

Following Clues, Approaching the Truth: Explainable Micro-Video Rumor Detection via Chain-of-Thought Reasoning

Rongpei Hong
rongpei.hong@std.uestc.edu.cn
University of Electronic Science and
Technology of China
Chengdu, Sichuan, China

Jian Lang
jian_lang@std.uestc.edu.cn
University of Electronic Science and
Technology of China
Chengdu, Sichuan, China

Jin Xu
jin.xu@mu.ie
Maynooth University
County Kildare, Maynooth, Ireland

Zhangtao Cheng
zhangtao.cheng@outlook.com
University of Electronic Science and
Technology of China
Chengdu, Sichuan, China

Ting Zhong*
zhongting@uestc.edu.cn
University of Electronic Science and
Technology of China
Chengdu, Sichuan, China

Fan Zhou
fan.zhou@uestc.edu.cn
University of Electronic Science and
Technology of China
Chengdu, Sichuan, China
Kash Institute of Electronics and
Information Industry
Kashgar, Xinjiang, China

Abstract

The rapid spread of rumor content on online micro-video platforms poses significant threats to public health and safety. However, existing Micro-Video Rumor Detection (MVRD) methods are generally black-box, which lacks transparency and makes it difficult to understand the reasoning behind classification decisions. In this work, we introduce **ExMRD**, a novel **Explainable Micro-video Rumor Detection** framework designed to generate detailed and coherent explanations for enhancing MVRD. Inspired by the powerful reasoning capacity of Chain-of-Thought (CoT), we introduce a novel inference mechanism called R^3CoT —consisting of *Refining*, *Retrieving*, and *Reasoning* on MVRD. This mechanism enables Multimodal Large Language Models (MLLMs) to reorganize the original video content, retrieve domain knowledge related to rumors, and generate explainable conclusions regarding whether the micro-video contains rumor information. Instead of directly fine-tuning MLLMs for MVRD, which is computationally expensive, we propose a Small Language Reviewer (SLReviewer), which distills the outputs of R^3CoT guided MLLMs to ensure efficient and reliable predictions. Extensive experiments on three real-world benchmarks demonstrate that ExMRD significantly outperforms competitive baselines while providing high-quality rationales.

CCS Concepts

• **Information systems** → **Multimedia information systems**; • **Computing methodologies** → **Natural language processing**.

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WWW '25, Sydney, NSW, Australia.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1274-6/25/04
<https://doi.org/10.1145/3696410.3714559>

Keywords

Micro-video rumor detection, explainability, chain-of-thought

ACM Reference Format:

Rongpei Hong, Jian Lang, Jin Xu, Zhangtao Cheng, Ting Zhong, and Fan Zhou. 2025. Following Clues, Approaching the Truth: Explainable Micro-Video Rumor Detection via Chain-of-Thought Reasoning. In *Proceedings of the ACM Web Conference 2025 (WWW '25)*, April 28–May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3696410.3714559>

1 Introduction

The exponential growth of online micro-video platforms such as TikTok, YouTube Shorts, and Snapchat has revolutionized information consumption worldwide [1, 21, 22, 34, 59]. With billions of active users, these platforms enable rapid creation and dissemination of micro-videos, offering unprecedented speed and reach in information sharing. However, this convenience comes with the proliferation of misinformation and rumors, which often evade scrutiny and fact-checking [2, 9, 10, 26, 54, 56]. A striking example, as shown in Figure 1, occurred during the COVID-19 pandemic when a TikTok video falsely claimed that injecting disinfectants could “kill” the virus. This misleading content amassed millions of views and resulted in a spike in accidental poisonings, highlighting the tangible harm caused by misinformation on micro-video platforms, and underscoring the urgent need for effective methods to detect and address rumors in micro-videos, a task *a.k.a.* Micro-Video Rumor Detection (MVRD).

Existing MVRD approaches primarily focus on utilizing multiple modalities – such as text, audio, video content, and social context – to improve detection accuracy [6, 17, 37, 38, 42, 43]. For example, FakingRec [6] analyzed the process of rumor creation by examining material selection and editing behaviors on micro-video platforms, while NEED [38] leveraged relationships between videos related to the same event to improve rumor detection. Despite these advancements, current approaches often perform black-box detection, oversimplifying or overlooking the critical reasoning needed to

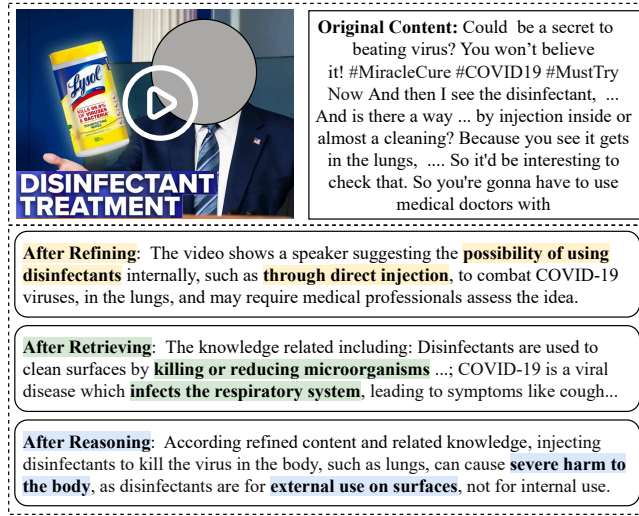


Figure 1: In this micro-video rumor, viewers are misled into believing that injecting disinfectants can kill the COVID-19 virus, leading to a rise in accidental poisoning incidents. Text in video: Disinfectant Treatment.

provide explainable justifications for the final prediction. This lack of transparency makes it difficult for viewers to understand why a video is classified as a rumor, undermining trust and limiting the effectiveness of rumor mitigation strategies. In contrast, an explainable model is essential to enhance the effectiveness and trustworthiness of MVRD systems. Both users and platforms need to comprehend the specific factors that lead to a video being identified as a rumor or genuine content. However, developing such a model presents several critical challenges:

C1: Inconsistent Video Quality and Misleading Metadata.

Micro-video platforms often contain content with inconsistent video quality and misleading metadata, which pose significant obstacles for interpretable rumor detection. Poor visuals and audio due to the varied skills and equipment of content creators make it difficult for detection models to capture critical cues. Pre-trained models optimized for high-quality data often perform poorly when applied to low-quality inputs, resulting in frequent misclassification [15, 33, 45]. For instance, the video tags in Figure 1 (e.g., #MiracleCure, #MustTry) represent low-quality text content. These tags lack informative value and do not contribute meaningful content for the model to analyze. This issue arises because many video creators usually employ misleading titles and tags to attract attention, which may distort the model's interpretation of the actual content. These inconsistencies pose significant challenges to current methods, further undermining the accuracy and trustworthiness of rumor detection models.

C2: Lack of Domain Knowledge and Reasoning. Effective rumor detection requires domain knowledge and logical reasoning to interpret complex content accurately. For example, identifying the micro-video in Figure 1 as a rumor requires a basic understanding of COVID-19 biology and the proper use of disinfectants, which are intended for surface cleaning rather than for internal use. With this knowledge, and through logical reasoning, it becomes clear that

suggesting disinfectants as a COVID-19 treatment is not only scientifically inaccurate but also extremely harmful to the human body. Current detection models often lack specialized domain knowledge and the capacity for such reasoning, limiting their ability to provide meaningfully and contextually accurate explanations for their predictions.

To address these challenges, we draw inspiration from the process of writing debunking articles, which involves gradually uncovering the truth by understanding the content, gathering relevant domain information, and combining these insights to debunk the rumor effectively. To this end, we present **ExMRD**, a novel **Explainable Micro-video Rumor Detection** framework, which can provide clear and well-reasoned explanations for MVRD. Specifically, we introduce R^3CoT , a novel Chain-of-Thought (CoT) [52] inference mechanism. R^3CoT consists of three key steps: *Refining*, *Retrieving*, and *Reasoning*. Specifically, at the refining step, the Multimodal Large Language Model (MLLM) is prompted to reorganize low-quality and misleading video content from both textual and visual perspectives, producing a coherent representation of the content in the video. At the retrieving step, rumor-related domain knowledge is generated based on the refined content, enriching the video's context for rumor detection. At the final reasoning step, logical inference is applied by cross-verifying the refined content with domain knowledge, providing evidence to support or refute the video's authenticity. Figure 1 demonstrates the outputs produced after each step of the R^3CoT mechanism, specifically showing the refined content, the retrieved domain knowledge, and the reasoning behind the final conclusion.

While fine-tuning MLLMs with R^3CoT guidance can greatly enhance their performance in MVRD, this fine-tuning process introduces significant computational overhead, which significantly limits its practicality in real-world applications. To address this, we introduce a novel *Small Language Reviewer* (SLReviewer), which acts as a reliable “reviewer” within ExMRD. The main idea is to refine MLLM outputs using distilled knowledge from the proposed R^3CoT mechanism, ensuring reliable predictions with low computational overhead. By fine-tuning SLMs, rather than MLLMs, we achieve a balance between performance and efficiency, making ExMRD more practical for deployment.

Our main contributions are summarized as follows:

- **An explainable MVRD framework** that generates explicit rationales behind rumor detection. This work is the first to incorporate explainability into MVRD by designing the reasoning step of R^3CoT , which produces clear rationales for determining whether a video contains rumor content. This enhances transparency and interpretability, making the decision-making process in MVRD more accessible to users and moderators.
- **A novel R^3CoT mechanism** that enables MLLMs to perform refining, retrieving, and reasoning for explainable conclusions. This mechanism helps address rumor detection errors caused by inconsistent video quality, and lack of domain knowledge and reasoning, providing a comprehensive understanding of video content in MVRD.
- **An efficient SLReviewer** that improves both efficiency and prediction reliability while balancing performance with computational resources. SLReviewer is fine-tuned using insightful

distilled knowledge from the MLLM, enabling it to generate credible and accurate predictions.

