

BLM-Guard: Explainable Multimodal Ad Moderation with Chain-of-Thought and Policy-Aligned Rewards

Yiran Yang^{1,2*†}, Zhaowei Liu^{1*}, Yuan Yuan¹, Yukun Song^{1,2}, Xiong Ma¹, Yinghao Song¹, Xiangji Zeng¹, Lu Sun¹, Yulu Wang¹, Hai Zhou¹, Shuai Cui^{1,3}, Zhaohan Gong¹, Jiefei Zhang^{1‡}

¹ Kuaishou Technology

² Beijing University of Posts and Telecommunications

³ Shandong University

yir2023110885@bupt.edu.cn, songyukun@bupt.edu.cn, 202315156@mail.sdu.edu.cn

{liuzhaowei03, yuanyuan06, maxiong, songyinghao03, zengxiangji, sunlu05, wangyulu05, zhouhai, gongzhaohan, zhangjiefei}@kuaishou.com

Abstract

Short-video platforms now host vast multimodal ads whose deceptive visuals, speech and subtitles demand finer-grained, policy-driven moderation than community safety filters. We present BLM-Guard, a content-audit framework for commercial ads that fuses Chain-of-Thought reasoning with rule-based policy principles and a critic-guided reward. A rule-driven ICoT data-synthesis pipeline jump-starts training by generating structured scene descriptions, reasoning chains and labels, cutting annotation costs. Reinforcement learning then refines the model using a composite reward balancing causal coherence with policy adherence. A multitask architecture models intra-modal manipulations (e.g., exaggerated imagery) and cross-modal mismatches (e.g., subtitle-speech drift), boosting robustness. Experiments on real short-video ads show BLM-Guard surpasses strong baselines in accuracy, consistency and generalization.

Introduction

With the rise of generative AI and the rapid growth of short-video platforms like *TikTok* and *Instagram Reels*, creating multimodal advertising content has become easier and cheaper than ever (Qi et al. 2022; Yeh et al. 2024). Brands and individuals can now swiftly produce videos that combine visuals, speech, and text with minimal cost. While this boosts monetization efficiency and creative diversity, it also poses new challenges for content moderation.

Unlike general community moderation—typically focused on coarse risks like violence or nudity (Helff et al. 2025; Inan et al. 2023)—ad moderation demands *fine-grained, policy-driven* compliance checks. Violations are often subtle or disguised, including exaggerated claims, misleading cues, or rule evasion (Goldstein et al. 2023), and frequently manifest as *modality misalignments*—e.g., visually truthful but verbally deceptive content, benign captions

with provocative imagery, or mismatched subtitles and audio (Guo et al. 2024).

However, existing moderation approaches—such as static rule-based filters or general-purpose vision-language models (VLMs)—struggle in high-risk, multimodal ad settings due to three key limitations: (i) limited cross-modal causal reasoning (Ahn et al. 2024); (ii) poor adaptability to policy drift (Lu et al. 2025); and (iii) lack of task-specific reasoning for nuanced commercial risks (Liu et al. 2025b). Most prior work targets generic harm detection (Qu et al. 2024), offering limited support for regulation-sensitive ad content.

To tackle these challenges, we propose **BLM-Guard**, a moderation framework for policy-sensitive, high-risk short-video ads. It introduces three key innovations: First, an *Interleaved-modal Chain-of-Thought (ICoT)* pipeline (Gao et al. 2025) interleaves visual grounding, region localization, and causal reasoning (Wang et al. 2023; Xia et al. 2025), enabling structured, cross-modal trace generation and explainable decisions. Second, self-adaptive GRPO reinforcement learning incorporates a hybrid reward that fuses rule correctness with *principle-based self-consistency* (Chen et al. 2025), enhancing adaptability to policy shifts. Third, a multi-task architecture jointly models intra-modal manipulations and cross-modal mismatches, boosting robustness and generalization across diverse ad violations.

To support systematic evaluation, we release the **BLM-Guard Benchmark**, a real-world dataset of short-video ads annotated with a three-level risk taxonomy aligned with platform policies: (i) *risk scenario* (e.g., illegal content, misleading marketing); (ii) *violation type* (e.g., income exaggeration, privacy exposure); (iii) *risk severity* (high, medium, low). This hierarchical design supports interpretable supervision and rigorous policy-aligned evaluation. A visual summary is shown in Figure 1.

Our main contributions are:

- We introduce **BLM-Guard Benchmark**, a real-world dataset for ad moderation, structured across seven risk scenarios and fine-grained violation types, enabling policy-grounded evaluation.

*These authors contributed equally.

†Work done during an internship at Kuaishou Technology.

