

MER-Bench: Benchmarking Multimodal Emotional Reappraisal for Image-Text Memes

Yiqi Nie^{1,2}, Fei Wang^{2,*}, Junjie Chen³, Kun Li⁴, Yudi Cai^{2,5},
Dan Guo^{2,3}, Chenglong Li¹, Meng Wang^{2,3,*}

¹ Anhui University

² Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

³ Hefei University of Technology ⁴ United Arab Emirates University

⁵ University of Science and Technology of China

nieyiqi5@gmail.com, wangfei@iai.ustc.edu.cn

Abstract

Multimodal Emotional Reappraisal (MER) aims to shift affective interpretation by rewriting and editing a situation without altering its underlying semantic context. Despite advances in controllable editing and sentiment transfer, MER remains challenging under coupled multimodal constraints, especially in internet memes where images, embedded text, and layout jointly encode affective meaning. To bridge this gap, we present **MER-Bench**, a systematic benchmark for MER on image-text memes, containing over 3,000 real-world memes with fine-grained annotations for emotion transformation and structure-preserving multimodal editing. MER-Bench spans diverse visual, textual, and structural settings and is paired with a structured MLLM-as-a-Judge evaluation suite, comprising eight fine-grained metrics and three aggregated objectives for emotional alignment, content fidelity, and transformation strength. Extensive evaluation of 14 advanced multimodal models reveals that: ❶ current models struggle to jointly balance affect, semantics, structure, and text; ❷ strong visual editing does not imply effective MER, with poor embedded text rendering and affect-aware rewriting remaining key bottlenecks. Our results identify MER as a critical capability gap in current multimodal generation systems and highlight affective image-text reasoning as a key challenge for future work. Project Page: <https://github.com/one-seven17/MER-Bench>.

1 Instruction

Emotional Reappraisal is a fundamental cognitive strategy for shifting affective interpretation by reinterpreting a situation without altering its underlying semantic context (Gross, 1998; Gross et al., 2014; Su et al., 2025). As illustrated in Fig. 1, Multimodal Emotional Reappraisal (MER)




Figure 1: Illustration of MER, which requires affective shift, semantic preservation, and structural fidelity when transforming negative memes into positive ones.

extends this idea to image-text settings, requiring models to transform emotional framing while jointly preserving semantic intent and structural coherence. Such reappraisal is particularly challenging in image-text memes, where affective meaning is jointly constructed through visual content, embedded text, layout structure, and culturally shared templates (Mortensen and Neumayer, 2021; Lin et al., 2026; Denisova, 2019). As a prevalent form of online affective expression, memes frequently convey frustration, anxiety, sarcasm, ridicule, and polarized attitudes through highly compressed multimodal signals (Brubaker et al., 2018), making them a challenging testbed for MER.

Recent studies (Lin et al., 2026; Sharma et al., 2024; Bai et al., 2025; Kiela et al., 2020; Yuan et al., 2025; Ouyang et al., 2025; Li et al., 2024) have substantially advanced meme understanding, harmful content detection, and template-aware gen-

*Corresponding author.

eration (Kiela et al., 2020; Lin et al., 2026; Bai et al., 2025; Chen et al., 2024b; Cai et al., 2025b; Wang et al., 2024b,a). Meanwhile, multimodal generation systems have demonstrated strong capabilities in controllable editing and affect-aware generation (Yang et al., 2025b, 2024, 2025a; Wang et al., 2026a). However, existing approaches primarily focus on recognition, generation, or coarse sentiment control, leaving emotional reappraisal under coupled multimodal constraints largely unexplored. Compared with conventional sentiment editing, MER introduces several unique challenges: (1) emotional meaning in memes is often implicit, culturally grounded, and structurally entangled across modalities, requiring multimodal reasoning over both explicit and implicit cues (Wiggins, 2019); (2) emotional transformation requires controllable modeling of directional affect shifts while preserving multimodal semantic consistency; and (3) evaluation is inherently multi-objective, since valid outputs must jointly satisfy emotional transformation, structural preservation, and content fidelity.

Towards this end, we present **MER-Bench**, a benchmark for MER on image-text memes. Grounded in cognitive reappraisal theory and Russell’s circumplex model of affect, MER-Bench operationalizes emotional reappraisal through four emotion-shift pairs, structured prompt-based rewriting protocols, and structure-preserving multimodal editing objectives. The benchmark contains over 3,000 real-world memes spanning diverse visual modalities, structural layouts, and affective conditions, with fine-grained annotations covering emotion transition, meme text rewriting, visual editing specifications, and structural taxonomy labels. These annotations enable systematic evaluation of whether models can achieve controllable affective shifts while preserving the semantic and structural identity of the original meme.

To systematically evaluate model capabilities under coupled multimodal constraints, we further introduce a structured **MLLM-as-a-Judge** evaluation suite. The framework includes eight fine-grained metrics: Visual Generation Quality (VGQ), Visual Emotion Alignment (VEA), Text Generation Quality (TGQ), Text Emotion Alignment (TEA), Layout Consistency (LC), Holistic Generation Quality (HGQ), Perceived Primary Emotion (PPE) and Perceived Emotion Shift (PES). To further capture the conjunctive nature of emotional reappraisal, we additionally introduce three aggregated objectives: Target Alignment Score (TAS),

Content Fidelity Score (CFS), and Shift Strength Score (SSS), which are unified into the final Reappraisal Fidelity Score (RFS). To validate the reliability of the proposed evaluation protocol, we conduct human evaluation and cross-configuration analysis under different prompting strategies. Four human experts score sampled outputs across multiple prompting configurations, providing reference annotations for cross-validation with MLLM-based judgments. The strong agreement across multiple statistical measures confirms the effectiveness, stability, and human alignment of our automatic evaluation framework. Applying this evaluation suite to 14 advanced multimodal generation models, we uncover two key findings: ❶ Current models often satisfy only part of MER’s coupled requirements, struggling to jointly preserve structure, maintain semantic identity, and perform target affective reinterpretation; ❷ MER is a multimodal reasoning problem rather than a visual editing or positive-generation task, as models must preserve the structural and pragmatic conditions of meme meaning while modifying affective framing.

Thus, our contributions are as follows:

- **Systematic Benchmark:** We introduce MER-Bench, a novel benchmark for MER in image-text memes, organized around three taxonomic axes and nine subcategories to evaluate controllable affective shifts under structure-preserving multimodal constraints.
- **Evaluation Suite:** We propose a structured MLLM-as-a-Judge evaluation suite for assessing MER under coupled multimodal constraints, and validate its reliability through human expert study and prompt ablation.
- **Empirical Analysis:** We evaluate 14 advanced multimodal models on MER-Bench, showing that effective MER requires coordinated affective transformation, semantic preservation, structural consistency, and embedded text rewriting, especially when meme meaning depends on rigid templates, sequential layouts, or ambiguous affective cues.

2 Related Works

2.1 Meme Emotion Understanding

Early benchmarks, such as the Facebook Hateful Memes challenge, formulate meme analysis as multimodal hate speech detection, advancing

Table 1: Benchmark comparison across emotion-related multimodal tasks.

Dataset	Modality	Emotion Recognition	Emotion Reasoning / Regulation	Emotion Generation	Social Media
Memotion 3 (Mishra et al., 2023)	Image-Text	✓	✗	✗	✓
MOOD (Sharma et al., 2024)	Image-Text	✓	✗	✗	✓
GOAT-Bench (Lin et al., 2026)	Image-Text	✓	✗	✗	✓
EmotionBench (Huang et al., 2023)	Text	✓	✓	✓	✗
EmoBench (Sabour et al., 2024)	Text	✓	✓	✓	✗
MOSABench (Song et al., 2024)	Image	✓	✗	✗	✗
MM-InstructEval (Yang et al., 2025c)	Image	✓	✗	✓	✗
MER-Bench (Ours)	Image-Text	✓	✓	✓	✓

image-text fusion and robustness to spurious correlations (Kiela et al., 2020). Subsequent datasets, including MemeInterpret (Park et al., 2025) and GOAT-Bench (Lin et al., 2026), introduce culturally grounded categories and fine-grained labels for real-world distribution shifts. Beyond classification, MemeReaCon (Zhao et al., 2025) and MET-Meme (Xu et al., 2022) require reasoning over social context, user comments, communicative intent, and multimodal metaphors. Sharma introduces MOOD, a six-emotion meme dataset, and ALFRED, an emotion-aware multimodal fusion model using visual cues and gated cross-modal fusion (Sharma et al., 2024). Memes have also been studied as socio-psychological signals of collective attitudes and public discourse (Brubaker et al., 2018; Wang et al., 2026b,a).

2.2 Meme Generation

Meme generation spans template-aware captioning, text-to-image generation, and punchline generation. Layout-aware systems XMeCap (Chen et al., 2024b) and IterMeme (Cai et al., 2025b) jointly model template structures and humorous text via template-conditioned or cloze-style generation. MemeBot (Sadasivam et al., 2020) frames meme creation as a translation-like process combining template images with captions. Large-scale resources such as Memeify (Vyalla and Udandara, 2020) support caption generation, while MemeCraft (Wang and Lee, 2024) leverages LLMs/VLMs for socially aligned meme synthesis. HUMOR adopts multi-path chain-of-thought reasoning with preference-based reinforcement learning (Li et al., 2025b), while theory-guided frameworks (Zhang et al., 2026b; Li et al., 2025a) integrate psychological humor theories.