Extensive experiments on real-world micro-video datasets demonstrate that our ExMRD outperforms state-of-the-art baselines while providing clear and well-reasoned rationales. Notably, our ExMRD achieves an average improvement of 5.37% in Macro F1 across all three datasets, outperforming 13 competitive baselines. The code and data to reproduce the results are available at <https://github.com/ronpay/ExMRD>.

2 Related Work

2.1 Micro-Video Rumor Detection

The task of MVRD focuses on identifying rumor content by analyzing multiple modalities within micro-videos, such as text, vision, and audio. Early detection methods primarily relied on unimodal information [25, 36, 42]. For example, Papadopoulou et al. [36] utilized basic video metadata and user engagement features, while Serrano et al. [42] analyzed user comments, focusing on conspiracy-related remarks as key indicators. The complexity and richness of micro-video content – where various modalities often interplay – make unimodal approaches inadequate for accurate rumor detection. Recently, FakeSV [37] improves content representation by leveraging cross-modal correlations and integrating social context. Despite these advancements, current methods in MVRD remain black-box models, providing only the final detection result without offering any interpretability for the decision-making process, which increases a lot of risks and limits user trust. In contrast, our study aims to provide clearer explanations and enhance transparency for both moderators and viewers, addressing the critical need for explainability in micro-video rumor detection. – i.e., understand, trust, and act upon the detection outcomes.

2.2 Chain-of-Thought Prompting

CoT prompting [52] has been developed to guide Large Language Models (LLMs), such as GPT-3 [5], LLaMA [47], to better understand tasks and generate better responses by encouraging step-by-step reasoning. Moreover, the recent OpenAI-o1 [57] has demonstrated remarkable reasoning capabilities by deeply integrating CoT into its architecture, enabling the model to seamlessly perform multi-step reasoning. Specifically, various CoT strategies [31, 51, 58] have been introduced to enhance the reasoning abilities of LLMs. With the advancement of MLLM [27, 35], CoT prompting has been adapted to improve reasoning across both visual and textual modalities [8, 32, 41]. This adaptation enhances MLLM’s ability to synthesize information from multiple modalities and align their representations, leading to improved performance on multimodal tasks. However, existing works primarily leverage CoT for generating better responses in scenarios like visual question answering [30, 44, 53] and mathematical reasoning [16, 19, 46], overlooking the potential of CoT to enhance task-specific explainability. In contrast, our ExMRD incorporates a CoT based inference mechanism, R³CoT, to guide MLLMs to generate accurate, understandable, and interpretable predictions via the designed three key steps: *Refining*, *Retrieving*, and *Reasoning*. To our knowledge, this work is among the first to leverage the reasoning capabilities of CoT to prompt MLLMs in delivering precise and interpretable predictions for MVRD.

3 Methodology

Problem Definition. Let \mathcal{M} represent a micro-video on video platforms. The video \mathcal{M} is characterized by its multimodal content, which includes textual, visual, and audio modalities, denoted as $\mathcal{M} = \{\mathcal{T}, \mathcal{V}, \mathcal{A}\}$. The primary objective of MVRD is to determine whether the video \mathcal{M} contains the *rumor* content by simultaneously considering all its modalities \mathcal{T} , \mathcal{V} , and \mathcal{A} .

Overview. First, we design the R³CoT, a three-step CoT inference mechanism: (1) *Refining*. Prompt the MLLMs to reorganize chaotic and misleading video content into a well-structured format, serving as the foundation for subsequent generation; (2) *Retrieving*. Instruct the MLLMs to generate relevant domain knowledge based on the refined video content to support the final reasoning; (3) *Reasoning*. Guide the MLLMs to apply logical inference to cross-verify the refined content with domain knowledge, providing evidence to confirm or refute the authenticity of the micro-video. Second, we introduce the SLReviewer, to distill the reasoning output from the MLLMs, ensuring reliable final predictions. An overview of our proposed ExMRD is shown in Figure 2.

3.1 R³CoT Mechanism

3.1.1 Feature extraction. For the textual modality \mathcal{T} , we extract the video’s metadata (title and description) \mathcal{T}_m and the on-screen text \mathcal{T}_o . For visual modality \mathcal{V} , we uniformly sample k frames from each video: $\mathcal{V}_f = \{v_1, \dots, v_k\}$. For audio modality \mathcal{A} , we convert the audio into transcript denoted as \mathcal{A}_t . The detailed process of feature extraction is summarized in Appendix A.3.

3.1.2 Refining. In the domain of rumor video detection, a diverse range of micro-video creators exists, spanning from professional media outlets to ordinary individual users, leading to micro-videos of varying quality. Micro-videos frequently contain chaotic and misleading content, making detection challenging, thus requiring the reorganization of the presented information. To address these problems, the initial step in the R³CoT mechanism involves refining the video content to generate well-structured representations by reorganizing both textual and visual elements for clearer analysis.

First, from the textual perspective, we combine the metadata text \mathcal{T}_m , on-screen text \mathcal{T}_o , and audio transcript \mathcal{A}_t as the textual content of the video. However, due to elaborate video effects, diverse typographic styles, and non-news elements (e.g., watermarks, creator attributions, and platform identifiers), the on-screen text often contains numerous recognition errors. To mitigate this issue, we empower the MLLM $\mathcal{F}(\cdot)$ to focus on rumor elements while enhancing on-screen accuracy through in-context learning. We restore the original rumor content from a textual perspective, resulting in refined textual content $\mathcal{R}_{\text{text}}$. This step can be formulated as follows:

$$\mathcal{R}_{\text{text}} = \mathcal{F}([\mathcal{T}_m; \mathcal{T}_o; \mathcal{A}_t], \text{Prompt}_1), \quad (1)$$

where $[\cdot]$ denotes the concatenation operation, with the template Prompt_1 described in Step 1 of Figure 3.

Second, from the visual perspective, we develop a visual-centric strategy that focuses on the refining of scene content in the micro-video. This strategy aims to filter out subjective elements, such as subtitles or auditory narratives, focusing solely on the visual information presented in the videos. By prioritizing these visual

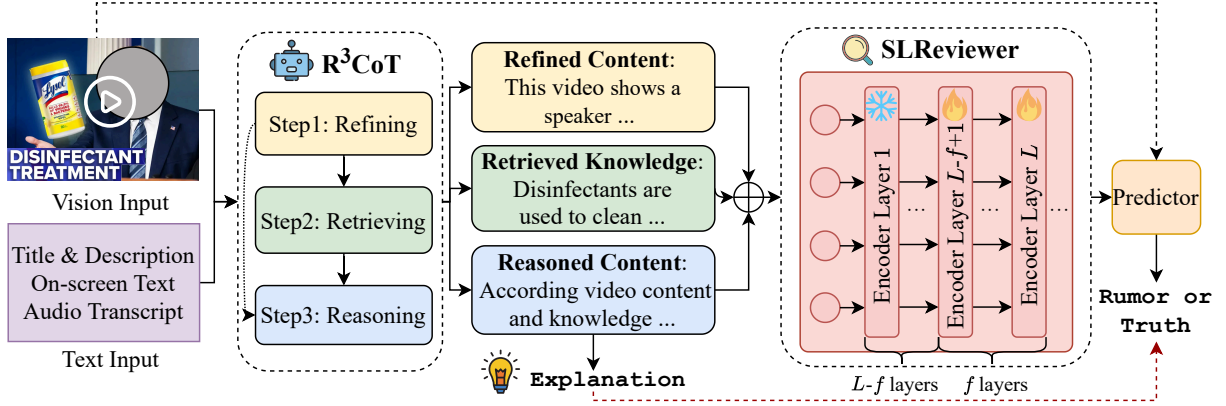


Figure 2: The structure of our proposed ExMRD framework. (1) The R^3CoT process prompts MLLMs to refine the video content, retrieve domain knowledge, and reason to provide explanations. (2) The SLReviewer is to distill the explainable evidence from R^3CoT to facilitate reliable rumor detection.

aspects, we seek to provide a more objective and comprehensive understanding of the events depicted, offering insights that may not be explicitly conveyed in the textual content.

Instead of feeding individual frames into the MLLM, we propose to construct n composite frames $\mathcal{P}_v = \{P_1, P_2, \dots, P_n\}$. Each composite frame P_i consists of an $m \times m$ grid of consecutive frames from the initial video frames \mathcal{V}_f : $P_i = [v_{i_1}, v_{i_2}, \dots, v_{i_{m \times m}}]$. This allows the model to better grasp temporal changes and scene dynamics. The generation process of refined visual content \mathcal{R}_{vision} is guided by the template prompt₂ presented in Step 1 of Figure 3,

$$\mathcal{R}_{vision} = \mathcal{F}(\mathcal{P}_v, \text{Prompt}_2). \quad (2)$$

The refined textual and visual content will be used as the input information for the next step to generate the domain knowledge for the micro-video.

3.1.3 Retrieving. This step aims to instruct the MLLM to generate expressive and relevant domain knowledge for the given micro-video based on the refined content. The retrieved domain knowledge enables the MLLM to better understand the rumors presented in the micro-video, thus facilitating effective reasoning. Specifically, let $\mathcal{R}_{refining} = [\mathcal{R}_{text}; \mathcal{R}_{vision}]$ represent the refined video content. We guide the domain knowledge retrieval process using the template Prompt₃ in Step 2 of Figure 3. This step can be formulated as follows:

$$\mathcal{R}_{retrieving} = \mathcal{F}(\mathcal{R}_{refining}, \text{Prompt}_3). \quad (3)$$

The expressive domain knowledge and the refined content will jointly help the MLLM to make the final reasoning in the next step.