‡Project Leader

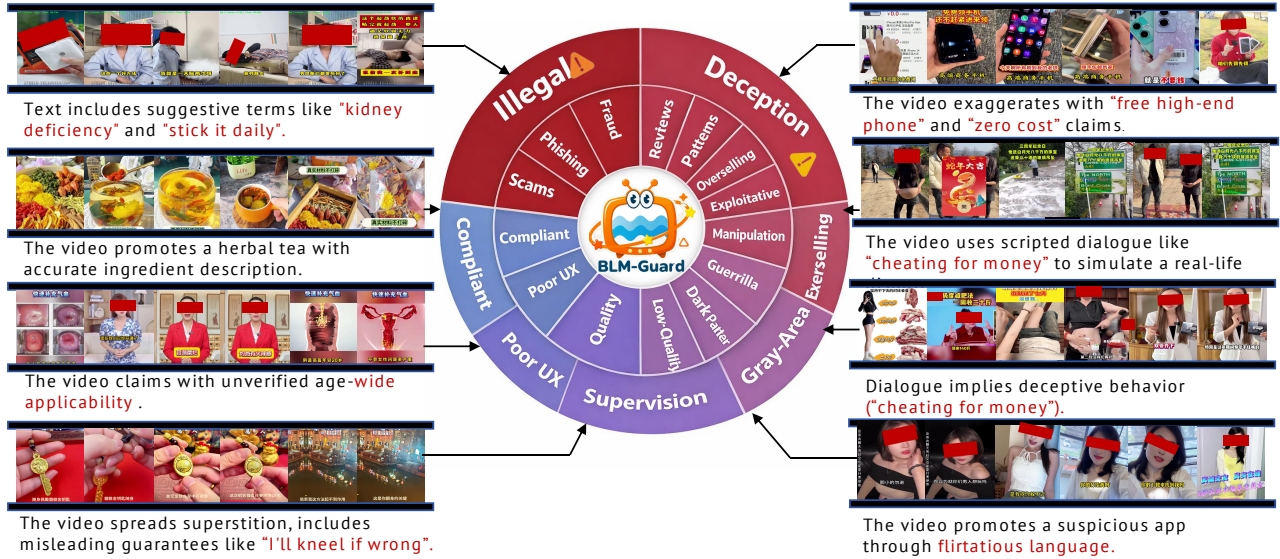


Figure 1: **BLM-Guard Benchmark Taxonomy.** Our benchmark organizes commercial short-video ads into a hierarchical risk taxonomy with seven core violation scenarios and fine-grained subtypes. Each node reflects a policy-sensitive violation type (e.g., income exaggeration, privacy leak, feudal superstition), and is further associated with a severity level (high, medium, low). This structure supports interpretable supervision and enables fine-grained evaluation of model performance on diverse and nuanced moderation cases. A dedicated “No Risk” category is also included to balance risk distribution.

- We propose **BLM-Guard**, a multimodal moderation framework combining rule-driven ICoT reasoning, consistency-aware reinforcement learning, and multi-task modeling to ensure policy-compliant and explainable decisions.

Task Formulation and BLM-Guard Benchmark

Task Definition

We formulate the task as policy-aligned moderation of commercial short-video ads. Given a video $\mathbf{v} = \langle f_1, f_2, \dots, f_T \rangle$ composed of multimodal frames (visuals and ASR transcripts), and a policy set $\mathcal{P} = \{\pi_1, \dots, \pi_n\}$ from platform guidelines, the goal is to determine for each policy π_i whether it is violated, and to generate an interpretable reasoning trace.

Formally, the moderation model G_{reasoner} outputs a decision $\hat{Y}_i \in \{0, 1\}$ and a reasoning chain R_i :

$$\{(\hat{Y}_i, R_i)\}_{i=1}^n = G_{\text{reasoner}}(\mathbf{v}, \mathcal{P}), \quad (1)$$

where $\hat{Y}_i = 1$ denotes a violation. Each R_i follows a structured path: *Scenario Identification* \rightarrow *Violation Typing* \rightarrow *Rule Matching* \rightarrow *Decision*.

This structure supports explainable moderation and facilitates reward modeling for reinforcement learning.

BLM-Guard Benchmark

To support structured moderation training, we introduce the **BLM-Guard Benchmark**, a short-video dataset with three-level policy labels (Figure 1).

- **Level 1: Severity** — High, Medium, or Low, reflecting potential legal, reputational, or user experience risks.
- **Level 2: Scenario** — e.g., illegal content, false marketing, misleading operations, etc.
- **Level 3: Violation Type** — e.g., income exaggeration, privacy leak, feudal superstition, and others.

This hierarchical structure enables fine-grained, policy-aligned supervision and supports nuanced detection of multimodal violations.

Data Construction

We collect 5–45s commercial ads from a short-video platform, covering domains like e-commerce, health, education, and lifestyle. The dataset includes:

1. **Expert-Audited Ads:** Videos labeled by professional moderators with verified risk types.
2. **Rule-Triggered Ads:** Samples flagged by rule-based moderation systems.
3. **High-Impression Compliant Ads:** Popular videos with low complaint rates to improve generalization.

Each video is annotated with its risk severity, scenario, and violation type. A subset contains structured reasoning traces to support both supervised learning and interpretable reward-based training.

Methodology

As shown in Fig. 2, we address the compliance and controllability challenges of short-video ad moderation via a two-stage pipeline. To overcome the limitations of single-stage

