEST	357	95	99	94	85	85	80	78	70	66	1109	307	3.61

4.1 Automatic Q&A Generation

Automatic benchmark creation consists of two steps shown in Fig. 1: Knowledge Extraction and Q&A Generation.

Knowledge Extraction (KE) The entire framework is modular and component-agnostic. We process one meme at a time, extracting knowledge per dimension using our Dimensions Framework (§3) coupled with the VLM (Qwen3-VL-8b-Instruct for M-QUEST-qwen, and Gemini3.6-Flash for M-QUEST-gemini). Each VLM call receives the dimension definition integrated in the prompt and the meme as a .png file, yielding one knowledge base entry per dimension. Examples are shown in Fig. 1 (KE Snippet). Prompts follow a fixed order, as in Table 5, Appendix A.3: basic dimensions first (Textual and Visual Material), then dependent ones (Scene), and finally those with complex inter-dimensional relations (Analogical Mapping, Semiotic Projection). At each step, previously extracted knowledge is included in the prompt, ensuring consistency and shared anchoring points across dimensions. The output is a comprehensive KB about the meme and its toxicity, according to the dimensions described in Section 3. A worked example is provided in Appendix A.3.

Candidate Q&A Generation Using the extracted knowledge, the VLM generates N questions per dimension, each with four answers: one correct, one plausible distractor, one implausible distractor, and a fixed “None of the others” option. Correct answer positions is balanced to avoid positional bias: for M-QUEST-qwen is A: 25.45%, B: 26.44%, C: 24.47%, D: 23.65%, while for M-QUEST-gemini is 25% for all answers, and a brief explanation accompanies each correct answer. Additional details are provided in Appendix A.4. The inputs to the VLM are the dimension definition, the previously extracted knowledge, and the meme image (Fig. 1).

Outcome We randomly selected 1,000 memes from the Hateful Memes dataset (Kiela et al., 2020). This dataset was selected for: 1. being widely adopted in previous harmful memes research, 2. including memes from different domains (e.g., Covid pandemics, elections) and 3. covering multiple types of toxicity (e.g., harmful, misogyny, etc.). Our KE step generates one KB per meme (statistics in Fig. 3). Creating one question per dimension yields 10,000 candidate Q&As in total (1 question $\times$ 10 dimensions $\times$ 1,000 memes).

4.2 Manual Data Validation

Human Validation Task Design Since our primary focus is toxicity, human judges validated the KBs and Q&A pairs for this dimension via a custom interface (Fig. 8). Per meme, judges evaluated up to six fields: (1) plausibility of the toxicity assessment; (2) if a toxicity-linked KB existed, plausibility of the associated element’s assessment; (3–4) whether the toxicity-derived question was on point and its correct answer plausible; (5–6) the same two judgments for the question derived from the toxicity-linked element. This results in a minimum of one annotation per meme (toxicity question plausibility) and a maximum of six.

Outcome Fourteen annotators (9 male, 5 female), all familiar with LLMs, KBs, and memes, validated 1,000 memes in total, generated by the Qwen VLM, focusing on the Toxicity dimension. In 307 cases, both the question and correct answer were positively validated, forming an intermediate dataset called ToxM-QUEST-qwen. To cover all dimensions, we sampled 50 additional questions per dimension from the 307 ToxM-QUEST-qwen memes and subjected them to human validation: 59.33% of memes received positive validation for both questions. Approval rates were highest for fundamental dimensions – Scene, Textual, and Visual Material (74–76%) – and lowest for Semiotic Projection and Target Community (32% and 40%, respectively). Inter Annotation Agreement

(IAA) is measured on 50-meme batches across two tasks: Toxicity Assessment (if the meme is toxic), and Q&A Plausibility (if the proposed question is pertinent, if the correct answer is actually correct, and if the distractors match the desired structure of plausible-not plausible-None of the above). For Toxicity Assessment: pairwise Cohen’s $\kappa \in [0.066, 0.387]$, Fleiss’ $\kappa = 0.158$, nominal $\alpha = 0.160$. For Answer Plausibility: pairwise $\kappa \in [0, 0.250]$, Fleiss’ $\kappa = 0.075$, nominal $\alpha = 0.060$. Further details about the human validation of M-QUEST-qwen are provided in Appendix A.6. Subsequently, the Q&A pairs generated by Gemini for the same 307 memes were also validated by one of the human judges following the same setup, resulting in M-QUEST-gemini.

4.3 M-QUEST

Key Statistics Table 2 summarizes the resulting benchmark M-QUEST, which comprises 1109 Q&A pairs for 307 memes – more than three questions per meme on average: one toxicity question is always present, while at least another reasoning question depends on the connection between the reasoning question and the toxicity one. The reasoning questions in M-QUEST-qwen are relatively balanced across the nine dimensions, ranging from 16 (Semiotic Projection) to 35 (Visual Material, Textual Material, Scene), while for M-QUEST-gemini they are perfectly balanced at 50 per dimension.

Release and Licensing Following Meta’s license on the Hateful Memes dataset, the M-QUEST HF repository includes generated graphs, human annotations, and Q&A pairs for 1,000 memes, but not the images themselves. The images can be requested from the Hateful Memes challenge repository,⁵ while the alignment is ensured via the original IDs.

5 Evaluation

Next, we conduct benchmarking experiments to gauge the performance of current state-of-the-art (SoTA) VLMs on M-QUEST. We present our experimental setup, followed by the results obtained and a qualitative analysis.

5.1 Experimental Setup

Evaluation Metrics We report Q&A *accuracy* (percentage of correctly answered questions), both

Table 3: Accuracy of VLMs on M-QUEST, comparing the Qwen-generated and Gemini-generated benchmark subsets: accuracy for All questions, grouped per meme (T+R), Toxicity accuracy, and combined Reasoning accuracy across other dimensions, along with the Macro-Average. Models are grouped by capability.

VLM	All	Group (T+R)	T only	R only	Macro Avg.
M-QUEST-qwen					
<i>Perception-focused</i>					
BLIP2-Flan-T5-xl	13.2	12.3	14.7	14.7	14.3
Pixtral-12B	70.3	60.6	72.0	54.7	57.4
<i>Instruction-tuned</i>					
InstructBLIP-Vicuna-7B	22.9	18.7	22.8	17.8	20.6
FastVLM 1.5B	32.5	14.0	31.1	37.7	34.3
FastVLM 7B	44.0	19.9	35.1	53.0	48.6
<i>Reasoning-oriented</i>					
LLaVA-v1.6-Vicuna-7B	47.3	46.8	50.2	40.2	41.5
LLaVA-v1.5-7B	56.8	53.4	54.7	70.4	68.3
<i>Instruction-tuned + reasoning</i>					
Qwen3-VL-8B	83.8	78.3	80.1	95.3	94.3
Qwen2.5-VL-7B	<u>86.4</u>	<u>83.3</u>	<u>86.6</u>	<u>94.7</u>	<u>93.3</u>
Qwen2-VL-7B	85.6	82.3	84.4	94.7	93.8
<i>Proprietary</i>					
Claude-Opus-5	90.5	88.0	90.6	90.3	90.3
Gemini-3.5-Flash-Lite	83.1	71.3	79.5	86.8	84.9
<i>Human baseline</i>					
Human annotation	82.8	78.3	87.0	79.2	82.9
M-QUEST-gemini					
<i>Perception-focused</i>					
BLIP2-Flan-T5-xl	82.4	70.7	94.0	81.1	82.4
Pixtral-12B	91.8	84.5	96.0	91.3	91.8
<i>Instruction-tuned</i>					
InstructBLIP-Vicuna-7B	36.0	14.2	42.0	35.3	36.0
FastVLM-1.5B	62.6	40.5	58.0	63.1	62.6
FastVLM-7B	80.8	65.9	72.0	81.8	80.8
<i>Reasoning-oriented</i>					
LLaVA-1.6-Vicuna-7B	45.4	26.7	38.0	46.2	45.4
LLaVA-1.5-7B	38.0	17.2	42.0	37.6	38.0
<i>Instruction-tuned + reasoning</i>					
Qwen3-VL-8B-Instruct	<u>95.2</u>	<u>90.5</u>	<u>98.0</u>	<u>94.9</u>	<u>95.2</u>
Qwen2.5-VL-7B-Instruct	81.2	65.9	90.0	80.2	81.2
Qwen2-VL-7B-Instruct	64.6	41.8	88.0	62.0	64.6
<i>Proprietary</i>					
Claude-Opus-5	98.8	97.4	100.0	98.7	98.8
Gemini-3.5-Flash-Lite	94.8	89.7	100.0	94.2	94.8
<i>Human baseline</i>					
Human annotation	89.0	87.8	100.0	87.8	89.0