2.3 Emotion-Controllable Generation

Emotion-controllable generation has been explored in both textual and visual domains. Text-based methods perform attribute-conditioned style trans-

fer while approximately preserving semantic content (Xu et al., 2025; Li et al., 2026), with extensions addressing finer cues such as sarcasm. Visual approaches control affect through emotion-conditioned generation or image manipulation, as in EmoGen (Yang et al., 2024), EmoCtrl (Yang et al., 2025b), and EmoEdit (Yang et al., 2025a). Diffusion-based models such as T2i-Adapter (Mou et al., 2024), IP-Adapter (Ye et al., 2023), and Instruct-Pix2Pix (Brooks et al., 2023) enable structured conditioning for layout and style control. Diffusion-DPO (Wallace et al., 2024) instead aligns diffusion models with human preferences through direct preference optimization. As summarized in Tab. 1, existing methods focus on coarse affect control and cannot jointly satisfy the semantic, structural, and reappraisal constraints required for meme transformation.

3 Task Definition

Let $\mathcal{M} = (I, T, \ell)$ denote an image-text meme, where I , T , and ℓ represent the visual content, overlaid text, and layout structure, respectively. We formulate MER as a constrained generation task grounded in emotion regulation theory (Gross, 1998; Gross et al., 2014). Given a source meme $\mathcal{M}^- = (I^-, T^-, \ell^-)$ with source affect a^- and a target affect a^+ , the goal is to generate a reappraised meme $\mathcal{M}^+ = (I^+, T^+, \ell^+)$ that shifts the perceived affect from a^- toward a^+ while preserving the original semantic situation and structure:

$$\mathcal{M}^+ = f(\mathcal{M}^-, a^+). \quad (1)$$

The generated meme should satisfy three coupled requirements:

$$a(\mathcal{M}^+) \approx a^+, \quad S(\mathcal{M}^+) \approx S(\mathcal{M}^-), \quad \ell^+ \approx \ell^-, \quad (2)$$

where $a(\cdot)$ denotes perceived affective framing, and $S(\cdot)$ denotes the underlying semantic situation, including core entities, events, and relations. Unlike conventional sentiment editing or open-ended




Figure 2: **Construction pipeline of MER-Bench.** (A) Psychological theories define the affective space and reappraisal objective for emotion transformation. (B) Iterative annotation, expert clarification, and annotator calibration refine labeling guidelines and resolve ambiguous cases. (C) LLM-assisted human rewriting performs emotion detection, reappraisal planning, and structure-preserving image-text meme generation with human verification.

meme generation, MER requires changing how a meme is emotionally interpreted without changing what it is about or how its multimodal content is structurally organized. Thus, successful MER requires jointly satisfying affective shift, semantic preservation, and structural fidelity.

4 Dataset Construction

To benchmark MER in a controlled and comparable setting, we construct **MER-Bench** from the **MOOD** meme corpus (Sharma et al., 2024). Starting with 10,005 memes, we retain 6,695 samples with clearly identifiable negative affective framing and then apply quality, duplication, safety, and reappraisal-feasibility filters, resulting in 3,117 high-quality image-text meme instances. Detailed filtering statistics are provided in Appendix A. Fig. 2 illustrates the overall construction pipeline.

4.1 Psychological Foundations

We define the reappraisal target space using Russell’s circumplex model of affect (Russell, 1980). As shown in Fig. 2 (A-1), MER-Bench covers four negative-to-constructive affective transitions: *Sad* → *Happy*, *Angry* → *Calm*, *Tense* → *Relaxed*, and *Bored* → *Excited*, with 1,026, 937, 1,020, and 134 samples (Fig. S7), respectively. Following Gross’s process model of emotion regulation (Gross, 1998; Gross et al., 2014), we design a reappraisal-oriented rewriting protocol that changes affective interpretation while preserving the underlying semantic situation and structure.

4.2 Construction Pipeline and Control

MER-Bench is constructed through an LLM-assisted human annotation workflow. Five trained annotators with graduate-level backgrounds participate in annotation, while four experts, including three computer science experts and one psychology expert, support rule refinement, ambiguity resolution, and quality control. Each candidate sample is checked by four annotators, with unresolved disagreements escalated to expert review.

After annotator calibration, we use a three-stage prompting workflow: (i) emotion detection and image description, (ii) reappraisal planning with target emotion and editing intent, and (iii) reappraised meme generation. The first two stages use Gemini-2.5-Flash¹, while the final multimodal generation stage uses Gemini-2.5-Flash-Image². LLMs are used only as annotation assistants; all retained samples are manually verified before inclusion. Detailed prompt templates are provided in Appendix A.3. Each sample is checked for source-emotion validity, target-emotion consistency, semantic preservation, and structural editability. Samples are revised or removed if the source affect is ambiguous, the target reappraisal is infeasible, the reappraised caption changes the core situation, or the generated meme fails to preserve essential layout and entity information. The complete annotation schema, human-control statistics, and inter-

¹<https://deepmind.google/technologies/gemini/flash/>

²<https://ai.google.dev/gemini-api/docs/models/gemini-2.5-flash-image>




Figure 3: MER-Bench provides a unified taxonomy and evaluation metrics for MER.

annotator agreement are provided in Appendix A.

4.3 Dataset Statistics

Fig. 3 summarizes the distribution of MER-Bench across three taxonomy axes and nine subcategories. The visual modality axis includes four categories: object/animal-centric, photographic human, illustrated/cartoon, and template-stylized memes. The structural layout axis distinguishes single-panel and multi-panel memes. The affective polarity axis describes the caption-level sentiment polarity of the original meme text, including negative, neutral, and positive captions. This taxonomy characterizes MER difficulty from complementary perspectives: visual grounding, layout complexity, and textual affective signal. Template-stylized memes form the largest visual category, reflecting the prevalence of template-based meme creation, while single-panel memes dominate the structural distribution, matching common meme formats on social media. Although MER-Bench is constructed with a negative holistic affective framing, some source captions are neutral or positive on the surface because negative affect can also be conveyed through visual context, layout, irony, or image-text incongruity.

5 Evaluation Protocol

5.1 MLLM-as-a-Judge

Evaluating MER is inherently multi-objective: a valid output must express the target emotion while preserving semantic content, meme structure, and textual readability. Single-axis metrics such as sentiment accuracy or semantic similarity cannot capture these coupled requirements (Sharma et al., 2024; Sharma and Patil, 2026; Lin et al., 2026). We adopt a structured *MLLM-as-a-Judge* paradigm (Chen et al., 2024a), using

Gemini-3-Pro³ as the judge. Given the source meme, edited meme, and target emotion, the judge returns a schema-constrained JSON output. To reduce free-form judgment bias, we impose three constraints: (i) *target-gated emotion scoring*, which limits alignment scores when the perceived emotion mismatches the target; (ii) *localized rationales*, which justify each score using criterion-specific evidence; and (iii) *explicit layout verification*, which checks panel structure, caption placement, and spatial consistency. The complete judge prompt and JSON schema are provided in Appendix C.

We define eight fine-grained metrics covering visual quality, text quality, emotional alignment, structural fidelity, and holistic reappraisal success. Except for *Perceived Primary Emotion (PPE)*, which records the dominant emotion from an eight-emotion label space plus other, all scalar metrics use a 1-5 Likert scale (Mellor and Moore, 2014): *Visual Generation Quality (VGQ)* assesses image quality and artifacts; *Visual Emotion Alignment (VEA)* measures visual match to the target emotion; *Text Generation Quality (TGQ)* evaluates caption legibility, formatting, fluency, and meme-style naturalness; *Text Emotion Alignment (TEA)* measures whether the caption conveys the target emotion while preserving the scenario; *Layout Consistency (LC)* checks panel structure, caption placement, and spatial organization; *Holistic Generation Quality (HGQ)* assesses overall visual-text coherence, readability, and content preservation; and *Perceived Emotion Shift (PES)* quantifies the strength of the target-consistent emotional shift. Together, these metrics disentangle quality, controllability, structure, and overall fidelity of reappraisal.

5.2 Prompt Sensitivity and Human Validation

Prompt design can affect the stability of MLLM-based evaluation. We study two orthogonal prompt factors: evidence order and reasoning order. For evidence order, we compare *Visual-First (VF)* and *Visual-Last (VL)* settings, which place visual evidence before or after textual task information. For reasoning order, we compare *Rationale-First (RF)* and *Rationale-Last (RL)* settings, which require localized rationales before or after scalar scores. Their combination yields four judge configurations: RF-VF, RF-VL, RL-VF, and RL-VL. Following expert-based affective evaluation practice in MME-Emotion (Zhang et al., 2026a), we conduct a hu-

³<https://deepmind.google/models/gemini/pro/>




Figure 4: Human validation and prompt ablation results for the MLLM-as-a-Judge evaluation protocol on 100 randomly sampled MER outputs.

man validation study to assess the reliability of our MLLM-as-a-Judge protocol. Four experts manually score 100 randomly sampled MER outputs using the same evaluation rubric as the automatic judge. These expert annotations serve as human reference scores for evaluating judge alignment across different prompt configurations.