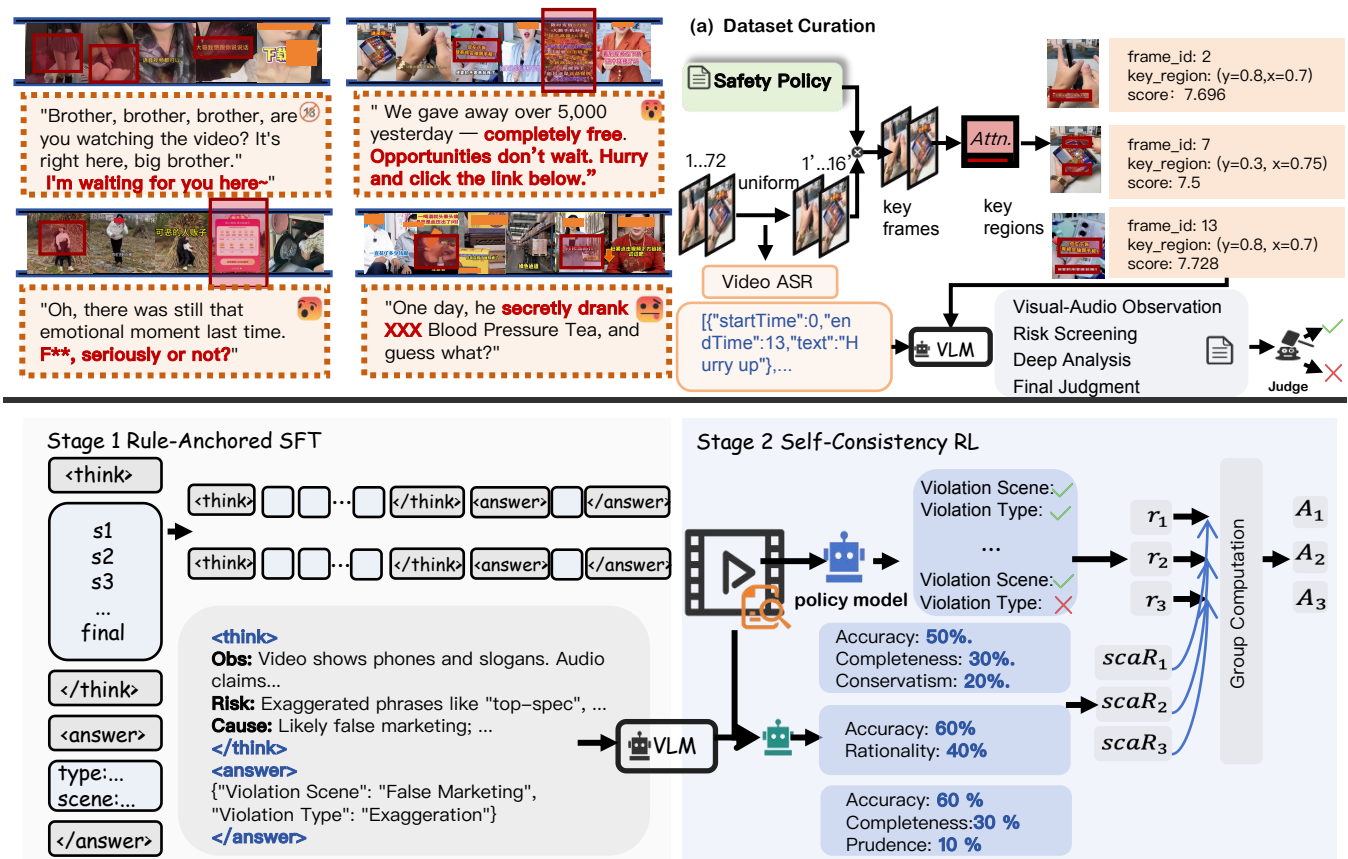


Figure 2: Our method adopts a progressive two-stage pipeline for policy-controllable content moderation. In Stage 1 (Rule-driven SFT Cold Start), we synthesize structured visual-language Chain-of-Thought (ICoT) data via keyframe selection and multi-step prompting using InternVL. This enables supervised fine-tuning with rule-anchored causal supervision. In Stage 2 (Self-adaptive GRPO Reinforcement Learning), we apply a safety-aware data curation strategy and propose a Self-Adaptive Critique Reward (SACR) to dynamically evaluate reasoning outputs. The model is optimized using a modified Group-wise Relative Policy Optimization (GRPO) algorithm with token-level normalization and dynamic sampling. This multi-stage process enables the model to first learn fine-grained compliance reasoning and then refine its moderation behavior through adaptive, reward-driven training.

tuning in policy coverage and interpretability, we decompose the training into:

- **Rule-guided SFT:** Injecting policy priors and reasoning skills into the VLM through Interleaved-modal Chain-of-Thought (ICoT) supervision.
- **Self-adaptive GRPO:** Optimizing behavior via Group-wise Relative Policy Optimization guided by dynamic rule-aware rewards.

This design enables the model to first acquire structured compliance reasoning and then adaptively refine moderation accuracy against evolving risk patterns.

Cold Start Strategy: Rule-Guided Causal Supervision

To provide an effective initialization for subsequent reinforcement learning, we propose a two-stage strategy that constructs structured visual reasoning data based on safety rules:

- **Stage-1:** Keyframe and Region Selection via Risk-Prompt Anchored Coarse Filtering
- **Stage-2:** Interleaved Multi-stage CoT Generation using Vision-Language Models

Keyframe and Region Extraction. We employ a two-step selection strategy to extract informative frames and salient regions that represent high-risk cues:

1. **CLIP-based Prompt Similarity.** For a video with $N = 16$ uniformly sampled frames $\{f_i\}_{i=1}^{16}$, we compute their semantic similarity to a set of $K = 7$ predefined risk prompts (e.g., “false marketing”, “illegal content”) using CLIP-ViT-L/14 (Radford et al. 2021). Denote the visual embedding of frame i as $\mathbf{v}_i \in \mathbb{R}^d$ and prompt embedding k as $\mathbf{t}_k \in \mathbb{R}^d$. The maximum similarity score of frame i is:

$$s_i = \max_k \left(\frac{\mathbf{v}_i^\top \mathbf{t}_k}{\|\mathbf{v}_i\| \cdot \|\mathbf{t}_k\|} \right) \quad (2)$$

This results in a similarity vector $\mathbf{s} = [s_1, s_2, \dots, s_{16}] \in \mathbb{R}^{16}$.

- Adaptive Keyframe Selection (AKS).** To ensure both temporal coverage and semantic diversity, we apply a hybrid *BIN+TOP* strategy (Tang et al. 2025) to select $m = 3$ keyframes. Specifically, the video is evenly divided into m temporal bins, and the most relevant frame is selected within each bin (BIN). If the number of selected frames is less than m , we supplement it with top-scoring frames across the entire video (TOP), ensuring complementary selection of globally salient frames.

$$\mathcal{K}_{\text{bin}} = \bigcup_{j=1}^m \left\{ \arg \max_{i \in \mathcal{S}_j} s_i \right\}, \quad \mathcal{S}_j = \left[\frac{(j-1) \cdot N}{m}, \frac{j \cdot N}{m} \right) \quad (3)$$

If $|\mathcal{K}_{\text{bin}}| < m$, we supplement it with top-ranked global scores not in \mathcal{K}_{bin} :

$$\mathcal{K}_{\text{final}} = \mathcal{K}_{\text{bin}} \cup \{\text{Top-}(m - |\mathcal{K}_{\text{bin}}|)\} \quad (4)$$

- Patch-level Region Extraction.** For each selected frame f_i , we feed it into InternViT-6B to extract patch-wise features $\mathbf{H}_i \in \mathbb{R}^{P \times d}$. The saliency score for patch p is computed via L2 norm:

$$\text{score}_{i,p} = \left\| \mathbf{H}_i^{(p)} \right\|_2 \quad (5)$$

The highest scoring patch is selected as the visual key region.

Combining the selected frames and regions, we define the structured visual clue as:

$$\mathcal{K} = \{(f_i, \text{region}_i)\}_{i=1}^3$$

Interleaved Multi-stage CoT Generation. Given the keyframe set \mathcal{K} , ASR transcripts, and a predefined safety rule schema, we perform a three-stage prompting process with a frozen InternVL-3-78B model to generate structured reasoning:

- Step-1 (Observation):** Describe the visual content, summarize ASR transcripts, and assess the consistency between modalities.
- Step-2+3 (Risk Screening + Causal Analysis):** Identify potential policy violations and analyze their underlying causes.
- Step-4 (Final Verdict):** Integrate multi-stage reasoning to reach a final compliance decision in accordance with the rule schema.

For each sample j , the model produces a visual-text-label triplet:

$$\left(K^{(j)}, T^{(j)}, y^{(j)} \right), \quad j = 1, 2, \dots, |\mathcal{D}|$$

where K denotes the selected keyframes, T the reasoning chain, and y the predicted compliance label.