overall and per dimension. For memes with both a toxicity and a reasoning question, we compute *group accuracy* – whether a model answers both correctly, i.e., whether model answers correctly about toxicity for the right reasons. Performance across dimensions is summarized via *macro-average* (unweighted mean). This setup enables both overall comparison and fine-grained identification of each VLM’s strengths and weaknesses.

Models We evaluate 10 open-source and 2 proprietary VLMs spanning four capability profiles: *Perception-focused* (no instruction tuning or ad-

⁵<https://ai.meta.com/tools/hatefulmemes/>




Figure 2: Accuracy of VLMs across dimensions of M-QUEST-qwen: BLIP2-Flan-T5-xl, InstructBLIP-Vicuna-7B, FastVLM-7B, FastVLM-1.5B, LLaVA-v1.6-Vicuna, LLaVA-v1.5, Pixtral-12B, Qwen3-VL-8B-Instruct, Qwen2-VL-7B-Instruct, Qwen2.5-VL-7B-Instruct, Claude-Opus-5, Gemini-3.5-Flash-Lite.

vanced reasoning): BLIP2-Flan-T5-xl (Li et al., 2023) and Pixtral-12B (Agrawal et al., 2024); *Instruction-tuned*: InstructBLIP-Vicuna-7B (Dai et al., 2023), FastVLM-1.5B/7B (Vasu et al., 2025); *Reasoning-oriented*: LLaVA-v1.5 (Liu et al., 2024a) and LLaVA-v1.6-Vicuna (Liu et al., 2024b); *Instruction-tuned with reasoning*: Qwen2-VL-7B (Wang et al., 2024), Qwen2.5-VL-7B (Bai et al., 2025b), Qwen3-VL-8B-Instruct (Bai et al., 2025a), Claude Opus 5 (Anthropic, 2026), and Gemini3.5-Flash-Lite (Google DeepMind, 2026).

Implementation Details All models are evaluated zero-shot with greedy decoding ($temperature=0$, $top-p=1.0$). For multiple-choice questions, the answer choice is generated first (stopping upon a valid option), followed by the CoT explanation (terminated naturally by the end-of-sequence token). All experiments were run on a single H100 GPU, taking approximately 10 hours in total.

5.2 Results

Performance Analysis As shown in Table 3, accuracy on M-QUEST-qwen ranges from below chance to 86% for open-weight models, 90% for proprietary models. Reasoning dimensions consistently score higher than toxicity-related ones, suggesting that explicit visual or textual reasoning, scene understanding, and background knowledge interpretation are easier than inferring latent toxic intent. Notably, Qwen models approach near-perfect accuracy on reasoning questions when toxicity is excluded. Still, performance drops when implicit inference is required as shown in Fig. 2, underscoring that *latent semantic understanding across modalities is systematically harder than surface-level cue matching*.

Substantial differences emerge across model families. BLIP2-Flan-T5-xl and InstructBLIP-

Vicuna-7B fall below random chance, reflecting limited instruction tuning. FastVLM-1.5B and 7B achieve moderate scores, capturing basic visual features but lacking the multimodal alignment needed for relational reasoning. Within the LLaVA family, v1.5 outperforms v1.6-Vicuna, showing that instruction tuning differences matter even when the vision encoder family is similar. Pixtral-12B benefits from a CLIP encoder’s semantically rich representations. Interestingly, Qwen2.5-VL-7B slightly surpasses Qwen3-VL-8B, suggesting that model size alone is not decisive; rather, *models optimized for complex multi-step reasoning may occasionally overreason, leading to slightly lower accuracy in this setting*.

These trends are explained by two key factors. *Instruction tuning* enables models to align textual cues with visual information, while *reasoning capability* supports multi-step inference over contextual relationships. Models lacking both (BLIP2) struggle the most; instruction tuning alone (InstructBLIP) yields only moderate gains; and models combining both (Qwen family) achieve the strongest overall performance. This suggests that while reasoning capability is essential for interpreting complex meme semantics, excessive reasoning may introduce inefficiencies when tasks require only shallow inference.

Table 3 also reports results on M-QUEST-gemini. Accuracy is consistently higher across nearly all VLMs than on M-QUEST-qwen, most notably on Toxicity, where several models and the human baseline approach or reach 100%, suggesting *item difficulty is sensitive to the model used during knowledge extraction*. Rankings shift slightly (e.g., Qwen3-VL-8B overtakes Qwen2.5-VL-7B), but proprietary models retain a comparable lead over open-weight ones in both settings, indicating this gap is a stable property of the models rather

than the benchmark variant.

In the rest of §5, we focus on M-QUEST-qwen, for more details on M-QUEST-gemini see Appendix A.5.

Performance Across Evaluation Dimensions

Group accuracy (correctly answering both the Reasoning and Toxicity questions for the same meme) reflects cross-dimensional consistency rather than isolated task performance. As shown in Table 3, group accuracy is generally lower than reasoning and toxicity accuracy, indicating that *aligning Reasoning outcomes with Toxicity judgments remains challenging*. Weak models (BLIP2, InstructBLIP) show low group accuracy due to limited multimodal alignment. Stronger models achieve higher group accuracy, suggesting improved cross-dimensional consistency when reasoning about the meme and evaluating its toxicity within the same context.

Per-dimension analysis (Fig. 2) reveals consistent patterns. Qwen models achieve near-perfect accuracy on perceptual dimensions (Scene, Visual Material) but weaker performance on latent relational dimensions (Toxicity, Overall Intent). LLaVA-v1.5 follows a similar pattern. BLIP2 and InstructBLIP represent the lower bound across most dimensions, failing to capture implicit vision-text relationships. Notably, Qwen-family models achieve near-perfect accuracy on Background Knowledge, suggesting that well-aligned instruction tuning enables effective external knowledge retrieval. Overall, *accurate perceptual recognition alone does not guarantee strong performance on latent relational dimensions; effective instruction tuning and multimodal reasoning remain essential for robust pragmatic inference*.

Implications The results highlight four key implications. First, *multimodal alignment and reasoning quality matter more than model size*. Second, models with stronger multimodal representations achieve consistently higher performance, highlighting the *importance of semantically rich visual-textual representations* for socially complex content: improved semantic grounding could substantially benefit toxicity analysis. Third, the gains from combining instruction tuning with reasoning training confirm that *toxic meme understanding is reasoning-intensive*: accurately inferring toxic intent requires integrating visual, textual, and contextual cues, which models lacking reasoning capability fail to capture. Each question is evaluated

in an isolated chat session, reasoning traces of advanced models suggest that answer selection relies on dimension-relevant meme understanding, rather than superficial cues. Finally, while current VLMs are approaching saturation on surface-level visual dimensions, *latent intent and toxicity remain challenging*: a critical gap for real-world deployment, where subtle harmful content remains most likely to be missed.

5.3 Qualitative Analysis

VLM Comparison on Toxicity Reasoning For the meme in Fig. 1, Box 1 presents a case study comparing four CoT-capable models with increasing performance on the same meme.

Box 1 – Q: This meme’s message targets which group?

