

The paradigm shift: A comprehensive survey on large vision language models for multimodal fake news detection

Wei Ai^a, Yilong Tan^a, Yuntao Shou^a, Tao Meng^{a,*}, Haowen Chen^b, Zhixiong He^c, Keqin Li^d

^a College of Computer and Mathematics, Central South University of Forestry and Technology, 410004, Hunan, Changsha, China

^b College of Computer Science and Electronic Engineering, Hunan University, 410004, Hunan, Changsha, China

^c College of Economics and Management, Central South University of Forestry and Technology, 410004, Hunan, Changsha, China

^d Department of Computer Science, State University of New York, New Paltz, New York, 12561, USA

ARTICLE INFO

Keywords:

Multimodal fake news detection
Large language models
Deep learning
Large vision-language models

ABSTRACT

In recent years, the rapid evolution of large vision–language models (LVLMs) has driven a paradigm shift in multimodal fake news detection (MFND), transforming it from traditional feature-engineering approaches to unified, end-to-end multimodal reasoning frameworks. Early methods primarily relied on shallow fusion techniques to capture correlations between text and images, but they struggled with high-level semantic understanding and complex cross-modal interactions. The emergence of LVLMs has fundamentally changed this landscape by enabling joint modeling of vision and language with powerful representation learning, thereby enhancing the ability to detect misinformation that leverages both textual narratives and visual content. Despite these advances, the field lacks a systematic survey that traces this transition and consolidates recent developments. To address this gap, this paper provides a comprehensive review of MFND through the lens of LVLMs. We first present a historical perspective, mapping the evolution from conventional multimodal detection pipelines to foundation model-driven paradigms. Next, we establish a structured taxonomy covering model architectures, datasets, and performance benchmarks. Furthermore, we analyze the remaining technical challenges, including interpretability, temporal reasoning, and domain generalization. Finally, we outline future research directions to guide the next stage of this paradigm shift. To the best of our knowledge, this is the first comprehensive survey to systematically document and analyze the transformative role of LVLMs in combating multimodal fake news. The summary of existing methods mentioned is in our Github: <https://github.com/Tan-YiLong/Overview-of-Fake-News-Detection>.

1. Introduction

The rapid proliferation of fake news across online platforms has emerged as a formidable societal challenge [1], undermining public trust [2], destabilizing democratic discourse, and exacerbating global crises such as the COVID-19 pandemic and geopolitical conflicts [3,4]. Unlike early waves of misinformation, which were predominantly textual, modern fake news campaigns increasingly exploit multimodal content combining manipulated images [5], misleading videos [6], and textually coherent yet semantically deceptive captions [7,8]. This growing reliance on cross-modal deception renders traditional unimodal detection techniques inadequate and motivates a shift toward Multimodal Fake News Detection (MFND) [9,10].

Multimodal fake news detection seeks to uncover inconsistencies between text, image, video, and other modalities to assess the veracity of online information [11]. However, the complex interplay between

modalities, such as subtle semantic mismatches [12], visual entailment contradictions, and temporal or contextual incoherence, presents formidable challenges [13,14]. Early MFND methods typically relied on late fusion architectures or shallow cross-modal alignment, which lacked the semantic depth and reasoning capacity to capture nuanced deception strategies [15].

Importantly, the notion of cross modal inconsistency in MFND is not ad hoc, but rooted in well documented limitations of vision–language pretraining (VLP) models. Prior work shows that even large-scale VLP models exhibit systematic weaknesses in fine grained semantic alignment, particularly for linguistic negation, attribute ownership, and spatial relations [16]. These limitations are especially consequential for fake news detection, where deceptive content often maintains surface level coherence while violating latent semantic constraints across modalities. As demonstrated by Wang et al. [16], such inconsistencies cannot

* Corresponding author.

Email address: mengtao@hnu.edu.cn (T. Meng).

be reliably identified through naive cross-modal similarity alone, revealing the brittleness of multimodal alignment without explicit linguistic grounding. This observation provides a theoretical basis for MFND, as many forms of multimodal misinformation intentionally exploit these alignment blind spots.

The emergence of Large Vision-Language Models (LVLMs) has fundamentally changed this landscape [17]. Built upon advances in large language models (LLMs) and pre-trained visual encoders [18,19], as shown in Fig. 1, LVLMs such as CLIP [20], BLIP-2 [21], Flamingo [22], Kosmos-1 [23], LLaVA [24], and GPT-4V [25] offer unified architectures capable of joint cross-modal representation learning and reasoning [26]. These models demonstrate impressive zero-shot and few-shot performance on tasks such as image-text matching [27], visual question answering [28], and multimodal entailment making them well-suited for high-level fake news detection that demands semantic grounding and cross-modal verification [29].

Nevertheless, the application of LVLMs to MFND remains fragmented and lacks a systematic understanding. Existing studies vary widely in how they integrate LVLMs, differing in architectural design, supervision strategy [30], task formulation [31], and training cost [32]. To unify this landscape, we propose a novel three branch taxonomy that categorizes current MFND approaches using LVLMs into three distinct paradigms: (1) **Parameter freezing applications**. These methods utilize pre-trained LVLMs without modifying their internal parameters [33]. Techniques such as in-context learning [34], prompt-based adaptation [35], or lightweight classification heads are employed on top of frozen backbones [15]. This paradigm is particularly attractive for resource constrained or real-time applications, offering fast deployment and high generalization at the cost of task specific adaptation [36]. (2) **Parameter tuning applications**. These approaches involve full or partial fine tuning of the LVLMs to enhance task alignment [11]. Techniques include full model fine tuning [37], adapter insertion [38], prefix tuning [39], or low rank adaptation (e.g., LoRA) [40]. Tuning allows the model to capture domain specific semantics [41], subtle modality inconsistencies [42], and contextual cues critical to accurate fake news detection, albeit with higher computational and data requirements [43]. (3) **Reasoning paradigm applications**. This new paradigm focuses on how LVLMs engage in structured, multi step reasoning to enhance the accuracy of fake news detection. Approaches under this paradigm emphasize the use of multi agent systems or explicit reasoning

pipelines that sequentially analyze multimodal evidence and the claim. Methods such as agent-based reasoning [44], which breaks down reasoning into distinct sub-tasks like evidence retrieval and contradiction analysis, and prompting-based reasoning [45], which guides the model through reasoning via carefully crafted prompts, are key examples. The reasoning paradigm enables a deeper understanding of complex misinformation, leveraging the strengths of both large pre-trained models and task-specific reasoning strategies. However, it typically incurs higher inference costs and complexity, requiring more elaborate model architectures and robust reasoning pipelines [46].

In this survey, we offer the first comprehensive and structured review of how large vision-language models are being utilized for multimodal fake news detection. The contributions made in this paper are summarized as follows:

- **New taxonomy:** We propose a novel three-branch taxonomy for applying large vision language models (LVLMs) to multimodal fake news detection (MFND), dividing existing methods into three main categories: (1) parameter-freezing applications, (2) parameter-tuning applications, and (3) reasoning-paradigm applications. This is the first comprehensive survey to systematically document and analyze the transformative role of LVLMs in combating multimodal fake news, incorporating the latest advancements in structured reasoning and agent-based strategies.
- **Comprehensive review:** We systematically review the evolution from unimodal to multimodal approaches, analyze representative LVLMs-based architectures, and compare their performance and design principles across visual-textual misinformation detection benchmarks.
- **Unified framework:** We build a structured analytical framework that encompasses model design, task formulation, training paradigms, and evaluation criteria. By aligning the strengths and limitations of parameter-freezing, parameter-tuning, and reasoning-based approaches, we provide guidance for selecting appropriate LVLM strategies under different resource and robustness constraints, helping to navigate the challenges of real-time and resource-constrained applications.
- **Benchmark and evaluation:** We summarize and compare publicly available datasets in MFND, highlighting their modality composition, annotation strategy, and real-world complexity. In addition, we examine current evaluation metrics and identify key gaps.

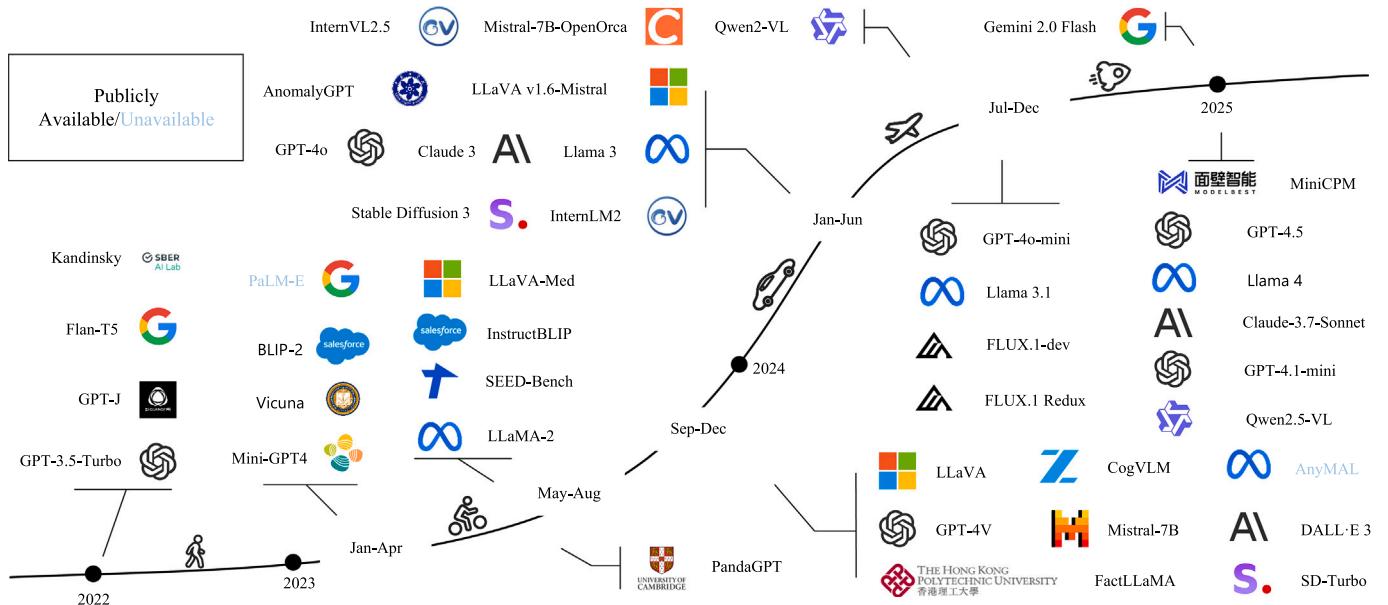


Fig. 1. A chronological overview of representative LVLMs is presented, highlighting the rapid growth of this field.

- **Research challenges and future directions:** We identify open problems in LVLM-based MFND, including multimodal hallucination, cross-modal bias, and data scarcity in misinformation domains. We further propose promising directions for future research, including counterfactual generation for cross-modal reasoning, and knowledge-enhanced LVLM adaptation.

This paper is organized as follows: Section 2 introduces the background and preliminary knowledge. Section 3 outlines the multimodal fake news detection paradigm based on LVLM. Section 4 introduces widely used benchmark datasets. Section 5 reviews evaluation metrics. Section 6 compares the experimental results of different methods. Section 7 discusses the advantages and challenges of LVLM in this field. Section 8 points out future research directions. Section 9 concludes the paper.

2. Preliminary information

2.1. Traditional single-modal fake news detection

Problem definition. Early approaches to fake news detection predominantly relied on a single modality [47,48], typically focusing on either textual or visual content [49,50]. In these methods, fake news is formulated as a binary classification problem over a single input space [51]. Let x_t denote textual features extracted from news articles [52], and x_v represent visual features derived from associated images [53]. The single-modal fake news classifier can generally be defined as $\mathcal{F}_{sm} : x \rightarrow \{0, 1\}$, where $x \in \{x_t, x_v\}$ and the output $\{0, 1\}$ corresponds to the labels real and fake [54].

Text-based methods. Traditional text-based approaches typically leverage linguistic and semantic features [55], such as word n -grams, syntactic patterns [56], or semantic embeddings [57]. Given a text sequence $T = \{w_1, w_2, \dots, w_n\}$, an embedding model $\phi(\cdot)$ maps words into a continuous space [58], yielding the representation as follows:

$$X_T = \phi(T) \in \mathbb{R}^d \quad (1)$$

where d is the embedding dimension. A classifier f_θ (e.g., logistic regression, SVM, or neural network) then predicts the probability of being fake as follows:

$$P(y = 1|T) = \sigma(f_\theta(X_T)) \quad (2)$$

with $\sigma(\cdot)$ denoting the sigmoid function. The model is trained using the standard binary cross-entropy loss as follows:

$$\mathcal{L}_{text} = - \sum_{i=1}^N [y_i \log P(y_i|T_i) + (1 - y_i) \log(1 - P(y_i|T_i))] \quad (3)$$

Image-based methods. Visual single-modal methods focus on detecting inconsistencies or manipulations in images [59,60]. A typical approach involves extracting features from an image I using a convolutional neural network (CNN) [3] or vision Transformer (ViT) [61] as follows:

$$X_V = \psi(I) \in \mathbb{R}^k \quad (4)$$

where $\psi(\cdot)$ is the feature extractor. The classifier g_θ then computes as follows:

$$P(y = 1|I) = \sigma(g_\theta(X_V)) \quad (5)$$

These methods aim to capture visual artifacts such as unnatural textures [62], illumination mismatches [63], or traces of generative adversarial networks (GANs) [64].

