

RAVEN: Robust Advertisement Video Violation Temporal Grounding via Reinforcement Reasoning

Deyi Ji^{1*} Yuekui Yang^{1,2*} Haiyang Wu¹ Shaoping Ma² Tianrun Chen³ Lanyun Zhu^{4†}

¹Tencent ²Department of Computer Science and Technology, Tsinghua University

³Zhejiang University ⁴Singapore University of Technology and Design

deyiji@tencent.com, yuekuiyang@tencent.com, gavinwu@tencent.com,
msp@tsinghua.edu.cn, tianrun.chen@zju.edu.cn, lanyun_zhu@mymail.sutd.edu.sg

Abstract

Advertisement (Ad) video violation detection is critical for ensuring platform compliance, but existing methods struggle with precise temporal grounding, noisy annotations, and limited generalization. We propose RAVEN, a novel framework that integrates curriculum reinforcement learning with multimodal large language models (MLLMs) to enhance reasoning and cognitive capabilities for violation detection. RAVEN employs a progressive training strategy, combining precisely and coarsely annotated data, and leverages Group Relative Policy Optimization (GRPO) to develop emergent reasoning abilities without explicit reasoning annotations. Multiple hierarchical sophisticated reward mechanism ensures precise temporal grounding and consistent category prediction. Experiments on industrial datasets and public benchmarks show that RAVEN achieves superior performances in violation category accuracy and temporal interval localization. We also design a pipeline to deploy the RAVEN on the online Ad services, and online A/B testing further validates its practical applicability, with significant improvements in precision and recall. RAVEN also demonstrates strong generalization, mitigating the catastrophic forgetting issue associated with supervised fine-tuning.

1 Introduction

In the modern digital landscape, advertisements play a pivotal role in sustaining the growth of internet platforms. To ensure compliance with local laws and regulations, promote sustainable development, and foster a user-friendly environment, platforms establish stringent guidelines to regulate the content uploaded by advertisers. Despite these efforts, violations of platform policies persist. Early

*The first two authors contribute equally to this work. We acknowledge Shaogang Tang for collaborating on data resources and application scenarios to validate and improve algorithm performance.

†Corresponding Author.

approaches relied on small-scale models (Dosovitskiy, 2020; He et al., 2016) to analyze and identify such violations, but these methods suffered from limited generalization capabilities. With the advent of large language models (LLMs) (Liu et al., 2023; Bai et al., 2023a), more advanced techniques have been increasingly adopted in practice to detect non-compliant content.

Among the various types of content, video advertisements present the most significant challenge for violation detection. In practice, it is not only necessary to predict the violation categories of a video but also to localize the specific sub-scenes corresponding to each category. A single video may contain multiple violation categories, each potentially associated with multiple temporal intervals. Existing methods typically follow a two-step process: (1) annotating each video with its violation categories and their corresponding temporal intervals, and (2) fine-tuning multimodal large language models (MLLMs) using supervised fine-tuning (SFT) techniques.

However, due to constraints in data volume, annotation costs, and the inherent difficulty of precise labeling, the annotated sub-scene intervals often contain natural errors or ambiguities. These inaccuracies can lead to unstable training or even misguided learning when using conventional SFT methods. As discussed in (Shao et al., 2024; Liu et al., 2025), SFT faces several limitations: its effectiveness notably diminishes in out-of-domain settings despite strong in-domain performance, it inherently causes catastrophic forgetting of general capabilities, and the absence of an explicit reasoning process limits its efficacy in complex scenarios. Recent research shows that pure reinforcement learning (RL) (Guo et al., 2025) fosters emergent reasoning during testing, underscoring the power of reward-driven optimization in boosting model reasoning. This method also tends to improve generalization, avoiding overfitting to specific datasets.