Internally, the model first generates intermediate reasoning tuples $(\text{obs}_i, \text{risk}_i, \text{cause}_i, \text{verdict}_i)_{i=1}^{16}$ for each frame, which are then aggregated into the final triplet (K, T, y) .

Rule-Anchored Supervised Fine-tuning (SFT) On top of the synthesized dataset, we perform supervised fine-tuning of the base VLM. The training objective consists of two components:

$$\mathcal{L} = \mathcal{L}_{\text{CE}}(\langle \text{answer} \rangle) + \lambda \text{KL}(p_{\text{think}} \parallel p_{\text{rule}}),$$

where the primary loss \mathcal{L}_{CE} computes the cross-entropy on the answer token, encouraging accurate prediction of the compliance label. The auxiliary KL term aligns the model’s reasoning distribution (think) with a rule-guided prior p_{rule} , which we construct by tokenizing the target violation scene and type (e.g., “false marketing”, “adult content”) into a keyword set. These keywords are normalized into a soft token-level target distribution, serving as an interpretable and structured prior to guide the model’s causal reasoning.

In summary, we construct a compact yet effective multi-modal CoT corpus and perform rule-anchored fine-tuning to establish a solid cold-start foundation for downstream reinforcement learning.

Reinforcement Learning: Principle-Guided Self-Consistent Optimization

Building on the cold-start model trained with multimodal CoT-style supervision, we apply online reinforcement learning (RL) to further enhance the model’s reasoning ability and policy alignment in content moderation scenarios. The overall training procedure is illustrated in **Algorithm 1**, which consists of three main components:

- Data Curation:** Constructing difficult and diverse samples to improve generalization.
- Reward Design:** Combining rule-based correctness, principle-guided critique, and format consistency.
- Policy Optimization:** Applying GRPO with token-level normalization and dynamic sampling.

To improve generalization and sample efficiency in RL, we construct a refined dataset \mathcal{D}_{RL} from the supervised set \mathcal{D} through the following strategies:

(1) Rejection Sampling. We sample from \mathcal{D} with high temperature for 4 passes, collecting all examples where the model fails consistently across generations as hard samples.

(2) Safety-Aware Concatenation. To increase multi-modal diversity and risk density, we concatenate semantically related prompts or risky ASR segments to simulate more complex decision contexts.

Note: We use only ASR since OCR in short videos is often noisy (e.g., overlays), which harms reasoning. ASR provides more stable guidance.

Reward Design We propose a hybrid reward mechanism that combines factual correctness, interpretable reasoning, and structured output format. The final reward used in RL is:

$$r = r_{\text{rule}} + r_{\text{format}} + r_{\text{scaR}}$$

(1) Rule-Based Reward. We assign a discrete correctness reward based on whether the predicted violation `scene` and `type` match the ground truth. Each correct field contributes 0.5:

$$r_{\text{rule}} = \begin{cases} 1.0, & \text{if } \hat{y}_{\text{scene}} = y_{\text{scene}} \text{ and } \hat{y}_{\text{type}} = y_{\text{type}} \\ 0.5, & \text{if } \hat{y}_{\text{scene}} = y_{\text{scene}} \text{ and } \hat{y}_{\text{type}} \neq y_{\text{type}} \\ 0.0, & \text{otherwise} \end{cases}$$

(2) Format-Aware Reward. To ensure outputs follow the required structure, we assign a binary reward based on the presence of both `<think>` and `<answer>` tags:

$$r_{\text{format}} = \mathbb{I}[\langle \text{think} \rangle \in \hat{y} \wedge \langle \text{answer} \rangle \in \hat{y}]$$

(3) SCA-R: Self-Consistency and Adaptive Reward. To address reward misalignment caused by policy drift, we propose **SCA-R**, a self-consistency reward guided by scenario-adaptive critique.

For each (input, response) pair, a guide model acts as a rating expert. It receives the `<think>` trace, ground-truth labels, and moderation guidelines, then dynamically constructs scoring principles $\mathcal{P} = \{p_k\}$ (e.g., causal clarity, risk attribution), each with weight w_k .

The model conducts step-by-step critique and assigns a score in $\{0, 0.5, 1\}$, where 1 indicates both reasoning and conclusion are aligned with policy.

$$r_{\text{scaR}} = \sum_{k=1}^K w_k \cdot \text{score}_{p_k}(\hat{y})$$

This adaptive reward enhances interpretability while maintaining alignment under evolving moderation standards.

Policy Optimization with GRPO We adopt Group Relative Policy Optimization (GRPO) and introduce several modifications for improved stability and sample efficiency:

Token-Level Normalization. Unlike standard GRPO that normalizes loss at the sequence level, we compute the objective at the token level to mitigate reward bias due to varying output lengths.

Group Advantage. For each input, we sample G responses and compute relative advantage:

$$A_i = \frac{r_i - \text{mean}(\{r_j\}_{j=1}^G)}{\text{std}(\{r_j\}_{j=1}^G)} \quad (6)$$

Clipped Policy Loss. The final RL objective is:

$$\mathcal{L}_{\text{RL}} = -\mathbb{E} \left[\frac{1}{G} \sum_{i=1}^G \min(K_i, \text{clip}(K_i, 1 - B, 1 + B)) \cdot A_i \right] \quad (7)$$

where the policy ratio is:

$$K_i = \frac{P_{\theta}(\hat{y}_i | x)}{P_{\theta_{\text{old}}}(\hat{y}_i | x)} \quad (8)$$

The clipping factor B is annealed with training step s :

$$B_s = \prod_{i=1}^s \left(\frac{s_{\text{total}} - i}{s_{\text{total}}} \right) \cdot \epsilon \quad (9)$$

Dynamic Sampling. To avoid gradient collapse when all G samples receive identical rewards, we skip such batches and continue sampling until variance is non-zero.

In summary, our hybrid reward design and GRPO optimization jointly enhance the model’s factual precision, causal reasoning clarity, and alignment with real-world moderation policies.