Limitations. Although single-modal detection has achieved early success, it suffers from several critical limitations [65]. Text-only methods often fail when the news content appears linguistically coherent

but is paired with misleading images [66]. Similarly, image-only methods cannot account for manipulative narratives conveyed by text [67]. Formally, if fake news involves a cross-modal inconsistency [68,69], i.e.,

$$(T, V) \in \mathcal{D}_{fake}, \quad \phi(T) \not\approx \psi(V)$$

then any function \mathcal{F}_{sm} defined on a single modality is inherently insufficient [70]. This limitation motivates the shift toward multimodal approaches, where joint reasoning over text and vision becomes essential [71].

2.2. Multimodal fake news detection

Problem definition. Let $\mathcal{M} = \{t, v\}$ denote the set of available modalities (e.g., text t , image/video v). A multimodal set is $x = \{x_m\}_{m \in \mathcal{M}}$, associated with a focal claim c and optional external context κ (e.g., knowledge bases, timelines, provenance) [72]. The goal of Multimodal Fake News Detection (MFND) is to infer a veracity label $y \in \{0, 1\}$ (with 1 meaning fake) together with a structured, referable rationale R grounding the decision in cross-modal evidence [73,74]. We formalize a reasoned prediction as follows:

$$(\hat{y}, \hat{R}) = \arg \max_{y \in \{0, 1\}, R \in \mathcal{R}} P_\theta(y, R | x, c, \kappa) \quad (6)$$

where θ parameterizes the large vision-language model (LVLM), and \mathcal{R} denotes the space of executable explanations (textual justifications with pointers to spans, frames, and regions).

2.3. Multimodal fake news reasoning

Principled dimensions of deception. We posit a principled, triadic decomposition of multimodal deception into three interacting yet distinct dimensions (i.e., Authenticity, Consistency, and Intent) that together determine veracity and the structure of model explanations [75]. Authenticity quantifies media-level integrity via forensic evidence and generative priors [76]. Consistency measures cross-modal semantic agreement between claim, text, image/video, and audio embeddings [59]. Intent captures manipulative or persuasive framing, decoupled from factuality, through stylistic, pragmatic, and source cues. These dimensions target complementary failure modes, including (i) deepfaked or altered media with otherwise aligned captions; (ii) truthful media paired with misleading text; and (iii) technically accurate content arranged to steer beliefs. Relying on any single dimension is insufficient; instead, we fuse calibrated posteriors within an interpretable surrogate to produce both the veracity label and a grounded rationale [77]. This decomposition provides a unifying lens across news articles, social posts, memes, and videos, and it aligns naturally with LVLM-based reasoning that constructs evidence graphs and referenceable explanations [78]. Below we will introduce the principles of each dimension.

Media authenticity. This dimension targets whether a medium in x has been synthetically generated or altered. Let p_{nat} and p_{alt} be (implicit) likelihoods under natural and altered media manifolds, respectively (instantiated by generative/forensic priors) [79]. A generic forgery score for a visual stream x_v is as follows:

$$\begin{aligned} S_F(x_v) &= -\log p_{nat}(x_v) + \log p_{alt}(x_v) \\ F &= \sigma(\phi_F(S_F(x_v))) \in [0, 1] \end{aligned} \quad (7)$$

where ϕ_F calibrates the score and σ is a sigmoid. F indicates tampering signals (e.g., seam/lighting inconsistencies, GAN fingerprints, lip-sync drift).

Cross-modal inconsistency. This dimension captures contradictions between modalities (e.g., the caption asserts “Category 5 hurricane” while the video depicts calm weather). Let $g_m : \mathcal{X}_m \rightarrow \mathbb{R}^d$ embed each modality into a shared semantic space [13]. For any pair $(m, n) \in \mathcal{M}$:

$$D_{m,n}(x) = 1 - \text{sim}(g_m(x_m), g_n(x_n)) \quad (8)$$

$$D = \max_{(m,n) \in \mathcal{M}} D_{m,n}(x) \in [0, 1] \quad (9)$$

where sim can be cosine similarity or an entailment-based score from a caption→NLI pipeline. Large D signals cross-modal contradiction or missing supporting evidence [80].

Manipulative intent. Distinct from factuality, this dimension evaluates whether content is crafted to steer beliefs via rhetorical devices (e.g., cherry-picked visuals, loaded language, fear appeals) [65]. We model intent as a latent probability as follows:

$$\begin{aligned} I &= P_\theta(\text{manipulative} \mid x, c, \kappa) \\ &= \sigma(\phi_I(\psi_t(x_t), \psi_v(x_v), \psi_a(x_a), u)) \in [0, 1] \end{aligned} \quad (10)$$

where ψ_\cdot extracts stance, sentiment, emotional arousal, propaganda patterns, and u encodes source-level priors.

Joint decision. The veracity decision should couple these dimensions while remaining explanation-seeking [81]. A factorized formulation is as follows:

$$\begin{aligned} P_\theta(y, F, D, I \mid x, c, \kappa) &\propto \psi_Y(y \mid x, c, \kappa) \psi_F(F \mid x) \\ \psi_D(D \mid x, c) \psi_I(I \mid x, c, \kappa) \psi_C(y, F, D, I) \end{aligned} \quad (11)$$

where ψ_C enforces structural compatibilities (e.g., severe tampering or strong cross-modal contradictions raise the prior for $y=1$, while high I alone is not conclusive) [26]. A transparent decision surrogate is as follows:

$$\begin{aligned} s(x, c, \kappa) &= h_\theta(x, c, \kappa) + \lambda_F F + \lambda_D D + \lambda_I I \\ \hat{y} &= \mathbb{1}[\sigma(s(x, c, \kappa)) > \tau] \end{aligned} \quad (12)$$

where h_θ captures claim-conditioning and external verification (e.g., retrieval-augmented grounding), and λ_\cdot balances the three dimensions.

Definition of reasoning in MFND. Reasoning is the construction of a *grounded evidence graph* $G = (V, E)$ and a minimal rationale $R \subseteq V$ such that a verifier V_{check} confirms sufficiency [81] and faithfulness as follows:

$$\begin{aligned} \underbrace{P_\theta(y \mid x, c, \kappa, R)}_{\text{sufficient}} - \underbrace{P_\theta(y \mid x, c, \kappa, \emptyset)}_{\text{baseline}} &\geq \epsilon \\ P(R \text{ used by } \theta) &\geq 1 - \delta \end{aligned} \quad (13)$$

where R consists of modality-aligned snippets (text spans, frames, regions) and explicit cross-modal relations (support/contradiction) that can be referenced and verified [9].

Learning objective. Supervision is naturally multi-task, promoting veracity accuracy, forensic robustness, semantic alignment, and intent recognition [82], while rewarding concise, grounded explanations as follows:

$$\mathcal{L} = \underbrace{\ell_{\text{cls}}(y, \hat{y})}_{\text{veracity}} + \alpha \underbrace{\ell_F(F, \hat{F})}_{\text{authenticity}} + \beta \underbrace{\ell_D(D, \hat{D})}_{\text{consistency}} + \gamma \underbrace{\ell_I(I, \hat{I})}_{\text{intent}} + \eta \underbrace{\ell_{\text{exp}}(R, \hat{R})}_{\text{grounded rationale}} \quad (14)$$

where α controls the relative importance of the content-level forgery detection loss ℓ_F , which focuses on identifying tampered or synthetic media, β weights the cross-modal consistency loss ℓ_D , encouraging the model to detect semantic misalignments between modalities [83], γ adjusts the manipulative intent recognition loss ℓ_I , which measures whether the content is strategically framed to mislead or persuade, η regulates the explanation generation loss ℓ_{exp} , ensuring that the model outputs interpretable and faithful rationales R grounded in multimodal evidence [84].

2.4. Positioning of large vision-language models

Unified parameterization. LVLMs provide a single probabilistic program to instantiate the factorization in Eq. (11) and the reasoned prediction [85] in Eq. (6). Let ϕ_t, ϕ_v be modality encoders and let $r_\kappa(c)$ denote a retrieval operator over external context κ conditioned on claim c . We define claim-aware token streams

$$\tilde{t} = \text{Pack}(x_t, c), \quad \tilde{v} = \text{Patch}(x_v), \quad \tilde{\kappa} = r_\kappa(c) \quad (15)$$

and a cross-modal aggregator \mathcal{A}_θ (multi-head cross-attention over concatenated tokens) that yields a joint representation as follows:

$$Z = \mathcal{A}_\theta(\phi_t(\tilde{t}), \phi_v(\tilde{v}), \phi_t(\tilde{\kappa})) \in \mathbb{R}^d \quad (16)$$

Concretely, the authenticity, consistency [86], and intent scores are realized by LVLM heads, which are mathematically expressed as follows:

$$\begin{aligned} \hat{F} &= \sigma(w_F^\top Z + b_F), \quad \hat{I} = \sigma(w_I^\top Z + b_I) \\ g_t(x_t) &= W_t \phi_t(\tilde{t}), \quad g_v(x_v) = W_v \phi_v(\tilde{v}) \\ \hat{D} &= 1 - \text{sim}(g_t(x_t), g_v(x_v)) \end{aligned} \quad (17)$$

The claim-conditioned verifier h_θ is implemented to fuse content evidence with retrieved context as follows:

$$h_\theta(x, c, \kappa) = w_Y^\top Z + b_Y \quad (18)$$

Cross-modal alignment pretraining. To make $\text{sim}(\cdot, \cdot)$ meaningful, LVLMs optimize contrastive alignment [87] between g_t and g_v :

$$\mathcal{L}_{\text{align}} = - \sum_{i=1}^N \log \frac{\exp(\langle g_t(x_t^{(i)}), g_v(x_v^{(i)}) \rangle / \tau)}{\sum_{j=1}^N \exp(\langle g_t(x_t^{(j)}), g_v(x_v^{(j)}) \rangle / \tau)} \quad (19)$$

This objective ensures that D faithfully reflects cross-modal consistency.

Claim-conditioned verification and retrieval grounding. External context κ is integrated by a differentiable retriever as follows:

$$r_\kappa(c) = \text{TopK}(\arg \max_{d \in \kappa} \text{sim}(q(c), k(d))) \quad (20)$$

with query encoder $q(\cdot)$ and key encoder $k(\cdot)$ sharing parameters with ϕ_t . The retrieved snippets are injected as memory tokens, allowing h_θ to realize the claim-conditioning term and to reduce spurious correlations in $\hat{F}, \hat{D}, \hat{I}$.

Evidence-consistent rationale generation. Let $R = (r_1, \dots, r_L)$ denote a textual rationale augmented [31] with pointers π to evidence units (text spans, frames, regions) from the evidence graph $G = (V, E)$. LVLMs generate R with an autoregressive head over Z as follows:

$$P_\theta(R \mid x, c, \kappa) = \prod_{t=1}^L P_\theta(r_t \mid r_{\leq t}, x, c, \kappa), \quad (21)$$

$$P_\theta(\pi_t = u \mid r_{\leq t}, x, c, \kappa) = \text{softmax}(\langle q_t, v(u) \rangle), \quad u \in V$$

where q_t is the decoder query at step t and $v(u)$ is the node embedding of the evidence unit u . By training Eqs. (21) jointly, the LVLM yields rationales that satisfy the sufficiency/faithfulness criteria via explicit grounding [88].

Decision surrogate and calibration. Plugging LVLM heads into Eq. (12) gives the following:

$$\begin{aligned} s(x, c, \kappa) &= h_\theta(x, c, \kappa) + \lambda_F \hat{F} + \lambda_D \hat{D} + \lambda_I \hat{I} \\ \hat{y} &= \mathbb{1}[\sigma(s) > \tau] \end{aligned} \quad (22)$$

where the weights $(\lambda_F, \lambda_D, \lambda_I)$ are learned by post-hoc calibration to respect the compatibility potential ψ_C in Eq. (11):

$$\min_{\lambda_F, \lambda_D, \lambda_I, \tau} \mathbb{E}[\ell_{\text{cls}}(y, \mathbb{1}[\sigma(s) > \tau])] + \Omega(\lambda_F, \lambda_D, \lambda_I), \quad (23)$$

with a regularizer Ω that discourages over-reliance on any single dimension.

Media authenticity with generative priors. To instantiate p_{nat} and p_{alt} , LVLMs either distill forensic experts into a head over Z , or use a learned energy surrogate [76] as follows:

$$S_F(x_v) \approx E_\theta(x_v) = -\log \sum_z \exp(-\mathcal{E}_\theta(x_v, z)) \quad (24)$$

$$\hat{F} = \sigma(\phi_F(E_\theta(x_v)))$$

where \mathcal{E}_θ is an amortized energy over latent codes z (e.g., visual artifacts, frequency cues). This provides a practical bridge to Eq. (7) within the LVLM.

Training objective. The multi-task learning is realized by augmenting it with alignment and instruction-following terms [89] as follows:

$$\mathcal{L}_{\text{LVLM}} = \underbrace{\ell_{\text{cls}}(y, \hat{y})}_{\text{veracity}} + \alpha \underbrace{\ell_F(F, \hat{F})}_{\text{authenticity}} + \beta \underbrace{\ell_D(D, \hat{D})}_{\text{consistency}} + \gamma \underbrace{\ell_I(I, \hat{I})}_{\text{intent}} + \eta \underbrace{\ell_{\text{exp}}(R, \hat{R})}_{\text{grounded rationale}} + \mathcal{L}_{\text{align}} + \mathcal{L}_{\text{LM}} + \mathcal{L}_{\text{ptr}} \quad (25)$$

where \mathcal{L}_{LM} is the next-token negative log-likelihood for rationale generation and \mathcal{L}_{ptr} is a cross-entropy over pointers π to evidence nodes [90].