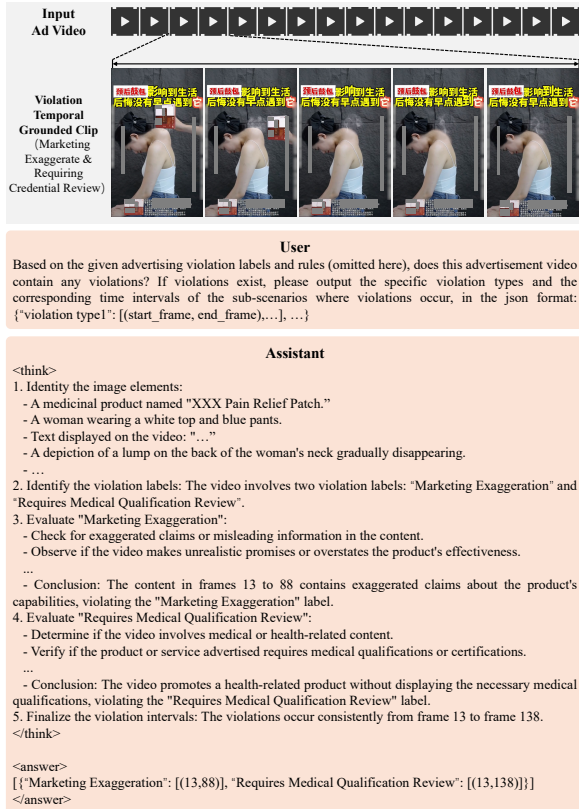


Figure 1: The example of RAVEN reasoning.

Building on these insights, we introduce RAVEN, a novel framework aimed at advancing reasoning and cognitive skills for detecting violation scenes in videos. The name RAVEN, symbolizing “keen insight”, reflects our aspiration for the system to detect violations with the sharpness of a raven. RAVEN is a structured reasoning MLLM that combines curriculum learning with hierarchical, multi-granular reinforcement. It employs GRPO (Group Relative Policy Optimization) (Shao et al., 2024; Guo et al., 2025) and structured thinking, eliminating the need for explicitly annotated reasoning process data. Instead, it leverages the self-evolution potential of MLLMs to develop reasoning capabilities from scratch. A significant advantage of RAVEN is its ability to robustly train on large-scale, noisy, coarsely annotated industrial data, achieving superior violation detection performance while preserving the strong generalization capabilities of MLLMs. To achieve this, we develop hierarchical sophisticated rewards mechanism comprising multiple types of rewards: format rewards, which enforce constraints on the structure of the reasoning process and violation sub-scene outputs, and accuracy rewards, which include primary rewards (e.g., IoU Reward), auxil-

ary rewards (e.g., Boundary Alignment Reward), and regularization rewards (e.g., Category Consistency Reward). As illustrated in Figure 1, RAVEN exhibits emergent test-time reasoning abilities, enabling it to handle complex instructions by breaking them down into sequential analytical steps, thus achieving precise localization of violation intervals. RAVEN demonstrates exceptional performance on both in-domain and out-of-domain data, significantly outperforming models trained via SFT.

To validate RAVEN, we conduct extensive experiments from both offline and online testing perspectives, using both publicly available datasets and proprietary industrial data. The results show that the RAVEN-7B model exhibits strong test-time reasoning capabilities and achieves superior generalization performance compared to models of the same scale. Our contributions are threefold: (1) We propose RAVEN, the novel architecture specifically designed for localizing violation scenes in advertisement content. Through its innovative design, RAVEN exhibits emergent reasoning abilities. (2) RAVEN is a practical system tailored for real-world industrial applications. It demonstrates remarkable robustness when trained on large-scale, noisy, coarsely annotated data, while retaining strong generalization capabilities. (3) Extensive experiments on both offline and online testing, using public datasets and proprietary industrial data, demonstrate that the RAVEN-7B model achieves superior reasoning and generalization performance compared to models of the same scale.

2 Related Work

2.1 Temporal Grounding in Videos

Temporal grounding aims to localize specific events or actions within a video. Prior work has focused on supervised learning with precise annotations (Gao et al., 2017). However, these methods struggle with noisy, coarsely annotated data, which is prevalent in industrial settings. Recent approaches like VSLNet (Zhang et al., 2020a) and 2D-TAN (Zhang et al., 2020b) have improved localization accuracy but lack robust reasoning capabilities for complex tasks like violation detection.

2.2 Multimodal Large Language Models

Multimodal Large Language Models (MLLMs) (Yin et al., 2023; Xu et al., 2024a; Maity et al., 2024), such as CLIP (Radford et al., 2021), Flamingo (Alayrac et al., 2022), and BLIP (Li et al.,

2022), have demonstrated remarkable capabilities in understanding and reasoning across modalities on various tasks (Wei et al., 2022a,b; Kojima et al., 2022; Ji et al., 2024b, 2023, 2024a, 2022, 2025; Liu et al., 2024; Zhu et al., 2024b). These models excel in tasks like image-text retrieval and video captioning but are often limited by their reliance on supervised fine-tuning (SFT), which can lead to catastrophic forgetting and poor generalization. Recent efforts like LLaVA (Liu et al., 2023; Xu et al., 2024b), Qwen (Bai et al., 2023a,b) and Video-ChatGPT (Maaz et al., 2024) have explored integrating reasoning into MLLMs, but they remain underutilized in temporal grounding tasks.