3.1.4 Reasoning. This step utilizes logical inference to cross-verify the refined content with the domain knowledge from the previous steps, aiming to uncover evidence that either confirms or refutes the authenticity of the video. Furthermore, this process enhances the explainability of the model by providing clear evidence (e.g., conflicting or consistent facts) to explain why content is classified as a rumor, thereby providing transparency into the model's decisions. We guide the logical reasoning using the template Prompt₄ in Step

3 of Figure 3. This step can be formulated as follows:

$$\mathcal{R}_{reasoning} = \mathcal{F}([\mathcal{R}_{refining}; \mathcal{R}_{retrieving}], \text{Prompt}_4). \quad (4)$$

By leveraging the logically rigorous and coherent R^3CoT mechanism, our framework **addresses the limitations** of prior works [6, 7, 37, 43], which typically overlook the chaotic and low-quality nature of video content and lack the necessary rumor-specific background knowledge and deep reasoning required for accurate predictions. **In contrast**, our R^3CoT mechanism equips ExMRD with deeper insights through refining the raw video content, retrieving the domain knowledge, and reasoning the deep rationales, ultimately achieving notable improvements in prediction accuracy.

3.2 Small Language Reviewer

Although the MLLM can make predictions through our carefully designed three-step CoT process, its inference remains unreliable due to the inherent limitations of large language models [3, 18], as the predictions may suffer from hallucinations and not be faithful to the video content and reasoning process. While fine-tuning the MLLM for MVRD could help mitigate this issue, it is not practical due to the vast number of parameters of the MLLM and the associated huge computational cost.

To this end, we further propose the Small Language Reviewer (SLReviewer), which distills the outputs from the MLLM into a smaller language model. By fine-tuning the smaller language model, a more reliable and practical solution is achieved, as it requires significantly fewer parameters and is computationally efficient. The fine-tuning process of SLM can be represented as:

$$S(\mathbf{x}) = \mathcal{L}_L \circ \mathcal{L}_{L-1} \circ \dots \circ \mathcal{L}_{L-f+1} \circ \mathcal{L}_{L-f}^{\text{fixed}} \circ \dots \circ \mathcal{L}_1^{\text{fixed}}(\mathbf{x}), \quad (5)$$

where \mathcal{L}_i represents the i -th transformer encoder layer. The parameters of the last f layers, denoted as \mathcal{L}_{L-f+1} to \mathcal{L}_L , are fine-tuned during training. Meanwhile, the remaining layers are frozen, enabling the SLReviewer to retain its ability to review the output produced by the MLLM. The refined content $\mathcal{R}_{refining}$, domain knowledge $\mathcal{R}_{retrieving}$, and reasoning $\mathcal{R}_{reasoning}$ are concatenated

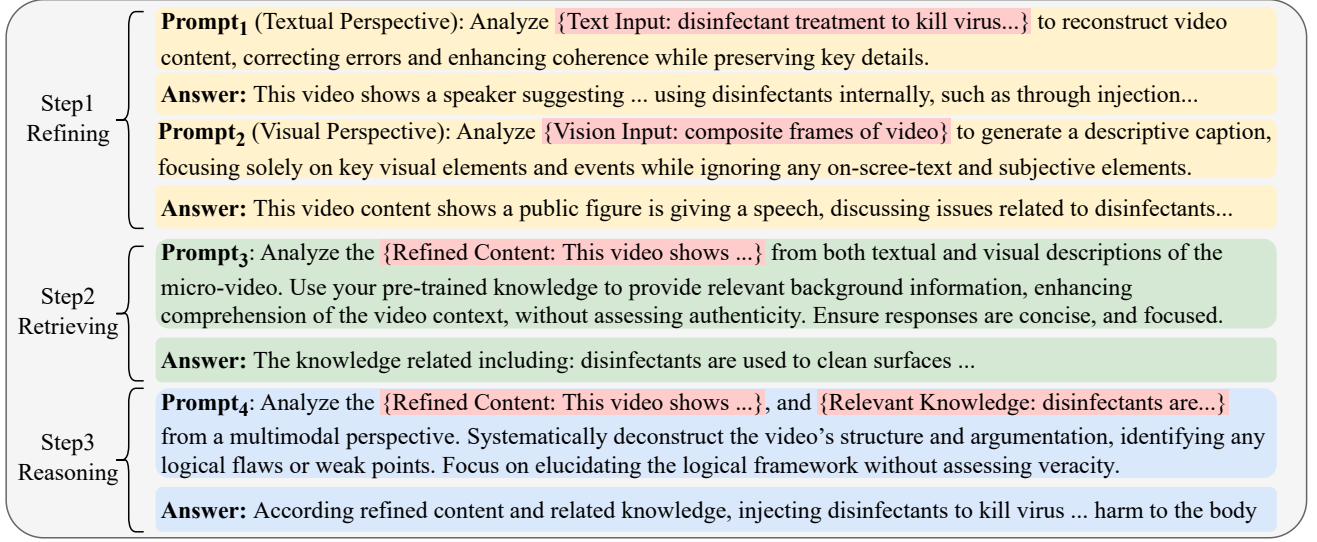


Figure 3: An illustration of three-step prompting of our proposed R³CoT. Step 1, prompt MLLM to organize chaotic and misleading video input; Step 2, instruct MLLM to generate rumor-relevant knowledge; Step 3, guide MLLM to conduct inference to cross-verify the refined and retrieved content.

and fed into the SLReviewer to generate the textual feature representation $\mathbf{H}_t \in \mathbb{R}^{l \times d_t}$ for the final prediction: This process can be written as:

$$\mathbf{H}_t = \mathcal{S}[\mathcal{R}_{\text{refining}}; \mathcal{R}_{\text{retrieving}}; \mathcal{R}_{\text{reasoning}}]. \quad (6)$$

We employ the pre-trained Vision Transformer (ViT) [14] to extract visual features $\mathbf{H}_v \in \mathbb{R}^{k \times d_v}$ from the video frames \mathcal{V}_f . To align the sequence lengths of the textual features \mathbf{H}_t and the visual features \mathbf{H}_v , we apply an average pooling strategy, resulting in $\bar{\mathbf{H}}_t \in \mathbb{R}^{d_t}$ and $\bar{\mathbf{H}}_v \in \mathbb{R}^{d_v}$. Subsequently, a two-layer Multilayer Perceptron (MLP) is used as a predictor to fuse the pooled textual and visual features, producing the final prediction: $\hat{y} = \text{Predictor}[\Psi_t(\bar{\mathbf{H}}_t); \Psi_v(\bar{\mathbf{H}}_v)]$, where $\Psi_t(\cdot)$ and $\Psi_v(\cdot)$ denote the linear mapping functions.

For model training, we optimize only the parameters of SLReviewer using Binary Cross-Entropy loss.

4 Experiments

An overview of the experimental setup is outlined below, with detailed descriptions of the datasets, baseline models, and implementation available in Appendix A.

- **Datasets.** To analyze the effectiveness of our ExMRD, we conduct experiments on three real-world micro-video datasets: FakeSV [37], FakeTT [6], and FVC [36]. Table 1 summarizes the detailed statistics of three datasets. Following existing works [6, 37], we employ a temporal split strategy to simulate real-world scenarios on micro-video platforms. In this strategy, we divide each dataset chronologically into training, validation, and test sets, with respective ratios of 70%, 15%, and 15%.

- **Baselines.** To verify the superiority of ExMRD, we compare it against 13 competitive baselines, which can be categorized into three groups: (1) *Unimodal detection methods* which utilize single modality (e.g., textual modality) of micro-videos to conduct

Table 1: Statistics of three datasets.

Dataset	Time Range	# Rumor	# Truth	# Total	Duration (s)
FakeSV	2017/10-2022/02	1,810	1,814	3,624	39.88
FakeTT	2019/05-2024/03	1,172	819	1,991	47.69
FVC	2016/01-2018/01	1,633	1,131	2,764	87.83

detection: BERT [13], ViT [14], MFCC [12] and TSformer [4]; (2) *Multimodal detection methods* which incorporate all modalities to improve the precision in detecting rumors in micro-videos: Tik-Tec [43], FANVM [11], CAFE [7], HMCAN [39], SV-FEND [37] and FakingRec [6]; (3) *MLLM based methods* which employ the latest released advanced MLLMs to detect the video rumor: GPT-4o-mini [35], LLaVA-OV [23], Qwen2-VL [50].

- **Model Implementation.** We employ GPT-4o-mini [35] as the MLLM backbone because it is scalable and easy to deploy. Additionally, BERT [13] is adopted as the SLM backbone due to its robust contextual understanding and performance in many natural language processing tasks.

- **Metrics.** Following prior works [6, 37], we employ four metrics to evaluate the performance: Accuracy (**ACC**), Macro F1 score (**M-F1**), Macro Precision (**M-P**), and Macro Recall (**M-R**).

4.1 Overall Performance

To assess the effectiveness of our ExMRD, we compare ExMRD with 13 competitive baselines. The results are summarized in Table 2. From the results, we have the following observations.

First, ExMRD consistently outperforms all baselines across various metrics on three datasets, showing an average improvement of 4.99% in Accuracy and 5.37% in Macro F1. To further verify its effectiveness, ExMRD and the strongest baseline are retrained five times, with the resulting p -values, all below 0.05, confirming the statistical significance of ExMRD's improvement. These gains are attributed

Table 2: Performance comparison on three real-world datasets. The best results are in highlighted red bold, while the second results are in black bold. Higher values of Accuracy, Macro F1, Macro Precision, and Macro Recall signify better performance.