InstructBLIP-Vicuna-7B CoT

The meme targets the group of people who are being accused of being terrorists, which is the group of Muslims or Middle Easterners. The meme is making a joke by saying “trust me, they’re not terrorists” and is poking fun at the stereotype that all camels are terrorists.

LLaVA-1.6-Vicuna CoT

The text above the camels reads “trust me... they’re all terrorists” which is a statement that is often used to stereotype and demonize Muslims or Middle Easterners. The humor in the meme is derived from the absurdity of camels being terrorists, which is a play on the idea that all Muslims are terrorists. The meme’s message targets a specific group, which is Muslims or Middle Easterners, as it uses a stereotype to make a joke. The humor is derived from the absurdity of the situation, but it is important to recognize that such stereotypes are harmful and perpetuate negative attitudes towards a specific group of people.

Pixtral-12B CoT ✓

The meme uses a stereotype about camels being associated with the Middle East or desert regions. The phrase “they’re all terrorists” is a generalization that implies all members of a certain group are terrorists, which is a derogatory and offensive statement. The context and the visual elements (camels in a desert) strongly suggest that the meme is targeting individuals from the Middle East or Muslims.

Qwen3-VL-8B-Instruct CoT ✓

The word “terrorists” is a highly charged term historically and politically associated with acts of violence, particularly those committed by individuals or groups from specific regions or religions; most notably, in Western discourse, often linked to Muslims or Middle Easterners. The meme uses camels, animals often culturally associated with the Middle East, as the speakers’ of this harmful generalization. [...] The target of the meme’s message is therefore the group that the word “terrorists” is being applied to in a sweeping, unjust manner, which, in this context, is Muslims or Middle Easterners, as that is the most common and historically charged association with the term in global discourse.

InstructBLIP correctly flags hateful group dynamics, but misses the symbolic camel-Muslim association, treating the text as directly describing the visual content. LLaVA-1.6 reaches the same conclusion, again reducing the meme to “camels as terrorists” without deeper inference. Pixtral-12B

draws the correct connection between the visual entity (camel) and the target group (Muslims) via background knowledge, and correctly identifies the meme as perpetuating a harmful stereotype. Qwen3 provides the most nuanced interpretation, contextualizing “terrorist” both semantically and culturally, and building a layered, consistent reading that correctly identifies the adoption of a harmful stereotype.

VLM Behavior Across Dimensions Fig. 10 illustrates how Qwen 2/2.5/3 reason differently across dimensions on the same meme. On the Scene dimension, all models correctly identify the salient visual elements (individuals in fruit costumes, people walking along a street) and infer the depicted activities. On the Toxicity dimension, however, all models incorrectly interpret the meme as implying fruit scarcity in Africa, missing its classification as false equivalence: the meme conflates dietary patterns with quality of life, reinforcing colonialist stereotypes.

These results reveal a recurring vulnerability: models performing adequately on unimodal or perceptual dimensions frequently fail when multimodal cue synthesis is required for toxicity assessment. LLaVA and Qwen3 show stronger recognition of implicit social critique and ironic framing, while earlier Qwen iterations exhibit a literal interpretation bias. These findings underscore the need for multi-dimensional evaluation to comprehensively assess VLMs’ capacity for nuanced multimodal reasoning.

6 Conclusions and Future Works

This paper studied the ability of state-of-the-art open-weight VLMs to assess toxicity and support their decisions by principled reasoning. To this end, we designed a semantic framework that consolidates 10 dimensions from prior work, including the understudied dimensions of analogical mapping and semiotic projection. Then, we designed a compositional benchmark, M-QUEST, with 307 memes and 1,109 questions, focusing on toxicity and its grounding in the interpretation of a second key dimension for each meme. Our experiments with 12 VLMs revealed substantial and systematic differences in the model performance. Models that incorporate instruction-tuning and reasoning significantly outperformed the others, especially on inference dimensions. This analysis revealed that meme toxicity assessment is an in-

herently knowledge-and reasoning-intensive task that requires rich multimodal representations, intent modeling, and commonsense reasoning. A key future direction is to extend the benchmark with dimensions that can further capture nuance in annotator judgments, namely, more granular background knowledge (Ilievski, 2024) and considerations of morality, culture, and legality (Forbes et al., 2020).

A key strength of M-QUEST is its structural composition into 10 literature-derived dimensions, extracted in a staged process; difficulty correlates with extraction order, suggesting a principled path for scaling benchmark difficulty without an ad hoc notion of “complexity”—future work can simply combine two or more existing dimensions within a single question, leveraging the compositional structure already present in the annotation schema.

We envision two complementary combination strategies. *Co-location* questions jointly draw on multiple dimensions without an explicit relationship between them, e.g., referencing both a visual element and an accompanying textual span. *Interaction* questions instead require reasoning across dimensions that are mutually dependent or spatially/semantically entangled, e.g., identifying an emotion expressed in a specific image region, or localizing toxicity arising from the interplay between an image region and its associated textual material. Co-location thus tests integration of independent signals, while interaction checks for deeper compositional and grounding capabilities, since the target phenomenon emerges only from the relationship between dimensions rather than from either alone. We envision the dimension-combination strategy as a promising direction for building increasingly challenging benchmark variants in an interpretable manner, since difficulty can be scaled and directly characterized by the number and types of dimensions combined.

Acknowledgments

This work was supported by Huawei Finland as part of the DReaMS Lab (WP5). All content represents the opinion of the authors, which is not necessarily shared nor endorsed by their respective employers and/or sponsors. We thank the fourteen annotators for their dedicated work. We thank Floris Den Hengst and Mohammadhossein Khojasteh for insightful discussions. We thank SURF (www.surf.nl) for the support and access to the Dutch national supercomputer Snellius.

7 Limitations

Benchmark Size and Languages M-QUEST comprises 1109 questions in total: 609 for M-QUEST-qwen (of which 307 on toxicity, 16–49 per other semantic dimension), and 500 for M-QUEST-gemini (of which 50 per each dimension). Its limited size may constrain the generalisability of the findings. Similarly, its findings are based on English memes and may not fully extend to other languages. Expanding both the scale and linguistic diversity of M-QUEST is a key future direction.

Annotator Effects Human validation is inherently influenced by annotators’ values, culture, and background. Our 14 judges are LLM-familiar scientists primarily of European origin, whose approval decisions may not generalize to other populations. Validation rates (Fig. 6) are further affected by cultural misalignment with the original memes, many of which reference US politics and everyday situations. Characterizing annotator (dis)agreement and its sources remains an important direction for future work.

Framework Design Our framework is built bottom-up by surveying existing semantic frameworks for memes. Consequently, dimensions such as target metadata (Tommasini et al., 2023), template configurations (Bates et al., 2025), and considerations of morality, ethics, culture, and legality (Forbes et al., 2020) are not included but warrant inclusion in future extensions. Alternative design approaches (e.g., combining top-down theoretical frameworks with bottom-up benchmark analysis) remain an important direction for future work.

8 Ethical Considerations

This work involves the analysis and automated processing of toxic internet memes, including content that may be harmful, offensive, or distressing. We acknowledge several ethical considerations. We further added a disclaimer after the Abstract.

Dataset and Content M-QUEST is derived from the Hateful Memes dataset (Kiela et al., 2020), which contains content that is intentionally hateful, offensive, or toxic. We do not redistribute the meme images, requiring users to obtain them directly from the original source in accordance with Meta’s license. Researchers accessing M-QUEST should be aware of the potentially distressing nature of the content.

Annotator Welfare Our 14 human annotators were exposed to toxic and hateful content as part of the validation process. Annotators were informed of the nature of the task prior to participation and provided consent. We recommend that future annotation efforts involving similar content implement structured exposure limits and access to psychological support.

Annotator Bias As noted in Section 7, our annotators are primarily of European origin with familiarity with LLMs and memes. Their judgments may not be representative of broader or more diverse populations, and cultural misalignment with the predominantly American cultural references in the dataset may have influenced validation rates.