Algorithm 1: BLM-Guard: Self-Adaptive GRPO Training

procedure TRAIN(π_{θ} , \mathcal{R} , \mathcal{D}_{RL} , S) {**Inputs:** π_{θ} = moderation model (from cold-start), \mathcal{R} = reward model (rule + format + SCA-R), \mathcal{D}_{RL} = reasoning dataset, S = training steps, G = group size per input, ϵ = clip scale}

for $s = 1$ to S **do**

for each x in batch from \mathcal{D}_{RL} **do**

Phase 1: Response Generation and Reward

Generate G responses $\{o_i\}_{i=1}^G \sim \pi_{\theta}(x)$

for $i = 1$ to G **do**

$r_i^{\text{rule}} \leftarrow$ rule match score

$r_i^{\text{format}} \leftarrow$ format compliance score

$r_i^{\text{scaR}} \leftarrow$ critique score from guide model

$r_i \leftarrow r_i^{\text{rule}} + r_i^{\text{format}} + r_i^{\text{scaR}}$

end for

if $\text{Var}(\{r_i\}) = 0$ **then**

continue {Skip zero-variance batch}

end if

$\mu, \sigma \leftarrow$ mean and std of $\{r_i\}$

Phase 2: GRPO Optimization

for $i = 1$ to G **do**

$A_i \leftarrow \frac{r_i - \mu}{\sigma}$ {Z-score advantage}

$K_i \leftarrow \frac{P_{\theta}(o_i)}{P_{\theta_{\text{old}}}(o_i)}$

end for

$B_s \leftarrow \epsilon \cdot \frac{S - s}{S}$ {Annealed clip factor}

Compute loss:

$$\mathcal{L}_{\text{RL}} = -\frac{1}{G} \sum_{i=1}^G \min(K_i A_i, \text{clip}(K_i, 1 - B_s, 1 + B_s) \cdot A_i)$$

$\theta \leftarrow \theta - \eta \cdot \nabla_{\theta} \mathcal{L}_{\text{RL}}$

$\theta_{\text{old}} \leftarrow \theta$ {Update reference policy}

end for

end for

return Optimized policy π_{θ}

end procedure

Experiments

Baselines We compare **BLM-Guard** against several state-of-the-art open-source multimodal models. These include guard-specific architectures **LlavaGuard-7B** and **QwenGuard-7B** (Helff et al. 2025), and strong general-purpose foundation models: **LLaVA-Next-Video-7B** (Zhang et al. 2024), **InternVL3-8B** (Zhu et al. 2025), and **Qwen2.5-VL-7B-Instruct** (Bai et al. 2025) (our backbone). We also evaluate **Kimi-VL-A3B-Thinking** (Team et al. 2025), which utilizes structured Chain-of-Thought reasoning. For fair comparison, all models are adapted to

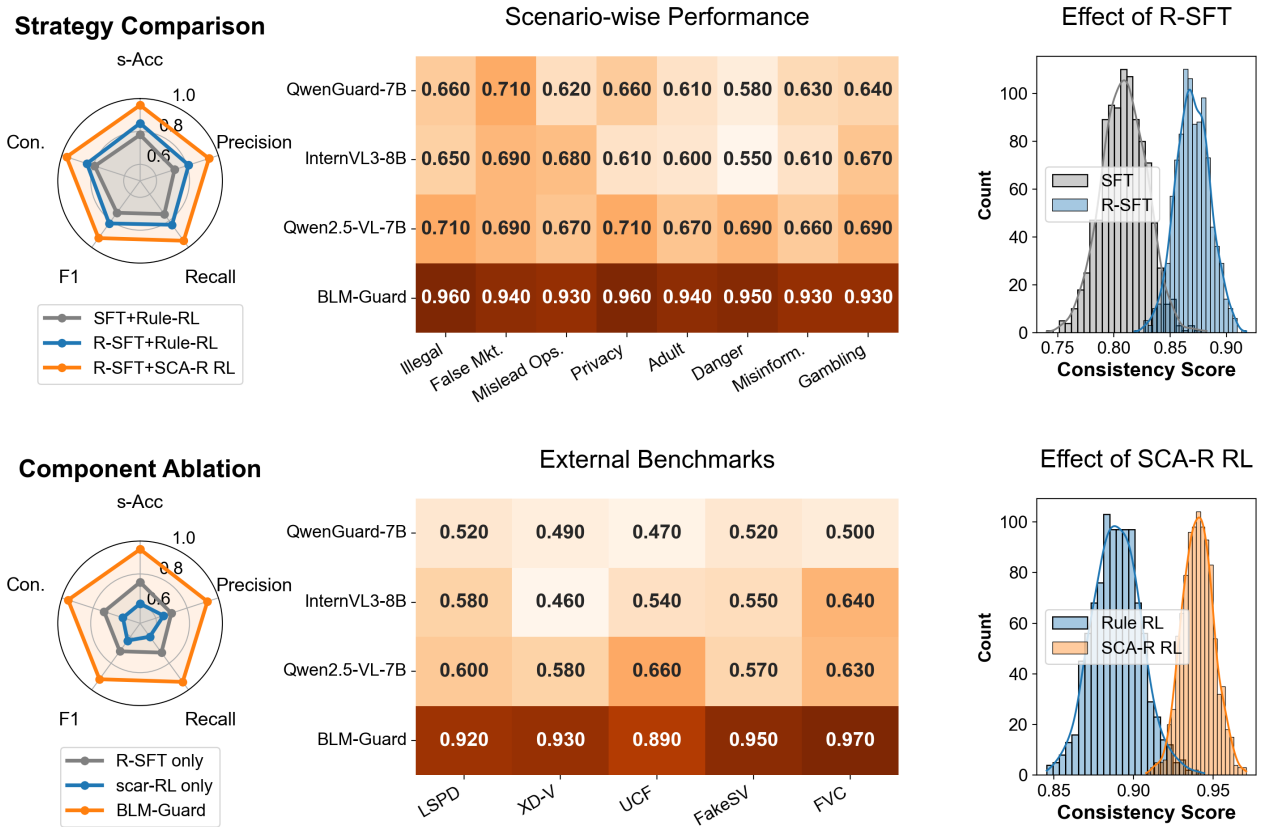


Figure 3: Comparison of model strategies and components. **Left:** Radar plots show gains in accuracy, precision, and consistency. **Middle:** Heatmaps display performance across risk scenarios and benchmarks. **Right:** Histograms illustrate consistency improvements from Rule-SFT and SCA-R reward learning.

video moderation via uniform frame sampling and ASR transcript concatenation.