We compare MLLM-based judgments with expert scores to examine whether the automatic evaluation captures human perception of multimodal quality, emotional alignment, and reappraisal success. In Fig. 4, RF-VF achieves the highest agreement with human judgments among the four configurations. This suggests that presenting visual evidence first and eliciting localized rationales before scalar scores better aligns MLLM-based evaluation with expert assessment. Overall, the results support for the reliability and human alignment of our schema-constrained MLLM-as-a-Judge protocol.

5.3 Aggregated Evaluation Score

Although the eight metrics provide fine-grained diagnostics, MER requires three objectives to hold jointly: target emotional alignment, content fidelity, and perceptible affective transformation. A model should not obtain a high overall score by excelling in only one dimension, such as visual quality, while failing to preserve structure or reach the target emotion. We therefore aggregate the metrics into three intermediate scores and a final conjunctive score. All scalar metrics are normalized to $[0, 1]$ before aggregation. Let $H \in \{0, 1\}$ be the target emotion hit indicator, where $H = 1$ if the perceived primary emotion matches the target emotion and $H = 0$ otherwise. Let E_v and E_t denote normalized visual and textual emotion alignment scores, Q_v , Q_t , and Q_o denote normalized visual, textual, and holistic generation quality scores, C denote normalized layout consistency, and S denote normalized perceived emotion shift.

Target Alignment Score (TAS). TAS measures whether the edited meme realizes the target emotion across visual and textual modalities:

$$\text{TAS} = H \cdot \text{HM}(E_v, E_t), \quad (3)$$

where $\text{HM}(x, y) = \frac{2xy}{x+y+\epsilon}$ is the harmonic mean. The target-hit gate suppresses the score when the perceived emotion misses the target, while the harmonic mean penalizes modality imbalance.

Content Fidelity Score (CFS). CFS measures semantic, structural, and generation-quality preservation:

$$\text{CFS} = C \cdot \text{GM}(Q_v, Q_t, Q_o), \quad (4)$$

where $\text{GM}(\cdot)$ denotes the geometric mean. The multiplicative layout term directly penalizes structural inconsistency.

Shift Strength Score (SSS). SSS quantifies the magnitude of target-consistent emotional transformation:

$$\text{SSS} = \frac{1 - \exp(-\alpha S)}{1 - \exp(-\alpha)}, \quad (5)$$

where $\alpha > 0$ controls the saturation rate. This normalized form maps $S = 0$ to 0 and $S = 1$ to 1.

Reappraisal Fidelity Score (RFS). Finally, we combine the three components using an ϵ -smoothed geometric conjunction:

$$\text{RFS} = [(TAS + \epsilon)(CFS + \epsilon)(SSS + \epsilon)]^{1/3} - \epsilon. \quad (6)$$

This soft conjunctive operator preserves the requirement that all three objectives should be satisfied, while avoiding numerical instability when one component is near zero. Thus, RFS reflects the logical structure of MER: strong reappraisal requires simultaneous target alignment, content fidelity, and affective transformation.

5.4 Evaluation Procedure

We evaluate 14 image-editing and multimodal-generation models spanning diffusion-based editors, instruction-guided frameworks, MLLM-conditioned generators, and large-scale image transformation systems. For each source meme and target emotion, a model generates one edited meme, which is independently evaluated using the MLLM-as-a-Judge protocol in Sec. 5.1. In total, $M = 14$ models are evaluated on $N = 3,117$ MER-Bench samples. Some models fail on certain inputs due to missing images, corrupted files,

Table 2: Model performance across evaluation metrics

Model	VGQ ↑	VEA ↑	TGQ ↑	TEA ↑	LC ↑	HGQ ↑	PPE ↑	PES ↑	TAS ↑	CFS ↑	SSS ↑	RFS ↑
Flux9B (Labs, 2025)	91.10 $\pm$ 1.16	98.12 $\pm$ 0.98	58.27 $\pm$ 2.00	92.47 $\pm$ 1.57	87.90 $\pm$ 2.33	68.77 $\pm$ 1.61	79.50 $\pm$ 2.99	99.70 $\pm$ 0.28	75.70 $\pm$ 3.07	62.81 $\pm$ 1.61	94.98 $\pm$ 0.04	76.78 $\pm$ 1.16
QwenEdit (Wu et al., 2025a)	84.08 $\pm$ 2.20	91.70 $\pm$ 2.33	50.68 $\pm$ 4.01	80.92 $\pm$ 2.57	85.80 $\pm$ 3.33	61.25 $\pm$ 3.21	75.10 $\pm$ 3.03	94.75 $\pm$ 2.06	64.57 $\pm$ 3.30	54.84 $\pm$ 3.99	94.16 $\pm$ 0.35	69.38 $\pm$ 2.77
Flux4B (Labs, 2025)	78.50 $\pm$ 1.53	95.98 $\pm$ 1.82	25.37 $\pm$ 1.49	74.42 $\pm$ 2.22	84.60 $\pm$ 4.03	47.35 $\pm$ 1.47	74.30 $\pm$ 4.32	98.85 $\pm$ 0.76	62.29 $\pm$ 4.05	38.49 $\pm$ 2.05	94.85 $\pm$ 0.12	61.06 $\pm$ 1.75
UniPic (Wei et al., 2026)	74.78 $\pm$ 1.24	61.72 $\pm$ 2.92	42.88 $\pm$ 3.35	62.85 $\pm$ 4.39	81.50 $\pm$ 1.72	40.27 $\pm$ 2.58	50.50 $\pm$ 4.74	70.13 $\pm$ 2.92	31.45 $\pm$ 3.76	41.17 $\pm$ 2.09	87.76 $\pm$ 1.10	48.43 $\pm$ 2.48
DreamO2 (Xia et al., 2025)	68.47 $\pm$ 2.54	46.25 $\pm$ 4.77	27.80 $\pm$ 2.70	48.30 $\pm$ 3.50	89.00 $\pm$ 2.94	25.45 $\pm$ 2.39	43.60 $\pm$ 4.84	55.45 $\pm$ 4.99	20.69 $\pm$ 3.70	32.37 $\pm$ 1.63	80.87 $\pm$ 2.69	37.79 $\pm$ 2.92
ZTurbo (Cai et al., 2025a)	85.10 $\pm$ 1.29	76.42 $\pm$ 4.38	70.38 $\pm$ 2.07	93.92 $\pm$ 1.79	9.00 $\pm$ 2.54	45.90 $\pm$ 2.06	65.70 $\pm$ 3.50	88.52 $\pm$ 2.62	55.40 $\pm$ 4.51	5.87 $\pm$ 1.71	92.96 $\pm$ 0.58	31.04 $\pm$ 3.34
LongCat (Team et al., 2025)	73.75 $\pm$ 1.80	83.97 $\pm$ 2.55	28.90 $\pm$ 5.11	5.58 $\pm$ 1.60	93.30 $\pm$ 1.83	29.23 $\pm$ 1.71	58.50 $\pm$ 4.35	67.88 $\pm$ 3.31	6.08 $\pm$ 1.66	36.89 $\pm$ 3.21	86.89 $\pm$ 1.32	26.74 $\pm$ 2.54
Step1X (Liu et al., 2025)	69.95 $\pm$ 2.39	71.20 $\pm$ 3.68	49.47 $\pm$ 5.06	5.20 $\pm$ 1.53	97.90 $\pm$ 1.52	28.55 $\pm$ 2.74	45.50 $\pm$ 4.81	54.65 $\pm$ 3.42	4.41 $\pm$ 1.32	45.22 $\pm$ 3.29	80.50 $\pm$ 2.03	25.13 $\pm$ 3.06
Bagel7B (Deng et al., 2025)	32.55 $\pm$ 3.62	23.07 $\pm$ 3.35	5.10 $\pm$ 2.37	0.58 $\pm$ 0.26	82.00 $\pm$ 3.27	4.55 $\pm$ 1.25	21.20 $\pm$ 3.36	23.40 $\pm$ 3.61	0.24 $\pm$ 0.11	7.29 $\pm$ 1.74	50.17 $\pm$ 5.75	4.86 $\pm$ 0.86
GoT (Fang et al., 2025)	26.77 $\pm$ 2.71	28.18 $\pm$ 4.20	34.88 $\pm$ 3.87	1.53 $\pm$ 0.89	65.70 $\pm$ 3.09	5.88 $\pm$ 1.33	24.10 $\pm$ 4.70	33.35 $\pm$ 4.57	0.74 $\pm$ 0.51	11.52 $\pm$ 1.46	62.92 $\pm$ 5.15	8.12 $\pm$ 1.98
IP2P (Brooks et al., 2023)	26.47 $\pm$ 2.64	17.88 $\pm$ 1.90	42.15 $\pm$ 4.50	1.05 $\pm$ 0.72	98.20 $\pm$ 1.87	4.47 $\pm$ 0.87	13.20 $\pm$ 2.74	23.03 $\pm$ 2.94	0.26 $\pm$ 0.21	16.72 $\pm$ 1.81	49.70 $\pm$ 4.58	6.37 $\pm$ 1.49
OmniG2 (Wu et al., 2025b)	22.00 $\pm$ 2.25	6.15 $\pm$ 2.91	42.37 $\pm$ 3.18	9.72 $\pm$ 2.99	84.40 $\pm$ 4.60	2.48 $\pm$ 1.17	9.40 $\pm$ 3.50	9.45 $\pm$ 3.22	0.78 $\pm$ 0.55	10.86 $\pm$ 1.48	24.37 $\pm$ 7.13	5.93 $\pm$ 2.13
ICEdit (Zhang et al., 2025)	28.82 $\pm$ 3.79	4.80 $\pm$ 0.70	10.85 $\pm$ 2.95	0.48 $\pm$ 0.48	89.70 $\pm$ 2.63	0.68 $\pm$ 0.44	5.40 $\pm$ 2.37	6.17 $\pm$ 1.05	0.04 $\pm$ 0.05	4.53 $\pm$ 2.57	16.87 $\pm$ 2.64	1.94 $\pm$ 0.82
SD (Rombach et al., 2022)	93.92 $\pm$ 1.01	84.88 $\pm$ 3.07	0.28 $\pm$ 0.25	1.20 $\pm$ 0.59	0.40 $\pm$ 0.52	18.93 $\pm$ 1.71	51.90 $\pm$ 7.19	91.00 $\pm$ 1.55	1.22 $\pm$ 0.63	0.03 $\pm$ 0.04	93.47 $\pm$ 0.31	2.32 $\pm$ 0.64