Dual-Use Risk While our framework is designed to advance toxicity detection, the dimensions and knowledge extraction pipeline could in principle be repurposed to assist in the *creation* of more effective toxic memes. We encourage responsible use of this framework strictly for detection, moderation, and research purposes.

Model Evaluation The VLMs evaluated in this work may themselves produce harmful or biased outputs when processing toxic content. Our evaluation is conducted in a controlled research setting and is not intended to endorse or promote any specific model for deployment in content moderation systems without further safety evaluation.

Scope of Toxicity Our operationalization of toxicity follows (Pandiani et al., 2025) and is inherently partial. Definitions of harmful content are culturally situated and may not generalise across

languages, communities, or legal jurisdictions. M-QUEST should not be treated as a definitive or universal standard for toxicity classification.

Use of Language Models Claude Sonnet 4.6 and Claude Sonnet 5 were used solely for English proofreading and language checking of this paper. All VLMs mentioned in Section 5.1 were used exclusively for the benchmark evaluation tasks described in the paper, and no generative model was used to produce the scientific content of this work.

9 Responsible NLP Research Checklist

A1. Did you describe the limitations of your work? Yes, in the dedicated Section 7 - Limitations.

A2. Did you discuss any potential risks of your work? Yes, in dedicated Section 8 - Ethical Concerns.

B1. Did you cite the creators of artifacts you used? We created a benchmark, which is provided as additional files. We do not disclose the name of the creators in order not to violate anonymity.

We reused the Hateful Memes existing dataset, which is properly cited in the paper, and, according to their material license, only a limited subset (max 100 images) of their data can be released publicly. We fully comply with their license.

B2. Did you discuss the license or terms for use and / or distribution of any artifacts? We discussed among us and we'll be release it under CC-BY 4.0 license when anonymity will be dropped, after the review period.

B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)? Yes, our use of existing artifacts is consistent.

Yes, we specify intended use and it is fully compatible with original resources.

B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it? This point does not apply to our work.

B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.? Yes, in Section 2, and Section 4.3.

B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created?

Yes, in Section 4.3, Section 5.1, and Section 5.3. Additionally further insights are provided in Appendix.

C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used? Yes, in Section 5.1.

C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values? Yes, in Section 5.1.

C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run? Yes, in Section 5.1, Section 5.3, and additionally in Appendix.

C4. If you used existing packages (e.g., for pre-processing, for normalization, or for evaluation, such as NLTK, Spacy, ROUGE, etc.), did you report the implementation, model, and parameter settings used? Yes, in Section 5.1 and Section 3.

D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.? Yes, it is described in Section 5.1, and screenshots are provided in Appendix.

D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)? Yes, annotators were volunteers, and we specify gender and collective provenance, maintaining anonymity.

D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? This point does not apply to our work.

D4. Was the data collection protocol approved (or determined exempt) by an ethics review board? This point does not apply to our work.

D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data? Yes, in Section 4.3.

E1. If you used any AI assistants, did you include information about your use? Yes, in Section 8.

References

- Siddhant Agarwal, Shivam Sharma, Preslav Nakov, and Tanmoy Chakraborty. 2024. Mememqa: multimodal question answering for memes via rationale-based inferencing. *arXiv preprint arXiv:2405.11215*.
- Pravesh Agrawal, Szymon Antoniak, Emma Bou Hanna, Baptiste Bout, Devendra Chaplot, Jessica Chudnovsky, Diogo Costa, Baudouin De Monicault, Saurabh Garg, Theophile Gervet, Soham Ghosh, Amélie Héliou, Paul Jacob, Albert Q. Jiang, Kartik Khandelwal, Timothée Lacroix, Guillaume Lample, Diego Las Casas, Thibaut Lavril, and 23 others. 2024. *Pixtral 12b*. Preprint, arXiv:2410.07073.
- Ahmad Alzu'bi, Lojin Bani Younis, Abdelrahman Abuarqoub, and Mohammad Hammoudeh. 2023. Multimodal deep learning with discriminant descriptors for offensive memes detection. *ACM Journal of Data and Information Quality*, 15(3):1–16.
- Anthropic. 2026. *Claude opus 5 system card*. Technical report, Anthropic. Accessed: 2026-08-31.
- Jamshid Bacha, Farman Ullah, Jebran Khan, Abdul Wasay Sardar, and Sungchang Lee. 2023. A deep learning-based framework for offensive text detection in unstructured data for heterogeneous social media. *IEEE Access*, 11:124484–124498.
- Jafar Badour and Joseph Alexander Brown. 2021. Hateful memes classification using machine learning. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–8. IEEE.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhi-fang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025a. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Luke Bates, Peter Ebert Christensen, Preslav Nakov, and Iryna Gurevych. 2025. A template is all you meme. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 10443–10475.
- Aashish Bhandari, Siddhant B Shah, Surendrabikram Thapa, Usman Naseem, and Mehwish Nasim. 2023. Crisishatemm: Multimodal analysis of directed and undirected hate speech in text-embedded images from russia-ukraine conflict. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1994–2003.
- Rajat Subhra Bhowmick, Isha Ganguli, Jayanta Paul, and Jaya Sil. 2021. A multimodal deep framework for derogatory social media post identification of a recognized person. *Transactions on Asian and Low-Resource Language Information Processing*, 21(1):1–19.
- Gabriella Bollici, Filip Ilievski, and Wouter van Atteveldt. 2026. Political memes as fragmented narratives: An interdisciplinary framework for meme analysis. In *The Annual International Communication Association (ICA) Conference*.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. InstructBLIP: Towards general-purpose vision-language models with instruction tuning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Mithun Das and Animesh Mukherjee. 2023. Banglaabusememe: A dataset for bengali abusive meme classification. *arXiv preprint arXiv:2310.11748*.
- Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021. *Detecting propaganda techniques in memes*. *arXiv preprint arXiv:2109.08013*.
- 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 *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 533–549.
- Maxwell Forbes, Jena D Hwang, Vered Shwartz, Maarten Sap, and Yejin Choi. 2020. Social chemistry 101: Learning to reason about social and moral norms. *arXiv preprint arXiv:2011.00620*.
- Kenneth D Forbus, Ronald W Ferguson, Andrew Lovett, and Dedre Gentner. 2017. Extending sme to handle large-scale cognitive modeling. *Cognitive Science*, 41(5):1152–1201.
- Francesca Gasparini, Giulia Rizzi, Aurora Saibene, and Elisabetta Fersini. 2022. Benchmark dataset of memes with text transcriptions for automatic detection of multi-modal misogynistic content. *Data in brief*, 44:108526.
- Google DeepMind. 2026. *Gemini 3.5 flash-lite: Model card*. Accessed: 2026-08-31.