Datasets. We evaluate **BLM-Guard** and baselines across diverse moderation tasks. Our custom benchmark, **BLM-Guard-Bench**, Each video is annotated with a three-level policy taxonomy: **Severity**, **Scenario**, and **Violation Type** (see Figure 1). A subset contains structured reasoning traces for supervised and reward-based learning. To ensure comparability, we further evaluate on five public moderation datasets: LSPD (Phan et al. 2022), XD-Violence (Wu et al. 2020), UCF-Crime (Sultani, Chen, and Shah 2018), FakeSV (Qi et al. 2023), and FVC (Papadopoulou, Papadopoulos, and Kompatsiaris 2018), spanning tasks such as violence detection, anomaly spotting, and misinformation detection. Inspired by SafeWatch (Chen et al. 2024), we additionally assess generalization to unseen policy domains.

Metrics We evaluate moderation along two dimensions:

(1) Violation Identification Accuracy. Based on structured `<answer>` outputs, we report *Wide Accuracy* (partial match of violation scene or type) and *Strict Accuracy* (full match of both). We also include macro/weighted precision, recall, and F1-score for harmful vs. harmless classification.

(2) Consistency Score. To assess reasoning quality in the `<think>` section, we use a GPT-4o-based guide model to score alignment with moderation principles. The score ranges from 0 to 1, with higher values indicating more consistent and interpretable reasoning.

Results on BLM-Guard Benchmark As Table 1 shows, **BLM-Guard** outperforms all baselines in severity classification, overall accuracy, binary risk detection, and reasoning consistency, producing more stable and principle-aligned outputs. This demonstrates the effectiveness of combining rule-guided supervision, structured reasoning, and reinforcement learning.

Generalization on Existing Datasets. **BLM-Guard** generalizes well to public benchmarks (LSPD, XD-Violence, UCF-Crime, FakeSV, FVC), particularly excelling in misinformation scenarios (FakeSV/FVC) where prior models struggle. While these datasets do not directly evaluate factual consistency, **BLM-Guard** effectively detects deceptive content and cross-modal inconsistencies, consistent with its internal benchmark performance (see Figure 3).

Model	Severity-wise Classification				Accuracy		Binary Detection					Con.
	High	Medium	Low	None	w-Acc.	s-Acc.	NR (B)	R (B)	P	R	F1 (R)	Con.
Qwen2.5-VL-7B	0.745	0.701	0.669	0.781	0.712	0.701	0.755	0.918	0.831	0.576	0.680	0.642
Qwen2.5-VL-32B	0.732	0.685	0.652	0.768	0.703	0.682	0.875	0.823	0.769	0.835	0.801	0.667
Kimi-VL-A3B-Thinking	0.728	0.684	0.651	0.624	0.529	0.511	0.887	0.561	0.588	0.898	0.711	0.701
InternVL3-14B	0.581	0.535	0.502	0.622	0.521	0.499	0.931	0.519	0.485	0.950	0.642	0.593
InternVL3-8B	0.695	0.645	0.607	0.701	0.389	0.631	0.921	0.299	0.461	0.961	0.623	0.611
QwenGuard-7B	0.733	0.631	0.599	0.693	0.582	0.512	0.823	0.695	0.642	0.785	0.706	0.598
LlavaGuard-7B	0.301	0.322	0.395	0.392	0.255	0.235	0.329	0.258	0.193	0.251	0.218	0.527
LLaVA-Next-Video-7B	0.321	0.345	0.435	0.443	0.342	0.312	0.348	0.3101	0.1540	0.1780	0.165	0.533
BLM-Guard (ours)	0.978	0.940	0.902	0.986	0.962	0.914	0.781	0.950	0.976	0.962	0.969	0.845

Table 1: Evaluation results of harmfulness detection for short-video advertisements. The left block reports severity-wise classification accuracy across *High*, *Medium*, *Low*, and *None* levels. The middle block presents overall accuracy, including lenient accuracy (**w-Acc.**) and strict accuracy (**s-Acc.**). The right block shows binary detection metrics: recall of non-risky (**NR (B)**) and risky (**R (B)**) classes, along with precision (**P**) and recall (**R**) for the risky class. “**Con.**” denotes reasoning–answer consistency scores assessed by GPT-4o (ranging from 0 to 1). Bold values indicate the best performance in each column; **BLM-Guard** achieves the strongest results across all metrics.

Ablation Study

We evaluate BLM-GUARD (Qwen2.5-VL-7B) across seven configurations to verify the impact of our training stages: (1) **Ans-SFT** and (2) **Think-SFT** (single-component tuning); (3) **Full-SFT (13k)** vs. (4) **Part-SFT (5k)** (data scaling); (5) **Rule-SFT (5k)** (rule-anchored joint tuning); (6) **+Rule-RL** (rule-based reward learning); and (7) **+SCA-R (Full)**, our final setup with SCA-R guided RL. The results, summarized in Table 2, highlight the synergy between rule-aligned SFT and SCA-R refinement.

Setting	s-Acc.	P	R	F1	Con.
Ans-SFT	0.648	0.765	0.702	0.732	0.412
Think-SFT	0.612	0.720	0.679	0.699	0.585
Full-SFT (13k)	0.7925	0.874	0.862	0.867	0.762
Part-SFT (5k)	0.768	0.780	0.840	0.809	0.762
Rule-SFT (5k)	0.783	0.882	0.853	0.867	0.776
+ Rule-RL	0.801	0.915	0.874	0.894	0.781
+ SCA-R (Full)	0.914	0.976	0.962	0.969	0.845

Table 2: Ablation on Qwen2.5-VL-7B. **Ans-SFT/Think-SFT** are single-supervision baselines. **Rule-SFT** uses rule-aligned tuning. **+Rule-RL/+SCA-R** introduce reward learning; SCA-R adds adaptive alignment. **Con.** denotes reasoning consistency.

Findings: (1) Rule-SFT outperforms single-stage SFT by improving both accuracy and interpretability. (2) RL enhances robustness; combining Rule-SFT with rule-based RL yields strong gains. (3) The full version (+SCA-R) achieves the best overall performance in reasoning quality and policy alignment.

Related Work

VLM Guardrails. Early guardrails like LLAMA-GUARD (Inan et al. 2023) and LLAVAGUARD (Helff et al. 2025) focused on text/image safety. While recent works extend to video (Chen et al. 2024) and legal alignment (Lu

et al. 2025), they often fail to capture nuanced ad violations like misleading visuals or exaggerated claims.