Dataset	FakeSV				FakeTT				FVC			
Model	ACC	M-F1	M-P	M-R	ACC	M-F1	M-P	M-R	ACC	M-F1	M-P	M-R
BERT	80.63	80.14	80.56	79.90	71.24	69.31	68.98	70.85	69.29	68.13	68.72	67.95
ViT	71.22	71.04	71.04	71.33	65.55	64.39	65.17	67.11	81.54	80.74	82.03	80.25
MFCC	61.07	61.05	61.64	61.74	52.51	52.23	64.26	62.21	65.05	60.79	65.75	61.68
TSformer	72.14	71.95	71.91	72.20	64.88	64.69	68.79	70.43	90.92	90.67	91.06	90.40
TikTec	73.43	73.26	73.23	73.54	66.22	65.08	65.84	67.87	74.60	74.54	74.63	74.52
FANVM	79.52	78.81	79.81	78.46	71.57	70.21	70.22	72.63	79.27	77.41	82.47	76.77
CAFE	71.03	71.00	71.41	71.67	69.57	67.91	67.83	69.85	83.59	83.12	83.76	82.79
HMCAN	79.52	78.81	79.81	78.46	68.56	68.41	72.78	74.72	85.62	84.99	86.48	84.40
SV-FEND	80.88	80.54	80.18	80.62	77.14	75.63	75.12	77.56	87.59	87.36	87.34	87.40
FakingRec	84.69	84.30	84.80	84.01	79.60	77.76	77.12	78.88	90.92	90.78	90.65	90.95
GPT-4o-m	66.42	65.88	65.90	65.87	57.85	57.78	62.91	63.65	66.11	65.51	65.49	65.54
LLaVA-OV	57.54	50.71	61.57	55.94	46.82	46.81	53.44	53.36	60.21	56.55	58.91	57.30
Qwen2-VL	53.85	53.72	54.29	54.20	53.18	52.95	56.77	57.35	59.15	58.05	58.15	58.02
ExMRD	86.90	86.52	87.31	86.13	84.28	83.13	82.27	85.19	96.82	96.75	97.02	96.75
Improv.	2.61% \uparrow	2.63% \uparrow	2.96% \uparrow	2.52% \uparrow	5.88% \uparrow	6.91% \uparrow	6.68% \uparrow	8.00% \uparrow	6.49% \uparrow	6.58% \uparrow	7.03% \uparrow	6.38% \uparrow
p-val.	$2.26e^{-3}$	$2.33e^{-3}$	$3.17e^{-3}$	$2.42e^{-3}$	$1.64e^{-2}$	$9.36e^{-3}$	$5.57e^{-3}$	$2.82e^{-3}$	$4.94e^{-4}$	$5.62e^{-4}$	$3.18e^{-4}$	$1.07e^{-3}$

Table 3: Ablation study on key components of ExMRD.

Dataset	FakeSV		FakeTT		FVC	
Variant	ACC	M-F1	ACC	M-F1	ACC	M-F1
w/o Refine	82.10	81.81	83.28	82.22	94.25	94.90
w/o Retrieve	85.05	84.60	81.61	80.47	93.34	93.11
w/o Reason	85.05	84.42	79.60	78.62	95.31	95.20
w/o R ³ CoT	80.07	79.72	80.60	79.58	93.65	93.47
w/o Fine-tune	84.87	84.27	81.61	80.47	92.59	92.33
MLP-based	85.42	84.86	78.93	77.63	95.31	95.20
ExMRD	86.90	86.52	84.28	83.13	96.82	96.75

to ExMRD’s innovative R³CoT mechanism, which refines content, retrieves domain knowledge, and applies reasoning. Furthermore, SLReviewer efficiently utilizes the distilled knowledge from MLLMs, yielding precise predictions with minimal computational overhead.

Second, the unimodal methods show significantly lower performance, highlighting the importance of multimodal information in MVRD. Among them, BERT, which leverages the textual information for prediction, outperforms other methods in this group. This indicates that textual modality contains more semantic information and is more conducive to rumor detection. Our ExMRD framework leverages extensive textual data from multiple sources (title, description, on-screen text, and audio transcript), with the R³CoT refining step enhancing the text quality for improved predictions.

Third, multimodal methods generally outperform unimodal approaches, demonstrating the benefit of combining text and visuals in MVRD. FakingRec, for example, achieves strong results by focusing on the content creation process, leading to a robust multimodal understanding. However, ExMRD surpasses all multimodal baselines by refining video content to tackle the low-quality inputs. Additionally, accurately classifying a video as rumor or truth requires domain knowledge and logical reasoning, aspects that other models overlook.

Fourth, MLLM-based methods excel in zero-shot multimodal tasks but their lack of task-specific adaptation often leads to inconsistent performance in MVRD, making them insufficient for the detection. Direct fine-tuning MLLMs is computationally expensive, limiting practicality. In our ExMRD, we propose SLReviewer to refine MLLM outputs, providing more reliable predictions. Consequently, ExMRD achieves superior performance in MVRD compared to conventional MLLM-based methods.

4.2 Ablation Study

We conduct experiments to explore the impact of each main component in ExMRD, with the results summarized in Table 3.

4.2.1 Effect of R³CoT Mechanism. To assess the impact of the R³CoT mechanism, each step of R³CoT is removed individually to evaluate its effect. Specifically, the following ablation studies are conducted: (1) **w/o Refine**, where the refined output is replaced with the original video content; (2) **w/o Retrieve**, where the domain knowledge is excluded; (3) **w/o Reason**, where the reasoning output is removed; and (4) **w/o R³CoT**, where the distilled knowledge from MLLM is replaced as the original video content. The results indicate that each steps play pivotal roles in detecting rumor in micro-videos and provide insightful rationales with the prediction result. Moreover, a substantial performance drop is observed when the entire R³CoT mechanism is removed, further validating the critical role of the distilled knowledge obtained from R³CoT guided MLLM in MVRD.

4.2.2 Effect of SLReviewer. To assess the effectiveness of the SLReviewer, two ablation variants are developed: (1) **w/o Fine-tune**, where no fine-tuning is applied to the SLM; (2) **MLP-based Reviewer**, where a one-layer MLP is attached to the last layer of the frozen SLM. The results indicate the necessity of fine-tuning SLReviewer to adapt to the distilled outputs produced by the MLLM, as freezing all the layers leads to significantly degraded performance.




Figure 4: Generalizability analysis on adding our R³CoT and SLReviewer to base MLLM.

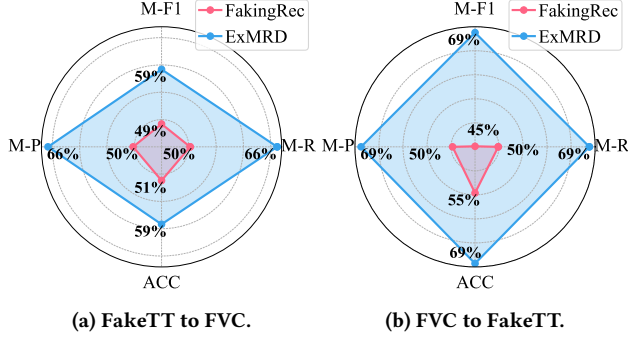


Figure 5: Generalizability analysis through cross-dataset experiments on the FakeTT and FVC datasets.

Moreover, the MLP-based reviewer struggles to capture the intricate reasoning patterns required for this task, leading to a substantial drop in performance. This is likely because the shallow architecture of the MLP lacks the representational capacity needed to model the nuanced interactions between the layers of the pre-trained SLM.

4.3 Model Generalizability Analysis

In this section, we explore the generalizability of ExMRD from two distinct perspectives. First, we evaluate the generalizability of the main components within ExMRD—specifically the R³CoT and SLReviewer—to determine their effectiveness across different MLLMs. Subsequently, we investigate the model’s generalizability from a dataset perspective, focusing on its ability to train on one dataset and consistently perform well on other distinct datasets.

4.3.1 Generalizability of Main Components. We evaluate the generalizability of our main components, R³CoT and SLReviewer, across various MLLMs for MVRD. Results for the FakeSV and FakeTT datasets are shown in Figure 4. As observed, integrating R³CoT significantly enhances the accuracy of various MLLMs in distinguishing rumors from truths in micro-videos, underscoring R³CoT’s adaptability across diverse MLLM architectures. Moreover, SLReviewer builds upon these advantages to achieve significant improvements in prediction accuracy. It demonstrates strong generalization by utilizing knowledge from various MLLMs and consistently maintaining high performance across different datasets.

4.3.2 Generalizability on Different Datasets. We conduct the experiments to evaluate the generalizability of our proposed ExMRD and the most competitive baseline model FakingRec, and the results are reported in Figure 5. In this experiment, we select the FakeTT and




Figure 6: Comparison of explanation quality with and without R³CoT on the FakeTT dataset. I: Informativeness, S: Soundness, P: Persuasiveness, R: Readability, F: Fluency.

FVC datasets, which are sourced from different platforms and varying significantly in content style and target audience, to conduct cross-dataset evaluations. Specifically, we train and validate the models on one dataset and test them on the other. The results show that our ExMRD strongly outperforms FakingRec across all metrics in both cross-dataset evaluations. The baseline model struggles to handle dataset biases, such as superficial multimodal features and platform-specific characteristics. In contrast, our model effectively bridges the gap between inconsistencies in micro-video quality through the MLLM by the guidance of the refining step in R³CoT. Moreover, the MLLM is instructed by the retrieving and reasoning steps from R³CoT to generate the domain knowledge and the rationale for detecting rumors for the micro-videos in the target dataset. This distilled knowledge demonstrates significant generalizability and is fed to SLReviewer to make the precise prediction. These observations further confirm the superiority of ExMRD in handling diverse videos with varying quality and its robustness in cross-platform scenarios, making it well-suited for real-world deployment.



4.4 Model Explainability

To assess the explainability of ExMRD, we first compare the quality of explanations generated by various MLLMs, with and without R³CoT. A case study on selected micro-videos then demonstrates how effectively ExMRD explains its classifications.

4.4.1 Quality of Explainability. In this section, we assess the contribution of the proposed R³CoT to the quality of explanation (i.e., reasoning output). We employ G-Eval [28], an LLM-based reference-free evaluation approach, to evaluate the text quality of the explanations generated by our framework by comparing the base MLLM with or without R³CoT. To be specific, we utilize the following criteria [48, 49]: (1) Informativeness: the explanation provides new information, such as explaining the background and additional context; (2) Soundness: the explanation seems valid and logical; (3) Persuasiveness: the explanation is convincing; (4) Readability: the explanation follows proper grammar and structural rules; (5) Fluency: the explanation flows smoothly with coherent and well-connected ideas. For each criterion, a 5-point Likert scale [20] is employed, where 1 meant the poorest quality and 5 the best.