- Ming Shan Hee, Wen-Haw Chong, and Roy Ka-Wei Lee. 2023. Decoding the underlying meaning of multimodal hateful memes. *arXiv preprint arXiv:2305.17678*.
- Keith J Holyoak and Paul Thagard. 1996. *Mental leaps: Analogy in creative thought*. MIT press.
- Filip Ilievski. 2024. *Human-Centric AI with Common Sense*. Springer.
- Bahrz Jabiyev, Jeremiah Onaolapo, Gianluca Stringhini, and Engin Kirda. 2021. e-game of fame: Automatic detection of fake memes. In *TTO*, pages 1–11.
- Prince Jha, Krishanu Maity, Raghav Jain, Apoorv Verma, Sriparna Saha, and Pushpak Bhattacharyya. 2024. Meme-ingful analysis: Enhanced understanding of cyberbullying in memes through multimodal explanations. *arXiv preprint arXiv:2401.09899*.
- Yifan Jiang, Kexuan Sun, Zhivar Sourati, Kian Ahra-bian, Kaixin Ma, Filip Ilievski, Jay Pujara, and 1 others. 2024. Marvel: Multidimensional abstraction and reasoning through visual evaluation and learning. *Advances in Neural Information Processing Systems*, 37:46567–46592.
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanu-j Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. The hateful memes challenge: Detecting hate speech in multimodal memes. *Advances in neural information processing systems*, 33:2611–2624.
- Gitanjali Kumari, Dibyanayan Bandyopadhyay, and Asif Ekbal. 2023. Emoffmeme: identifying offensive memes by leveraging underlying emotions. *Multime-dia Tools and Applications*, 82(29):45061–45096.
- Gitanjali Kumari, Anubhav Sinha, Asif Ekbal, Arindam Chatterjee, and Vinutha B N. 2024. Enhancing the fairness of offensive memes detection models by mit-igating unintended political bias. *Journal of Intelli-gent Information Systems*, 62(3):735–763.
- George Lakoff and Mark Johnson. 2024. *Metaphors we live by*. University of Chicago press.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large lan-guage models. In *ICML*.
- Hongzhan Lin, Ziyang Luo, Bo Wang, Ruichao Yang, and Jing Ma. 2024. Goat-bench: Safety insights to large multimodal models through meme-based social abuse. *ACM Transactions on Intelligent Systems and Technology*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024a. Improved baselines with visual instruc-tion tuning. In *Proceedings of the IEEE/CVF Con-ference on Computer Vision and Pattern Recognition (CVPR)*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024b. *Llava-next: Improved reasoning, ocr, and world knowledge*.
- Krishanu Maity, Prince Jha, Sriparna Saha, and Push-pak Bhattacharyya. 2022. A multitask framework for sentiment, emotion and sarcasm aware cyberbullying detection from multi-modal code-mixed memes. In *Proceedings of the 45th international ACM SIGIR conference on research and development in informa-tion retrieval*, pages 1739–1749.
- Delfina S Martinez Pandiani, Erik Tjong Kim Sang, and Davide Ceolin. 2025. Ontoxkg: An ontology-based knowledge graph of toxic symbols and their manifestations. In *International Conference on Web Engineering*, pages 119–127. Springer.
- Lambert Mathias, Shaoliang Nie, Aida Mostafazadeh Davani, Douwe Kiela, Vinodkumar Prabhakaran, Bertie Vidgen, and Zeerak Talat. 2021. Findings of the woah 5 shared task on fine grained hateful memes detection. In *Proceedings of the 5th Work-shop on Online Abuse and Harms (WOAH 2021)*, pages 201–206.
- Ryan M Milner. 2018. *The world made meme: Public conversations and participatory media*. mit Press.
- Karthick Myilvahanan, B Shashank, Tushar Raj, Chi-raanth Attanti, and Shivam Sahay. 2023. A study on deep learning based classification and identification of offensive memes. In *2023 Third International Conference on Artificial Intelligence and Smart En-ergy (ICAIS)*, pages 1552–1556. IEEE.
- Khoi PN Nguyen, Terrence Li, Derek Lou Zhou, Gabriel Xiong, Pranav Balu, Nandhan Alahari, Alan Huang, Tanush Chauhan, Harshavardhan Bala, Emre Guzelordu, and 1 others. 2025. Memeqa: Holistic evaluation for meme understanding. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 18926–18946.
- Thanh Tin Nguyen, Nhat Truong Pham, Ngoc Duy Nguyen, Hai Nguyen, Long H Nguyen, and Yong-Guk Kim. 2022. Hcilab at memotion 2.0 2022: Anal-ysis of sentiment, emotion and intensity of emotion classes from meme images using single and multi modalities (short paper). In *DE-FACTIFY@ AAIL*.
- Carles Onielfa, Carles Casacuberta, and Sergio Escalera. 2022. Influence in social networks through visual analysis of image memes. In *Artificial Intelligence Research and Development*, pages 71–80. IOS Press.
- Delfina S Martinez Pandiani, Erik Tjong Kim Sang, and Davide Ceolin. 2025. ‘toxic’memes: A survey of computational perspectives on the detection and explanation of meme toxicities. *Online Social Net-works and Media*, 47:100317.
- Kim Ngan Phan, Guee-Sang Lee, Hyung-Jeong Yang, and Soo-Hyung Kim. 2022. Little flower at memo-tion 2.0 2022: Ensemble of multi-modal model using

- attention mechanism in memotion analysis. In *Proceedings of De-Factify: Workshop on Multimodal Fact Checking and Hate Speech Detection, CEUR*.
- Shraman Pramanick, Dimitar Dimitrov, Rituparna Mukherjee, Shivam Sharma, Md Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. 2021a. Detecting harmful memes and their targets. *arXiv preprint arXiv:2110.00413*.
- Shraman Pramanick, Shivam Sharma, Dimitar Dimitrov, Md Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. 2021b. Momena: A multimodal framework for detecting harmful memes and their targets. *arXiv preprint arXiv:2109.05184*.
- Jingnong Qu, Liunian Harold Li, Jieyu Zhao, Sunipa Dev, and Kai-Wei Chang. 2022. Disinfomeme: A multimodal dataset for detecting meme intentionally spreading out disinformation. *arXiv preprint arXiv:2205.12617*.
- Yiting Qu, Xinlei He, Shannon Pierson, Michael Backes, Yang Zhang, and Savvas Zannettou. 2023. On the evolution of (hateful) memes by means of multimodal contrastive learning. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 293–310. IEEE.
- Kshitij Rajput, Raghav Kapoor, Kaushal Rai, and Preeti Kaur. 2022. Hate me not: detecting hate inducing memes in code switched languages. *arXiv preprint arXiv:2204.11356*.
- Sathyanarayanan Ramamoorthy, Nethra Gunti, Shreyash Mishra, S Suryavardan, Aishwarya Reganti, Parth Patwa, Amitava DaS, Tanmoy Chakraborty, Amit Sheth, Asif Ekbil, and 1 others. 2022. Memotion 2: Dataset on sentiment and emotion analysis of memes. In *Proceedings of De-Factify: workshop on multimodal fact checking and hate speech detection, CEUR*, volume 17.
- Benet Oriol Sabat, Cristian Canton Ferrer, and Xavier Giro-i Nieto. 2019. Hate speech in pixels: Detection of offensive memes towards automatic moderation. *arXiv preprint arXiv:1910.02334*.
- Lanyu Shang, Christina Youn, Yuheng Zha, Yang Zhang, and Dong Wang. 2021a. Knowmeme: A knowledge-enriched graph neural network solution to offensive meme detection. In *2021 IEEE 17th International Conference on eScience (eScience)*, pages 186–195. IEEE.
- Lanyu Shang, Yang Zhang, Yuheng Zha, Yingxi Chen, Christina Youn, and Dong Wang. 2021b. Aomd: An analogy-aware approach to offensive meme detection on social media. *Information Processing & Management*, 58(5):102664.
- Chhavi Sharma, Deepesh Bhageria, William Scott, Srinivas Pykl, Amitava Das, Tanmoy Chakraborty, Viswanath Pulabaigari, and Bjorn Gambäck. 2020. Semeval-2020 task 8: Memotion analysis—the visuo-lingual metaphor! *arXiv preprint arXiv:2008.03781*.
- Shivam Sharma, Md Shad Akhtar, Preslav Nakov, and Tanmoy Chakraborty. 2022a. Disarm: Detecting the victims targeted by harmful memes. *arXiv preprint arXiv:2205.05738*.
- Shivam Sharma, Tharun Suresh, Atharva Kulkarni, Himanshi Mathur, Preslav Nakov, Md Shad Akhtar, and Tanmoy Chakraborty. 2022b. Findings of the constraint 2022 shared task on detecting the hero, the villain, and the victim in memes. In *Proceedings of the Workshop on Combating Online Hostile Posts in Regional Languages during Emergency Situations*, pages 1–11.
- Limor Shifman. 2014. The cultural logic of photo-based meme genres. *Journal of visual culture*, 13(3):340–358.
- Limor Shifman. 2019. Internet memes and the twofold articulation of values. *Society and the internet: How networks of information and communication are changing our lives*, pages 43–57.
- Shardul Suryawanshi, Bharathi Raja Chakravarthi, Mihael Arcan, and Paul Buitelaar. 2020a. Multimodal meme dataset (multioff) for identifying offensive content in image and text. In *Proceedings of the second workshop on trolling, aggression and cyberbullying*, pages 32–41.
- Shardul Suryawanshi, Bharathi Raja Chakravarthi, Mihael Arcan, and Paul Buitelaar. 2023. Trollswithopinion: A taxonomy and dataset for predicting domain-specific opinion manipulation in troll memes. *Multi-media Tools and Applications*, 82(6):9137–9171.
- Shardul Suryawanshi, Bharathi Raja Chakravarthi, Pranav Verma, Mihael Arcan, John Philip McCrae, and Paul Buitelaar. 2020b. A dataset for troll classification of tamilmemes. In *Proceedings of the WILDRE5–5th workshop on indian language data: resources and evaluation*, pages 7–13.
- Leo Tolstoy. 1878. *Anna karenina*. The Russian Messenger. First published in Russian 1873–1877.
- Riccardo Tommasini, Filip Ilievski, and Thilini Wijesiriwardene. 2023. Imkg: The internet meme knowledge graph. In *European Semantic Web Conference*, pages 354–371. Springer.
- Pavan Kumar Anasosalu Vasu, Fartash Faghri, Chun-Liang Li, Cem Koc, Nate True, Albert Antony, Gokula Santhanam, James Gabriel, Peter Gräsch, Oncel Tuzel, and 1 others. 2025. Fastvlm: Efficient vision encoding for vision language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19769–19780.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.