Notes: Results are reported on the comparable subset with $N = 2,711$ memes. For each of the 10 runs, we randomly sample 100 memes from this subset and evaluate all models on the same sampled instances. All values are normalized to $[0, 100]$ and reported as mean $\pm$ standard deviation over the 10 runs. Darker shading indicates better performance within each metric.

Table 3: Subcategory level performance across evaluation metrics

	VGQ ↑	VEA ↑	TGQ ↑	TEA ↑	LC ↑	HGQ ↑	PPE ↑	PES ↑	TAS ↑	CFS ↑	SSS ↑	RFS ↑
Visual Modality												
Object/Animal-Centric	65.65 $\pm$ 0.00	59.25 $\pm$ 0.00	38.75 $\pm$ 0.00	37.20 $\pm$ 0.00	74.00 $\pm$ 5.66	30.25 $\pm$ 0.00	49.60 $\pm$ 8.68	61.20 $\pm$ 0.00	22.65 $\pm$ 4.13	31.42 $\pm$ 3.34	83.99 $\pm$ 1.50	38.96 $\pm$ 2.59
Photographic Human	61.25 $\pm$ 0.00	60.05 $\pm$ 0.00	36.80 $\pm$ 0.00	35.75 $\pm$ 0.00	77.80 $\pm$ 2.90	29.60 $\pm$ 0.00	46.80 $\pm$ 6.68	61.65 $\pm$ 0.00	21.03 $\pm$ 3.88	31.48 $\pm$ 2.06	84.23 $\pm$ 1.17	38.14 $\pm$ 2.90
Illustrated/Cartoon	61.75 $\pm$ 0.00	55.30 $\pm$ 0.00	36.15 $\pm$ 0.00	37.20 $\pm$ 0.00	77.40 $\pm$ 2.99	28.60 $\pm$ 0.00	48.20 $\pm$ 5.20	57.25 $\pm$ 0.00	21.48 $\pm$ 3.35	30.90 $\pm$ 2.06	81.93 $\pm$ 2.15	37.86 $\pm$ 2.64
Template Stylized	60.05 $\pm$ 0.00	58.15 $\pm$ 0.00	30.75 $\pm$ 0.00	32.25 $\pm$ 0.00	74.20 $\pm$ 3.71	26.10 $\pm$ 0.00	45.00 $\pm$ 7.50	60.45 $\pm$ 0.00	18.60 $\pm$ 3.18	26.88 $\pm$ 1.26	83.61 $\pm$ 1.75	34.67 $\pm$ 2.28
Structural Layout												
Single-Panel	61.95 $\pm$ 0.00	56.55 $\pm$ 0.00	35.95 $\pm$ 0.00	34.65 $\pm$ 0.00	75.40 $\pm$ 3.13	28.15 $\pm$ 0.00	46.80 $\pm$ 5.01	58.65 $\pm$ 0.00	20.10 $\pm$ 3.29	29.82 $\pm$ 2.78	82.70 $\pm$ 1.91	36.71 $\pm$ 3.18
Multi-Panel	57.05 $\pm$ 0.00	54.65 $\pm$ 0.00	30.80 $\pm$ 0.00	33.80 $\pm$ 0.00	72.60 $\pm$ 5.74	23.30 $\pm$ 0.00	46.60 $\pm$ 7.78	56.20 $\pm$ 0.00	19.62 $\pm$ 4.79	24.93 $\pm$ 2.43	81.37 $\pm$ 2.11	34.00 $\pm$ 3.65
Affective Polarity												
Positive Valence	64.43 $\pm$ 0.00	62.07 $\pm$ 0.00	34.93 $\pm$ 0.00	39.89 $\pm$ 0.00	74.94 $\pm$ 6.51	31.42 $\pm$ 0.00	54.92 $\pm$ 5.44	66.68 $\pm$ 0.00	26.58 $\pm$ 3.20	30.92 $\pm$ 3.92	86.38 $\pm$ 1.61	41.35 $\pm$ 2.86
Neutral Valence	60.10 $\pm$ 0.00	55.95 $\pm$ 0.00	32.80 $\pm$ 0.00	34.30 $\pm$ 0.00	72.60 $\pm$ 4.90	26.90 $\pm$ 0.00	44.00 $\pm$ 7.18	59.40 $\pm$ 0.00	18.83 $\pm$ 4.30	27.23 $\pm$ 3.35	82.89 $\pm$ 3.38	34.73 $\pm$ 3.45
Negative Valence	62.80 $\pm$ 0.00	59.60 $\pm$ 0.00	34.60 $\pm$ 0.00	33.25 $\pm$ 0.00	75.80 $\pm$ 4.26	27.10 $\pm$ 0.00	49.40 $\pm$ 9.71	59.85 $\pm$ 0.00	21.16 $\pm$ 5.21	29.47 $\pm$ 3.44	83.25 $\pm$ 2.28	37.16 $\pm$ 4.22

Notes: Results report category-wise performance on the comparable subset with $N = 2,711$ memes. For each subcategory, we randomly sample up to 50 memes from the comparable subset and evaluate all models on the same sampled instances. This process is repeated 10 times, and results are reported as mean $\pm$ standard deviation over the 10 runs. All values are normalized to $[0, 100]$. Darker shading indicates better performance within each metric.

unreadable or missing meme text, safety refusals, or severe format violations. We treat only such technical or format failures as invalid; low-quality but scorable outputs are retained and penalized by the metrics. For fair comparison, we retain the intersection subset where all models produce valid outputs, yielding $\tilde{N} = 2,711$ comparable memes. We adopt a repeated randomized subset evaluation protocol, with scores computed at the output level and aggregated by model, taxonomy category, and overall performance. Unless otherwise specified, all main tables use this comparable subset to avoid sample-composition shifts across models.

6 Quantitative Results

6.1 Model-wise Performance Analysis

Table 2 reports metric-level results for 14 models detailed in Table S8, while Fig. 5(a)-(n) shows subcategory-wise RFS scores. Flux9B (Labs, 2025) achieves the best overall RFS of 76.78, followed by QwenEdit (Wu et al., 2025a) with 69.38 and Flux4B (Labs, 2025) with 61.06. These top models maintain relatively balanced performance across emotional alignment, structural preserva-

tion, and generation quality. Notably, although SD attains the highest VGQ, its near-zero text and layout scores result in a very low RFS, indicating that visual quality alone is insufficient for MER. The middle tier, including UniPic (Wei et al., 2026), DreamO2 (Xia et al., 2025), and ZTurbo (Cai et al., 2025a), exhibits strong but uneven capabilities. For example, ZTurbo achieves high text quality and emotional alignment, but its extremely low LC indicates severe structural disruption. Conversely, instruction-based editors such as Step1X (Liu et al., 2025) and IP2P (Brooks et al., 2023) preserve layout well but show limited emotional transformation. Traditional diffusion baselines and smaller MLLMs, including Bagel7B (Deng et al., 2025), GoT (Fang et al., 2025), OmniG2 (Wu et al., 2025b), and ICDit (Zhang et al., 2025), generally fail to satisfy the joint requirements of MER, especially in text generation and affective alignment.