2.3 Reinforcement Learning for Video Understanding

Reinforcement learning (RL) (Guo et al., 2025; Kaelbling et al., 1996; Christiano et al., 2017; Zhu et al., 2024a; Rafailov et al., 2024; Song et al., 2024; Liu et al.) has been applied to video understanding tasks, such as action segmentation and event detection. Methods like SM-RL (Wang et al., 2019a,b) and RLPP (Li et al., 2018) use RL to optimize temporal localization but are limited by their inability to handle multimodal inputs or perform complex reasoning. Curriculum reinforcement learning (Narvekar et al., 2020; Bengio et al., 2009) has shown promise in improving RL’s robustness and generalization, but its application to temporal grounding remains unexplored.

2.4 Advertisement Video Violation Detection

Existing methods for advertisement video violation detection rely heavily on rule-based systems or supervised learning with precise annotations. These approaches are effective in controlled environments but fail to generalize to large-scale, noisy industrial datasets. Recent works (Wang et al., 2024b; Lu et al., 2024) have explored using MLLMs for content moderation, but these methods lack the temporal grounding and reasoning capabilities required for precise violation detection. Our work bridges these gaps by introducing RAVEN, a curriculum reinforcement learning framework that integrates MLLMs with sophisticated reward mechanisms and structured reasoning for robust and precise advertisement video violation detection. By leveraging both precisely and coarsely annotated data, RAVEN addresses the limitations of existing methods and sets a new benchmark for temporal grounding in industrial applications.

3 Methodology

3.1 Problem Overview

Given an input video V , a predefined list of violation labels T , and a prompt P , the Advertisement Video Violation Temporal Grounding task aims to output: (1) The violation labels associated with the video. (2) The temporal intervals of the sub-scenes corresponding to each violation label. Note that a single video may contain multiple violation labels, and each label may correspond to multiple sub-scenes. This requires the model to perform reasoning to accurately identify the most relevant frame fragments. Inspired by recent advancements in the reasoning capabilities of large models, we leverage this ability to develop a pipeline for reasoning-based violative sub-scene temporal grounding.

We first employ reinforcement learning (RL) on a Multimodal Large Language Model (MLLM) to activate its reasoning ability, enabling it to generate a reasoning process and predict all violation categories $\mathcal{C} = \{c_1, c_2, \dots, c_n\}$ and their corresponding accurate sub-scene locations $\mathcal{X}_c = (t_c^l, t_c^r)$ for each category c . Here, t_c^l and t_c^r denote the start and end times of the sub-scene, respectively.

However, the manually annotated results $\mathcal{Y}_c = (y_c^l, y_c^r)$ often deviate from the ground truth $\mathcal{Z}_c = (z_c^l, z_c^r)$ due to annotation errors or ambiguities. To prevent supervised fine-tuning (SFT) from forcing the model to fit \mathcal{Y}_c , which could lead to significant deviations from \mathcal{Z}_c , we instead use RL for training. Additionally, to enhance the accuracy of the reasoning process, we follow DeepSeek (Dai et al., 2024) and employ explicit structured thinking tags ‘<think>’ for chained reasoning.

3.2 Data Construction

In real-world scenarios, for each advertisement video V , when a violation is found, we annotate the precise violation category c and the corresponding temporal sub-interval $\mathcal{Y}_c = (y_c^l, y_c^r)$ where the violation occurs. However, due to limitations in annotation resources, cost constraints, and inherent ambiguity in many videos, we can only maintain relatively accurate violation categories, while the annotated temporal intervals \mathcal{Y}_c often exhibit some degree of deviation from the ground truth $\mathcal{Z}_c = (z_c^l, z_c^r)$. To address this, we organize the data based on a curriculum learning approach. Specifically, we select a subset of data with precisely annotated temporal intervals for the early stages of curriculum learning, while the remaining

coarsely annotated data is used in the later stages. Additionally, it is important to note that for the reasoning training of RAVEN, we do NOT need to generate any offline reasoning data, meaning that RAVEN’s reasoning does not require a cold-start training process.