Bo Xu, Tingting Li, Junzhe Zheng, Mehdi Naseriparsa, Zhehuan Zhao, Hongfei Lin, and Feng Xia. 2022. Met-meme: A multimodal meme dataset rich in metaphors. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, pages 2887–2899.

Ning Zhang, Feng Xuan, Gu Tianlong, and Chang Liang. 2023. Mvlp: Multi-perspective vision-language pre-training model for ethically aligned meme detection. *TechRxiv arXiv:10.36227/techrxiv.23689437.v1*.

A Appendix

A.1 Overview of Toxicity Benchmarks

Table 4 provides an overview of toxicity work through the lens of the dimensions in our framework. Note that all of these dimensions, except for the last one, are covered by our methodological framework. The final dimension is template structure, The *Template* feature is focused on representing the meme image’s structural organization. Some example of Template patterns are: N-panels, vertical or horizontal progression, layout templates for X vs Y, dialogical organization, and structural patterns. Several previous works focus on multi-faceted notions of meme templates (Sharma et al., 2022a; Mathias et al., 2021; Pramanick et al., 2021a,b). This dimension is not covered in the current version of our framework; we detail possible related research paths in the Conclusions and Future Works Section.

A.2 Additional Related Work: Theoretical Frameworks Beyond Toxicity

Other research beyond toxicity benchmarks has focused on capturing the semantics of internet memes. Here, the Internet Meme Knowledge Graph (IMKG) (Tommasini et al., 2023) models textual, visual, and external knowledge about memes by connecting three sources. A key resource in IMKG is KnowYourMeme,⁶ which catalogs internet meme templates and instances together with background knowledge on their emergence. OnToxKG (Martinez Pandiani et al., 2025) contributes an ontology-grounded knowledge graph of nearly 800 symbols, and shows its impact on downstream analysis of toxicity. Bollici et al. (2026) operationalize the notion of narratives in memes into 1) entities, roles, and social identities; 2) legitimization strategies; and 3) emerging narratives. Pandiani et al. (2025) organize prior work on toxicity in memes, yielding a taxonomy of toxicity categories, and defining three core dimensions of toxicity: intent (why), tactic (how), and target (who). Our semantic framework and associated benchmark M-QUEST builds on these theoretical insights to provide a practical tool for characterizing hate speech via underlying semantic dimensions and for evaluating the ability of VLMs to discover such links automatically.

⁶<https://knowyourmeme.com/>

A.3 Running Example for our Extraction Procedure

In the running example shown in Fig. 1, meme 05864.png from the Hateful Memes dataset depicts three camels on in a desert-like scenario, with upper text: `trust me...`, followed by lower text: `...they’re all terrorists`.

The meme is passed to the KE module, where Qwen3-VL-8b-Instruct is prompted to extract the dimensions sequentially. The prompt order is shown in Table 5: Textual and visual dimensions are extracted first, and the knowledge material is stored and passed in the prompts for subsequent dimensions. The output is a knowledge base. A snippet of the generated Knowledge Base is shown in Fig. 1. Among others, three dimensions are displayed: Toxicity, Overall Intent, and Background Knowledge. The Toxicity is described as such: *The meme equates all camels (symbolizing Middle Easterners or Muslims) with terrorism, which is a direct attack on a protected group based on ethnicity and religion. This is a clear example of Islamophobia and hate speech.*

Further knowledge extracted for Overall Intent and Background Knowledge dimensions is shown below Toxicity in Fig. 1.

The knowledge base is stored, and salient knowledge material is passed to the Q&A Generation component. For the Toxicity dimension, the above mentioned chunk of text is passed, successfully resulting in the question: This meme’s message targets which group?, with the correct answer (*a4*: Muslims or Middle Easterners) and distractors (*a1*: Camel owners only; *a2*: Animal lovers everywhere; and *a3*: None of the others). Finally, the provided explanation justifies answer *a4* as correct, identifying a mapping between camels and the term “terrorists,” with background knowledge reference to the stereotype of Muslims being terrorists (possibly following the 9/11 terrorist attack).

Modularity and Component Agnosticism Note that the entire framework is modular and component-agnostic: the knowledge base can be enriched with additional modules, such as refined modules focused on emotions, adopting a theory-driven approach and reusing well-established emotion theories as in (Maity et al., 2022). The more refined dimension definitions would then be used to prompt the VLM. The neural component—in this case a VLM—could be replaced with other spe-

Table 4: **Memes dimensions have been sparsely and non-systematically analysed.** Here the full table of the 40 works grounding our framework used to build M-QUEST benchmark.