Key Insights. Model-wise results show that MER is not bottlenecked by low-level visual generation, but by the ability to coordinate affective transformation with semantic and structural preservation. Current systems often optimize only one




Figure 5: Subcategory-wise and overall analysis on Multimodal Emotional Reappraisal. In each run, 700 memes are sampled as evenly as possible from the edited outputs of different models, and scores are averaged over 10 runs. The bar colors in (o) are consistent with the metric-specific color scheme used in Tables 2 and 3.

side of this trade-off: ZTurbo improves affective and textual cues at the cost of layout consistency, while Step1X and IP2P preserve templates without sufficiently changing affective framing. The strong performance of Flux9B and QwenEdit suggests that effective Emotional Reappraisal requires integrated multimodal reasoning rather than isolated image editing or text rewriting.

6.2 Category-wise Performance Analysis

Tab. 3 summarizes performance across dataset subcategories, and Fig. 5(a)-(n) shows the corresponding RFS distributions. The results reveal that MER difficulty varies substantially with visual modality, layout structure, and affective polarity. Across visual modalities, object/animal-centric memes achieve the highest RFS of 38.96, followed by photographic human (38.14) and illustrated/cartoon memes (37.86), while template-stylized memes perform worst (34.67). This drop is mainly reflected in TGQ and TEA, suggesting that fixed templates and stylized captions make text rewriting more error-prone and semantically fragile. For structural layout, single-panel memes outperform multi-panel memes (36.71 vs. 34.00), as multi-panel memes require preserving both spatial organization and cross-panel narrative coherence. Affective polarity also affects performance: positive-valence memes

achieve the highest RFS (41.35), followed by negative (37.16) and neutral memes (34.73), indicating that ambiguous or weakly expressed affect makes controlled reappraisal harder.

Key Insights. Category-wise results show that MER becomes difficult when meme meaning is tied to rigid templates, sequential layouts, or ambiguous affective cues. These findings suggest that the main challenge is not generating positive content, but preserving the structural and pragmatic conditions under which the original meme conveys meaning.

6.3 Overall Benchmark Analysis

Fig. 5(o) summarizes the overall metric distribution across the filtered evaluation subset. The global results reveal a clear gap between low-level generation quality and higher-level reappraisal capability. Visual quality and basic structural metrics, such as VGQ and LC, remain relatively strong across models, suggesting that current systems can often produce visually coherent outputs and preserve coarse layouts. In contrast, text- and affect-related metrics, including TGQ, TEA, and TAS, are substantially weaker, indicating persistent difficulty in rewriting embedded text and aligning the edited meme with the target affect. Notably, the high SSS but low TAS, CFS, and RFS suggest that models can often induce visible emotional changes, but struggle to



Figure 6: Qualitative comparison of MER results across models on MER-Bench.

make them target-aligned and content-faithful.

Key Insights. Overall results show that MER is not primarily limited by image synthesis quality, but by the ability to coordinate textual rewriting, affective transformation, and semantic preservation under structural constraints. This confirms that MER is a multimodal reasoning problem rather than a purely visual editing task.

7 Qualitative Analysis

Fig. 6 shows examples of common failure modes that are consistent with quantitative trends. Models often preserve the template and layout but fail to induce an affective shift toward the target emotion. This pattern appears in Bage17B (Deng et al., 2025), GoT (Fang et al., 2025), SD (Rombach et al., 2022), and IP2P (Brooks et al., 2023), where outputs remain close to the source meme yet weakly align with the target emotion. A contrasting failure occurs when models achieve the desired affective shift at the cost of structural preservation. For example, ZTurbo (Cai et al., 2025a) often produces emotionally plausible edits but disrupts the meme template, caption placement, or layout. Multi-panel memes are harder, as they require preserving spatial organization, panel order, and cross-panel narrative coherence. Many models distort inter-panel relations or produce inconsistent content, leading to semantic drift despite plausible visuals.

Key Insights. The qualitative results reveal that current models often fail at the interface between structure, text, and affect. Preserving layouts alone does not guarantee emotional transformation, while affective edits frequently introduce caption rendering errors, layout distortion, or semantic drift. This confirms that MER is not merely an image-editing problem but a coordinated multimodal reasoning challenge involving visual structure, embedded text, semantic intent, and emotional framing.

8 Conclusion

We study **Multimodal Emotional Reappraisal (MER)** as a structure-preserving affective transformation problem for image-text memes, and introduce **MER-Bench** for evaluation. MER-Bench provides annotations and a human-validated MLLM-as-a-Judge evaluation suite. Experiments on 14 multimodal generation systems show that models still struggle to jointly achieve emotional transformation, semantic preservation, and structural fidelity. These findings position MER as a challenging multimodal reasoning problem beyond conventional image editing or sentiment transfer, and we hope MER-Bench will support future research on controllable affect-aware generation.

Limitations

MER-Bench focuses on memes, leaving dynamic formats such as GIFs, short videos, and audio-visual memes unexplored. It also does not explicitly address multilingual or cross-cultural variation, despite the influence of language, culture, and community conventions on meme interpretation and emotional reappraisal. Finally, although expert-validated, our MLLM-as-a-Judge protocol may still inherit biases from the underlying judge model. Future work will extend MER toward multilingual, multicultural, and temporally grounded settings.

Ethical Considerations

MER-Bench is built from available meme data for research on controllable, affect-aware multimodal generation. It focuses on transforming negative affect into more positive or constructive interpretations rather than generating harmful, deceptive, or manipulative content. Unsafe or infeasible samples are filtered, and retained annotations are manually verified to reduce misuse. Although downstream risks cannot be eliminated, MER-Bench is designed to promote safer affective transformation.

Acknowledgment

This work was supported by the Yangtze River Delta Science and Technology Innovation Community Joint Research (Basic Research) Project (2025CSJZN01600), the New Generation of Artificial Intelligence National Science and Technology Major Project (2025ZD0123303).

References

- Zewen Bai, Liang Yang, Shengdi Yin, Junyu Lu, Jingjie Zeng, Haozhao Zhu, Yuanyuan Sun, and Hongfei Lin. 2025. State toxic: A benchmark for span-level target-aware toxicity extraction in chinese hate speech detection. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 10206–10219.
- Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18392–18402.
- Pamela Jo Brubaker, Scott Haden Church, Jared Hansen, Steven Pelham, and Alison Ostler. 2018. One does not simply meme about organizations: Exploring the content creation strategies of user-generated memes on imgur. *Public relations review*, 44(5):741–751.
- Huanqia Cai, Sihan Cao, Ruoyi Du, Peng Gao, Steven Hoi, Zhaohui Hou, Shijie Huang, Dengyang Jiang, Xin Jin, Liangchen Li, and 1 others. 2025a. Z-image: An efficient image generation foundation model with single-stream diffusion transformer. *arXiv preprint arXiv:2511.22699*.
- Yaqi Cai, Shancheng Fang, Yadong Qu, Xiaorui Wang, Meng Shao, and Hongtao Xie. 2025b. Itermeme: Expert-guided multimodal llm for interactive meme creation with layout-aware generation. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pages 720–728.
- Dongping Chen, Ruoxi Chen, Shilin Zhang, Yaochen Wang, Yinuo Liu, Huichi Zhou, Qihui Zhang, Yao Wan, Pan Zhou, and Lichao Sun. 2024a. MLLM-as-a-judge: Assessing multimodal LLM-as-a-judge with vision-language benchmark. In *Proceedings of the 41st International Conference on Machine Learning*, pages 6562–6595.
- Yuyan Chen, Songzhou Yan, Zhihong Zhu, Zhixu Li, and Yanghua Xiao. 2024b. Xmeap: Meme caption generation with sub-image adaptability. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 3352–3361.
- Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, and 1 others. 2025. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*.
- Anastasia Denisova. 2019. *Internet memes and society: Social, cultural, and political contexts*.
- Rongyao Fang, Chengqi Duan, Kun Wang, and 1 others. 2025. Got: Unleashing reasoning capability of multimodal large language model for visual generation and editing. *arXiv preprint arXiv:2503.10639*.
- James J Gross. 1998. The emerging field of emotion regulation: An integrative review. *Review of general psychology*, 2(3):271–299.
- James J Gross and 1 others. 2014. Emotion regulation: Conceptual and empirical foundations. *Handbook of emotion regulation*, 2(1):3–20.
- Jen-tse Huang, Man Ho Lam, and 1 others. 2023. Emotionally numb or empathetic? evaluating how llms feel using emotionbench. *arXiv preprint arXiv:2308.03656*.
- EunJeong Hwang and Vered Shwartz. 2023. Meme-cap: A dataset for captioning and interpreting memes. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1433–1445.
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj 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.
- Black Forest Labs. 2025. FLUX.2: Frontier Visual Intelligence.
- Jiankai Li, Yunhong Wang, and Weixin Li. 2024. Mhrn: A multimodal hierarchical reasoning network for topic detection. *IEEE Transactions on Multimedia*, 26:6968–6980.
- Kun Li, Jihao Gu, Fei Wang, Zhiliang Wu, Hehe Fan, and Dan Guo. 2026. Ma-bench: Towards fine-grained micro-action understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20118–20128.
- Kun Li, Pengyu Liu, Dan Guo, Fei Wang, Zhiliang Wu, Hehe Fan, and Meng Wang. 2025a. Mmad: Multi-label micro-action detection in videos. In *2025 IEEE/CVF International Conference on Computer Vision*, pages 13225–13236.
- Xueyan Li, Yingyi Xue, Mengjie Jiang, Qingzi Zhu, and Yazhe Niu. 2025b. From perception to punchline: Empowering vlm with the art of in-the-wild meme. *arXiv preprint arXiv:2512.24555*.
- Hongzhan Lin, Ziyang Luo, Bo Wang, Ruichao Yang, and Jing Ma. 2026. Goat-bench: Safety insights to large multimodal models through meme-based social abuse. *ACM Transactions on Intelligent Systems and Technology*, 17(4):1–25.