3. Taxonomy

We classify multimodal fake news detection methods into three paradigms, namely Parameter-Frozen Paradigm, Parameter-Tuning

Paradigm, and Reasoning Paradigm. The Parameter-Frozen Paradigm relies on zero-shot or few-shot prompting without updating the model's parameters, ensuring that the backbone model remains fixed. In contrast, the Parameter-Tuning Paradigm involves adapting the model's parameters through full fine-tuning or parameter-efficient techniques such as adapters, prefix tuning, and Low Rank Adaptation (LORA). Finally, the Reasoning Paradigm categorizes methods based on the type of reasoning employed, with prompting-based reasoning conducting implicit reasoning within a single inference pass and agent-based reasoning involving explicit, structured, multi step interactions for veracity reasoning. This classification is illustrated in Fig. 2, providing a comprehensive framework for understanding the distinct approaches in multimodal fake news detection.

3.1. Parameter-frozen paradigm

In the parameter-frozen paradigm, the backbone LVLM remains unchanged, and task adaptation is realized purely through prompting and in-context specification [156]. This setting is particularly attractive for multimodal fake news detection because it avoids expensive fine-tuning on large models, and it enables rapid transfer across datasets and platforms with heterogeneous news formats [75]. Let $X = \{x_t, x_i, x_v\}$ denote text, image, and video inputs after modality-specific encoding or serialization into model-acceptable inputs, and let I denote a task instruction template that defines the goal (e.g., “Determine whether the given news is fake or real and justify the decision with cross-modal evidence.”) [157]. The model produces a distribution over label strings given the constructed

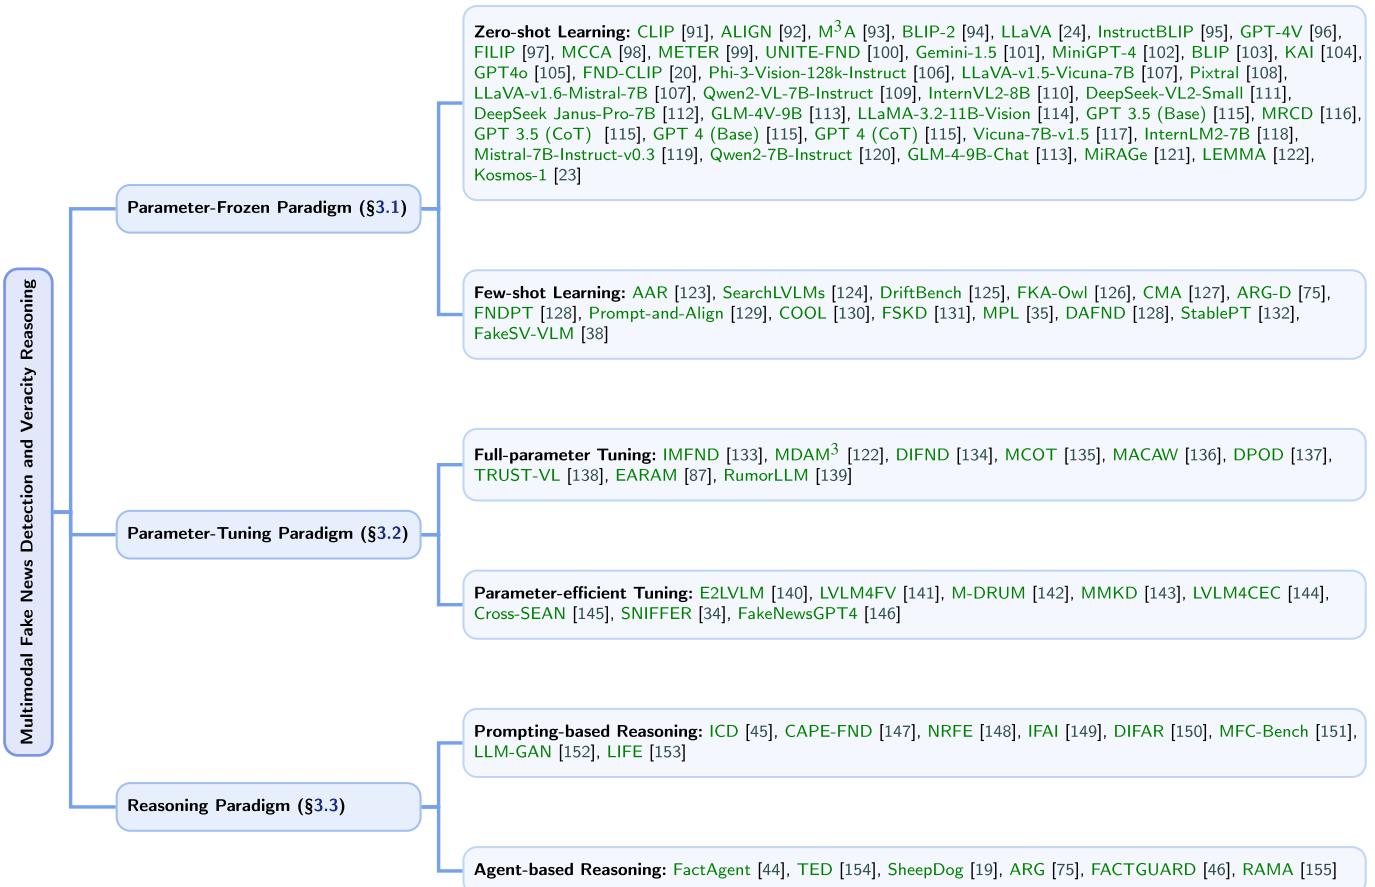


Fig. 2. Taxonomy of multimodal fake news detection and veracity reasoning. We systematically categorize multimodal fake news detection methods according to their parameter adaptation paradigms, including full fine-tuning, parameter-efficient tuning, prompting-based inference, and agent-based reasoning. This taxonomy provides a structured overview of state-of-the-art approaches and clarifies how different parameter interaction strategies are leveraged for multimodal veracity assessment [91,92,97–103,105–108,111–114,117–120,127,128,130–132,135,137,139,146,155].

prompt $\Pi(X, I)$ as follows:

$$P_{\Theta}(y | X, I) = \sum_{s \in v(y)} P_{\Theta}(s | \Pi(X, I)) \quad (26)$$

where Θ are frozen parameters and $v(y)$ is a verbalizer mapping each class y (e.g., *fake*, *real*) to one or more label strings. The final decision is $\hat{y} = \arg \max_y P_{\Theta}(y | X, I)$.

Mechanistic view under frozen parameters. From a mechanistic perspective, parameter-frozen LVLMs perform multimodal fake news detection by integrating visual tokens and textual instructions into a unified autoregressive reasoning process, rather than relying on explicit cross modal classifiers. Since model parameters remain unchanged, prompts act as the primary interface for controlling how multimodal evidence is attended to, decomposed, and compared. Different prompting strategies induce distinct reasoning behaviors [129]. Direct classification prompts typically encourage holistic judgments, which may obscure fine grained semantic inconsistencies such as mismatched object attributes, quantities, or spatial relations. In contrast, structured prompts that explicitly request visual entity extraction, attribute comparison, or step by step reasoning promote finer grained alignment between visual evidence and textual claims [158]. Chain-of-thought prompting further externalizes intermediate reasoning states, reducing over reliance on language priors and improving sensitivity to subtle cross-modal discrepancies. This mechanism distinguishes parameter-frozen LVLMs from traditional multimodal pipelines, where inconsistency detection is enforced through explicit similarity objectives or supervised heads. In frozen LVLMs, fine grained semantic inconsistency detection emerges implicitly from prompt guided reasoning trajectories, highlighting the central role of inference time design in determining detection robustness.

Zero-shot learning. Zero-shot multimodal fake news detection relies solely on natural-language instructions without task-specific examples [159]. As shown in Fig. 3, a well-engineered prompt serializes heterogeneous evidence into a structured context that highlights cross-modal credibility cues while minimizing spurious correlations [160]. In LLM-based pipelines, non-text modalities are first converted into textual evidence snippets, e.g., image captions, entity-relation triples, and video scene summaries through auxiliary encoders, and concatenated with raw news articles using role tags (e.g., <Text>, <Image-Caption>, <Video-Scene>)). In LVLM-based pipelines, raw embeddings are passed alongside textual tokens, but the model is still queried through instructions

[65]; hence Eq. (26) applies with $\Pi(\cdot)$ injecting modality tokens or connector-produced embeddings. To reduce prompt sensitivity and improve robustness, zero-shot systems commonly employ instruction variants and prompt ensembling with majority voting or probability averaging, self-consistency with reasoning where latent rationales r are sampled and marginalized as follows:

$$P_{\Theta}(y | X, I) \approx \sum_{r \in R} P_{\Theta}(y | r, X, I) P_{\Theta}(r | X, I) \quad (27)$$

and contextual calibration subtracts a prior estimated from a content-free prompt to alleviate label-word frequency bias. Zero-shot decoding can output both a veracity label (fake or real) and an explanation, providing weak but valuable interpretability for auditing fake news decisions [161].

Representative methods have instantiated these principles in diverse ways. For instance, MRCD [116] introduces a zero-shot multi round collaborative framework that iteratively generates hypotheses, retrieves external evidence, and refines predictions without supervised fine tuning. This approach effectively supports the detection of emerging fake news while enhancing decision explainability. Additionally, MRCD leverages collaborative evidence gathering, which improves robustness by facilitating cross source information validation and hypothesis refinement. LEMMA [122] adapts this paradigm by leveraging exemplar driven fusion pipelines in a zero-shot setting, showing that multimodal alignment and reasoning can be improved without task specific supervision. This model introduces an innovative approach by using exemplar based learning to guide reasoning, enabling better performance in resource constrained settings. MiRAGE [121] improves cross modal alignment in zero-shot scenarios through synthetic data augmentation and augmented connectors, enhancing multimodal reasoning. By synthesizing additional data, MiRAGE reduces the dependency on manually curated datasets, making it highly effective for scaling fake news detection across diverse domains. Similarly, M³A [93] integrates text, image, and domain adaptation features via large scale pre-training, achieving transferability across diverse disinformation benchmarks without task specific fine tuning. This model demonstrates the power of large scale pre-training, allowing it to generalize effectively across multiple fake news detection challenges. KAI [104] proposes a knowledge enhanced interpretable network that leverages a large language model to generate target specific analytical perspectives, which are integrated through a bidirectional knowledge guided neural generation system.

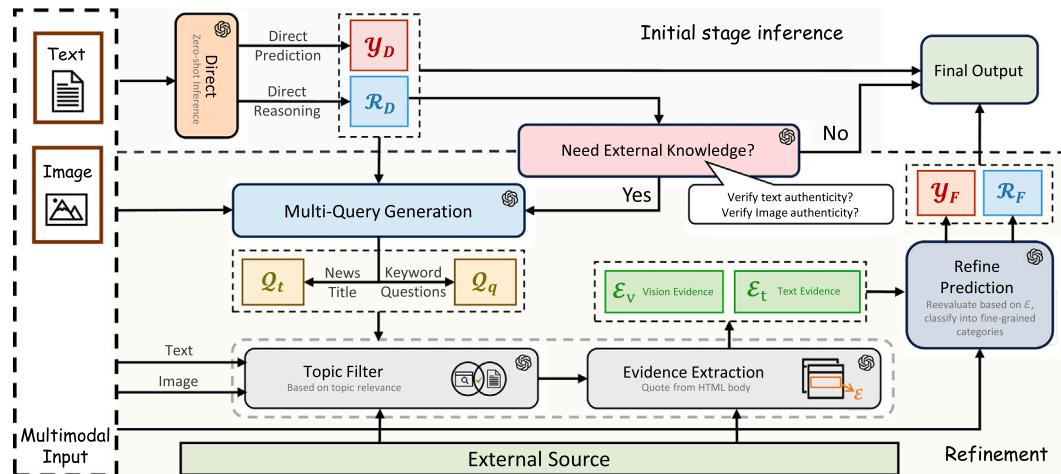


Fig. 3. Zero-shot learning framework for multimodal fake news detection. The architecture integrates direct prediction and reasoning mechanisms to process text and image inputs. It employs a multi-query generation module to formulate news-related queries based on the input's title and keywords. These queries are filtered for topical relevance and used to extract evidence from external sources. If external knowledge is deemed necessary, the system verifies the authenticity of both text and image inputs. The extracted evidence, along with the initial predictions, is refined through a reasoning process to produce the final output. This end-to-end approach enables the detection of fake news without prior training on labeled data, leveraging the complementary strengths of direct reasoning and external evidence integration [122].

This enables zero-shot stance detection with improved interpretability, making KAI particularly suitable for real world applications that demand transparency and explainability in decision-making processes.

Few-shot learning. Few-shot multimodal fake news detection augments the instruction with k demonstration pairs $\{(X_i, y_i)\}_{i=1}^k$ injected into the prompt as in-context exemplars. The predictive distribution is as follows:

$$P_{\Theta}(y | X^*, \mathcal{I}, \mathcal{D}_k) = \sum_{s \in v(y)} P_{\Theta}(s | \Pi(X^*, \mathcal{I}, \mathcal{D}_k)) \quad (28)$$

where the selection of \mathcal{D}_k is crucial. Effective strategies balance relevance to the query and diversity across news topics and modalities [6,162], often via embedding-based retrieval with determinantal or max-min objectives to avoid redundancy [163]. Demonstrations should expose the model to prototypical cross-modal inconsistencies (e.g., text-image mismatch, fabricated video snippets, or misleading multi-modal narratives) and discourse contexts (e.g., satire, conspiracy, and propaganda) that are underrepresented in pre-training [32]. Ordering also matters: placing task definition first, then structured exemplars (“context \rightarrow rationale \rightarrow label”), and finally the query tends to reduce hallucination [148]. For classification with free-form generation, a rationale-then-verbalizer template that first elicits a short explanation r and then constrains the final answer to $v(y)$ often improves calibration and stability [74]. When label spaces differ across datasets, dynamic verbalizers provide a principled bridge by defining synonyms per class and aggregating token probabilities [55].