Paper Reference	Toxicity	Textual Material	Visual Material	Scene	Background Knowledge	Overall Intent	Emotion	Target Community	Semiotic Projection	Analogical Mapping
(Hee et al., 2023)	✓	✓	✓	✗	✓	✓	✓	✓	✗	✗
(Bhowmick et al., 2021)	✓	✓	✓	✗	✓	✗	✓	✓	✗	✗
(Nguyen et al., 2022)	✗	✓	✓	✗	✗	✗	✓	✗	✗	✗
(Shang et al., 2021b)	✓	✓	✓	✓	✗	✗	✓	✗	✗	✓
(Qu et al., 2022)	✓	✓	✓	✓	✓	✓	✓	✗	✗	✓
(Lin et al., 2024)	✓	✓	✓	✗	✗	✗	✗	✓	✗	✗
(Rajput et al., 2022)	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
(Sabat et al., 2019)	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
(Badour and Brown, 2021)	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
(Shang et al., 2021a)	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗
(Myilvahanan et al., 2023)	✗	✓	✓	✓	✓	✗	✗	✗	✗	✓
(Das and Mukherjee, 2023)	✓	✓	✓	✗	✗	✓	✓	✓	✗	✗
(Gasparini et al., 2022)	✓	✓	✗	✗	✗	✓	✗	✓	✗	✗
(Kumari et al., 2023)	✓	✓	✓	✗	✗	✗	✓	✓	✗	✗
(Kumari et al., 2024)	✗	✓	✓	✗	✗	✓	✗	✓	✗	✗
(Sharma et al., 2022b)	✗	✓	✓	✓	✓	✗	✓	✗	✗	✗
(Mathias et al., 2021)	✓	✓	✓	✗	✗	✓	✓	✓	✗	✗
(Phan et al., 2022)	✓	✓	✓	✗	✗	✓	✓	✗	✗	✗
(Jha et al., 2024)	✓	✓	✓	✗	✗	✓	✓	✗	✗	✓
(Ramamoorthy et al., 2022)	✗	✓	✓	✓	✗	✓	✓	✗	✗	✗
(Xu et al., 2022)	✗	✗	✗	✗	✗	✓	✓	✗	✗	✓
(Pramanick et al., 2021b)	✓	✓	✓	✗	✓	✗	✗	✓	✗	✗
(Alzu’bi et al., 2023)	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
(Suryawanshi et al., 2020a)	✓	✓	✓	✗	✗	✗	✗	✗	✗	✓
(Zhang et al., 2023)	✓	✓	✓	✗	✓	✗	✗	✗	✗	✗
(Kiela et al., 2020)	✓	✓	✓	✗	✗	✓	✗	✓	✗	✗
(Sharma et al., 2020)	✓	✓	✓	✗	✗	✓	✓	✗	✗	✗
(Hee et al., 2023)	✓	✓	✓	✗	✓	✗	✓	✓	✗	✗
(Suryawanshi et al., 2020b)	✓	✓	✓	✗	✓	✓	✗	✗	✗	✗
(Bacha et al., 2023)	✓	✓	✓	✗	✓	✗	✓	✗	✗	✗
(Bhowmick et al., 2021)	✓	✓	✓	✗	✓	✓	✗	✗	✗	✗
(Maity et al., 2022)	✓	✓	✓	✗	✗	✗	✓	✗	✗	✗
(Fersini et al., 2022)	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
(Suryawanshi et al., 2023)	✓	✓	✓	✗	✗	✓	✗	✗	✗	✗
(Bhandari et al., 2023)	✓	✓	✓	✗	✓	✓	✗	✗	✗	✗
(Pramanick et al., 2021a)	✓	✓	✓	✗	✓	✓	✗	✓	✗	✗
(Dimitrov et al., 2021)	✓	✓	✓	✗	✗	✓	✓	✗	✗	✗
(Sharma et al., 2022a)	✓	✓	✓	✓	✓	✗	✗	✓	✗	✗
(Qu et al., 2022)	✗	✓	✓	✗	✓	✓	✗	✗	✗	✗
(Nguyen et al., 2025)	✗	✓	✓	✗	✓	✓	✓	✓	✗	✗




Figure 3: Total number of knowledge facts generated for 1,000 memes, and the number of knowledge facts per dimension.

cialized components for specific tasks (e.g., emotion extraction). Finally, the KB used for retrieval-augmented generation can be substituted with one (or more) focused on a more specific topic. We used one focused on toxicity assessment here because it aligns with our scope.

Regarding the KE component, as shown in Fig. 3, the number of KBs generated per 1,000 memes is considerable, with at least one individual per dimension per meme, with notable exceptions being Emotion and Toxicity. As expected, Background Knowledge - the broadest dimension, encompassing potentially any kind of object - is the most prolific, generating five times as many individuals compared to other dimensions.

A.4 Additional Benchmark Construction Details

Table 6 details the composition of M-QUEST and in particular of M-QUEST-qwen.

Answer “None of the above” placement is also balanced: A: 23.15% | B: 25.45% | C: 23.48% | D: 27.91%. By design, “None of the above” was meant to be a distractor only, being the questions generated starting from - and grounded in - the knowledge extraction process.

A.5 M-QUEST-gemini Details

Fig. 4 details this models’ gap per dimension. Bigger models perform better, while the Qwen family models remain consistently high-performing despite their small size.

The per-dimension heatmap in Fig. 4 reveals that model capability tiers are not uniformly separated across all evaluation dimensions, but rather diverge most sharply on dimensions requiring la-

tent or pragmatic inference. *Overall Intent* emerges as the most uniformly difficult dimension: even top-performing models show their largest within-model drop here (Claude-Opus-5 gets to 96% and Gemini-3.5-Flash-Lite to 82%, their lowest scores across all ten dimensions), while weaker models collapse most severely (InstructBLIP-Vicuna-7B at 20%, LLaVA-1.5-7B at 26%). By contrast, *Textual Material* shows the widest overall range (from 14% for LLaVA-1.5-7B to 100% for both proprietary models), making it the single most discriminative dimension for separating weak from strong models, though it is not uniformly hard: several mid-tier open models such as Pixtral-12B and Qwen3-VL-8B-Instruct achieve near-ceiling accuracy on it.

A second pattern concerns capability-tier consistency rather than raw score. Within the *Instruction-tuned + reasoning* group (see Table 2), Qwen3-VL-8B-Instruct dominates its smaller Qwen-family models on perceptual dimensions (*Visual Material*, *Background Knowledge*: both 100%) but the gap narrows or reverses on *Scene*, where Qwen2.5-VL-7B-Instruct (62%) and Qwen2-VL-7B-Instruct (42%) trail Qwen3 by a wide margin, suggesting that scene-level grounding benefits disproportionately from the newer backbone rather than scaling alone. The *Reasoning-oriented* group (LLaVA-1.5-7B, LLaVA-1.6-Vicuna-7B) is the only tier that fails to show a clear internal hierarchy, performing inconsistently across dimensions (e.g., LLaVA-1.6 outperforms LLaVA-1.5 on *Semiotic Projection* and *Analogical Mapping* but underperforms it on *Scene* and *Target Community*), indicating that its instruction-tuning does not generalize evenly across dimension types. Finally, the two proprietary models track each other closely across most dimensions, with Gemini-3.5-Flash-Lite trailing Claude-Opus-5 primarily on dimensions requiring integrative judgment (*Overall Intent*, *Semiotic Projection*) rather than on perceptual ones, where the two are statistically indistinguishable (both at 100% on *Toxicity*, *Textual Material*, and *Visual Material*).

Overall, the qualitative trends from M-QUEST-qwen, namely, the difficulty of latent toxicity relative to explicit reasoning, and the benefit of combining instruction tuning with reasoning, hold on M-QUEST-gemini as well, though their magnitude varies with the generating model.

Table 5: Framework sequential extraction order and context dependencies across dimensions.

Step	Dimension	Receives Context From
1	OverallIntent	None (extracted first)
2	TextualMaterial	OverallIntent graph
3	VisualMaterial	OverallIntent graph
4	Scene	OverallIntent graph + VisualMaterial entities
5	BackgroundKnowledge	OverallIntent graph + VisualMaterial, TextualMaterial, Scene entities
6	EmotionExpression	OverallIntent graph + VisualMaterial, TextualMaterial, Scene, BackgroundKnowledge entities
7	AnalogicalMapping	OverallIntent graph + VisualMaterial, TextualMaterial, Scene, BackgroundKnowledge, EmotionExpression entities
8	SemioticProjection	OverallIntent graph + VisualMaterial, TextualMaterial, Scene, BackgroundKnowledge, AnalogicalMapping entities
9	ToxicityAssessment	OverallIntent graph + VisualMaterial, TextualMaterial, Scene, BackgroundKnowledge, EmotionExpression, AnalogicalMapping, SemioticProjection entities
10	TargetCommunity	OverallIntent graph + VisualMaterial, TextualMaterial, Scene, BackgroundKnowledge, AnalogicalMapping, ToxicityAssessment entities

Table 6: Q&As distribution across dimensions for different benchmark subsets. M-QUEST-qwen is in turn divided in Tox-M-QUEST-qwen, focused on toxicity, and Gen-M-QUEST-qwen, about all the other dimensions. Being focused on Toxicity, Tox-M-QUEST-qwen is the set providing memes to all other M-QUESTpartitions: it consists of 307 memes, the number 180 for Gen-M-QUEST-qwen and 232 for M-QUEST-gemini in the “Unique Memes” column, should be interpreted as a subsets of the 307 memes in the ToxM-QUEST-qwen row.