- Shiyu Liu, Yucheng Han, Peng Xing, and 1 others. 2025. Step1x-edit: A practical framework for general image editing. *arXiv preprint arXiv:2504.17761*.
- David Mellor and Kathleen A Moore. 2014. The use of likert scales with children. *Journal of pediatric psychology*, 39(3):369–379.
- Shreyash Mishra, S Suryavardan, and 1 others. 2023. Memotion 3: Dataset on sentiment and emotion analysis of codemixed hindi-english memes. *arXiv preprint arXiv:2303.09892*.
- Mette Mortensen and Christina Neumayer. 2021. The playful politics of memes.
- Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. 2024. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 4296–4304.
- Kun Ouyang, Liqiang Jing, Xueming Song, Meng Liu, Yupeng Hu, and Liqiang Nie. 2025. Sentiment-enhanced graph-based sarcasm explanation in dialogue. *IEEE Transactions on Multimedia*.
- Jeongsik Park, Khoi PN Nguyen, Jihyung Park, Minseok Kim, Jaeheon Lee, Jae Won Choi, Kalyani Ganta, Phalgun Ashrit Kasu, Rohan Sarakinti, Sanjana Vipplerla, and 1 others. 2025. Memeinterpret: Towards an all-in-one dataset for meme understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 16073–16087.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.
- James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology*, 39(6):1161.
- Sahand Sabour, Siyang Liu, Zheyuan Zhang, June Liu, Jinfeng Zhou, Alvionna Sunaryo, Tatia Lee, Rada Mihalcea, and Minlie Huang. 2024. Emobench: Evaluating the emotional intelligence of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, pages 5986–6004.
- Aadhavan Sadasivam, Kausic Gunasekar, Hasan Davulcu, and Yezhou Yang. 2020. Memebot: Towards automatic image meme generation. *arXiv preprint arXiv:2004.14571*.
- Akshit Sharma and Prashant W Patil. 2026. Memetag: Keyword-driven meme classification through tag embedding reconstruction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7679–7688.
- Shivam Sharma, S Ramaneswaran, Md Shad Akhtar, and Tanmoy Chakraborty. 2024. Emotion-aware multimodal fusion for meme emotion detection. *IEEE Transactions on Affective Computing*, 15(3):1800–1811.
- Shezheng Song, Chengxiang He, and 1 others. 2024. Mosabench: Multi-object sentiment analysis benchmark for evaluating multimodal large language models understanding of complex image. *arXiv preprint arXiv:2412.00060*.
- Xuanyu Su, Yansong Li, Diana Inkpen, and Nathalie Japkowicz. 2025. A context-aware contrastive learning framework for hateful meme detection and segmentation. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 5201–5215.
- Meituan LongCat Team, Hanghang Ma, Haoxian Tan, Jiale Huang, and 1 others. 2025. Longcat-image technical report. *arXiv preprint arXiv:2512.07584*.
- Suryatej Reddy Vyalla and Vishaal Udandaraao. 2020. Memeify: A large-scale meme generation system. In *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD*, pages 307–311.
- Bram Wallace, Meihua Dang, Rafael Rafailov, and 1 others. 2024. Diffusion model alignment using direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8228–8238.
- Fei Wang, Dan Guo, Kun Li, and Meng Wang. 2024a. Eulermormer: Robust eulerian motion magnification via dynamic filtering within transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 5345–5353.
- Fei Wang, Dan Guo, Kun Li, Zhun Zhong, and Meng Wang. 2024b. Frequency decoupling for motion magnification via multi-level isomorphic architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18984–18994.
- Fei Wang, Jiangnan Yang, Junjie Chen, Yuxin Liu, Kun Li, Yanyan Wei, Dan Guo, and Meng Wang. 2026a. Xinsight: Integrative stage-consistent psychological counseling support agents for digital well-being. *arXiv preprint arXiv:2603.06583*.
- Fei Wang, Xinye Zheng, Kun Li, Yanyan Wei, Yuxin Liu, Ganpeng Hu, Tong Bao, and Jingwen Yang. 2026b. Multimodal protein language models for enzyme kinetic parameters: From substrate recognition to conformational adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15829–15839.
- Han Wang and Roy Ka-Wei Lee. 2024. Memecraft: Contextual and stance-driven multimodal meme generation. In *Proceedings of the ACM Web Conference 2024*, pages 4642–4652.

- Hongyang Wei, Hongbo Liu, Zidong Wang, Yi Peng, Baixin Xu, and 1 others. 2026. Skywork unipic 3.0: Unified multi-image composition via sequence modeling. *arXiv preprint arXiv:2601.15664*.
- Bradley E Wiggins. 2019. *The discursive power of memes in digital culture: Ideology, semiotics, and intertextuality*.
- Chenfei Wu, Jiahao Li, Jingren Zhou, and 1 others. 2025a. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*.
- Chenyuan Wu, Pengfei Zheng, Ruiran Yan, and 1 others. 2025b. Omnigen2: Exploration to advanced multimodal generation. *arXiv preprint arXiv:2506.18871*.
- Bin Xia, Bohao Peng, Yuechen Zhang, and 1 others. 2025. Dreamomni2: Multimodal instruction-based editing and generation. *arXiv preprint arXiv:2510.06679*.
- 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.
- Yangyang Xu, Zhuoer Zhao, Xiao Sun, and Xun Yang. 2025. Prompt learning with multiperspective cues for emotional support conversation systems. *IEEE Transactions on Computational Social Systems*, 12(5):2648–2660.
- Jingyuan Yang, Jiawei Feng, and Hui Huang. 2024. Emogen: Emotional image content generation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6358–6368.
- Jingyuan Yang, Jiawei Feng, Weibin Luo, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. 2025a. Emoedit: Evoking emotions through image manipulation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24690–24699.
- Jingyuan Yang, Weibin Luo, and Hui Huang. 2025b. Emoctrl: Controllable emotional image content generation. *arXiv preprint arXiv:2512.22437*.
- Xiaocui Yang, Wenfang Wu, Shi Feng, Ming Wang, Daling Wang, Yang Li, Qi Sun, Yifei Zhang, Xiaoming Fu, and Soujanya Poria. 2025c. Mm-instructeval: Zero-shot evaluation of (multimodal) large language models on multimodal reasoning tasks. *Information Fusion*, 122:103204.
- Hu Ye, Jun Zhang, Sibao Liu, Xiao Han, and Wei Yang. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Shaozu Yuan, Yiwei Wei, Hengyang Zhou, Qinfu Xu, Meng Chen, and Xiaodong He. 2025. Enhancing semantic awareness by sentimental constraint with automatic outlier masking for multimodal sarcasm detection. *IEEE Transactions on Multimedia*.
- Fan Zhang, Zebang Cheng, Chong Deng, Haoxuan Li, Zheng Lian, Qian Chen, Huadai Liu, Wen Wang, YiFan Zhang, Renrui Zhang, and 1 others. 2026a. Mme-emotion: A holistic evaluation benchmark for emotional intelligence in multimodal large language models. In *International Conference on Learning Representations*, volume 2026, pages 48767–48807.
- Jiajun Zhang, Shijia Luo, Ruikang Zhang, and Qi Su. 2026b. Humorchain: Theory-guided multi-stage reasoning for interpretable multimodal humor generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19176–19185.
- Yuanmin Zhang, Kexin Xu, Chunzhi Xie, and Zhisheng Gao. 2024. Emotion inference in conversations based on commonsense enhancement and graph structures. *PloS One*, 19(12):e0315039.
- Zechuan Zhang, Ji Xie, Yu Lu, Zongxin Yang, and Yi Yang. 2025. Enabling instructional image editing with in-context generation in large scale diffusion transformer. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Zhengyi Zhao, Shubo Zhang, Yuxi Zhang, Yanxi Zhao, Yifan Zhang, Zezhong Wang, Huimin Wang, Yutian Zhao, Bin Liang, Yefeng Zheng, and 1 others. 2025. Memereacon: Probing contextual meme understanding in large vision-language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 3559–3582.