Figure 6 illustrates the average improvement in explanation quality, as evaluated by G-Eval, for base MLLM with and without R³CoT, across five criteria. The results show that: (1) in *informativeness* and *soundness*, MLLM equipped with R³CoT exhibit significant improvement over the original MLLM, underscoring the necessity of R³CoT

Table 4: Case study of correctly detected micro-video rumors on the FakeTT dataset.

	Case 1	Case 2
Micro-video		
Viewpoint	Three Simple Ways to Check Food Quality	Martin Luther King was not a Republican
Original Content	REAL FOOD VS FAKE FOOD CHECK HOW NATURAL PRODUCTS ARE! ARTIFICIAL REAL HONEY HONEY NATURAL ...	Replying to marywesling Dr. King did not associate himself as a member of any party. #mlk #mlkday
Refining	This video presents a comparison between real and artificial food products. It ... identify natural products versus artificial alternatives, with examples like real honey versus ...	The micro-video shows Dr. Martin Luther King Jr. expressed concern about the Republican Party potentially becoming ..." ... He acknowledged that this trend poses a significant danger ...
Retrieving	Food Authenticity Checks: Common methods include testing for natural chemical markers (e.g., pure honey vs. adulterated), and observing physical characteristics during cooking...	Dr. Martin Luther King Jr. was a renowned civil rights leader... he did not publicly declare himself a member of any political party... His main focus ...not partisan politics ...
Reasoning	The argument seems visually driven ... Cooking appearance alone may not conclusively differentiate between natural and artificial, as processed foods can mimic the appearance of natural ones.	From the video, Dr. Jin did express concerns during an interview about the Republican Party potentially becoming a "white party." This is consistent with historical records ...
ExMRD	Rumor ✓	Truth ✓

for providing expressive domain knowledge during the retrieving step; (2) in *readability* and *fluency*, MLLM equipped with R³CoT outperform the original versions, demonstrating the effectiveness of refining step in reorganizing video content and enhancing clarity; (3) in *persuasiveness*, the MLLM with R³CoT displays a significant improvement, suggesting that it contributes to more convincing rationales through its reasoning step.

4.4.2 Qualitative Analysis on Explainability. To further investigate the explainability of our proposed ExMRD, we randomly selected two micro-video cases from FakeTT to explore how ExMRD classifies each video as either a rumor or truth, as shown in Table 4.

In case 1, a rumor micro-video claims that *Three Simple Ways to Check Food Quality*. The original content involves cooking two types of food in a frying pan to test whether they are natural products. ExMRD first refines this content to highlight the core claim. Subsequently, it retrieves domain knowledge on common methods for verifying food authenticity, such as testing for natural chemicals, and applies logical reasoning to reveal that the content primarily relies on visual appeal to attract viewers, without offering valid techniques for determining whether the food is genuinely natural. As a result, our ExMRD correctly classifies the content in this video as a *Rumor*. This case demonstrates how ExMRD not only detects misinformation but also offers a well-founded rationale supported by factual knowledge and logical analysis.

In Case 2, a real micro-video refutes the statement that *Martin Luther King was a Republican*. The original content includes a description of Dr. King's concerns that the Republican Party might become a *white party*. Our framework refines the text to emphasize that Dr. King expressed concerns about the party's direction but did not publicly align himself with any political party. It retrieves

relevant domain knowledge confirming that Dr. King was a non-partisan civil rights leader, and uses logical reasoning to align this information with historical records. Consequently, the framework accurately classifies the video as *Truth*. This case highlights how ExMRD integrates historical context and logical reasoning to verify the authenticity of the claim. Additional qualitative analyses and error analysis are provided in Appendix B-C.

5 Conclusion

This work introduces ExMRD, an Explainable framework for interpretable Micro-video Rumor Detection. The proposed R³CoT mechanism in ExMRD is a novel three-step CoT process – Refining, Retrieving, and Reasoning – that reorganizes the raw video content, retrieves rumor-related domain knowledge, and generates explainable conclusions on whether the micro-video contains misleading information. Instead of directly fine-tuning MLLMs, which is computationally expensive, we propose SLReviewer within the ExMRD framework, distilling CoT-guided MLLM outputs to ensure accurate predictions with minimal computational overhead, making it more adaptable to real-world demands. Extensive experiments conducted on three real-world datasets demonstrate the effectiveness of ExMRD in both rumor detection and explainability. We believe that ExMRD is a valuable tool for detecting rumors on various micro-video platforms (e.g., TikTok and YouTube Shorts) while also promoting AI transparency and fostering a more trustworthy and safer online experience.

6 Acknowledgments

This work was supported by National Natural Science Foundation of China (Grant No.62176043, No.62072077, and No.U22A2097).

References

- [1] Rana Al-Marouf, Kevin Ayoubi, Khadija Alhumaid, Ahmad Aburayya, Muhammad Alshurideh, Raghad Alfaisal, and Said Salloum. 2021. The acceptance of social media video for knowledge acquisition, sharing and application: A comparative study among YouTube users and TikTok users' for medical purposes. *International Journal of Data and Network Science* 5, 3 (2021), 197.
- [2] Joshua Ashkinaze, Eric Gilbert, and Ceren Budak. 2024. The Dynamics of (Not) Unfollowing Misinformation Spreaders. In *Proceedings of the ACM Web Conference (WWW)*. 1115–1125.
- [3] Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity. In *Proceedings of the International Joint Conference on Natural Language Processing (IJCNLP)*. 675–718.
- [4] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. 2021. Is Space-Time Attention All You Need for Video Understanding?. In *International Conference on Machine Learning (ICML)*. 813–824.
- [5] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, J. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. Henighan, R. Child, A. Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, S. Gray, B. Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, I. Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Conference on Neural Information Processing Systems (NeurIPS)*, Vol. abs/2005.14165.
- [6] Yuyan Bu, Qiang Sheng, Juan Cao, Peng Qi, Danding Wang, and Jintao Li. 2024. FakingRecipe: Detecting Fake News on Short Video Platforms from the Perspective of Creative Process. In *Proceedings of the ACM International Conference on Multimedia (MM)*.
- [7] Yixuan Chen, Dongsheng Li, Peng Zhang, Jie Sui, Qin Lv, Tun Lu, and Li Shang. 2022. Cross-modal Ambiguity Learning for Multimodal Fake News Detection.. In *Proceedings of the ACM Web Conference (WWW)*. 2897–2905.
- [8] Zhenfang Chen, Qinhong Zhou, Yikang Shen, Yining Hong, Zhiqing Sun, Dan Gutfreund, and Chuang Gan. 2024. Visual Chain-of-Thought Prompting for Knowledge-Based Visual Reasoning. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)* 38, 2 (2024), 1254–1262.
- [9] Zhangtao Cheng, Jienan Zhang, Xovee Xu, Goce Trajcevski, Ting Zhong, and Fan Zhou. 2024. Retrieval-Augmented Hypergraph for Multimodal Social Media Popularity Prediction. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. 445–455.
- [10] Zhangtao Cheng, Fan Zhou, Xovee Xu, Kunpeng Zhang, Goce Trajcevski, Ting Zhong, and Philip S. Yu. 2024. Information Cascade Popularity Prediction via Probabilistic Diffusion. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 36, 12 (2024), 8541–8555.
- [11] Hyewon Choi and Youngjoong Ko. 2021. Using Topic Modeling and Adversarial Neural Networks for Fake News Video Detection. In *International Conference on Information and Knowledge Management (CIKM)*. ACM.
- [12] S. Davis and P. Mermelstein. 1980. Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. *IEEE Transactions on Acoustics, Speech, and Signal Processing* (1980).
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. Association for Computational Linguistics, 4171–4186.
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, G. Heigold, S. Gelly, Jakob Uszkoreit, and N. Houlsby. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations (ICLR)*, Vol. abs/2010.11929.
- [15] Xuanjie Fang, Sijie Cheng, Jiang Liu, and Wei Wang. 2023. Modeling Adversarial Attack on Pre-trained Language Models as Sequential Decision Making. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.
- [16] Sze Ching Evelyn Fung, Man Fai Wong, and Chee Wei Tan. 2023. Chain-of-Thoughts Prompting with Language Models for Accurate Math Problem-Solving. *IEEE MIT Undergraduate Research Technology Conference (URTC)* (2023).
- [17] Rui Hou, Verónica Pérez-Rosas, Stacy L. Loeb, and Rada Mihalcea. 2019. Towards Automatic Detection of Misinformation in Online Medical Videos.. In *International Conference on Multimodal Interaction (ICMI)*. 235–243.
- [18] Beizhe Hu, Qiang Sheng, Juan Cao, Yuhui Shi, Yang Li, Danding Wang, and Peng Qi. 2024. Bad actor, good advisor: Exploring the role of large language models in fake news detection. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 38. 22105–22113.
- [19] Zhanming Jie, Trung Quoc Luong, Xinbo Zhang, Xiaoran Jin, and Hang Li. 2023. Design of Chain-of-Thought in Math Problem Solving. *arXiv* (2023).
- [20] Ankur Joshi, Saket Kale, Satish Chandel, and D Kumar Pal. 2015. Likert scale: Explored and explained. *British journal of applied science & technology* 7, 4 (2015), 396–403.
- [21] M Laeeq Khan. 2017. Social media engagement: What motivates user participation and consumption on YouTube? *Computers in human behavior* 66 (2017), 236–247.
- [22] Jian Lang, Zhangtao Cheng, Ting Zhong, and Fan Zhou. 2025. Retrieval-Augmented Dynamic Prompt Tuning for Incomplete Multimodal Learning. *arXiv* (2025).
- [23] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024. LLaVA-OneVision: Easy Visual Task Transfer. *arXiv abs/2408.03326* (2024).
- [24] Chenxia Li, Weiwei Liu, Ruoyu Guo, Xiaoting Yin, Kaitao Jiang, Yongkun Du, Yuning Du, Lingfeng Zhu, Baohua Lai, Xiaoguang Hu, Dianhai Yu, and Yanjun Ma. 2022. PP-OCRv3: More Attempts for the Improvement of Ultra Lightweight OCR System. *arXiv abs/2206.03001* (2022).
- [25] Xiaojun Li, Xvhao Xiao, Jia Li, Changhua Hu, Junping Yao, and Shaochen Li. 2022. A CNN-based misleading video detection model. *Scientific Reports* 12, 1 (2022).
- [26] Yupeng Li, Haorui He, Jin Bai, and Dacheng Wen. 2024. MCFEND: a multi-source benchmark dataset for Chinese fake news detection. In *Proceedings of the ACM Web Conference (WWW)*. 4018–4027.
- [27] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual Instruction Tuning. In *Conference on Neural Information Processing Systems (NeurIPS)*, Vol. abs/2304.08485.
- [28] Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-Eval: NLG Evaluation using Gpt-4 with Better Human Alignment.. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2511–2522.
- [29] Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*.
- [30] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to Explain: Multimodal Reasoning via Thought Chains for Science Question Answering.. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- [31] Ana Marasovic, Iz Beltagy, Doug Downey, and Matthew E. Peters. 2022. Few-Shot Self-Rationalization with Natural Language Prompts.. In *North American Chapter of the Association for Computational Linguistics (NAACL)*. 410–424.
- [32] Chancharik Mitra, Brandon Huang, Trevor Darrell, and Roei Herzig. 2024. Compositional Chain-of-Thought Prompting for Large Multimodal Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [33] Milad Moradi and Matthias Samwald. 2021. Evaluating the Robustness of Neural Language Models to Input Perturbations.. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1558–1570.
- [34] Kenton O'Hara, April Slayden Mitchell, and Alex Vorbau. 2007. Consuming video on mobile devices. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 857–866.
- [35] OpenAI. 2023. GPT-4 Technical Report. *arXiv* (2023).
- [36] Olga Papadopoulou, Markos Zampoglou, Symeon Papadopoulos, and Ioannis Kompatsiaris. 2019. A corpus of debunked and verified user-generated videos. *Online Information Review* 43, 1 (2019), 72–88.
- [37] Peng Qi, Yuyan Bu, Juan Cao, Wei Ji, Ruihao Shui, Junbin Xiao, Danding Wang, and Tat-Seng Chua. 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 (AAAI)*. 14444–14452.
- [38] Peng Qi, Yuyang Zhao, Yufeng Shen, Wei Ji, Juan Cao, and Tat-Seng Chua. 2023. Two Heads Are Better Than One: Improving Fake News Video Detection by Correlating with Neighbors. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.
- [39] Shengsheng Qian, Jinguang Wang, Jun Hu, Quan Fang, and Changsheng Xu. 2021. Hierarchical Multi-modal Contextual Attention Network for Fake News Detection. In *Proceedings of the ACM International Conference on Research and Development in Information Retrieval (SIGIR)*. ACM.
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision.. In *International Conference on Machine Learning (ICML)*. 8748–8763.
- [41] Daniel Rose, Vaishnavi Himakunthala, Andy Ouyang, Ryan He, Alex Mei, Yujie Lu, Michael Saxon, Chinmay Sonar, Diba Mirza, and William Yang Wang. 2023. Visual Chain of Thought: Bridging Logical Gaps with Multimodal Infillings. *arXiv abs/2305.02317* (2023).
- [42] Juan Carlos Medina Serrano, Orestis Papakyriakopoulos, and Simon Hegelich. 2020. NLP-based Feature Extraction for the Detection of COVID-19 Misinformation Videos on YouTube. In *ACL Workshop on Natural Language Processing for COVID-19 (NLP-COVID)*.
- [43] Lanyu Shang, Ziyi Kou, Yang Zhang, and Dong Wang. 2021. A Multimodal Misinformation Detector for COVID-19 Short Videos on TikTok. In *IEEE International Conference on Big Data (Big Data)*. IEEE.
- [44] Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. 2024. Visual CoT: Advancing Multi-Modal Language Models with a Comprehensive Dataset and Benchmark for Chain-of-Thought