CoT and RL Alignment. CoT interpretability (Tanneru et al. 2024; Xu et al. 2024) and multi-step reasoning (e.g., Tree-of-Thoughts (Yao et al. 2023), GuardAgent (Xiang et al. 2025)) are foundational to our approach. In multi-modal safety, GuardReasoner (Liu et al. 2025a) and MM-RLHF (Zhang et al. 2025) explore preference-based RL. We extend these by incorporating policy-specific, dynamic principles to boost alignment.

Rule-Guided Data & ICoT. Instruction tuning (e.g., Self-Instruct (Wang et al. 2023)) and policy-grounded generation (e.g., MSR-Align (Xia et al. 2025)) minimize annotation costs. Unlike these, we introduce Interleaved CoT (ICoT) and adaptive keyframe sampling specifically designed for structured reasoning in multimodal ad contexts.

Conclusion

We propose **BLM-Guard**, an explainable moderation framework designed for short-video advertisements. It integrates *Chain-of-Thought* (CoT) reasoning, rule-based policy alignment, and critic-guided reinforcement learning to enable accurate and transparent decision-making. Built upon a risk-tiered benchmark reflecting real-world regulatory concerns, BLM-Guard leverages rule-annotated supervision for policy-grounded alignment in both classification and reasoning tasks. It adopts a multi-task design that captures intra-modal manipulations (e.g., exaggerated visuals) and cross-modal inconsistencies (e.g., subtitle–speech mismatch), which are common in deceptive or policy-violating content. Extensive experiments show that BLM-Guard outperforms strong open- and closed-source baselines in accuracy and consistency, and generalizes well across datasets and unseen policy categories, demonstrating robustness, interpretability, and practical value in real-world moderation systems.

References

- Ahn, D.; Choi, Y.; Yu, Y.; Kang, D.; and Choi, J. 2024. Tuning Large Multimodal Models for Videos using Reinforcement Learning from AI Feedback. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 923–940. Bangkok, Thailand: Association for Computational Linguistics.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; Zhong, H.; Zhu, Y.; Yang, M.; Li, Z.; Wan, J.; Wang, P.; Ding, W.; Fu, Z.; Xu, Y.; Ye, J.; Zhang, X.; Xie, T.; Cheng, Z.; Zhang, H.; Yang, Z.; Xu, H.; and Lin, J. 2025. Qwen2.5-VL Technical Report. arXiv:2502.13923.
- Chen, Y.; Ge, Y.; Wang, R.; Ge, Y.; Cheng, J.; Shan, Y.; and Liu, X. 2025. GRPO-CARE: Consistency-Aware Reinforcement Learning for Multimodal Reasoning. arXiv:2506.16141.
- Chen, Z.; Pinto, F.; Pan, M.; and Li, B. 2024. SafeWatch: An Efficient Safety-Policy Following Video Guardrail Model with Transparent Explanations. arXiv:2412.06878.
- Gao, J.; Li, Y.; Cao, Z.; and Li, W. 2025. Interleaved-Modal Chain-of-Thought. arXiv:2411.19488.
- Goldstein, J. A.; Sastry, G.; Musser, M.; DiResta, R.; Gentzel, M.; and Sedova, K. 2023. Generative Language Models and Automated Influence Operations: Emerging Threats and Potential Mitigations. arXiv:2301.04246.
- Guo, K.; Utkarsh, A.; Ding, W.; Ondracek, I.; Zhao, Z.; Freeman, G.; Vishwamitra, N.; and Hu, H. 2024. Moderating Illicit Online Image Promotion for Unsafe User-Generated Content Games Using Large Vision-Language Models. arXiv:2403.18957.
- Helff, L.; Friedrich, F.; Brack, M.; Kersting, K.; and Schramowski, P. 2025. LlavaGuard: An Open VLM-based Framework for Safeguarding Vision Datasets and Models. arXiv:2406.05113.
- Inan, H.; Upasani, K.; Chi, J.; Rungta, R.; Iyer, K.; Mao, Y.; Tontchev, M.; Hu, Q.; Fuller, B.; Testuggine, D.; and Khabisa, M. 2023. Llama Guard: LLM-based Input-Output Safeguard for Human-AI Conversations. arXiv:2312.06674.
- Liu, Y.; Gao, H.; Zhai, S.; Xia, J.; Wu, T.; Xue, Z.; Chen, Y.; Kawaguchi, K.; Zhang, J.; and Hooi, B. 2025a. GuardReasoner: Towards Reasoning-based LLM Safeguards. arXiv:2501.18492.
- Liu, Y.; Zhai, S.; Du, M.; Chen, Y.; Cao, T.; Gao, H.; Wang, C.; Li, X.; Wang, K.; Fang, J.; Zhang, J.; and Hooi, B. 2025b. GuardReasoner-VL: Safeguarding VLMs via Reinforced Reasoning. arXiv:2505.11049.
- Lu, X.; Zhang, T.; Meng, C.; Wang, X.; Wang, J.; Zhang, Y.; Tang, S.; Liu, C.; Ding, H.; Jiang, K.; Tang, K.; Wen, B.; Zheng, H.-T.; Yang, F.; Gao, T.; Zhang, D.; and Gai, K. 2025. VLM as Policy: Common-Law Content Moderation Framework for Short Video Platform. arXiv:2504.14904.
- Papadopoulou, O.; Papadopoulos, S.; and Kompatsiaris, Y. 2018. A two-level verification approach for multimedia content with few available ground truth labels. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1982–1986. IEEE.
- Phan, D.; Nguyen, T. T.; Nguyen, Q. H.; Tran, H. L.; Nguyen, K. N. K.; and Vu, D. L. 2022. LSPD: A Large-Scale Pornographic Dataset for Detection and Classification. *International Journal of Intelligent Engineering Systems*, 15(1): 198–210.
- Qi, P.; Bu, Y.; Cao, J.; Ji, W.; Shui, R.; Xiao, J.; Wang, D.; and Chua, T. 2023. FakeSV: A Multimodal Benchmark with Rich Social Context for Fake News Detection on Short Video Platforms. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Qi, P.; Bu, Y.; Cao, J.; Ji, W.; Shui, R.; Xiao, J.; Wang, D.; and Chua, T.-S. 2022. FakeSV: A Multimodal Benchmark with Rich Social Context for Fake News Detection on Short Video Platforms. arXiv:2211.10973.
- Qu, Y.; Shen, X.; Wu, Y.; Backes, M.; Zannettou, S.; and Zhang, Y. 2024. UnsafeBench: Benchmarking Image Safety Classifiers on Real-World and AI-Generated Images. arXiv:2405.03486.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. arXiv:2103.00020.
- Sultani, W.; Chen, C.; and Shah, M. 2018. Real-World Anomaly Detection in Surveillance Videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tang, X.; Qiu, J.; Xie, L.; Tian, Y.; Jiao, J.; and Ye, Q. 2025. Adaptive Keyframe Sampling for Long Video Understanding. arXiv:2502.21271.
- Tanneru, S. H.; Ley, D.; Agarwal, C.; and Lakkaraju, H. 2024. On the Hardness of Faithful Chain-of-Thought Reasoning in Large Language Models. arXiv:2406.10625.
- Team, K.; Du, A.; Yin, B.; Xing, B.; Qu, B.; Wang, B.; Chen, C.; Zhang, C.; Du, C.; Wei, C.; Wang, C.; Zhang, D.; Du, D.; Wang, D.; Yuan, E.; Lu, E.; Li, F.; Sung, F.; Wei, G.; Lai, G.; Zhu, H.; Ding, H.; Hu, H.; Yang, H.; Zhang, H.; Wu, H.; Yao, H.; Lu, H.; Wang, H.; Gao, H.; Zheng, H.; Li, J.; Su, J.; Wang, J.; Deng, J.; Qiu, J.; Xie, J.; Wang, J.; Liu, J.; Yan, J.; Ouyang, K.; Chen, L.; Sui, L.; Yu, L.; Dong, M.; Dong, M.; Xu, N.; Cheng, P.; Gu, Q.; Zhou, R.; Liu, S.; Cao, S.; Yu, T.; Song, T.; Bai, T.; Song, W.; He, W.; Huang, W.; Xu, W.; Yuan, X.; Yao, X.; Wu, X.; Li, X.; Zu, X.; Zhou, X.; Wang, X.; Charles, Y.; Zhong, Y.; Li, Y.; Hu, Y.; Chen, Y.; Wang, Y.; Liu, Y.; Miao, Y.; Qin, Y.; Chen, Y.; Bao, Y.; Wang, Y.; Kang, Y.; Liu, Y.; Dong, Y.; Du, Y.; Wu, Y.; Wang, Y.; Yan, Y.; Zhou, Z.; Li, Z.; Jiang, Z.; Zhang, Z.; Yang, Z.; Huang, Z.; Huang, Z.; Zhao, Z.; Chen, Z.; and Lin, Z. 2025. Kimi-VL Technical Report. arXiv:2504.07491.
- Wang, Y.; Kordi, Y.; Mishra, S.; Liu, A.; Smith, N. A.; Khushabi, D.; and Hajishirzi, H. 2023. Self-Instruct: Aligning Language Models with Self-Generated Instructions. arXiv:2212.10560.
- Wu, Z.; Li, D.; Liu, B.; Su, H.; Zhu, J.; and Tan, T. 2020. Not only look, but also listen: Learning multimodal violence