Supplementary Material

Contents

A Dataset Construction Details	13
A.1 Filtering Statistics	13
A.2 Annotation Schema	13
A.3 Prompt Templates	13
A.4 Human Control over LLM Assistance	14
A.5 Inter-annotator Agreement	14
A.6 Dataset Statistics	14
B Model Details	15
C Judge Prompt and JSON Schema	15
C.1 Judge Input	15
C.2 System Prompt	15
C.3 User Prompt	16
C.4 Evaluation Rubric	16
C.5 JSON Output Schema	17
C.6 Metric Mapping	18
C.7 Prompt Configuration Ablation . .	18
D Additional Qualitative Results	18
E Discussion	18
E.1 Application Scenario	18

A Dataset Construction Details

This appendix provides additional details on MER-Bench construction, including filtering statistics, annotation schema, LLM-assisted human control, inter-annotator agreement, and emotion-shift distribution.

A.1 Filtering Statistics

MER-Bench is constructed from the *MOOD* meme corpus (Sharma et al., 2024). As shown in Table S4, we start from 10,005 original memes and retain 6,695 samples with clearly identifiable negative affective framing. After quality filtering, deduplication, safety filtering, and reappraisal-feasibility filtering, the final benchmark contains 3,117 image-text memes.

A.2 Annotation Schema

Each retained sample contains structured annotations for source understanding, target reappraisal, and multimodal editing. Table S5 summarizes the annotation schema.

A.3 Prompt Templates

We use three prompt templates for LLM-assisted human rewriting. The templates are normalized here using the annotation field names in Table S5. The first two prompts are used with Gemini-2.5-Flash, and the third prompt is used with Gemini-2.5-Flash-Image. All outputs are manually checked before inclusion in MER-Bench.

Prompt I: Emotion Detection and Image Description.

For the given meme, consisting of an image and overlaid text, complete the following tasks.

Emotion labeling. Choose one source emotion label from: sad, angry, tense, or bored.

Use the following emotion definitions: sad = loss or disappointment; angry = frustration or unfairness; tense = pressure or worry; bored = lack of interest.

When multiple people appear, focus on the main subject whose expression, posture, or role best conveys the meme’s emotion. If the visual expression and the overlaid text conflict, follow the emotion implied by the text. Always output only one source emotion.

Image description. Briefly describe the meme using the following information: visual cues, including facial expression, gesture, color, scene, and layout; textual cues, including tone and phrasing; and how these cues support the selected source emotion.

Output format. Return a JSON object with the following fields: src_emo and img_desc.

Prompt II: Reappraisal Planning.

Rewrite the meme toward its target reappraisal emotion while preserving the original topic, entities, and core situation.

Inputs. src_emo: one of sad, angry, tense, or bored; img_desc: a concise description of the meme’s visual and textual cues; txt: the original overlaid meme text.

Emotion mapping. sad → happy; angry → calm; tense → relaxed; bored → excited.

Tasks. First, set tgt_emo to the mapped target emotion. Second, write rewritten_caption, a reappraised meme caption that expresses the target emotion while keeping the same topic and context. Do not change the core situation, entities, or event relations. Avoid adding hostility, anxiety, fear, sexual content, violent content, or unrelated negativity. Third, write edit_spec, a concise visual editing specification describing changes in facial expression, posture, gesture, atmosphere, and visual tone.

Tone guide. happy: warm, optimistic, joyful; calm: peaceful, balanced, stable; relaxed: easy-going, comfortable, natural; excited: energetic, playful, enthusiastic.

Output format. Return a JSON object with the following fields: tgt_emo, rewritten_caption, and edit_spec.

Field	Description	Purpose
id	Original meme identifier inherited from MOOD	Trace the source sample
img_desc	Brief image description covering scene, entities, expressions, and poses	Support visual understanding and editing
txt	Original overlaid meme text	Preserve the source textual content
src_emo	Source emotion label	Identify the original affective framing
tgt_emo	Target emotion label	Define the intended reappraisal direction
rewritten_caption	Reappraised meme text	Provide the target textual framing
edit_spec	Visual editing specification	Guide structure-preserving multimodal editing
gen_id	Generated meme identifier	Track the generated output
notes	Optional annotation or quality-control notes	Record ambiguity, revision, or failure cases

Table S5: Annotation schema of MER-Bench.

Outcome	# Samples	Ratio
Accepted without modification	1,126	16.8%
Accepted after minor revision	1,421	21.2%
Accepted after substantial rewriting	570	8.5%
Removed after human verification	3,578	53.4%
Total checked	6,695	100.0%

Table S6: Human control statistics for LLM-assisted annotation.

Annotation Task	Fleiss' κ
Source emotion labeling	0.78
Target emotion verification	0.81
Visual modality labeling	0.74
Layout structure labeling	0.76
Affective polarity labeling	0.72
Quality-control decision	0.69
Average	0.75

Table S7: Inter-annotator agreement for MER-Bench annotations.

distribution, and word clouds for source and reappraised captions. The benchmark covers diverse meme formats, with single-panel memes dominating the structural distribution and template-stylized memes forming the largest visual modality category. The source emotion distribution corresponds to the four reappraisal directions used in MER-Bench: *Sad* $\rightarrow$ *Happy*, *Angry* $\rightarrow$ *Calm*, *Tense* $\rightarrow$ *Relaxed*, and *Bored* $\rightarrow$ *Excited*. The *Bored* $\rightarrow$ *Excited* category is smaller because bored memes with clear image-text grounding and feasible reappraisal targets are less frequent after filtering.

B Model Details

Table S8 lists the 14 models evaluated in MER-Bench, including their full model names, short names, model categories, and editing paradigms.

C Judge Prompt and JSON Schema

This appendix provides the prompt design and structured output schema used by our MLLM-as-a-Judge evaluation protocol. For each evaluation

instance, the judge receives the source meme, the edited meme, the source emotion, the target emotion, and optional reference information, including the source caption, target caption, and editing instruction. The judge is required to return a schema-constrained JSON object without any free-form text outside the JSON output.

C.1 Judge Input

Each evaluation sample is represented as:

$$x = \{I_{\text{src}}, I_{\text{edit}}, a_{\text{src}}, a_{\text{tgt}}, T_{\text{src}}, T_{\text{tgt}}, e\}, \quad (7)$$

where I_{src} and I_{edit} denote the source and edited meme images, a_{src} and a_{tgt} denote the source and target emotions, T_{src} and T_{tgt} denote optional source and target captions, and e denotes the optional editing instruction. The judge compares the edited meme against the source meme under the target emotion constraint.

C.2 System Prompt

The system prompt defines the role of the judge and enforces four global constraints:

System Prompt

You are an expert evaluator for the task: **Meme Reappraisal**.

You will be given a **SOURCE** meme image, an **EDITED** meme image, reference labels for **SOURCE** emotion and **TARGET** emotion, and optional reference captions and editing instructions.

Your job is to compare **SOURCE** vs. **EDITED** and output only a JSON object that strictly follows the provided JSON schema.

Global constraints:

1. **Content preservation.**

The **EDITED** meme must preserve the same situation, entities, identity, and core meaning as the **SOURCE** meme.

Penalize outputs where the topic changes, new events are introduced, or the implied situation differs.

2. **Layout and structure preservation.**

If the **SOURCE** meme is single-panel, the **EDITED** meme should remain single-panel.

If the **SOURCE** meme is multi-panel or grid-based, the **EDITED** meme should preserve the panel structure, panel order, and caption regions.

Any mismatch should be marked as **layout inconsistency**.

3. **Target-gated emotion scoring.**




Figure S8: Full qualitative comparison across evaluated models. This figure extends the qualitative examples in Fig. 6 by showing outputs from all evaluated models for representative source memes. The examples illustrate typical MER behaviors, including successful reappraisal, insufficient affective shift, layout disruption, text-rendering errors, and semantic drift.

Emotion–alignment scores must first check whether the ****EDITED meme**** matches the given ****TARGET emotion****. If the ****EDITED meme**** does not match the ****TARGET emotion****, the emotion–alignment score must not exceed ****2**** on the ****1—5**** scale.

4. ****Localized rationale.****
Each score must have its own rationale.

The rationale for each block should only justify that specific score or decision and should not be copied across blocks. Each rationale should mention target–emotion alignment, scenario preservation, and text–visual semantic alignment when applicable.

Return ****JSON only**** and do not add extra keys outside the schema.

SOURCE emotion: `{src\_emotion}`

TARGET emotion: `{tgt\_emotion}`

Optional references:

– SOURCE caption: `{src\_caption}`
– TARGET caption: `{tgt\_caption}`
– Editing instruction: `{edit\_instruction}`

Evaluate the ****EDITED meme**** according to the following criteria: visual assessment, text assessment, layout consistency, overall emotion classification, perceived emotion shift, and overall generation quality.

Each score should follow its own rubric and rationale.

C.3 User Prompt

The user prompt provides the task instance and the criterion-specific rubrics. The following template is used:

User Prompt

Task: Evaluate the ****EDITED meme**** for ****Meme Reappraisal**** given the ****SOURCE meme**** and ****TARGET emotion****.

Allowed emotion labels:
`angry`, `sad`, `tense`, `bored`, `calm`, `happy`, `excited`, `relaxed`, and `other`.