3.3 RAVEN Model

We use Qwen2.5-VL (Bai et al., 2023b) as the reasoning model F_{reason} in RAVEN. Although Qwen2.5-VL demonstrates some temporal grounding capabilities on public video understanding datasets, it struggles with accurate localization in real-world industrial applications. A straightforward approach would be to use precisely annotated temporal grounding data for SFT. However, acquiring large-scale, precisely annotated data is challenging and costly, especially for frame-level localization, which requires significant effort from annotators.

Instead, we opt for coarse-grained annotations, which are faster and more cost-effective to produce. During the reinforcement learning stage, format rewards are employed to ensure the model generates structured outputs. This process can be formulated as:

$$\mathcal{C}, \mathcal{X} = F_{\text{reason}}(V, T, P), \quad (1)$$

where \mathcal{C} represents the predicted violation categories, and \mathcal{X} denotes the corresponding temporal intervals.

Reasoning is a critical component in temporal grounding tasks. Inspired by DeepSeek-R1-Zero (Dai et al., 2024), we intentionally avoid using any explicit Chain-of-Thought (CoT) (Wei et al., 2022a) data to teach RAVEN reasoning skills. Instead, we aim to activate its reasoning capabilities from scratch, enabling the model to autonomously generate a logical CoT before producing the final answer. To achieve this, we design a structured user prompt and hierarchical sophisticated rewards that guides the reasoning model to follow specific instructions. As shown in Figure 1, the user prompt instructs RAVEN to analyze and compare objects in the video, beginning by generating a reasoning process within ‘<think>’ tags, followed by the final answer in a predefined format enclosed in ‘<answer>’ tags.

3.4 Reward Functions Design

Reward functions play a pivotal role in RL, as they determine the optimization direction of the model.

We manually design the following reward functions for RL:

3.4.1 Thinking Format Reward

The reward mechanism is designed to facilitate a structured cognitive process within the model (Shao et al., 2024; Guo et al., 2025). Specifically, it directs the model to articulate its reasoning steps within the designated <think> and </think> tags, while the final output is to be presented between the <answer> and </answer> tags.

3.4.2 Grounding Format Reward

Our framework incorporates two levels of temporal grounding format rewards: soft and strict (Shao et al., 2024; Guo et al., 2025). The soft approach validates the format if temporal coordinates are included in the answer, regardless of their organization. The strict approach, however, mandates that the model follows the predefined structure exactly, utilizing specific keywords like "temporal start" and "temporal end" to achieve correctness.

3.4.3 Temporal IoU Reward

As the primary reward, the Temporal IoU Reward evaluates the overlap between the predicted subscene intervals \mathcal{X}_c and the annotated intervals \mathcal{Y}_c . To maintain robustness against annotation noise, we binarize the IoU value using a threshold:

$$R_{\text{IoU}} = \begin{cases} 1 & \text{if } \text{IoU}(\mathcal{X}_c, \mathcal{Y}_c) > 0.5, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

3.4.4 Temporal Boundary Alignment Reward

Building on the IoU Reward, the Temporal Boundary Alignment Reward encourages the predicted interval boundaries (t_c^l, t_c^r) to align closely with the annotated boundaries (y_c^l, y_c^r) . This reward is continuous and serves as an auxiliary reward with a smaller weight:

$$R_{\text{Boundary}} = \exp\left(-\sigma^2 \left[(t_c^l - y_c^l)^2 + (t_c^r - y_c^r)^2 \right]\right), \quad (3)$$

where σ is a scaling factor.

3.4.5 Violation Category Consistency Reward

The Violation Category Consistency Reward ensures the predicted violation category c_p matches the annotated category c_g . This reward is binary:

$$R_{\text{Category}} = \begin{cases} 1 & \text{if } c_p = c_g, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

where c_p and c_g indicates the prediction and groundtruth respectively.

3.5 Curriculum Reasoning with Hierarchical Rewards

RAVEN does not require a cold-start reasoning training process. We initiate training directly from the pre-trained Qwen2.5-VL model, utilizing the aforementioned rewards and applying the GRPO (Shao et al., 2024) algorithm in the subsequent curriculum reinforcement training process.

We utilize the Curriculum GRPO with hierarchical rewards, which leverages a combination of precisely annotated and coarsely annotated data, progressively refining the model’s ability to predict both the temporal intervals and the associated violation categories. The training process is divided into three stages, each designed to optimize specific aspects of the model’s performance.