Subset	Toxicity	Textual Material	Visual Material	Scene	Background Knowledge	Overall Intent	Emotion	Analogical Mapping	Target Community	Semiotic Projection	Unique Memes	Avg. Q&As per Meme
M-QUEST-qwen												
ToxM-QUEST-qwen	307	10	14	9	9	9	—	—	—	—	307	1.16
GenM-QUEST-qwen	—	35	35	35	26	26	30	28	20	16	180	1.48
M-QUEST-gemini												
M-QUEST-gemini	50	50	50	50	50	50	50	50	50	50	232	2.16
M-QUEST	357	95	99	94	85	85	80	78	70	66	307	3.61

A.6 Additional Results of the Human Validation on M-QUEST-qwen

As illustrated in Fig. 6, the automatic extraction achieves particularly high validation rates for fundamental dimensions, with Scene (76.0%), Textual Material (76.0%), and Visual Material (74.0%) all substantially exceeding the average of 59.33%. These results are noteworthy given that these basic dimensions serve as building blocks for more complex cognitive assessments. Mid-level dimensions such as Overall Intent (64.0%) and Emotion (60.0%) maintain solid validation rates, while the validation of the more abstract and cognitively demanding dimensions like Analogical Mapping (56.0%) is near the average, reflecting the inherent complexity of these reasoning tasks. Surprisingly, Background Knowledge, although a more fundamental dimension, performs near-average (56.0%), possibly again due to the broadness of the dimension, which creates confusion in the model when generating Q&As. The relatively lower validation

rates for Target Community (40.0%) and Semiotic Projection (32.0%) can be attributed to: (i) their highly contextual and interpretative nature, which often requires deeper cultural knowledge and subjective judgment, (ii) their inconsistent presence in memes, and (iii) possible definitional confusion.

Inter Annotator Agreement is measurable but limited, stronger for toxicity and weaker for Q&A plausibility, likely reflecting subjectivity and label-prevalence effects. However, the human baseline is high (83%), attributable to selective inclusion criteria: unanimous annotator agreement on Q&A Plausibility yields 91.7% benchmark inclusion; two-of-three agreement yields 60.7%.

This gradient suggests that the benchmark construction process, achieved via filtering for those memes that received positive validation on Toxicity Assessment, and Answer Plausibility, while not explicitly using IAA as a filter, was biased toward items with higher rater consensus, ensuring high-quality data.




Figure 4: Heatmap per dimension, per model, for M-QUEST-gemini.

A.7 Additional CoT Examples

A Chain-of-Thought example for a Q&A in M-QUEST-qwen is shown in Fig. 5.

A.8 Prompt Template

A template of the prompt we use for zero-shot VLM evaluation is shown in Fig. 7.

A.9 Annotation Interface

Our human annotation task interface used both for M-QUEST-qwen and M-QUEST-gemini is shown in Fig. 8.

A.10 Model Generation Examples

An example M-QUEST-qwen Q&A pair with its explanation is shown in Fig. 9. An example with Q&A with different dimensions is shown in Fig. 10.



Figure 5: Illustration of the diverse reasoning patterns of VLMs in detecting toxicity assessment dimension in the same meme image. Qwen2 and Qwen2.5 exhibit more straightforward reasoning patterns and are less capable of reasoning about latent relationships compared to the three other VLMs/VLM families. Meanwhile, Qwen3 provides the most sophisticated interpretation.




Figure 6: M-QUEST-qwen human validation per each dimension.

```

{
  "messages": [
    {
      "role": "user",
      "content": [
        {"type": "image", "image_path": "image.jpg"},
        {"type": "text", "text": "{QUESTION}\n\n
          A. {option_A}, B. {option_B}, C. {option_C}, D. {option_D}\n\n
          First, analyze the meme image carefully and explain your reasoning
          (Chain of Thought). Then provide your answer by selecting one of the
          options above."}
      ]
    }
  ]
}

```

Figure 7: Prompt template for zero-shot VLM evaluation.



Toxicity

The meme directly attacks a religious group based on a false and degrading stereotype, associating Muslims with sexual violence and predatory behavior, which is a clear example of Islamophobia.

Is the above statement about the meme on point?

☒ YES ☐ NO ☐ MAYBE

this Toxicity is related to:

BackgroundKnowledge

The meme relies on the harmful stereotype that Muslim men are inherently sexually violent or uncontrollable, which is a common trope in Islamophobic discourse. This stereotype is not visible in the image but is essential to understanding the hateful intent of the meme's text.

Is the above statement about the meme on point?

☒ YES ☐ NO ☐ MAYBE

Q&As

ToxicityAssessment

Question: What type of stereotype does this meme promote?

Answers:

- A: None of the others
- B: Misplaced blame for economic issues
- C: Islamophobic sexual violence trope ✓ (Correct)
- D: Generalized animal cruelty claim

Does it make sense to ask this question about this meme?

☒ YES ☐ NO ☐ MAYBE

Is the proposed correct answer plausible?

☒ YES ☐ NO ☐ MAYBE

BackgroundKnowledge

Question: What stereotype does the meme's text rely on?

Answers:

- A: Sheep as religious symbols
- B: Muslims as sexually aggressive ✓ (Correct)
- C: Racing as a religious ritual
- D: None of the others

Does it make sense to ask this question about this meme?

☒ YES ☐ NO ☐ MAYBE

Is the proposed correct answer plausible?

☒ YES ☐ NO ☐ MAYBE

Next Meme

I don't understand / I'm not sure about this meme

Figure 8: Human annotation task user interface. Top: Toxicity assessment and the related toxic element (Background Knowledge). Bottom: Q&As generated for both the Toxicity and the related toxic element (Background Knowledge).

Question:

What does 'the fresh prince' imply about Meghan Markle?

Options:

- (A) She is an outsider in royal family
- (B) None of the others
- (C) She is secretly a spy
- (D) She is too wealthy for royals

Answer: A



Explanation

The term 'fresh prince' devalues Meghan by suggesting she does not belong, which is harmful and undermines her legitimacy.

Question:

What does the meme imply about historical justice?

Options:

- (A) It celebrates military victories only
- (B) It promotes slavery as necessary
- (C) It critiques hypocrisy in liberation movements
- (D) None of the others

Answer: C



Explanation

The meme highlights the irony of fighting for freedom while neglecting to free enslaved people, pointing to moral inconsistency.

Figure 9: Example of M-QUEST-qwen question-answer pairs along with their explanations.



Question: What are the people in the scene doing?

Options:

- (A) Walking in fruit costumes
- (B) None of the others
- (C) Dancing in a parade
- (D) Running across a field

Question: Which stereotype does this meme reinforce?

Options:

- (A) Cultural superiority of Western food
- (B) False equivalence about African diets
- (C) None of the others
- (D) Africa lacks all fruits and veggies

Dimension: Scene

The image features a group of five people. They are dressed in large, colorful, full-body mascots or costumes. Specifically, from left to right, they are dressed as a banana, a pea pod, a carrot, a bunch of grapes, and a red fruit or vegetable. While some are technically vegetables, they fall under the general category of produce or "fruit" in a colloquial sense. The people are positioned on a paved city street. Their body language—legs apart in mid-stride and holding hands—indicates forward motion. They appear to be walking, skipping, or running together in a line.

Dimension: Tox. Eval.

The meme plays on Western misconceptions about Africa being resource-poor. It suggests Africans are so deprived that they transform into the food itself. This perpetuates false stereotypes about the entire African continent lacking basic agricultural resources.

Figure 10: This figure presents the results of the Qwen models accurately reasoning over explicit visual cues but failing to capture latent contextual meaning, such as the misrepresentation of African diets in the meme.