Representative few-shot systems in fake news detection operationalize these principles through diverse mechanisms. For instance, AAR [123] demonstrates that incorporating adaptive exemplars enables models to align retrieved evidence with news claims under limited supervision. By dynamically selecting relevant exemplars, AAR allows for effective alignment of multimodal evidence, even in settings with scarce labeled data, thus improving robustness in detecting inconsistencies. SearchLVLMs [124] show that few-shot multimodal demonstrations can guide large vision-language models to detect subtle inconsistencies between text and images by leveraging retrieval augmented exemplars. The integration of retrieval augmented learning in SearchLVLMs allows for more precise detection of fine grained inconsistencies, effectively enhancing model sensitivity to cross-modal discrepancies, such as conflicting visual cues or ambiguous textual claims. DriftBench [125] extends few-shot evaluation by systematically assessing model robustness under limited exemplars and distribution shifts, providing insights into adaptability against evolving disinformation patterns. This framework facilitates a deeper understanding of how models perform under real world conditions where training data may vary, highlighting the importance of model adaptability to shifting distributions in fake news contexts. FKA-Owl [126] further enhances robustness by integrating structured attention over curated demonstrations, revealing that hierarchical exemplar design mitigates the impact of noisy or adversarial content. This model’s hierarchical attention mechanism helps filter out noise and adversarial influences, ensuring that models remain reliable even when presented with imperfect or misleading evidence. At a broader scale, FakeSV-VLM [38] exemplifies few-shot adaptability by exploiting cross source and cross event multimodal exemplars, improving generalization without necessitating full model fine tuning. By leveraging cross source and event specific data, FakeSV-VLM improves the model’s transferability across diverse contexts, ensuring more effective detection across different types of disinformation. Collectively, these systems highlight that few-shot learning strikes a pragmatic balance between efficiency and adaptability, making it a compelling paradigm for fake news detection in scenarios where annotated data is limited but multimodal evidence must be effectively utilized.

3.2. Parameter-tuning paradigm

Unlike the parameter-frozen paradigm that keeps Θ fixed, the parameter-tuning paradigm updates part or all of the LVLM parameters

Θ to enhance domain specialization [164], multimodal calibration [59], and interpretable reasoning [165]. Let each modality encoder be ϕ_t, ϕ_v for text, vision inputs, and let \mathcal{A}_{θ} denote the cross-modal aggregator. Given a multimodal dataset as follows:

$$\mathcal{D} = \{(x_i, c_i, y_i, R_i)\}_{i=1}^N \quad (29)$$

where $x_i = \{x_t, x_v\}$ are paired modalities [18], c_i is the claim, $y_i \in \{0, 1\}$ is the veracity label [41], and R_i the annotated rationale, parameter tuning seeks the optimal Θ that minimizes a task-aligned multi-objective loss as follows:

$$\mathcal{L}_{tune}(\Theta) = \mathcal{L}_{LVLM} + \lambda_{reg} \|\Theta - \Theta_0\|_2^2 \quad (30)$$

where \mathcal{L}_{LVLM} follows Eq. (25) and Θ_0 is the pretrained initialization, acting as a regularizer to prevent catastrophic forgetting.

Mechanistic view under parameter-tuning paradigm. From a mechanistic perspective, the parameter-tuning paradigm enhances the multimodal reasoning capabilities of LVLMs by fine tuning a subset of model parameters, allowing for a more flexible integration of text and image modalities. Unlike traditional models that treat text and images as separate inputs, LVLMs in this paradigm jointly optimize the parameters of both the textual and visual encoders through cross modal aggregators. This enables the model to better capture fine grained semantic inconsistencies between text and image, such as mismatched attributes or spatial relations. The effectiveness of this mechanism is largely driven by the choice of prompting strategy. Direct classification prompts encourage holistic judgments, which may overlook subtle mismatches, whereas structured prompts (e.g., entity extraction, attribute comparison, or step-by-step reasoning) guide the model to attend to specific details, improving the alignment between the visual and textual data. Furthermore, chain-of-thought prompting externalizes intermediate reasoning steps, reducing over reliance on language priors and enhancing the model’s sensitivity to cross modal discrepancies. Through the fine tuning of model parameters and strategic prompting, the parameter-tuning paradigm allows LVLMs to implicitly detect fine grained semantic inconsistencies, offering a more adaptive and robust approach compared to traditional methods.

Full-parameter tuning. Full fine-tuning unfreezes all LVLM components [166], including modality encoders ϕ_t, ϕ_v , aggregator \mathcal{A}_{θ} , and decision heads (w_F, w_D, w_I, w_Y) . The optimization jointly refines the authenticity [167], consistency [86], and intent estimators as follows:

$$\begin{aligned} \hat{F} &= \sigma(w_F^T Z + b_F) \\ \hat{D} &= 1 - \text{sim}(g_t(x_t), g_v(x_v)) \\ \hat{I} &= \sigma(w_I^T Z + b_I) \\ h_{\theta}(x, c, \xi) &= w_Y^T Z + b_Y \end{aligned} \quad (31)$$

and the calibrated decision surrogate remains as follows:

$$s(x, c, \xi) = h_{\theta}(x, c, \xi) + \alpha_F \hat{F} + \alpha_D \hat{D} + \alpha_I \hat{I} \quad (32)$$

During training, all parameters receive gradients [168] from \mathcal{L}_{LVLM} , enabling domain-specific adaptation [169] as follows:

$$\begin{aligned} \mathcal{L}_{LVLM} &= l_{cls}(y, \hat{y}) + \lambda_F l_F(F, \hat{F}) + \lambda_D l_D(D, \hat{D}) + \lambda_I l_I(I, \hat{I}) \\ &+ \lambda_R l_{exp}(R, \hat{R}) + \mathcal{L}_{align} + \mathcal{L}_{LM} + \mathcal{L}_{ptr} \end{aligned} \quad (33)$$

To stabilize optimization [145], practical implementations adopt stage-wise unfreezing, mixed-precision gradient checkpointing, and curriculum fine-tuning (progressing from single-claim to multi-evidence reasoning) [170]. Full tuning allows LVLMs to internalize complex cross-modal contradictions and high-level intents but requires significant computational and data resources [26].

Recent full-tuning approaches in multimodal fake news detection have demonstrated both the strengths and limitations of this paradigm

[171]. For instance, IMFND [133] established the foundation by fully fine tuning pretrained multimodal transformers on paired text-image datasets [172], achieving strong results in capturing cross modal inconsistencies in news content. By optimizing all parameters, IMFND improves the model's ability to learn fine grained interactions between text and images, significantly enhancing its capacity to identify misleading content. Building on these successes, video oriented systems such as MDAM³ [122] extended full parameter tuning to spatiotemporal news verification, showing that large scale video pretraining coupled with end-to-end optimization yields significant gains in detecting deepfake based misinformation [173]. MDAM³ pushes the boundaries of video based fake news detection by effectively incorporating both temporal and spatial cues, improving the model's robustness against dynamic manipulations often found in deepfake content. More recent models such as DIFND [134] and TRUST-VL [138] scale this strategy by integrating multimodal evidence with reasoning rich supervision, demonstrating that fully updating all parameters enhances both interpretability and robustness against adversarial news [165]. These models exemplify how reasoning guided updates allow for better disambiguation of complex multimodal evidence, improving both the model's generalization and its ability to handle adversarial inputs in fake news scenarios. Similarly, MACAW [136] leverages parameter intensive optimization to jointly model discourse level attention and multimodal alignment [174], highlighting the effectiveness of full-tuning for handling long context and linguistically nuanced misinformation [55]. By modeling discourse level attention, MACAW improves the model's ability to capture long range dependencies and subtle language nuances, which are critical in understanding and detecting misinformation embedded in lengthy articles or complex narratives.

While these systems highlight the versatility and strong performance of full-parameter tuning in fake news detection, they also expose its substantial computational costs [162], which motivates ongoing exploration [175] of more efficient alternatives such as parameter-efficient tuning.

Parameter-efficient tuning. To alleviate the prohibitive cost of full fine-tuning while preserving most of its performance benefits [176], parameter-efficient tuning (PET) has become a practical strategy for multimodal fake news detection [168]. In this paradigm, the backbone of large vision-language or video-language models is kept frozen, while lightweight modules, such as Adapter-tuning [177], LoRA [40], Prefix/Prompt-tuning [178,179], or gating mechanisms (e.g., IA³ [180], BitFit [181]) are optimized to capture domain-specific patterns. For multimodal news involving text, images, and videos, PET allows detectors to emphasize subtle lexical manipulations, visual tampering artifacts, and temporal inconsistencies without incurring the high computational overhead of updating the entire model. Recent advances such as QLoRA [182] further improve scalability by quantizing the frozen backbone to 4-bit while training LoRA adapters in higher precision, enabling efficient deployment at scale. Moreover, PET is often combined with connector tuning or selective unfreezing of higher Transformer layers to enhance discourse-level reasoning [183] and cross-modal alignment under constrained resources [184]. Specifically, parameter-efficient tuning keeps most of Θ frozen and introduces a small trainable subset Θ' through lightweight modules [185] as follows:

$$\Theta' = \{A_{\text{adapter}}, P_{\text{prefix}}, L_{\text{LoRA}}\} \quad (34)$$

Each module injects learnable low-rank or prefix parameters into the transformer blocks of \mathcal{A}_θ or modality encoders as follows:

$$Z' = \mathcal{A}_{\theta+\Theta'}(\phi_t(\tilde{t}), \phi_v(\tilde{v}), \phi_t(\tilde{\xi})) \quad (35)$$

The downstream heads (w_F, w_D, w_I, w_Y) are fine-tuned on Z' , while the frozen backbone retains general cross-modal alignment from pre-training. The loss function remains Eq. (33) but with gradient flow restricted to Θ' . Regularization such as orthogonality or low-rank

constraints on LoRA matrices (rank $r \ll d$) ensures stable adaptation:

$$\min_{\Theta'} \mathcal{L}_{\text{LVLM}}(\Theta') + \beta \|A_{\text{LoRA}} A_{\text{LoRA}}^\top - I\|_F^2 \quad (36)$$

Recent PET-based systems have demonstrated the practicality of parameter-efficient tuning in multimodal fake news detection [19]. For instance, E2LVLM [140] employs selective adaptation to integrate retrieved textual and visual evidence, significantly reducing training overhead while maintaining cross domain generalization. E2LVLM's selective adaptation approach allows the model to adapt only relevant parts of the pretrained model, ensuring efficient use of resources while still achieving strong performance across diverse domains and multimodal data. LVLM4FV [141] and M-DRUM [142] extend PET strategies to evidence driven verification, incorporating lightweight adapters into frozen LVLMs to efficiently capture cross modal reasoning under limited supervision. These models demonstrate that by incorporating lightweight adapters, it is possible to enhance model efficiency and generalization without the need for extensive fine tuning, offering a scalable solution for multimodal fake news detection. As shown in Fig. 4, MMKD [143] illustrates how PET can unify multimodal representation learning with knowledge guided reasoning, showing that adapter based modules can approximate the gains of full tuning while remaining computationally efficient. MMKD's approach of combining knowledge guided reasoning with PET allows the system to maintain interpretability and efficiency, providing a lightweight yet powerful solution for multimodal reasoning. Domain specific adaptations, such as LVLM4CEC [144], demonstrate that parameter efficient designs can preserve sensitivity to subtle context event correlations when annotated multimodal corpora are scarce. LVLM4CEC's ability to adapt to domain specific nuances without relying heavily on large annotated datasets highlights its effectiveness in specialized fake news detection tasks, particularly in resource constrained environments. Cross-SEAN [145] illustrates how adapter based architectures can be tuned alongside frozen backbones to support cross lingual and rationale grounded debunking. Cross-SEAN's cross lingual capabilities enhance its versatility in addressing fake news in multiple languages, while rationale-grounded debunking adds a layer of interpretability by explaining the model's reasoning process. Collectively, these systems exemplify how PET techniques balance the high performance of full parameter tuning with the scalability required for real-world deployment, underscoring their growing role in advancing multimodal fake news detection [169].

3.3. Reasoning paradigm

At both the system and cognitive levels, multimodal fake news detection methods employ distinct reasoning strategies. Prompting-based reasoning operates at the system level by implicitly guiding the model's reasoning process through instructions or prompts. In contrast, agent-based reasoning functions at the cognitive level by structuring reasoning as an explicit sequence of coordinated sub-tasks, simulating a more deliberate and human like approach to decision-making.