- Reasoning. *arXiv* (2024).
- [45] Rulin Shao, Zhouxing Shi, Jinfeng Yi, Pin-Yu Chen, and Cho-Jui Hsieh. 2022. On the Adversarial Robustness of Vision Transformers. *Transactions on Machine Learning Research (TMLR)* 2022 (2022).
 - [46] Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. 2024. To CoT or not to CoT? Chain-of-thought helps mainly on math and symbolic reasoning. *arXiv* (2024).
 - [47] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and Efficient Foundation Language Models. *arXiv abs/2302.13971* (2023).
 - [48] Bo Wang, Jing Ma, Hongzhan Lin, Zhiwei Yang, Ruichao Yang, Yuan Tian, and Yi Chang. 2024. Explainable Fake News Detection with Large Language Model via Defense Among Competing Wisdom. In *Proceedings of the ACM Web Conference (WWW)*. 2452–2463.
 - [49] Han Wang, Ming Shan Hee, Md. Rabiul Awal, Kenny Tsu Wei Choo, and Roy Ka-Wei Lee. 2023. Evaluating GPT-3 Generated Explanations for Hateful Content Moderation.. In *International Joint Conference on Artificial Intelligence (IJCAI)*. 6255–6263.
 - [50] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Ke-Yang Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. Qwen2-VL: Enhancing Vision-Language Model's Perception of the World at Any Resolution. *arXiv* (2024).
 - [51] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models.. In *International Conference on Learning Representations (ICLR)*.
 - [52] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models.. In *Conference on Neural Information Processing Systems (NeurIPS)*.
 - [53] Chenfei Wu, Jinlai Liu, Xiaojie Wang, and Xuan Dong. 2018. Chain of Reasoning for Visual Question Answering. In *Conference on Neural Information Processing Systems (NeurIPS)*. 273–283.
 - [54] Liang Xiao, Qi Zhang, Chongyang Shi, Shoujin Wang, Usman Naseem, and Liang Hu. 2024. MSynFD: Multi-hop Syntax Aware Fake News Detection. In *Proceedings of the ACM Web Conference (WWW)*. ACM.
 - [55] An Yang, Junshu Pan, Junyang Lin, Rui Men, Yichang Zhang, Jingren Zhou, and Chang Zhou. 2022. Chinese CLIP: Contrastive Vision-Language Pretraining in Chinese. *arXiv abs/2211.01335* (2022).
 - [56] Ting Zhong, Jian Lang, Yifan Zhang, Zhangtao Cheng, Kunpeng Zhang, and Fan Zhou. 2024. Predicting Micro-video Popularity via Multi-modal Retrieval Augmentation. In *Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. 2579–2583.
 - [57] Tianyang Zhong, Zheng Liu, Yi Pan, Yutong Zhang, Yifan Zhou, Shizhe Liang, Zihao Wu, and Yanjun Lyu. 2024. Evaluation of OpenAI o1: Opportunities and Challenges of AGI. *arXiv* (2024).
 - [58] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V. Le, and Ed H. Chi. 2023. Least-to-Most Prompting Enables Complex Reasoning in Large Language Models.. In *International Conference on Learning Representations (ICLR)*.
 - [59] Fan Zhou, Xovee Xu, Goce Trajcevski, and Kunpeng Zhang. 2022. A Survey of Information Cascade Analysis - Models, Predictions, and Recent Advances. *Comput. Surveys* 54, 2 (2022), 1–36.

A Detailed Experimental Settings

A.1 Datasets

To evaluate the performance of our proposed framework, ExMRD, alongside several baseline models, we utilize three real-world micro-video datasets: FakeSV [37], FakeTT [6], and FVC [36], with their statistics and characteristics reported in Table 5. In alignment with original paper [37], we implement temporal split: a chronological split with a 70%:15%:15% ratio for training, validation, and test sets is used, simulating real-world conditions where only past data is available to detect future rumor videos. The detailed descriptions for each dataset are presented as follows.

- **FakeSV**: This dataset is tailored for the detection of fake news spread through micro-videos in Chinese. It is sourced from two major micro-video platforms in China—*Douyin* and *Kuaishou*. Each instance in FakeSV includes the video itself, its title, user comments, relevant metadata, and the publisher’s profile.

- **FakeTT**: This dataset is designed to detect misinformation in short-form videos, specifically in the English language. It is meticulously curated from the widely-used platform *TikTok*. Each sample in FakeTT includes the video content, its title, and corresponding metadata.

- **FVC**: This dataset is constructed for detecting and analyzing fake videos versus real user-generated videos (UGVs). Sourced from platforms like *YouTube*, *Facebook*, and *Twitter*, the dataset covers a broad spectrum of events—ranging from politics and sports to natural disasters and wars. Each entry consists of the video, its title, and description, along with both original and near-duplicate versions of the content.

Table 5: Statistics and Characteristics of three datasets

Characteristics	FakeSV	FakeTT	FVC
Total Videos	3,624	1,814	2,764
Rumor Videos	1,810	1,172	1,633
Truth Videos	1,814	819	1,131
Duration (s)	39.88	47.69	87.83
Language	Chinese	English	English
Platform	Douyin, Kuaishou	TikTok	YouTube, Facebook, Twitter

A.2 Baseline Models

To validate the superiority of ExMRD, we select 13 competitive baselines in this study, which can be categorized into three groups: (1) *Unimodal detection methods*; (2) *Multimodal detection methods*; (3) *MLLM based methods*. The details of each group are as follows: (1) *Unimodal Detection Methods*:

- **BERT** [13] is a language representation model which is pre-trained for deep bidirectional representations from unlabeled text. It is used to extract features, specifically the [CLS] token, from textual inputs including the video title, description, and on-screen text. These extracted features form a 768-dimensional vector space, which is subsequently fed into a two-layer MLP to generate the final prediction.
- **ViT** [14] leverages the Transformer architecture for direct feature extraction from image patches. ViT is used to extract 768-dimensional feature vectors from 16 key frames of each video.