3.5.1 Stage 1: Training on Precisely Annotated Data

In the initial stage, the model is trained on a subset of data where the temporal intervals $\mathcal{Y}_c = (y_c^l, y_c^r)$ are precisely annotated. The reward function for this stage is designed to ensure the model learns the overall position of the interval while also improving boundary precision and category consistency. The total reward R_{Total} is defined as:

$$R_{\text{Total}} = R_{\text{IoU}} + \alpha_1 \cdot R_{\text{Boundary}} + R_{\text{Category}}, \quad (5)$$

where R_{IoU} measures the overlap between the predicted interval \mathcal{X}_c and the annotated interval \mathcal{Y}_c , binarized to ensure robustness against annotation noise. R_{Boundary} encourages precise alignment of the predicted boundaries (t_c^l, t_c^r) with the annotated boundaries (y_c^l, y_c^r) . R_{Category} ensures the predicted violation category c_p matches the annotated category c_g . α_1 is the reward weight. This stage focuses on establishing a strong foundation for interval prediction by prioritizing overall position (via R_{IoU}) while gradually refining boundary precision (via R_{Boundary}) and ensuring category consistency (via R_{Category}).

3.5.2 Stage 2: Training on the Large-Scale Coarsely Annotated Data

In the second stage, the model is trained on data where the temporal intervals are coarsely annotated. Here, the reward function is simplified to focus on overall position and boundary alignment, as the

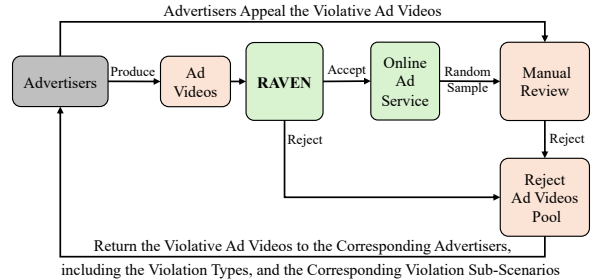


Figure 2: The deployment of RAVEN.

imprecise nature of the annotations makes category consistency less reliable. The total reward R_{Total} is defined as:

$$R_{\text{Total}} = R_{\text{IoU}} + \alpha_2 \cdot R_{\text{Boundary}}. \quad (6)$$

where α_2 is the reward weight. By retaining R_{IoU} and R_{Boundary} , the model learns to predict approximately correct intervals even with noisy annotations, while still improving boundary precision.

3.5.3 Stage 3: Fine-Tuning on Full Dataset

In the final stage, the model is fine-tuned on the full dataset, combining both precisely and coarsely annotated data. The reward function is adjusted to balance overall position, boundary precision, and category consistency:

$$R_{\text{Total}} = \alpha_3 \cdot R_{\text{IoU}} + \alpha_4 \cdot R_{\text{Boundary}} + \alpha_5 \cdot R_{\text{Category}}, \quad (7)$$

where α_3 , α_4 , and α_5 are the reward weights. This stage ensures the model achieves a robust balance between interval prediction and category identification, leveraging the strengths of both precise and coarse annotations.

4 Deployment

We design a pipeline to deploy the RAVEN on the online Ad services in Figure 2, which include 3 parts: (1) RAVEN Review: It is the core of the entire pipeline, handling the primary review functions. (2) Advertisers Appeal: It provides a channel for advertisers to appeal if they believe their ad is not violative. (3) Manual Review: It is primarily applied in two scenarios. (a) Random Sampling Review: For Ads already published on the platform, random samples are reviewed to identify potential violations. This helps to: (i) address cases missed by the review model, and (ii) quickly detect new types of violations, providing decision-making references for subsequent model optimization. (b)

Method	Marketing Exaggerate		Discomforting Content		Vulgar Content		Requiring Credential Review		Prohibited Goods/Services		Average	
	Cate.(P/R)	Gro.	Cate.(P/R)	Gro.	Cate.(P/R)	Gro.	Cate.(P/R)	Gro.	Cate.(P/R)	Gro.	Cate.(P/R)	Gro.
Small Models	0.681/0.532	-	0.707/0.679	-	0.667/0.654	-	0.711/0.687	-	0.721/0.734	-	0.697/0.657	-
LLaVA-v1.5-SFT	0.796/0.756	0.398	0.798/0.772	0.385	0.771/0.799	0.400	0.754/0.701	0.432	0.789/0.761	0.567	0.782/0.758	0.436
Qwen2.5-VL-7B-SFT	0.832/0.787	0.424	0.821/0.798	0.402	0.800/0.810	0.411	0.773/0.702	0.461	0.797/0.771	0.580	0.805/0.774	0.456
RAVEN	0.851/0.801	0.521	0.843/0.812	0.477	0.810/0.831	0.565	0.802/0.713	0.541	0.825/0.784	0.669	0.826/0.788	0.555

Table 1: Performance of Violation Category (Precision/Recall) and Violation Temporal Grounding (mIoU) on Industrial Dataset. ‘‘Cate.’’ indicates ‘‘Category’’, and ‘‘Gro.’’ indicates ‘‘Grounding’’.