Mechanistic view under reasoning paradigm. From a mechanistic perspective, the reasoning paradigm in LVLMs offers a significant advancement over traditional models in detecting fine grained semantic inconsistencies between text and images. In prompting-based reasoning, LVLMs use task-specific prompts to guide the model's reasoning in a single inference pass, which contrasts with traditional methods that typically rely on separate pipelines for text and image processing. However, while prompting-based methods enable efficient, zero-shot generalization, they may obscure subtle mismatches between textual and visual evidence, especially when using direct classification prompts. In contrast, structured prompts such as those for entity extraction or step-by-step reasoning allow LVLMs to more explicitly align and detect discrepancies in object attributes, quantities, and spatial relationships, improving the model's sensitivity to fine grained inconsistencies. On the other hand, agent-based reasoning decomposes the detection process into a sequence of explicit tasks performed by specialized agents,

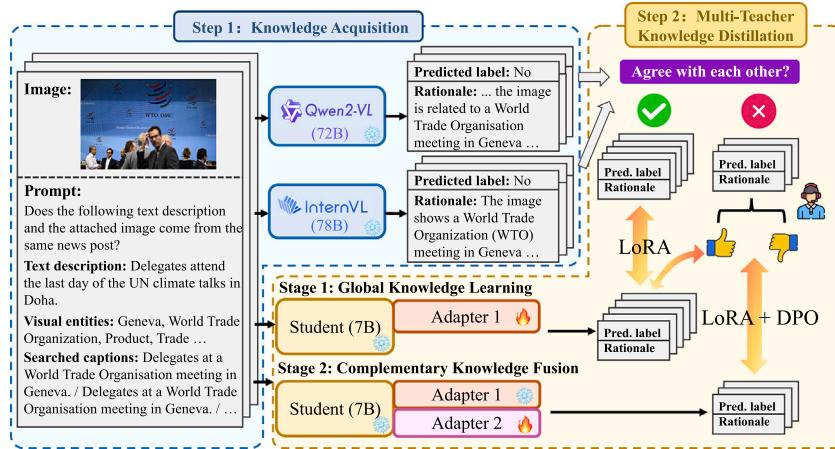


Fig. 4. Parameter-efficient tuning architecture for multimodal fake news detection. The framework operates in two stages: global knowledge learning and complementary knowledge fusion. In the first stage, a student model with 7 billion parameters is trained alongside two adapters, processing visual and textual inputs to identify relevant entities and search options. The second stage integrates the student model with the adapters to enhance reasoning. Knowledge acquisition is achieved through two vision-language models, Qwen2-VL and InternVL, which provide predictions and rationales for the input image. Multi-teacher knowledge distillation follows, where the student model and adapters align their predictions and rationales, facilitated by LoRA and DPO techniques to refine the model's outputs [143].

which interact to refine the reasoning process. This approach provides greater transparency and interpretability compared to prompting-based reasoning, allowing for more deliberate, human-like decision-making. It also facilitates error localization and enables strategies like counter argumentation, which are essential for handling complex, long context misinformation. By combining these two paradigms, LVLMs are able to surpass traditional models in robustness and interpretability, providing a more flexible and effective approach to multimodal fake news detection.

Prompting-based reasoning. Prompting-based reasoning methods treat multimodal fake news detection as a single pass inference problem, where reasoning is implicitly induced through carefully designed prompts, instructions, or constraints, without the need for explicit intermediate decision states or tool driven interactions. Given multimodal evidence $x = \{x_t, x_v\}$ and a claim c , the model directly produces a veracity prediction:

$$\hat{y} = f_{\Theta}(x_t, x_v, c; \mathcal{P}) \quad (37)$$

where x_t and x_v represent the textual and visual modalities of the multimodal evidence, respectively, and c denotes the claim whose veracity is being assessed. \mathcal{P} is a task specific prompt encoding reasoning cues, such as cross modal consistency, credibility, or relevance of the evidence.

This paradigm leverages the emergent reasoning capabilities of large vision-language models (LVLMs) or language models, using prompt engineering to guide the model toward detecting semantic inconsistencies, fabricated visual evidence, or misleading narratives. Representative approaches include ICD [45] and CAPE-FND [147], which mitigate multimodal hallucination and calibration errors through instruction level constraints, ensuring that models provide more accurate interpretations by imposing explicit guidance on reasoning pathways. These systems enhance the reliability of LVLMs by constraining the model's outputs to adhere more strictly to logical reasoning patterns, reducing the occurrence of hallucinations in multimodal content. NRFE [148] and IFAI [149], which employ confidence aware or feedback enhanced prompting to improve robustness, further address the challenge of multimodal inconsistency by adjusting the confidence levels in the model's output. NRFE utilizes feedback loops to refine reasoning accuracy, while IFAI focuses on boosting robustness in video based fake news detection through the strategic use of feedback enhanced prompts that guide the model's attention toward more reliable features. Specifically, IFAI addresses challenges in fake news video detection by constructing multimodal prompts through prompt engineering to semantically understand news videos and generate auxiliary information from the perspectives

of video style, content, and information matching. The video information interactor enables small models to effectively learn supplementary knowledge from LVLMs, while the key information selector evaluates the importance of inference rationales, thus improving the efficiency of utilizing knowledge from LVLMs. This mechanism significantly improves the model's capacity to discern subtle discrepancies between video content and textual narratives, enhancing detection performance in dynamic, real-world contexts. Recent works, such as DIFAR [150] and LIFE [153], demonstrate that structured prompts can replicate reasoning behaviors typically associated with explicit reasoning pipelines. DIFAR introduces a robust framework for prompt guided reasoning, leveraging structured prompts to replicate the reasoning steps commonly used in traditional, explicit pipelines, while LIFE focuses on long term effectiveness by applying structured prompt templates that align multimodal evidence over extended interactions. Benchmarks like MFC-Bench [151] and generative frameworks such as LLM-GAN [152] highlight both the potential and limitations of prompt induced reasoning, especially under adversarial or distribution shifted conditions. These benchmarks demonstrate that, while prompt based methods show significant promise in handling a wide range of multimodal scenarios, challenges remain in maintaining consistent reasoning under adversarial settings or when the distribution of data shifts significantly from the training conditions.

Despite their efficiency and strong zero-shot generalization, prompting-based methods inherently combine perception, reasoning, and decision-making within a single inference pass. This monolithic approach limits controllability and interpretability, particularly when dealing with complex, multi-hop misinformation.

Agent-based reasoning. Unlike prompting-based reasoning, which relies on implicit reasoning in a single inference step, agent-based reasoning explicitly decomposes the multimodal fake news detection process into a sequence of structured, multi-step tasks. This paradigm involves multiple specialized agents, each performing a distinct sub-task such as evidence retrieval, cross modal verification, contradiction detection, and final adjudication. These agents interact with one another to refine the reasoning process, ensuring that the detection system can handle complex scenarios requiring in-depth analysis. Formally, the inference process is modeled as a sequence of agent interactions:

$$S_{t+1} = \mathcal{A}_k(S_t, x, c), \quad k \in \{1, \dots, K\} \quad (38)$$

where S_t represents the intermediate reasoning state at time step t , and \mathcal{A}_k denotes a task-specific agent responsible for executing a particular sub-task.

Systems like FactAgent [44] and TED [154] exemplify the agent-based reasoning paradigm by coordinating multiple agents that iteratively verify claims using multimodal evidence, allowing for explicit control over the reasoning order and the selection of relevant evidence at each step. This coordination of agents enables more flexible and interpretable reasoning processes, as each agent can specialize in evaluating a particular type of evidence, whether textual, visual, or audio, and can adaptively select the most relevant information for the task at hand. These systems enable better traceability of the reasoning process, ensuring transparency and accountability in decision-making by offering a detailed log of the steps each agent took in reaching its conclusion. Furthermore, agent-based frameworks like SheepDog [19] and ARG [75] take this approach further by introducing adversarial or debate-style agents, which can engage in counter-argumentation, helping the system improve robustness against false or manipulated narratives. The adversarial agents engage in a structured exchange of opposing viewpoints, forcing the system to critically evaluate evidence from different perspectives and refine its reasoning, thus enhancing its ability to detect inconsistencies or biases in the evidence. By simulating a back-and-forth debate, these agents can identify subtle inconsistencies in the evidence, improving the system's ability to handle misleading or deceptive content. More recent architectures, such as FACTGUARD [46], have integrated tool-augmented agents that not only reason about the evidence but also maintain structured memory and rationale supervision. FACTGUARD introduces a memory mechanism that allows agents to retain and reference previous reasoning steps, making it possible to build a more coherent and structured argument over time. This approach enhances interpretability by providing clear, human readable reasoning traces that explain how conclusions are reached. Furthermore, these systems demonstrate improved generalization under complex, long-context misinformation scenarios, where the model must synthesize and analyze large volumes of data over extended reasoning chains. By incorporating memory and rationale supervision, these systems can maintain contextual continuity across longer interactions, improving the accuracy of fake news detection in scenarios where misinformation evolves gradually.

While agent-based reasoning systems are computationally more demanding and involve greater system complexity, they provide several key advantages. These include enhanced transparency, the ability to modularize reasoning processes, and improved error localization, all of which make agent-based methods highly suited for high stakes applications such as judicial decision-making, legal analysis, and evidence intensive fake news verification. By ensuring that each step in the reasoning process is explicitly documented and accountable, agent-based systems offer a level of trustworthiness and flexibility that is critical in real-world verification tasks.

4. Popular benchmark dataset

In recent years, the research progress in multimodal fake news detection has benefited greatly from the continuous construction and availability of high-quality datasets. Datasets not only provide a basis for model training and evaluation, but also largely define the boundaries and difficulty of research problems. Starting from early social media rumor detection datasets such as Twitter15 [186] and Twitter16 [187], researchers have gradually expanded from single text or image modalities to multimodal resources integrating text, images, videos and social context, and have continuously improved data scale, modality diversity and annotation precision, thereby promoting the development of more complex detection tasks that are closer to real-world scenarios. As shown in Fig. 5, looking back at the related work from 2015 to 2025, we can observe that dataset construction demonstrates the following four development trends as follows:

- **Data scale expansion:** Early datasets such as LIAR [66], Weibo [188], and PHEME [189] typically contained only a few thousand to tens of thousands of examples and were primarily used for training and validating small-scale models. As research deepened, subsequent datasets (such as M³A [93], Fakeddit [190], and NewsCLIPpings [191]) expanded to hundreds of thousands or even millions of examples, facilitating the application of deep learning and large-scale pre-trained models, and supporting the systematic evaluation of complex models at varying data scales.
- **Modal diversity:** Initially, data was primarily text-based. Subsequently, images (such as Weibo [188] and GoodNews [192]), image-text alignment data (such as NeuralNews [193] and COSMOS [7]), and multilingual content (such as MM-COVID [194] and CHECKED [195]) were introduced, and further expanded to include audio, video, and social interaction information (such as M³A [93] and MDAM³-DB [196]). This increased modal diversity not only provides the model with richer learning signals but also better reflects the complex spread of disinformation on real social platforms.
- **Task dimension expansion:** Early research focused primarily on binary classification tasks (authenticity and fake identification), but recent datasets have gradually expanded to support multi-class labeling (such as LIAR [66] and Fakeddit [190]), tampered region localization (such as DGM⁴ [88] and MFND [197]), image-text inconsistency detection (such as COSMOS [7] and NewsCLIPpings [191]), and AI-generated content recognition (such as MiRAGE News [121] and MMFakeBench). This expansion of task dimensions has greatly increased the challenge and application value of detection, enabling research to cover diverse objectives, from coarse-grained to fine-grained, and from discriminative to generative.

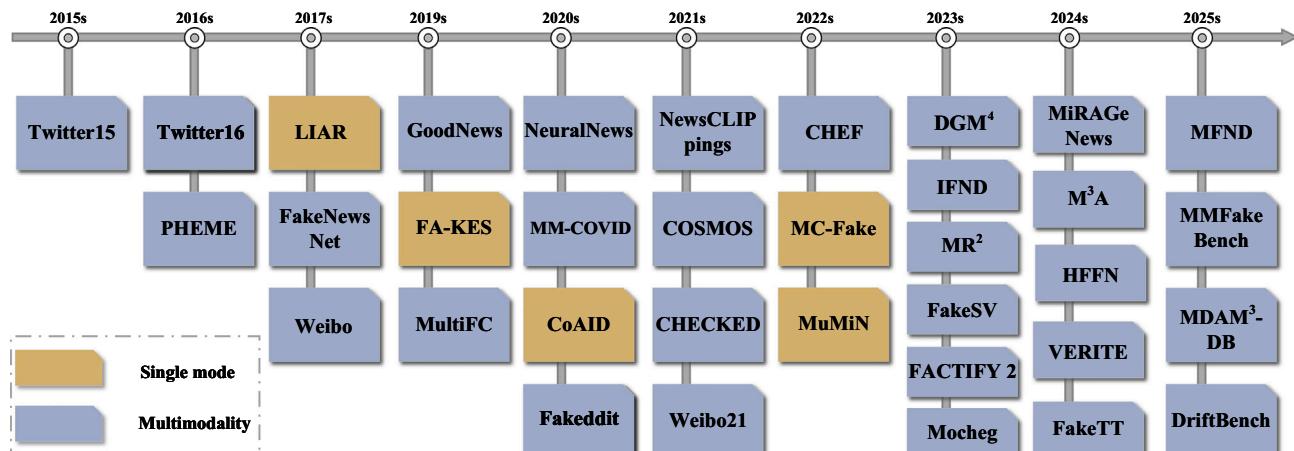


Fig. 5. Timeline of multimodal fake news detection datasets.

Table 1

Publicly available benchmark datasets in multimodal fake news detection.