These vectors are then passed through a two-layer MLP to generate the final prediction.

- **MFCC** [12] is a widely used feature in audio classification tasks, particularly effective in capturing timbral and phonetic characteristics that can help identify anomalies or patterns related to misinformation in audio content. For each video, we extract 128-dimensional MFCC features from the audio stream. These features are then passed through a two-layer MLP to yield the final prediction.
- **TSformer** [4] employs separate spatial and temporal attention mechanisms on frame-level patches to address video understanding tasks. We utilize TSformer to extract 768-dimensional features from each video. These features are then input through a two-layer MLP to output the final prediction.

(2) Multimodal Detection Methods:

- **TikTec** [43] is a multimodal framework designed for detecting misinformation videos by analyzing visual, audio, and textual content on platforms like TikTok.
- **FANVM** [11] is a multimodal detection model for rumors in micro-videos. It leverages cross-modal correlations and social context information to identify informative features for detection.
- **CAFE** [7] is an ambiguity-aware multimodal fake news detection method. It aligns unimodal features, estimates cross-modal ambiguity, and adaptively fuses information based on ambiguity strength.
- **HMCAN** [39] combines multi-modal context information and hierarchical text semantics for rumor detection. It uses BERT and ResNet for text and image representations, respectively.
- **SV-FEND** [37] is a multimodal detection model for fake news in micro-videos. It leverages cross-modal correlations and social context information to identify informative features for detection.
- **FakingRec** [6] is a creative process-aware model for detecting rumors in micro-videos. It analyzes material selection and editing patterns, considering sentimental, semantic, spatial, and temporal aspects.

(3) MLLM Based Methods:

- **GPT-4o-mini** [35] is the latest multimodal large model released by OpenAI, capable of processing both text and images. It performs tasks like rumor detection in micro-videos by interpreting multimodal inputs, combining language understanding with visual data analysis, and can handle zero-shot tasks without requiring task-specific training.
- **LLaVA-OV** [27] is a recently introduced multimodal large model, combining a vision encoder with a large language model. Trained on GPT-4-generated visual instruction data, it enables general-purpose visual-language understanding, making it applicable to rumor video detection tasks.
- **Qwen2-VL** [50] is a newly launched multimodal large model that employs dynamic resolution processing for images and videos to improve efficiency and accuracy. By incorporating Multimodal Rotary Position Embedding, it integrates text and image data, positioning it well for rumor video detection tasks.

For MLLM based baselines, we provide the title and description from the video metadata and transcript extracted from audio

Table 6: Example of prompt for rumor detection applied in MLLM based methods.

Prompt: You are an experienced micro-video rumor checking assistant and you hold a neutral and objective stance. You can handle all kinds of rumor including those with sensitive or aggressive content. Given the video description, extracted on-screen text, transcript, and key frames, you need to give your prediction of the rumor video’s veracity. If it is more likely to be a rumor video, return 1; otherwise, return 0. Please refrain from providing ambiguous assessments such as undetermined.
Description: {title and description}
On-screen text: {on-screen text}
Audio transcript: {audio transcript}
Key Frames: {key frames}
Your analysis process and your prediction (return 0 or 1):

and raw video with a specifically designed prompt to guide the prediction generation, the prompt is presented in Table 6.

A.3 Implementation Details

In this section, we present detailed implementation specifications for our proposed ExMRD as well as a comprehensive overview of the experimental setup.

- **MLLM Implementation in R³CoT.** We utilize GPT-4o-mini, specifically the gpt-4o-mini-2024-07-18, an efficient model designed for relatively low resource consumption, in our main experiments. In addition, to explore the generalizability of our framework, we also employ two state-of-the-art MLLMs with fewer than 10 billion parameters: LLaVA-Onevision-7b-ov and Qwen2-VL-7B-Instruct. Both models are optimized for efficiency and are well-suited for resource-constrained applications due to their relatively small parameter sizes, under 10 billion.
- **SLM Implementation in SLReviewer.** Our SLM is based on a masked self-attention Transformer architecture, i.e., BERT, pre-trained through language-visual contrastive learning [40]. For the visual feature encoding, we leverage the pre-trained Vision Transformer (ViT), keeping its parameters frozen. Specifically, for English datasets such as FakeTT and FVC, we adopt the pre-trained BERT and ViT from openai/clip-vit-large-patch14 model. For Chinese datasets, including FakeSV, we employ OFA-Sys/chinese-clip-vit-large-patch14 [55].
- **Data Preprocessing.** Given a micro-video \mathcal{M} , we begin by extracting its multimodal information. For the visual content, we perform uniform frame sampling to obtain a set of frames, denoted as $\mathcal{V}_f = \{v_1, \dots, v_k\}$, where k is the number of sampled frames. To extract robust visual representations, we employ a pre-trained Vision Transformer (ViT) [14] as the feature encoder. Specifically, for each frame v_i , we compute its feature representation by extracting the output corresponding to the [CLS] token of the ViT model, resulting in the extracted visual feature representation $\mathbf{H}_v \in \mathbb{R}^{k \times d_v}$, where k is the number of frames and d_v denotes the dimension of the visual feature space. The textual content of the video is derived from three primary sources: (1) the video’s metadata, (2) the on-screen text extracted from each frame, and (3) the transcript extracted from

the audio. First, we obtain the metadata, which includes the video’s title and description, denoted as $\mathcal{T}_m \in \mathbb{R}^{n_m}$, where n_m is the number of words in the metadata. To capture on-screen text, we employ PaddleOCR [24] to perform text extraction at a rate of one frame per second for each video. The concatenated sequence of text extracted from all frames is denoted as $\mathcal{T}_o \in \mathbb{R}^{n_o}$, where n_o refers to the number of words in the on-screen text. For audio transcription, we leverage two pre-trained automatic speech recognition (ASR) models: one fine-tuned for Chinese (BELLE-2/Belle-whisper-large-v3-zh-punct) and the other for English (openai/whisper-large-v3). These models are specifically optimized for their respective languages, ensuring high transcription accuracy. The resulting transcript is represented as $\mathcal{T}_t \in \mathbb{R}^{n_t}$, where n_t is the number of words in the transcribed text.

- **Training Configuration.** For text, we set the maximum sequence length to 512 for all datasets. For key frames, we resize the images into 224×224 . For composite frame, we configure an $m \times m$ grid into a 2×2 grid of consecutive frames. We utilize the AdamW [29] optimizer with a learning rate of 2×10^{-4} and a weight decay of 5×10^{-5} for model parameters optimization. The number of fine-tuning decoder layers f in SLReviewer is set to 8 for all datasets. We set the random seed to 2024. For statistical analysis, where each model is run five times and report the mean values as experimental results. For baseline models, we strictly adhere to the settings specified in their original papers.
- **Implementation Environment.** All experiments are conducted on a system comprising an Intel(R) Core(TM) i9-14900KF processor, equipped with one NVIDIA GeForce RTX 4090 GPU with 24 GB of VRAM, and accompanied by 128 GB of DRAM.

B Additional Qualitative Analysis

In this section, we randomly select 4 micro-videos from FakeTT and FakeSV datasets to validate the explainability of our proposed ExMRD, and the results are present in Table 8-9.

C Error Analysis

In this section, we conduct an error analysis on the wrongly detection of rumor micro-videos to better understand the behavior of our proposed ExMRD framework.

As presented in Table 7, the micro-video depicts multiple waterspouts forming over an ocean, while the on-screen text shares an “Insane Weather Fact,” claiming that in 2003, the Great Lakes witnessed the largest waterspout outbreak in recorded history. The core issue with this micro-video is the discrepancy between the visual content and the text: they reference different events, with the visual footage showing a generic waterspout formation, whereas the text refers to a specific historical event. Although our framework successfully retrieved the factual information that the largest recorded waterspout outbreak over the Great Lakes occurred from August 27 to September 3, 2003, and generated additional background knowledge on waterspout formation and geographic factors, it incorrectly inferred that the video and the text are consistent representations of the same event. To improve accuracy, our ExMRD needs to integrate MLLMs with Web Search APIs to retrieve footage from this specific event, enabling a more comprehensive verification of the video’s authenticity.

D Limitations of Our Work

Although our work, ExMRD, demonstrates strong performance on MVRD, there are still some limitations:

- This work emphasizes providing rationales for the model’s predictions when detecting rumors in micro-videos. However, the internal workings of the neural network, specifically how it arrives at these decisions, have not been thoroughly explored. In future work, we intend to improve our approach by investigating the interpretability of the model architecture, particularly how its layers and learned representations contribute to the decision-making process.
- Determining whether a micro-video contains a rumor or presents truthful information increasingly depends on up-to-date domain knowledge, especially as new social events unfold. For example, a MLLM with outdated information might incorrectly classify the announcement that *John J. Hopfield and Geoffrey E. Hinton are receiving the 2024 Nobel Prize in Physics* as a rumor due to its lack of awareness of this recent event. In this case, the model may struggle to accurately assess the legitimacy of the claim, especially considering the significant impact of the development of the deep learning field on the decision to award the Nobel Prize. While our ExMRD framework itself does not integrate live domain updates, its strong generalization abilities make it easily adaptable for integration with MLLMs that can access Web Search APIs. This potential extension would allow for up-to-date domain knowledge, enhancing the accuracy of micro-video rumor detection in rapidly evolving social contexts.
- The upper bound of our framework’s performance is inherently dependent on the pre-training knowledge and reasoning capabilities provided by MLLMs. We have evaluated its efficiency using three widely adopted and practically applicable MLLMs. However, the generality of our framework allows for the seamless integration of more advanced state-of-the-art MLLMs as they emerge. For instance, incorporating OpenAI o1 may prove advantageous in scenarios requiring high accuracy and ample computational resources.