The source and edited meme images are provided as image inputs. In the Visual-First setting, the images are placed before the textual instruction. In the Visual-Last setting, the textual instruction is placed before the images.

C.4 Evaluation Rubric

The judge evaluates six structured blocks, from which the eight fine-grained metrics are derived.

Visual Assessment. The judge scores:

Table S8: Details of models evaluated on MER-Bench.

Model Name (Short Name)	Size	Category	Editing Paradigm
black-forest-labs/FLUX.2-klein-9B (Flux9B) (Labs, 2025)	9B	Diffusion editing model	Instruction-based image editing
Qwen/Qwen-Image-Edit-2509 (QwenEdit) (Wu et al., 2025a)	–	Vision-language editing model	Multimodal instruction editing
black-forest-labs/FLUX.2-klein-4B (Flux4B) (Labs, 2025)	4B	Diffusion editing model	Instruction-based image editing
Skywork/Unipic3 (UniPic) (Wei et al., 2026)	1.5B	Vision-language editing model	Multimodal instruction editing
xiabs/DreamOmni2 (DreamO2) (Xia et al., 2025)	–	Vision-language editing model	Multimodal instruction editing
Tongyi-MAI/Z-Image-Turbo (ZTurbo) (Cai et al., 2025a)	6B	Vision-language editing model	Multimodal image editing
meituan-longcat/LongCat-Image (LongCat) (Team et al., 2025)	6B	Vision-language editing model	Multimodal image editing
stepfun-ai/Step1X-Edit-v1p2 (Step1X) (Liu et al., 2025)	–	Diffusion editing model	Instruction-based image editing
ByteDance-Seed/BAGEL-7B-MoT (Bagel7B) (Deng et al., 2025)	7B	Vision-language editing model	Multimodal instruction editing
GoT (GoT) (Fang et al., 2025)	–	Vision-language editing model	Multimodal instruction editing
timbrooks/instruct-pix2pix (IP2P) (Brooks et al., 2023)	–	Diffusion editing model	Text-guided image-to-image
OmniGen2/OmniGen2 (OmniG2) (Wu et al., 2025b)	–	Diffusion editing model	Text-guided image-to-image
In-Context Edit (ICEdit) (Zhang et al., 2025)	–	In-context editing baseline	Prompt-only editing
stable_diffusion (SD) (Rombach et al., 2022)	–	Diffusion generation model	Text-guided image-to-image

Notes: The table lists the full model names used in our benchmark and the short names used throughout the paper. The size column reports the approximate number of parameters when publicly available; otherwise, we mark it as “–”. Model categories are grouped by their primary editing or generation mechanism.

- `generation_quality_1_5`: visual cleanliness, artifact level, realism or style coherence, and readability of visual details.
- `emotion_accuracy_1_5`: whether visual affective cues match the target emotion, including facial expression, posture, lighting, color tone, and tension or relaxation.

If the visual affect does not match the target emotion, the visual emotion score is capped at 2.

Text Assessment. The judge scores:

- `generation_quality_1_5`: text legibility, absence of garbling or cropping, formatting consistency, and natural meme-style placement.
- `emotion_accuracy_1_5`: whether the rewritten text expresses the target emotion while preserving the same scenario.

If the text does not convey the target emotion, the text emotion score is capped at 2.

Layout Consistency. The judge returns a Boolean value:

- `consistent`: true only if the edited meme preserves the source layout type, panel structure, panel order, and caption regions.

Overall Emotion Classification. The judge predicts the primary perceived emotion of the edited meme:

- `label`: one of the eight emotions or other.
- `other_emotion`: a specific emotion name if and only if `label` is other; otherwise null.

Perceived Emotion Shift. The judge scores:

- `magnitude_1_5`: the perceived strength of emotional change from the source meme to the edited meme.

Overall Generation Quality. The judge scores:

- `overall_generation_quality_1_5`: holistic quality considering target-emotion success, content preservation, layout preservation, and readability.

C.5 JSON Output Schema

The judge output follows the strict JSON schema below. The order of the rationale field depends on the prompt configuration: in Rationale-First settings, each block places rationale before scalar fields; in Rationale-Last settings, each block places rationale after scalar fields.

```
{
  "visual_assessment": {
    "rationale": string,
    "generation_quality_1_5": integer in [1, 5],
    "emotion_accuracy_1_5": integer in [1, 5]
  },
  "text_assessment": {
    "rationale": string,
    "generation_quality_1_5": integer in [1, 5],
    "emotion_accuracy_1_5": integer in [1, 5]
  },
  "layout_consistency": {
    "rationale": string,
    "consistent": boolean
  },
  "overall_emotion_classification": {
    "rationale": string,
    "label": one of [
      "angry", "sad", "tense", "bored",
      "calm", "happy", "excited", "relaxed", "other"
    ],
    "other_emotion": string or null
  },
  "perceived_emotion_shift": {
```

Metric	JSON Field
VGQ	visual_assessment.generation_quality_1_5
VEA	visual_assessment.emotion_accuracy_1_5
TGQ	text_assessment.generation_quality_1_5
TEA	text_assessment.emotion_accuracy_1_5
LC	layout_consistency.consistent
HGQ	overall_generation_quality. overall_generation_quality_1_5
PPE	overall_emotion_classification. label
PES	perceived_emotion_shift.magnitude_1_5

Table S9: Mapping between judge JSON fields and evaluation metrics.

```

"rationale": string,
"magnitude_1_5": integer in [1, 5]
},
"overall_generation_quality": {
  "rationale": string,
  "overall_generation_quality_1_5": integer in [1, 5]
}
}

```

We additionally apply post-validation to ensure that other_emotion is non-empty when label is other, and null otherwise.

C.6 Metric Mapping

The JSON fields are mapped to the eight metrics used in the main paper as follows:

Scalar 1–5 scores are normalized to $[0, 1]$ for aggregated scoring. For LC, true is mapped to 1 and false is mapped to 0. For PPE, we compute the target emotion hit indicator by checking whether the predicted label matches the target emotion.

C.7 Prompt Configuration Ablation

We evaluate four prompt configurations by varying two orthogonal factors:

- **Evidence order:** Visual-First (VF) places image evidence before textual instructions, while Visual-Last (VL) places image evidence after textual instructions.
- **Reasoning order:** Rationale-First (RF) places the rationale before scalar scores in each JSON block, while Rationale-Last (RL) places the rationale after scalar scores.

Their Cartesian product yields four configurations: RF-VF, RF-VL, RL-VF, and RL-VL. The main experiments use the configuration that best aligns with human expert judgments.

D Additional Qualitative Results

Fig. S8 provides the full model-wise qualitative comparison for the examples in the main text. Due to space constraints, the main paper shows only a subset of representative outputs, while the appendix includes results for all evaluated models.

E Discussion

Our results suggest that Emotional Reappraisal is a distinct capability beyond meme understanding, meme generation, and generic emotion-controllable editing. Compared with meme-understanding and safety-oriented benchmarks such as Hateful Memes and GOAT-Bench (Kielbaso et al., 2020; Lin et al., 2026), MER reveals a gap between recognizing affective meaning and transforming it under structural constraints. Identifying the emotion of a meme is substantially easier than reinterpreting it while preserving its visual layout, embedded text, and semantic context.

Relative to meme generation methods such as MemeCap, XMeCap, and IterMeme (Hwang and Schwartz, 2023; Chen et al., 2024b; Cai et al., 2025b), MER introduces an invariance requirement: the edited meme should preserve the original identity, layout, and core situation while changing its affective interpretation. This shifts the problem from open-ended generation to constrained transformation. Empirically, models that are strong in image synthesis, layout preservation, or text rewriting alone do not reliably achieve successful reappraisal when these objectives must be satisfied jointly.

MER also extends prior work on emotion-controllable generation (Zhang et al., 2024; Yang et al., 2025b, 2024; Brooks et al., 2023). While existing methods demonstrate affect control in generic settings, our results show that emotion manipulation becomes substantially harder when affect is encoded through tightly coupled visual content, embedded text, layout structure, and meme conventions. This multimodal interdependence creates a structural-affective trade-off: current models often either preserve meme structure without changing emotional framing, or modify affective cues while disrupting layout, text rendering, or semantic identity.

E.1 Application Scenario

Fig. S9 illustrates a potential application scenario of MER. A user-posted negatively framed meme can be reinterpreted into a more constructive ver-




Figure S9: Example application scenario of Meme Reappraisal. A user-shared negative meme (left) is reinterpreted into a positive version (right) while preserving the original situation and meme structure. This illustrates its potential for constructive reframing in emotionally aware online communities.

sion while keeping the same situation and meme structure. This example highlights the practical motivation for studying structure-preserving affect transformation, but also underscores the need for careful evaluation: effective reappraisal should not simply make a meme “positive”, but should preserve what the meme is about while changing how its situation is emotionally framed.

Overall, these findings indicate that structure-preserving affect transformation constitutes a distinct evaluation axis for multimodal systems. MER-Bench, therefore, complements existing benchmarks by testing whether models can coordinate visual editing, text rewriting, layout preservation, and affective reinterpretation within a single constrained multimodal task.