Datasets	Year	Modality	Available at
Twitter15 [186]	2015	Text,Image	https://github.com/majingCUHK/Rumor_RvNN?tab=readme-ov-file
Twitter16 [187]	2016	Text,Image	https://github.com/majingCUHK/Rumor_RvNN?tab=readme-ov-file
PHEME [189]	2016	Text,Image	https://www.kaggle.com/datasets/usharengaraju/pHEME-dataset
LIAR [66]	2017	Text	https://www.kaggle.com/datasets/yuktibishambu/liar-dataset-labeled
FakeNewsNet [198]	2017	Text,Image	https://github.com/KaiDMML/FakeNewsNet
Weibo [188]	2017	Text,Image	https://github.com/plw-study/MRML?tab=readme-ov-file
GoodNews [192]	2019	Text,Image	https://github.com/furkanbiten/GoodNews
FA-KES [199]	2019	Text	https://www.kaggle.com/datasets/mohamadalhasan
MultiFC [199]	2019	Text,Image	https://huggingface.co/datasets/pszemraj/multi_fc
Fakeddit [190]	2020	Text,Image	https://github.com/entitiz/fakeddit
NeuralNews [193]	2020	Text,Image	https://drive.google.com/file/d/1vD4DtyJ0IjRzchPtCQu-KPrUjgTiWSmo/view
MM-COVID [194]	2020	Text,Image	https://drive.google.com/drive/folders/1gd4AvT6BxPRtymNd9Z7ukyaVhae5s7U
CoAID [200]	2020	Text	https://github.com/cuilingm/CoAID
NewsCLIPpings [191]	2021	Text,Image	https://github.com/g-luo/news_clippings?tab=readme-ov-file
COSMOS [7]	2021	Text,Image	https://shivangi-aneja.github.io/projects/cosmos/
CHECKED [195]	2021	Text,Image	https://github.com/cyang03/CHECKED/tree/master/dataset
Weibo21 [164]	2021	Text,Image	https://github.com/kennqiang/MDFEND-Weibo21
CHEF [201]	2022	Text	https://github.com/THU-BPM/CHEF?tab=readme-ov-file
MC-Fake [202]	2022	Text	https://github.com/qwerfdsapking/MC-Fake
MuMin [203]	2022	Text,Image	https://mumin-dataset.github.io/
DGM ⁴ [88]	2023	Text,Image	https://huggingface.co/datasets/rshaojimmy/DGM4
IFND [204]	2023	Text,Image	https://www.kaggle.com/datasets/sonalgarg174/ifnd-dataset
MR ² [205]	2023	Text,Image	https://github.com/THU-BPM/MR2
MocheG [206]	2023	Text,Image	https://github.com/PLUM-Lab/MocheG
FakeSV [207]	2023	Text,Vedio	https://github.com/ICTMCG/FakeSV
FACTIFY 2 [208]	2023	Text,Image	https://github.com/surya1701/Factify-2.0?tab=readme-ov-file
MiRAGENews [121]	2024	Text,Image	https://huggingface.co/datasets/anson-huang/mirage-news
M ³ A [93]	2024	Text,Image,Audio,Vedio	https://github.com/FinalYou/M3A?tab=readme-ov-file
HFFN [142]	2024	Text,Image	–
VERITE [209]	2024	Text,Image	https://github.com/stevejpapad/image-text-verification
FakeTT [210]	2024	Text,Vedio	https://github.com/ICTMCG/FakingRecipe/tree/main?tab=readme-ov-file
MFND [197]	2025	Text,Image	https://github.com/yunan-wang33/sdml
MMFakeBench [211]	2025	Text,Image	https://huggingface.co/datasets/liuxuannan/MMFakeBench
MDAM ³ -DB [196]	2025	Text,Image,Audio,Vedio	–
DriftBench [125]	2025	Text,Image	–

- Cross-language and cross-domain adaptation:** With the widespread global spread of disinformation, the need for cross-language and cross-domain detection has become increasingly prominent. Some datasets (such as MM-COVID [194], Weibo21 [164], and CHEF) focus on multilingual, cross-topic, and cross-domain authenticity detection. These datasets not only enrich the training corpus but also provide important support for research on model transferability and cross-cultural adaptability.

Overall, dataset evolution over the past decade has seen a gradual transition from early rumor detection to complex cross-modal disinformation identification. These datasets have continuously improved in terms of authenticity annotation granularity, modal coverage, cross-lingual adaptability, and generative content detection. These datasets not only provide a solid experimental benchmark for the research community but also drive the field's progress from traditional classification tasks to more complex multi-task, multi-modal, and multi-lingual approaches, forming a relatively complete evolutionary trajectory and research ecosystem. **Table 1** lists 32 fake news detection benchmark datasets. We analyze the release time, modality, and open-source URL of each dataset. Next, we briefly introduce each dataset.

4.1. Overview of benchmark datasets

Twitter15. The Twitter15 dataset [186], created by crawling two rumor-tracking websites, Snopes¹ and Emergent,² collects 2299 news items published up to March 2015. After screening, it contains 94 true news items and 446 false news items. To obtain tweets related to an

event, a keyword-based query was constructed, and a web crawler was used to obtain the complete history. This was then sampled and cross-checked by researchers. Furthermore, real events were collected using Twitter's free data stream and a clustering algorithm. The final dataset contains 421 true events and 421 false events.

Twitter16. The Twitter16 dataset [187], based on Snopes,³ an online rumor-debunking service, collects 778 events from March to December 2015, 64% of which are rumors. Researchers extracted and optimized keywords to obtain relevant tweets, supplemented with public datasets, ultimately creating a balanced dataset containing 498 rumors and 494 non-rumors, which is widely used in rumor detection research.

PHEME. The PHEME dataset [189], constructed through a combination of automated crawling and manual verification, covers tweets, comments, and user interaction data, such as retweets and likes, related to specific events. Its multi-dimensional information structure and high-quality annotations, including the differentiation of rumor types, provide a rich resource for rumor detection, information dissemination analysis, and user behavior research, supporting machine learning model training and cross-event comparative research.

LIAR. LIAR [66] is a multimodal fake news detection dataset constructed in 2017. It contains 12,836 manually annotated sentences from PolitiFact, covering a variety of contexts including news, speeches, interviews, advertisements, and social media. The dataset provides fine-grained authenticity annotations (six-category labels) and rich metadata, including speaker identity, context, and historical credibility. It is suitable for research in disinformation detection, credibility analysis, and multimodal modeling, providing an important benchmark resource in the field.

¹ snopes.com

² emergent.info

³ snopes.com

FakeNewsNet. The FakeNewsNet dataset [198] is a multidimensional data repository containing two fact-checking datasets based on Politifact and Gossipcop, covering news content, social context, and spatiotemporal information. Constructed by the FakeNewsTracker system, this dataset aims to advance open research questions in the field of fake news research.

Weibo. The Weibo dataset [188], built from Weibo's official rumor-busting system and Xinhua News Agency news sources, covers verified rumor and non-rumor posts from 2012–2016. The dataset contains approximately 40,000 tweets with images, including text, images, and social context. The dataset is deduplicated and quality-screened using the LSH algorithm, and the training/test set is split into an 8:2 ratio based on event clustering to minimize data leakage. This dataset provides an authoritative and high-quality benchmark for multimodal rumor detection.

GoodNews. The GoodNews dataset [192], sourced from The New York Times, contains over 460,000 news images, corresponding articles, and headlines, making it one of the largest multimodal news resources available. The average article length exceeds 650 words, 97% of headlines contain named entities, and 68% contain human names, fully demonstrating the complexity and semantic relevance of news text. This dataset not only supports news image and text generation and alignment tasks but also provides a core data source for building multimodal fake news detection benchmarks such as NeuralNews. Because approximately half of the news texts exceed 512 tokens, GoodNews presents significant challenges for long-text modeling, multimodal alignment, and named entity recognition.

FA-KES. The FA-KES dataset [199], constructed for diverse media coverage of the Syrian war, utilizes a semi-supervised annotation process and fact-checking mechanism to collect approximately 804 English news articles (with a nearly balanced distribution of true and false news). Its rich structured information (including title, date, source, etc.) and reliable label generation process provide a solid foundation for fake news detection, especially for research on meta-learning, weak supervision, and multi-feature fusion models in few-shot scenarios. Its significant generalizability makes it suitable for dataset construction and algorithm validation in other military conflict scenarios.

MultiFC. The MultiFC dataset [199], collected from 26 English-language fact-checking websites, covers 34,918 naturally occurring factual claims, accompanied by supporting evidence, context, and rich metadata, all annotated by professional journalists. Its core features include the entities involved in the claims, contextual information, and multi-dimensional metadata. These additional attributes significantly improve model performance in automated claim verification tasks. MultiFC provides a solid foundation for the development and evaluation of fact-checking models.

Fakeddit. The Fakeddit dataset [190], constructed from 22 Reddit subreddits, covers topics ranging from politics to everyday life and contains over one million posts. It undergoes a multi-stage review process and employs distant supervision to provide 2/3/6 classification labels. The data includes text, images, metadata, and comments, with approximately 64% of the samples being multimodal, supporting research on multimodal and hierarchical fake news detection.

NeuralNews. The NeuralNews dataset [193], built on the GoodNews dataset, covers approximately 128,000 real and machine-generated news articles, divided into four categories (combinations of real/generated articles and headlines). The machine-generated content, generated by models such as Grover and entity-aware image caption models, includes text, images, and headline information. Its multimodal and fine-grained design provides a more realistic and challenging benchmark for detecting machine-generated news, and is of great significance for advancing research in multimodal fake news detection.

MM-COVID. MM-COVID [194] is a multilingual, multimodal dataset for COVID-19 fake news detection. It covers six languages and contains 3981 fake news items and 7192 real news samples. The data was collected from February to July 2020 and comes from sources including

social media, traditional media, blogs, and fact-checking organizations. This dataset integrates multimodal content, preserves user interactions, and retains timestamp information. It supports cutting-edge tasks such as cross-language detection, multimodal fusion, and social and spatiotemporal feature analysis, providing key support for building efficient and generalizable fake news detection models.

CoAID. CoAID [200] is a multimodal dataset for detecting COVID-19 health disinformation. It covers 4251 news articles and statements, 296,000 user tweets and replies, and 926 social media posts collected from real-world online environments between December 2019 and July 2020. CoAID supports a variety of tasks, including multimodal fake news detection and social contextual propagation analysis, providing a critical data foundation for building efficient and generalizable disinformation identification models in real-world scenarios.

Weibo21. The Weibo 21 dataset [164] was released in 2021 and covers 4488 fake news and 4640 real news from 9 different fields (i.e., Science, Military, Education, Disasters, Politics, Health, Finance, Entertainment, Society). The dataset was created by collecting fake news and real news on Sina Weibo between December 2014 and March 2021.

NewsCLIPpings. The NewsCLIPpings dataset [191], built on VisualNews, addresses the problem of mismatched images and misleading headlines. By generating forged samples using multiple strategies and combining them with CLIP filtering, we have generated 1111,828 balanced image–text pairs. This dataset is challenging and avoids unimodal bias, making it a popular tool for evaluating the performance of multimodal models in image–text inconsistency detection.

COSMOS. COSMOS [7] is a typical multimodal dataset focused on detecting contextual inconsistencies between images and captions to support research on disinformation detection. The dataset consists of 160,000 training samples, 40,000 validation samples, and 1700 test samples. Each image is assigned up to 10 bounding boxes, and the data is primarily sourced from news websites and the Snopes platform. Unlike traditional annotation methods, COSMOS does not directly use contextual misuse annotations during training, but only incorporates them during evaluation. This significantly increases the challenge of the task and enhances the generalization capabilities of the model. This dataset provides a high-quality benchmark for multimodal disinformation detection, and is of great significance for comparative studies and method evaluation.

CHECKED. CHECKED [195] is the first Chinese multimodal dataset for COVID-19-related fake news detection, covering 2104 samples collected from Weibo between December 2019 and August 2020, including 344 fake news and 1760 real news. The dataset integrates multimodal information such as text, images, and videos, and provides large-scale social context data, including more than 1.86 million reposts, 1.18 million comments, and 56.85 million likes. CHECKED provides important benchmark experimental support for multimodal fake news detection, propagation behavior analysis, and early identification.

CHEF. CHEF [201] is a multi-domain Chinese dataset for evidence-based fact-checking, covering politics, public health, science, society, and culture. The dataset contains 10,000 manually verified true claims, each accompanied by manually curated and annotated evidence collected from the internet, including text and some images and videos, to ensure the reliability and accuracy of the annotations. The data source comes from six Chinese fact-checking and news websites (e.g., Piya, TFC, MyGoPen, Jiaozhen, and Cnews), and is divided into a training set of 5754 claims, a validation set of 666 claims, and a test set of 666 claims. Multiple rounds of annotation and verification ensure annotation consistency, with a Fleiss Kappa of 0.74.

MC-Fake. The MC-Fake dataset [202], collected from Twitter, covers 28,334 news events across five major themes: politics, entertainment, health, COVID-19, and the Syrian War. This dataset is unique in that it provides both news text and rich social context, including tweets, retweets, replies, user attributes, and their social connections. Compared to existing content-based fake news benchmarks, MC-Fake complements

existing content-based fake news benchmarks by incorporating social context and relationship networks. This provides a new experimental platform for studying the cross-modal characteristics of fake information and holds significant potential for open research.

MuMiN. The MuMiN dataset [203], built on Twitter, covers 41 languages. It integrates multimodal information, including tweets, replies, users, images, articles, and hashtags, spanning over a decade and covering a wide range of topics and events. This dataset offers significant advantages in scale, diversity, and cross-linguality, providing a large-scale benchmark for studying cross-modal modeling, cross-lingual propagation, and the long-term evolution of disinformation.

DGM⁴. The DGM⁴ dataset ([88]) is a large-scale dataset for detecting and localizing multimodal media manipulation. Built on VisualNews, it covers real-world sources from The Guardian, BBC, USA Today, and The Washington Post. It contains 230,000 news image and text samples, including 77,426 pairs of original images, 152,574 pairs of manipulated images (e.g., face swaps, attribute edits, and text replacements), and 32,693 pairs of hybrid images. The dataset is constructed with a distribution of nine manipulation categories. All samples are annotated with fine-grained and sentimentally balanced annotations, making it more challenging than existing datasets and suitable for multimodal forgery detection and localization research.

IFND. The IFND dataset [204] covers multimodal news related to India between 2013 and 2021, totaling 56,868 text and image samples. Real news comes from mainstream media outlets such as Times Now News and The Indian Express, while false news is collected by authoritative fact-checking platforms such as Alt News and Boom Live and is manually verified for labeling accuracy. To alleviate the category imbalance, researchers introduced an intelligent enhancement algorithm to generate semantically plausible false statements and, based on LDA topic modeling, categorized the news into five major categories: elections, politics, COVID-19, violence, and others. IFND is widely used for performance evaluation of machine learning and deep learning models, expanding the research boundaries of multimodal disinformation detection.