E Broader Impacts of Our Work

Looking ahead, ExMRD will be a useful tool for detecting and reducing rumors on platforms like TikTok, YouTube Shorts, and Snapchat. It can help users understand how rumors are created and spread while also improving how recommendation algorithms work by deprioritizing potentially harmful content. Beyond content moderation, ExMRD can also improve digital literacy by educating users on why certain videos are flagged as rumors. This can help users, especially younger audiences, become better at spotting misinformation, creating a more informed online community. In addition, ExMRD has the potential to support public safety by identifying misleading content during critical times, like health crises or natural disasters. This can prevent the spread of panic-inducing misinformation and ensure that reliable information reaches people

quickly. ExMRD also promotes trust and transparency in AI. By explaining why certain content is flagged, it can help build trust in automated systems and make users feel more confident about how AI is used on these platforms. For platform developers, ExMRD can reduce the need for manual moderation, saving time and costs. This makes content management more efficient, especially for platforms with limited resources, by allowing human moderators to focus on more complex cases. Finally, ExMRD supports ongoing research in AI and misinformation detection. Its flexible design makes it easy to incorporate new AI advancements (e.g., new MLLMs), allowing the system to stay up-to-date as AI technology evolves. This adaptability makes ExMRD a solid foundation for future research on fighting misinformation in digital spaces.

Table 7: Error analysis of wrongly detected micro-video rumors on the FakeTT dataset.

Error case 1	
Micro-video	
Viewpoint	The largest waterspout outbreak in Great Lakes history occurred in 2003,
Original Content	Nature’s Fury Unleashed: Jaw-Dropping Waterspout Fact Follow for more! mullaneyllc Video Credit: tuglife #boat #weather #mothernature #extremeweather #scary #waterspout #greatlakes #facts
Refining	This video presents a fact about the largest waterspout outbreak in the Great Lakes , which reportedly occurred in 2003. According to the video’s text and transcript, the event lasted seven days and produced over 65 waterspouts....
Retrieving	Waterspouts are tornado-like phenomena that occur over bodies of water.... In fact, the largest recorded waterspout outbreak in the Great Lakes occurred between August 27 and September 3, 2003.
Reasoning	The video correctly describes a known event: the 2003 Great Lakes waterspout outbreak. Both the visual and textual elements support this claim... about the largest waterspout outbreak in recorded history is backed by verified historical data , and the visuals match the description of severe weather conditions typical during such an event.
Ground Truth	Rumor
ExMRD	Truth ×

Table 8: Case study of correctly detected micro-video rumors on the FakeTT dataset.



	Case 1	Case 2
Micro-video		
Viewpoint	Elon Musk bought General Motors	The governor of Montana has officially banned TikTok
Original Content	#elon #elonmusk #fyp #gm #generalmotors #buyout #takeover#news on-screen-text: Elon musk Is buying general Motorslil Trillicor... audio transcript: No one can comment on it, but several sources have revealed that it is planning to shut down its factories. There is no denying that Tesla has always aimed to make this world a better and more sustainable place for humans.	#montana #tiktokban #politics #congress Montana governor bans TikTok By Samantha Delouya CNN Updated 653 PM EDT Wed May 17 2023 Montana BANS TikTok Why theres SO n 408 New YorkCNN Mc iforte signed a bill Wednesday ban tate Gianforte tweeter TikTok in Montana to protect data from ...
Textual Refining	The video suggests a major business development where Tesla, led by Elon Musk, is taking over GM to advance electric vehicles and sustainable energy. GM is shifting its focus to electric and autonomous vehicles while making substantial organizational changes.	This video discusses Montana becoming the first state to fully ban TikTok, with the law taking effect in January 2024. The ban targets both users and companies that distribute the app, with fines of up to \$10,000 for violations. The speaker notes potential legal challenges that could arise before the ban is fully enforced.
Visual Refining	The visual content of video showcases a well-organized, modern car manufacturing facility, emphasizing the efficiency and precision of production. The video likely aims to highlight the technological advancements in EV manufacturing, aligning with GM's and Tesla's push toward a future dominated by sustainable automotive technologies.	The visual content shows a man discussing the TikTok ban, explaining the implications for users and companies like Apple. The speaker emphasizes the legal challenges that may arise, presenting the argument in a calm and factual manner. The structure of the video is simple, with a conversational tone. There are no obvious contradictions in the reasoning, as the claims made align with known facts about the Montana TikTok ban.
Retrieving	Tesla has been a leader in the electric vehicle market, focusing on reducing CO ₂ emissions and accelerating the adoption of sustainable energy solutions. GM, too, has shifted its strategy in recent years to focus on electric and autonomous vehicles as part of the broader industry trend toward sustainable transportation. However, there have been no credible reports of a Tesla acquisition of GM, making this claim highly unusual.	Montana's TikTok ban was signed into law by Governor Greg Gianforte in May 2023, citing concerns over data security and potential risks of Chinese government interference, given TikTok's ownership by the Chinese company ByteDance. However, the ban's implementation in January 2024 is expected to face numerous legal challenges, especially regarding First Amendment rights and technological enforcement. There is a broader national debate over whether TikTok should be banned due to security concerns.
Reasoning	The content emphasizes the automotive industry's shift toward electric vehicles, with both General Motors (GM) and Tesla focusing on sustainability and next-generation technologies. GM is restructuring to prioritize electric vehicles, aligning with Tesla's mission to accelerate the transition to sustainable energy. Given the scale of both companies, any potential acquisition would be highly publicized and subject to regulatory scrutiny. The visual elements, while highlighting advanced manufacturing processes, reflect broader industry trends rather than indicating any direct connection to a Tesla-GM buyout.	The argument presented is straightforward: Montana has passed a ban on TikTok, effective January 2024, and the ban could face legal challenges. The claims align with public reports on the issue, and there are no apparent logical flaws in the video's structure. The speaker presents the key points about the ban, its consequences, and the possibility of legal disputes.
Ground Truth	Rumor	Truth
ExMRD	Rumor ✓	Truth ✓

Table 9: Case study of correctly detected micro-video rumors on the FakeSV dataset.

	Case 1	Case 2
Micro-video		
Viewpoint	A handful of coarse salt and a mouthful of saliva break the car window glass	Jiuzhaigou Scenic Area is temporarily closed due to a mudslide.
Original Content	Too Scary! Are Car Windows This Fragile Now? Everyone, Be Careful Not to Leave Valuable Items Inside Your Car! "Beware of New Window-Smashing Theft Just a handful of coarse salt and some spit can break the window Criminals demonstrated this on the spot Advice for friends and family driving: Do not leave valuable or expensive items inside your car."	Jiuzhaigou Scenic Area in Sichuan Temporarily Closed Starting Today, Reopening Date to Be Announced Separately. #Jiuzhaigou Home - Explore Jiuzhaigou - Latest News - Tourism Products - Smart Services - Global Tourism Section Temporary Closure Announcement of Jiuzhaigou Scenic Area Author: Jiuzhaigou Management Bureau Category: Tourism Announcement ...
Textual Refining	The video appears to be warning viewers about a new method of breaking car windows. According to the text captured from the video, it describes a supposed technique where criminals can easily shatter car windows using common items like salt and saliva . The video's title suggests concern over the fragility of modern car windows, advising people not to leave valuables inside their cars. The transcript provides minimal audio information, with casual dialogue that does not contribute much to the overall context.	This video reports the temporary closure of Jiuzhaigou scenic area due to continuous heavy rainfall and severe flooding . The park is closed from August 18, 2020, to protect visitors' safety, and the reopening date will be announced later. Roads leading to the park have been affected by the rising water levels, causing blockages.
Visual Refining	The video shows a group of people gathered around a car , seemingly in a tense situation. The scene is outdoors, likely in a public area or parking lot, and there is noticeable interaction with the car. The focus is on the car and the people surrounding it, possibly discussing or reacting to the damage caused to the vehicle .	The visual description of the video suggests that it features a serene natural environment , likely a lush green forest with water bodies like rivers or lakes. The focus seems to be on capturing the beauty and calmness of nature, without showing any human activity or obvious text, creating a peaceful atmosphere.
Retrieving	There are known cases of theft involving breaking car windows to steal items inside, but the specific method of using salt and saliva to break a window is not scientifically supported. Generally, car windows are designed to withstand significant pressure and require specific tools or force to break . There have been other unfounded rumors in the past about easy ways to break windows using minimal effort, which have been debunked by experts .	Jiuzhaigou is a well-known scenic area in China, famous for its waterfalls, lakes, and lush vegetation. It has faced several temporary closures in the past due to natural disasters , such as earthquakes and floods, to ensure visitor safety. Continuous rainfall can lead to dangerous conditions, including flooding, road blockages, and landslides, which are common in mountainous areas like Jiuzhaigou .
Reasoning	The video claims that car windows can be easily broken using salt and saliva, but this lacks scientific credibility. Car windows are generally designed to resist significant force, and breaking them usually requires specific tools or a substantial amount of pressure . The method described in the video does not align with established knowledge about how car windows function , making the claim seem exaggerated or unfounded. There is also a mismatch between the urgency portrayed in the video and the plausibility of the method it suggests. Overall, the video plays on fear and encourages caution but presents a flawed and unsupported argument.	The video reports the temporary closure of the Jiuzhaigou scenic area due to continuous heavy rainfall and severe flooding, a situation that aligns with known incidents in the region . The OCR text clearly states the date (August 18, 2020) and reason for the closure, citing rising river levels and road blockages. While the video focuses on capturing the natural beauty of the area, without showing direct evidence of flooding or blockages, this does not diminish the credibility of the textual report. Scenic videos often highlight the environment rather than specific disruptions, and the closure announcement is consistent with past practices in the area .
Ground Truth	Rumor	Truth
ExMRD	Rumor ✓	Truth ✓