Method	Average	
	Cate. (P/R)	Gro.
LLaVA-v1.5-SFT	0.509/0.501	0.370
Qwen2.5-VL-7B-SFT	0.537/0.517	0.384
RAVEN	0.551/0.530	0.435

Table 2: Performance of Violation Category (Precision/Recall) and Violation Temporal Grounding (mIoU) on Public MultiHateClip Dataset.

Appeal Review: For cases that are appealed by advertisers, manual review provides the final decision. (3) Model Iteration: Based on the continuously increasing volume and variety of online violation data, including (a) new types of violations, (b) more violation data, (c) difficult negative samples misidentified by the model, and (d) difficult positive samples missed by the model, we continuously iterate and optimize the RAVEN.

5 Experiments and Results

To comprehensively evaluate the performance of RAVEN, we conduct extensive experiments from both offline testing and online testing perspectives, utilizing both public dataset and practical industrial dataset.

5.1 Datasets

To validate RAVEN’s performance in real-world industrial scenarios, we construct a dataset comprising approximately 38,000 training videos, which include both precisely annotated and coarsely annotated data, and 5,000 precisely annotated test videos. The use of a precisely annotated test set ensures reliability in evaluation. The annotations cover six major violation categories (‘‘Discomforting Content’’, ‘‘Marketing Exaggeration’’, ‘‘Requiring Credential Review’’, ‘‘Vulgar Content’’, ‘‘Pro-

Model	Online Sample Average	
	Cate.(P/R)	Gro.
Small Models	0.711/0.668	-
Qwen2.5-VL-7B-SFT	0.800/0.787	0.478
RAVEN	0.821/0.803	0.563

Table 3: A/B Test on the Online Serving.

hibited Goods/Services’’, and ‘‘Normal’’) and the corresponding temporal intervals. The definitions of these major categories are inspired by both existing works (Wang et al., 2024b,a; Lu et al., 2023) and the actual platform management rules. These major classes are further divided into multiple sub-categories, forming a hierarchical and structured labeling system. In all experiments, we primarily focus on the major class labels to evaluate the model’s performance and robustness in high-level violation classification tasks.

MultiHateClip (Wang et al., 2024a) is a publicly available dataset for hateful and offensive content detection on platforms like YouTube and Bilibili, featuring annotations for ‘‘hateful’’, ‘‘offensive’’, and ‘‘normal’’ content. Due to the unavailability of some videos, we conduct experiments on a downloadable subset of Bilibili, and manually annotate the temporal intervals.

5.2 Offline Testing

We compare RAVEN against several baseline models, including LLaVA-v1.5 (Liu et al., 2023), Qwen2-VL-7B (Bai et al., 2023b), and Qwen2.5-VL-7B (Bai et al., 2023b), as well as their fine-tuned versions (SFT). The results in Table 1 and Table 2 demonstrate that RAVEN significantly outperforms both the base pretrained models and the SFT models in ‘‘violation category accuracy’’ and ‘‘temporal grounding precision’’. Specifically, RAVEN

Model	Average	
	Cate.(P/R)	Gro.
Qwen2.5-VL-7B-SFT	0.805/0.774	0.456
RAVEN(w/o Structured Thinking)	0.810/0.779	0.537
RAVEN	0.826/0.788	0.555

Table 4: Study on the Structured Thinking.

achieves superior accuracy in sub-scene interval localization, highlighting the effectiveness of its curriculum reinforcement learning approach in enhancing the robustness of MLLMs.

5.3 Online A/B Testing

We conduct day-long online A/B testing on a practical business platform, allocating 20% of the overall traffic for evaluation. RAVEN is compared against a small legacy model and Qwen2.5-VL-7B-SFT. The results in Table 3 show that RAVEN significantly improves violative video identification, achieving both higher precision and recall in category detection compared to the legacy model. Additionally, RAVEN outperforms the Qwen2.5-VL-7B-SFT model by 8.5% in temporal interval localization accuracy.