MR². The MR² dataset [205] consists of two subsets, Weibo and Twitter, covering both text and image news, and provides external evidence retrieved from the internet for both modalities. The dataset supports both Chinese and English, covers multiple fields including politics, society, technology, and entertainment, and contains a rich collection of text, images, and webpage information, enabling a more realistic reflection of the cross-platform rumor propagation and verification process. With its bilingual nature and multimodal design, MR² provides an important benchmark for the training and evaluation of multimodal rumor detection models, making it particularly suitable for lightweight CNN experiments and content moderation scenarios.

Mocheg. The Mocheg dataset [206], constructed from fact-checked claims from PolitiFact and Snopes, contains 15,601 truth-labeled claims, along with 33,880 text paragraphs and 12,112 images as evidence. This dataset is unique in that it provides both textual and image evidence, supporting, for the first time, evidence retrieval, multimodal fact-checking, and explanation generation within a single, end-to-end framework. As a key benchmark for multimodal fact-checking, Mocheg is not only suitable for validating model detection capabilities but also provides a systematic evaluation platform for interpretability research.

FakeSV. The FakeSV dataset [207] is the largest Chinese short-video benchmark for multimodal fake news detection, containing rich social context information. It includes news video content, user comments, and publisher profiles, providing a comprehensive view of the dissemination environment of short video news. FakeSV enables multimodal analysis by jointly leveraging visual, textual, and social modalities, addressing the limited exploitation of multimodal correlations in prior works. In addition to fake/real labels, the dataset supports exploratory analysis of fake news propagation characteristics across content and social dimensions. To establish baseline performance, the authors further introduce a multimodal detection model named SV-FEND, which adaptively exploits

cross-modal correlations and social contextual cues to enhance detection accuracy. Overall, FakeSV bridges the gap between multimodal content understanding and social behavior analysis, offering a valuable resource for future research on fake news detection in short video platforms.

FACTIFY 2. The FACTIFY 2 dataset [208] is a large-scale multimodal fact-checking benchmark that extends FACTIFY 1 by incorporating new data sources and adding satire news articles, resulting in 50,000 additional data instances. Each sample consists of textual claims paired with visual content, enabling research on multimodal verification. Similar to its predecessor, FACTIFY 2 categorizes samples into three broad labels—support, no-evidence, and refute—with sub-categories that reflect the entailment between text and image. The dataset provides a foundation for developing and evaluating models that jointly reason over textual and visual information. Baseline experiments using BERT and Vision Transformer architectures achieve a test F1 score of 65%, highlighting the challenge of multimodal fact verification.

MiRAGENews. The MiRAGENews dataset [121] is a multimodal benchmark designed specifically for AI-generated fake news detection. It contains 12,500 pairs of real and generated image-headline examples. Real examples are sourced from The New York Times articles in the TARA dataset, while the fake ones are generated using GPT-4 headlines and Midjourney V5.2 images, ensuring the content is highly realistic and misleading. This dataset includes training, validation, and cross-generator/publisher test sets, emphasizing the model's generalization capabilities.

HFFN. Guided by the core principles of “human-centeredness” and “factual relevance,” the HFFN dataset [142] comprises multimodal samples consisting of image-text pairs covering entertainment, sports, politics, and other fields. These samples are generated through image manipulation, text manipulation, and fact manipulation, and are accompanied by detailed human annotations. This significantly enhances the research value of fake news detection models in both authenticity judgment and manipulation reasoning.

VERITE. VERITE [209] is a dataset designed specifically for multimodal disinformation detection, aiming to address the unimodal bias commonly found in existing datasets. Constructed from real-world news and social media image and text pairs, VERITE ensures that detection relies on cross-modal information rather than unimodal shortcuts by balancing modalities and eliminating asymmetric multimodal disinformation.

M³A. M³A [93] is a large-scale multimodal disinformation dataset covering text, images, audio, and video. Collected from 60 leading news outlets worldwide, the dataset contains 708,425 real news items and 6,566,386 fake news items. It provides multi-category, fine-grained topic and sentiment annotations and serves as a unified benchmark for various disinformation detection tasks, such as out-of-context detection and deepfake detection. This dataset aims to promote the development of robust multimodal disinformation analysis techniques.

FakeTT. The FakeTT dataset [210] is an English-language benchmark specifically designed for fake news detection on short video platforms. It contains a large collection of short videos paired with textual descriptions, user comments, and video metadata, capturing the rich but heterogeneous multimodal information inherent in short-form content. Unlike previous datasets that focus mainly on content analysis, FakeTT emphasizes the creative process behind video production, providing insights into material selection and editing patterns commonly found in fake news videos. In addition to real/fake labels, FakeTT supports research on creative-process-aware detection, enabling models to learn from sentimental, semantic, spatial, and temporal cues in video content. Together with the existing Chinese FakeSV dataset, FakeTT provides a comprehensive resource for developing and evaluating multimodal models that consider both content and production characteristics in fake news detection on short video platforms.

MFND. The MFND dataset [197] (released on May 11, 2025) contains 125,000 multimodal news samples across four combinations: real image-real text (RIRT), fake image-real text (FIRT), real image-fake text

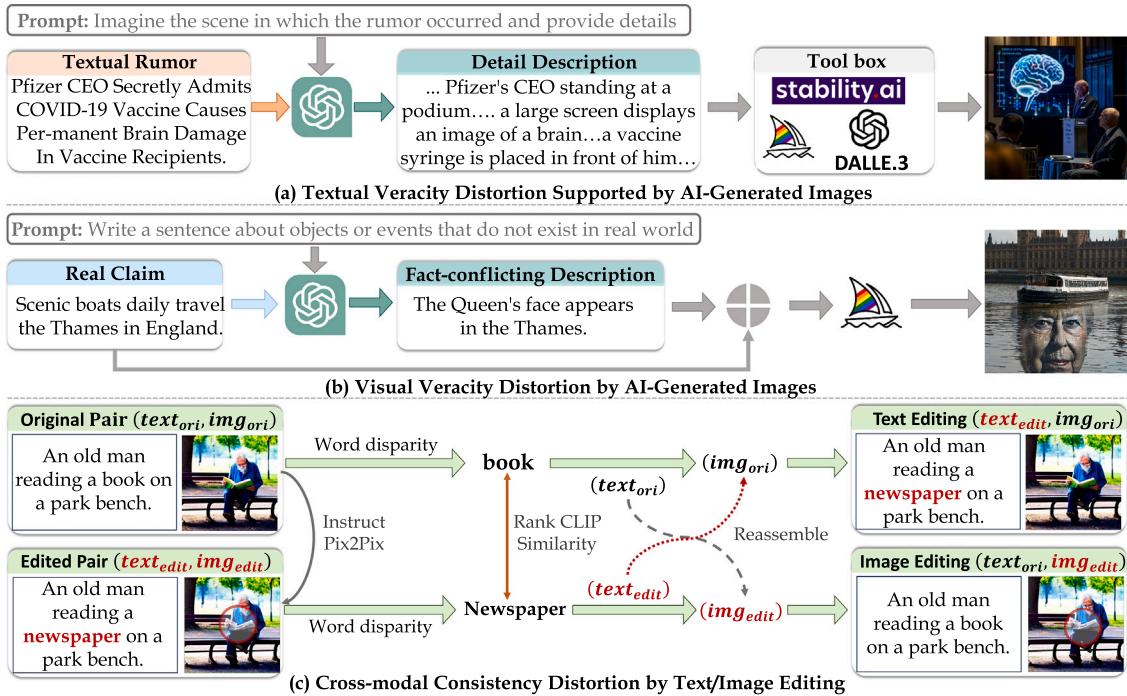


Fig. 6. Instructions for constructing the MMFakeBench dataset. (a) AI-generated images support text authenticity distortion. (b) Visual authenticity distortion is caused by generating images that conflict with facts. (c) Cross-modal consistency distortion is caused by text/image editing [211].

(RIFT), and fake image-fake text (FIFT). The fabricated content is constructed using 11 advanced generative techniques, including StyleGAN3 for synthesized images and multimodal large models (LVLMs) for text generation. In addition to true and fake labels, MFND also provides labels for image and text manipulation detection and precise annotation of manipulated image regions. Compared to existing datasets, MFND is more authentic to real-world communication scenarios and supports multi-task research such as detection and localization.

MMFakeBench. As shown in Fig. 6, the MMFakeBench dataset [211] (released in 2025) collects multimodal disinformation samples from various sources, covering both image and text modalities. The dataset provides complete annotations for each sample, including text content, image path, source information, and binary and multi-class labels. As a mixed-source multimodal disinformation detection benchmark, MMFakeBench is particularly suitable for evaluating the performance of large-scale vision-language models in multimodal fake news detection.

MDAM³-DB. MDAM³-DB [196] is a comprehensive multimodal disinformation detection benchmark consisting of 90,000 text, image, video, and audio samples. This dataset not only supports disinformation detection and modal inconsistency modeling across multimodal inputs, but also covers various deception scenarios, including AI-generated content, factual conflicts, and cross-modal mismatches. MDAM³-DB was built by integrating an interpretable analysis module driven by a large-scale visual-language model (LVLM). Its usability and practical performance were validated through systematic user studies. MDAM³-DB provides a solid foundation for promoting more comprehensive, transparent, and trustworthy disinformation detection research.

DriftBench. The DriftBench dataset [125] contains 16,000 news articles with images and text, covering six major topics. It focuses on simulating two typical scenarios: genre drift and evidence drift. Its tasks cover authenticity verification, adversarial evidence detection, and cross-variant consistency reasoning. DriftBench not only reveals the vulnerabilities of LVLM in the GenAI era but also provides a key benchmark for research on robustness and generalization methods.

4.2. Comparison and taxonomy of misinformation datasets

To comprehensively understand the landscape of misinformation datasets, Table 2 categorizes 32 representative benchmarks across three major axes of veracity distortion: (i) Textual Veracity Distortion, (ii) Visual Veracity Distortion, and (iii) Cross-modal Consistency Distortion. Each axis captures different aspects of misinformation generation, ranging from textual fabrication to visual manipulation and semantic inconsistency between modalities. Specifically:

- Textual veracity distortion.** This category involves falsified or manipulated textual content that misrepresents facts. Datasets such as Twitter15/16, LIAR, and FakeNewsNet primarily focus on this dimension, providing annotated claims and news statements labeled as true, false, or unverified. Some recent datasets, including FA-KES and GoodNews, extend this by incorporating AI-generated or repurposed text samples to simulate emerging misinformation patterns. These resources enable fine-grained studies on linguistic cues, contextual framing, and source reliability.
- Visual veracity distortion.** The second dimension focuses on visual manipulations, including photo editing, compositional blending, and AI-synthesized imagery. Datasets such as Fakeddit, NeuralNews, and MM-COVID explicitly mark images that have been Photoshop-edited (PS-edited) or generated by diffusion and GAN-based models. Unlike textual datasets, these visual corpora are often multimodal, pairing each image with a corresponding textual caption or headline, thus supporting research in multimodal forgery detection and image-text alignment. This line of datasets becomes increasingly important with the rise of generative models producing hyper-realistic yet fabricated content.
- Cross-modal consistency distortion.** Cross-modal distortion refers to semantic or factual mismatches between modalities, for example, when the text claims an event that the image does not depict. Datasets such as Weibo, FakeNewsNet, and COSMOS explicitly annotate image-text repurposing or editing

Table 2

Comparison of misinformation datasets across three categories of distortion: textual veracity, visual veracity, and cross-modal consistency.

Dataset	Textual Veracity Distortion			Visual Veracity Distortion			Cross-modal Consistency Distortion	
	Text (Rumor)	Supporting Image		Text (Veracity)	Fact-conflicting Image		Image/Text Repurposing	Image/Text Editing
		Repurposed	AI-generated		PS-edited	AI-generated		
Twitter15 [186]	✓	✗	✗	✗	✗	✗	✗	✗
Twitter16 [187]	✓	✗	✗	✗	✗	✗	✗	✗
PHEME [189]	✓	✗	✗	✗	✗	✗	✗	✗
LIAR [66]	✓	✗	✗	✗	✗	✗	✗	✗
FakeNewsNet [198]	✓	✗	✗	✗	✗	✗	✗	✗
Weibo [188]	✓	✗	✗	✗	✗	✗	✗	✗
GoodNews [192]	✓	✗	✗	✗	✗	✗	✗	✗
FA-KES [199]	✓	✗	✗	✗	✗	✗	✗	✗
MultiFC [199]	✓	✗	✗	✗	✗	✗	✗	✗
Fakeddit [190]	✗	✗	✗	✗	✗	✗	✗	✗
NeuralNews [193]	✓	✗	✗	✗	✗	✗	✗	✗
MM-COVID [194]	✓	✗	✗	✗	✗	✗	✗	✗
CoAID [200]	✓	✗	✗	✗	✗	✗	✗	✗
NewsCLIPpings [191]	✗	✗	✗	✗	✗	✗	✓	✗
COSMOS [7]	✗	✗	✗	✗	✗	✗	✗	✗
CHECKED [195]	✓	✗	✗	✗	✗	✗	✗	✗
Weibo21 [164]	✓	✗	✗	✗	✗	✗	✗	✗
CHEF [201]	✓	✗	✗	✗	✗	✗	✗	✗
MC-Fake [202]	✓	✗	✗	✗	✗	✗	✗	✗
MuMiN [203]	✓	✗	✗	✗	✗	✗	✗	✗
DGM ⁴ [88]	✗	✗	✗	✗	✗	✗	✗	✓
IFND [204]	✓	✗	✗	✗	✗	✗	✗	✓
MR ² [205]	✓	✗	✗	✗	✗	✗	✗	✗
Mocheg [206]	✓	✗	✗	✗	✗	✗	✗	✗
FakeSV [207]	✓	✗	✗	✗	✗	✗	✗	✗
FACTIFY 2 [208]	✓	✗	✗	✗	✗	✗	✗	✗
MiRAGENews [121]	✗	✓	✓	✓	✓	✓	✓	✓
M ³ A [93]	✗	✓	✓	✓	✓	✓	✓	✓
HFFN [142]	✗	✓	✓	✓	✓	✓	✓	✓
VERITE [209]	✓	✗	✗	✗	✗	✗	✗	✗
FakeTT [210]	✓	✗	✗	✗	✗	✗	✗	✗
MFND [197]	✓	✓	✓	✓	✓	✓	✓	✓
MMFakeBench [211]	✓	✓	✓	✓	✓	✓	✓	✓
MDAM ³ -DB [196]	✓	✓	✓	✓	✓	✓	✓	✓
DriftBench [125]	✓	✓	✓	✓	✓	✓	✓	✓

inconsistencies, making them valuable for multimodal reasoning tasks. More recent datasets like M3FD and VERITE further introduce controlled repurposing and image–text editing tasks to support explainable evaluation.