detection under weak supervision. In *European Conference on Computer Vision*, 322–339. Springer.

Xia, Y.; Jiang, Y.; Tan, Y.; Zhu, X.; Yue, X.; and Zheng, B. 2025. MSR-Align: Policy-Grounded Multimodal Alignment for Safety-Aware Reasoning in Vision-Language Models. arXiv:2506.19257.

Xiang, Z.; Zheng, L.; Li, Y.; Hong, J.; Li, Q.; Xie, H.; Zhang, J.; Xiong, Z.; Xie, C.; Yang, C.; Song, D.; and Li, B. 2025. GuardAgent: Safeguard LLM Agents by a Guard Agent via Knowledge-Enabled Reasoning. arXiv:2406.09187.

Xu, J.; Fei, H.; Pan, L.; Liu, Q.; Lee, M.-L.; and Hsu, W. 2024. Faithful Logical Reasoning via Symbolic Chain-of-Thought. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 13326–13365. Bangkok, Thailand: Association for Computational Linguistics.

Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T. L.; Cao, Y.; and Narasimhan, K. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. arXiv:2305.10601.

Yeh, C.; Chang, Y.-M.; Chiu, W.-C.; and Yu, N. 2024. T2Vs Meet VLMs: A Scalable Multimodal Dataset for Visual Harmfulness Recognition. arXiv:2409.19734.

Zhang, Y.; Li, B.; Liu, h.; Lee, Y. j.; Gui, L.; Fu, D.; Feng, J.; Liu, Z.; and Li, C. 2024. LLaVA-NeXT: A Strong Zero-shot Video Understanding Model.

Zhang, Y.-F.; Yu, T.; Tian, H.; Fu, C.; Li, P.; Zeng, J.; Xie, W.; Shi, Y.; Zhang, H.; Wu, J.; Wang, X.; Hu, Y.; Wen, B.; Yang, F.; Zhang, Z.; Gao, T.; Zhang, D.; Wang, L.; Jin, R.; and Tan, T. 2025. MM-RLHF: The Next Step Forward in Multimodal LLM Alignment. arXiv:2502.10391.

Zhu, J.; Wang, W.; Chen, Z.; Liu, Z.; Ye, S.; Gu, L.; Tian, H.; Duan, Y.; Su, W.; Shao, J.; Gao, Z.; Cui, E.; Wang, X.; Cao, Y.; Liu, Y.; Wei, X.; Zhang, H.; Wang, H.; Xu, W.; Li, H.; Wang, J.; Deng, N.; Li, S.; He, Y.; Jiang, T.; Luo, J.; Wang, Y.; He, C.; Shi, B.; Zhang, X.; Shao, W.; He, J.; Xiong, Y.; Qu, W.; Sun, P.; Jiao, P.; Lv, H.; Wu, L.; Zhang, K.; Deng, H.; Ge, J.; Chen, K.; Wang, L.; Dou, M.; Lu, L.; Zhu, X.; Lu, T.; Lin, D.; Qiao, Y.; Dai, J.; and Wang, W. 2025. InternVL3: Exploring Advanced Training and Test-Time Recipes for Open-Source Multimodal Models. arXiv:2504.10479.