5.4 Study on Generalization Capabilities

As discussed in Section 1, SFT often leads to catastrophic forgetting of general capabilities, while RL enhances the generalization of MLLMs. To validate this claim, we conduct experiments on the Industrial dataset. Specifically, we train RAVEN on three in-domain categories (Discomforting Content, Marketing Exaggeration, Requiring Credential Review) and test it on the remaining two out-of-domain categories (Vulgar Content, Prohibited Goods/Services). The results in Table 5 demonstrate that RAVEN, trained with RL, achieves higher accuracy and better generalization compared to the Qwen2.5-VL SFT model.

5.5 Study on Structured Thinking

We further investigate the impact of reasoning training of structured thinking in RAVEN. Table 4 shows that both w/o and w/ structured thinking outperform the SFT baseline, indicating that RL effectively boosts the model’s capabilities. However, RAVEN with structured thinking demonstrates even better performance, highlighting the importance of the reasoning process in handling complex video samples.

Method	In-Domain (Average Gro.)	Out-of-Domain (Average Gro.)
	Qwen2.5-VL-7B-SFT	0.433
RAVEN	0.546	0.408

Table 5: Study on Generalization Capabilities.

Temporal Boundary Alignment Reward	Grounding Format Reward	Curriculum Reinforcement Learning	Gro.
✗	strict	✓	0.540
✓	soft	✓	0.547
✓	strict	✗	0.508
✓	strict	✓	0.555

Table 6: Study on Reward Functions and Curriculum Reinforcement Learning.

5.6 Study on Reward Functions

To validate the effectiveness of our reward function design, we conduct ablation studies on the format reward and temporal boundary alignment reward on the Industrial dataset. The results in Table 6 demonstrate the effectiveness of the two reward functions.

5.7 Study on Curriculum Reinforcement Learning

To evaluate the effectiveness of the curriculum reinforcement learning strategy in RAVEN, we also conduct an ablation study on the Industrial dataset. As shown in Table 6, when remove the progressive curriculum learning, the results shown in a significant drop in performance, with temporal interval localization (mIoU) dropping by 4.7%, highlighting the importance of leveraging multi-stage training for robust learning.

6 Conclusion

RAVEN is a novel framework for advertisement video violation detection, integrating curriculum reinforcement learning with multimodal large language models (MLLMs) to address challenges in temporal grounding and noisy annotations. Its progressive training strategy and hierarchical reward mechanism ensure precise localization and consistent category prediction. Experiments and online A/B testing demonstrate superior performance in accuracy, precision, and recall, while mitigating catastrophic forgetting. RAVEN establishes a promising methodological approach for practical violation detection, offering significant potential for advancing the field and addressing real-world challenges.

7 Ethical Statement

Our research adheres to ethical principles and prioritizes user rights. The dataset samples are for scientific analysis only and do not reflect the authors' views. All resources are intended for scientific research purposes only, contributing to the development of more secure and reliable digital platforms.

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023a. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023b. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Damai Dai, Chengqi Deng, Chenggang Zhao, RX Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y Wu, et al. 2024. Deepseek-moe: Towards ultimate expert specialization in mixture-of-experts language models. *arXiv preprint arXiv:2401.06066*.
- Alexey Dosovitskiy. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. 2017. Tall: Temporal activity localization via language query. In *Proceedings of the IEEE international conference on computer vision*, pages 5267–5275.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Deyi Ji, Haoran Wang, Mingyuan Tao, Jianqiang Huang, Xian-Sheng Hua, and Hongtao Lu. 2022. Structural and statistical texture knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16876–16885.
- Deyi Ji, Feng Zhao, Hongtao Lu, Mingyuan Tao, and Jieping Ye. 2023. Ultra-high resolution segmentation with ultra-rich context: A novel benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23621–23630.
- Deyi Ji, Feng Zhao, Hongtao Lu, Feng Wu, and Jieping Ye. 2025. Structural and statistical texture knowledge distillation and learning for segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Deyi Ji, Feng Zhao, Lanyun Zhu, Wenwei Jin, Hongtao Lu, and Jieping Ye. 2024a. Discrete latent perspective learning for segmentation and detection. In *International Conference on Machine Learning*, pages 21719–21730. PMLR.
- Deyi Ji, Lanyun Zhu, Siqi Gao, Peng Xu, Hongtao Lu, Jieping Ye, and Feng Zhao. 2024b. Tree-of-table: Unleashing the power of llms for enhanced large-scale table understanding. *arXiv preprint arXiv:2411.08516*.
- Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR.
- Shuang Li, Shuai Xiao, Shixiang Zhu, Nan Du, Yao Xie, and Le Song. 2018. Learning temporal point processes via reinforcement learning. *Advances in neural information processing systems*, 31.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916.

- Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J Liu, and Jialu Liu. Statistical rejection sampling improves preference optimization. In *The Twelfth International Conference on Learning Representations*.
- Tongxuan Liu, Wenjiang Xu, Weizhe Huang, Xingyu Wang, Jiaying Wang, Hailong Yang, and Jing Li. 2024. Logic-of-thought: Injecting logic into contexts for full reasoning in large language models. *arXiv preprint arXiv:2409.17539*.
- Yuqi Liu, Bohao Peng, Zhisheng Zhong, Zihao Yue, Fanbin Lu, Bei Yu, and Jiaya Jia. 2025. Seg-zero: Reasoning-chain guided segmentation via cognitive reinforcement. *arXiv preprint arXiv:2503.06520*.
- Junyu Lu, Bo Xu, Xiaokun Zhang, Changrong Min, Liang Yang, and Hongfei Lin. 2023. Facilitating fine-grained detection of chinese toxic language: Hierarchical taxonomy, resources, and benchmarks. *arXiv preprint arXiv:2305.04446*.
- Junyu Lu, Bo Xu, Xiaokun Zhang, WangHongbo, Hao-hao Zhu, Dongyu Zhang, Liang Yang, and Hongfei Lin. 2024. Towards comprehensive detection of chinese harmful memes. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. 2024. Video-chatgpt: Towards detailed video understanding via large vision and language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12585–12602.
- Krishanu Maity, Poornash Sangeetha, Sriparna Saha, and Pushpak Bhattacharyya. 2024. Toxvidlm: A multimodal framework for toxicity detection in code-mixed videos. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 11130–11142.
- Sanmit Narvekar, Bei Peng, Matteo Leonetti, Jivko Sinapov, Matthew E Taylor, and Peter Stone. 2020. Curriculum learning for reinforcement learning domains: A framework and survey. *Journal of Machine Learning Research*, 21(181):1–50.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Feifan Song, Bowen Yu, Minghao Li, Haiyang Yu, Fei Huang, Yongbin Li, and Houfeng Wang. 2024. Preference ranking optimization for human alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18990–18998.
- Han Wang, Tan Rui Yang, Usman Naseem, and Roy Ka-Wei Lee. 2024a. Multihateclip: A multilingual benchmark dataset for hateful video detection on youtube and bilibili. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 7493–7502.
- Weining Wang, Yan Huang, and Liang Wang. 2019a. Language-driven temporal activity localization: A semantic matching reinforcement learning model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 334–343.
- Weining Wang, Yan Huang, and Liang Wang. 2019b. Language-driven temporal activity localization: A semantic matching reinforcement learning model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 334–343.
- Zhenting Wang, Shuming Hu, Shiyu Zhao, Xiaowen Lin, Felix Juefei-Xu, Zhuowei Li, Ligong Han, Harihar Subramanyam, Li Chen, Jianfa Chen, et al. 2024b. Mllm-as-a-judge for image safety without human labeling. *arXiv preprint arXiv:2501.00192*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022a. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, Yang Wang, and Enhong Chen. 2024a. Large language models for generative information extraction: A survey. *Frontiers of Computer Science*, 18(6):186357.
- Guowei Xu, Peng Jin, Li Hao, Yibing Song, Lichao Sun, and Li Yuan. 2024b. Llava-o1: Let vision language models reason step-by-step. *arXiv preprint arXiv:2411.10440*.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. 2023. A survey on multimodal large language models. *arXiv preprint arXiv:2306.13549*.
- Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. 2020a. Span-based localizing network for natural language video localization. In *Proceedings of the 58th*

Annual Meeting of the Association for Computational Linguistics, pages 6543–6554.

Songyang Zhang, Houwen Peng, Jianlong Fu, and Jiebo Luo. 2020b. Learning 2d temporal adjacent networks for moment localization with natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 12870–12877.

Lanyun Zhu, Tianrun Chen, Deyi Ji, Jieping Ye, and Jun Liu. 2024a. Llafs: When large language models meet few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3065–3075.

Lanyun Zhu, Deyi Ji, Tianrun Chen, Peng Xu, Jieping Ye, and Jun Liu. 2024b. Ibd: Alleviating hallucinations in large vision-language models via image-biased decoding. *arXiv preprint arXiv:2402.18476*.