In summary, as summarized in **Table 2**, most early rumor datasets (e.g., Twitter15/16, PHEME, LIAR) emphasize textual veracity classification, whereas later benchmarks gradually incorporate visual and cross-modal distortions to reflect real-world misinformation complexity. Notably, datasets such as MM-COVID, M3FD, and DriftBench cover all three dimensions, providing a holistic platform for evaluating multimodal reasoning and robustness. The trend indicates a paradigm shift from single-modality rumor detection toward comprehensive multimodal misinformation understanding.

5. Evaluation metrics

In multimodal fake news detection, evaluation metrics not only serve as fundamental tools to assess model performance but also guide research progress. Unlike traditional text classification tasks, fake news detection involves high-risk, cross-modal, and adversarial characteristics. Therefore, its evaluation framework must systematically

consider classification performance, robustness and generalization, cross-modal consistency, and interpretability and user perception. The following subsections provide a structured overview of commonly used metrics, along with mathematical definitions.

In the evaluation of multimodal fake news detection systems, particularly those framed as classification problems, standard performance metrics provide essential insights into model behavior across diverse and often imbalanced datasets. Given the binary or multi-class nature of fake news categorization (e.g., real vs. fake, or fine-grained labels such as satire, misleading, fabricated), the following metrics are widely adopted to quantify predictive efficacy.

Accuracy measures the proportion of correctly classified instances among all predictions.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (39)$$

where TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively. More specifically, TP is the number of samples that are actually “fake news” and are correctly judged as “fake news” by the model, TN is the number of samples that are actually “real news” and are correctly judged as “real news” by the model,

FP is the number of samples that are actually “real news” but are mistakenly judged as “fake news” by the model, FN is the number of samples that are actually “fake news” but are mistakenly judged as “real news” by the model. While intuitive, accuracy can be misleading under class imbalance, a common scenario in fake news datasets, where genuine news often dominates.

To address this limitation, Precision and Recall offer class-specific perspectives as follows:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad \text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (40)$$

Precision reflects the reliability of positive predictions (i.e., how many flagged items are truly fake), whereas recall indicates coverage (i.e., how many actual fake instances are detected). In safety-critical applications like misinformation mitigation, high recall is often prioritized to minimize undetected disinformation.

F1 Score harmonizes precision and recall via their harmonic mean as follows:

$$\text{F1} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (41)$$

This single-value metric is especially informative for imbalanced settings, as it penalizes extreme disparities between precision and recall.

For multi-class fake news taxonomies (e.g., distinguishing clickbait, deepfake, conspiracy), Macro-F1 and Micro-F1 extend the F1 score to aggregate performance across all classes. **Macro-F1** computes the unweighted mean of per-class F1 scores, treating all classes equally regardless of size as follows:

$$\text{Macro-F1} = \frac{1}{C} \sum_{i=1}^C \text{F1}_i \quad (42)$$

where C is the number of classes. This metric highlights model performance on minority classes, which is crucial when rare but high-impact misinformation types must be identified.

Conversely, **Micro-F1** aggregates TP, FP, and FN globally before computing precision and recall, effectively weighting each class by its frequency:

$$\text{Micro-F1} = \frac{2 \cdot \sum_{i=1}^C \text{TP}_i}{2 \cdot \sum_{i=1}^C \text{TP}_i + \sum_{i=1}^C \text{FP}_i + \sum_{i=1}^C \text{FN}_i} \quad (43)$$

Micro-F1 thus reflects overall instance-level performance and aligns closely with accuracy in balanced scenarios, but remains robust to label skew through its grounding in contingency counts.

6. Experimental performance

Due to the heterogeneous nature of multimodal misinformation detection benchmarks, this survey follows the original evaluation protocols, baseline selections, and metrics defined in the corresponding benchmark papers, rather than enforcing a unified setting across all datasets. The use of different proprietary LVLM baselines (e.g., GPT-4V in Tables 3–5 and GPT-4o-mini in Table 7) reflects dataset-specific evaluation objectives and historical contexts. In particular, GPT-4V is adopted in earlier benchmarks as a representative high-capacity LVLM, whereas GPT-4o-mini is used in DriftBench to reflect deployment-oriented evaluation under controlled diversity and distribution shift, consistent with the original benchmark design. Each table therefore constitutes a self-contained evaluation environment, where all compared methods share the same reference baseline, ensuring fair within-benchmark comparison. Similarly, evaluation metrics are retained as defined by each dataset to account for differences in task formulation, class imbalance, and annotation granularity. Metrics such as macro-F1, accuracy, AUC,

Table 3

Binary overall results of different models on the MM-FakeBench validation and test set with the comparison of standard prompting (Standard) and proposed MMD-Agent framework. The best results are **bolded**.

Model name	Language model	Prompt method	Validation (1000)				Test (10000)			
			F1	Precision	Recall	ACC	F1	Precision	Recall	ACC
Human evaluation			54.9	56.6	57.8	56.8	–	–	–	–
LVLMs with 7B parameter										
Otter-Image [212]	MPT-7B	Standard	7.9	4.1	4.5	7.9	8.6	32.4	5.0	8.6
MiniGPT4 [213]	Vicuna-7B	Standard	40.4	38.2	45.7	63.1	41.7	41.0	47.4	65.2
InstructBLIP [95]	Vicuna-7B	Standard	14.7	30.8	13.2	8.1	16.1	40.5	14.2	8.8
Qwen-VL [214]	Qwen-7B	Standard	43.6	50.6	44.9	60.3	44.0	51.6	45.2	60.5
VILA [215]	LLaMA2-7B	Standard	41.2	35.0	50.0	70.0	41.2	35.0	50.0	70.0
PandaGPT [216]	Vicuna-7B	Standard	24.6	60.6	50.5	30.9	24.1	61.7	50.4	30.6
mPLUG-Owl2 [217]	LLaMA2-7B	Standard	47.2	64.9	52.3	70.6	48.7	71.1	53.3	71.4
BLIP2 [94]	FlanT5-XL	Standard	41.2	35.0	50.0	70.0	41.2	35.0	50.0	70.0
LLaVA-1.6 [218]	Vicuna-7B	Standard	48.1	48.2	48.5	59.5	52.5	53.0	52.6	62.5
LVLMs with 13B parameter										
		Standard	41.1	35.0	50.0	70.0	41.1	35.0	50.0	70.0
VILA [215]	LLaMA2-13B	MMD-Agent	56.5	62.2	56.9	70.3	56.6	64.3	57.2	71.2
		Standard	41.1	35.0	49.9	69.9	41.1	35.0	49.9	69.8
InstructBLIP [95]	Vicuna-13B	MMD-Agent	51.3	53.4	54.0	53.1	47.9	50.1	50.1	49.9
		Standard	31.6	63.4	53.6	35.5	30.6	64.9	53.4	34.9
BLIP2 [94]	FlanT5-XXL	MMD-Agent	51.5	53.4	54.0	53.6	51.8	54.0	54.7	53.5
		Standard	41.1	35.0	50.0	69.7	42.3	57.3	50.1	69.5
LLaVA-1.6 [218]	Vicuna-13B	MMD-Agent	51.8	66.7	54.6	71.4	50.2	67.3	53.9	71.3
LVLMs with 34B parameter										
		Standard	62.9	67.1	70.0	63.4	64.3	68.8	71.7	64.8
LLaVA-1.6 [218]	Nous-Hermes-2 -Yi-34B	MMD-Agent	67.2	70.4	66.0	75.1	68.1	71.1	67.0	75.6
Proprietary LVLMs										
		Standard	72.3	72.1	72.8	75.6	74.2	73.5	76.9	76.4
GPT-4V [96]	ChatGPT	MMD-Agent	74.0	73.4	75.5	76.8	72.8	72.4	75.4	75.0

Table 4

Performance comparison of baseline methods on the Twitter and Fakeddit dataset. We show the results of eight different baseline methods. Additionally, we present the results of two ablation studies: one without initial-stage inference, and the other without resource distillation and evidence extraction. The best two results are **bolded** and underlined.

Dataset	Method	Accuracy	Rumor			Non-Rumor		
			Precision	Recall	F1	Precision	Recall	F1
Twitter	Direct (LLaVA [218])	0.605	0.688	0.590	0.635	0.522	0.626	0.569
	CoT (LLaVA [218])	0.468	0.563	0.231	0.635	0.441	0.765	0.560
	Direct (InstructBLIP [95])	0.494	0.751	0.171	0.277	0.443	0.902	0.599
	CoT (InstructBLIP [95])	0.455	0.813	0.067	0.112	0.428	0.921	0.596
	Direct (GPT-4 [115])	0.637	0.747	0.578	0.651	0.529	0.421	0.469
	CoT (GPT-4 [115])	0.667	0.899	0.508	0.649	0.545	0.911	0.682
	FactTool (GPT-4 [115])	0.548	0.585	0.857	0.696	0.273	0.082	0.125
	Direct (GPT-4V [96])	0.757	0.866	0.670	0.756	0.673	0.867	0.758
	CoT (GPT-4V [96])	0.678	0.927	0.485	0.637	0.567	<u>0.946</u>	0.709
	LEMMA [122]	0.824	0.943	<u>0.741</u>	0.830	0.721	0.937	0.816
Fakeddit	w/o initial-stage infer	<u>0.809</u>	0.932	0.736	<u>0.823</u>	<u>0.699</u>	0.919	<u>0.794</u>
	w/o visual retrieval	0.781	<u>0.953</u>	0.672	0.788	0.652	0.949	0.773
	Direct (LLaVA)	0.663	0.588	<u>0.797</u>	0.677	0.777	0.558	0.649
	CoT (LLaVA [218])	0.673	0.612	0.400	0.484	0.694	0.843	0.761
	Direct (InstructBLIP [95])	0.726	0.760	0.489	0.595	0.715	0.892	0.793
	CoT (InstructBLIP [95])	0.610	0.685	0.190	0.202	0.604	0.901	0.742
	Direct (GPT-4 [115])	0.677	0.598	0.771	0.674	0.776	0.606	0.680
	CoT (GPT-4 [115])	0.691	0.662	0.573	0.614	0.708	0.779	0.742
	FactTool (GPT-4 [115])	0.506	0.476	0.834	0.606	0.624	0.232	0.339
	Direct (GPT-4V [96])	0.734	0.673	0.723	0.697	0.771	0.742	0.764
LEMMA [122]	<u>0.828</u>	0.881	0.706	0.784	0.800	<u>0.925</u>	0.857	
	w/o initial-stage infer	<u>0.803</u>	0.857	0.692	<u>0.766</u>	<u>0.786</u>	0.891	0.830
	w/o visual retrieval	0.792	0.818	0.675	0.740	0.778	0.883	0.854

Table 5

Performance metrics for different misinformation types in MDAM³-DB, comparing direct prompt queries with results obtained using the proposed MDAM³ framework. The best two results are **bolded** and underlined.

Model	Process	Fact-conflicting			AI-generated			Offensive			OOC		
		Acc	AUC	AP	Acc	AUC	AP	Acc	AUC	AP	Acc	AUC	AP
InstructBLIP [95]	Direct	0.481	0.492	0.498	0.360	0.389	0.385	0.612	0.639	0.634	0.493	0.496	0.496
	MDAM ³	0.580	0.644	0.638	0.766	0.798	0.797	0.746	0.783	0.779	0.607	0.654	0.652
BLIP2 [94]	Direct	0.467	0.484	0.479	0.355	0.367	0.359	0.621	0.688	0.685	0.472	0.484	0.480
	MDAM ³	0.589	0.651	0.646	0.764	0.823	0.814	0.747	0.778	0.768	0.617	0.662	0.654
LLaVA [24]	Direct	0.503	0.512	0.509	0.398	0.402	0.399	0.631	0.696	0.688	0.622	0.639	0.635
	MDAM ³	0.712	0.761	0.760	0.787	0.852	0.849	0.795	0.848	0.839	0.728	0.773	0.761
VILA [215]	Direct	0.484	0.490	0.489	0.323	0.365	0.358	0.611	0.656	0.652	0.615	0.626	0.624
	MDAM ³	0.730	0.756	0.755	0.754	0.796	0.774	0.750	0.768	0.764	0.712	0.769	0.758
GPT-4V [96]	Direct	0.611	0.684	0.667	0.483	0.491	0.488	0.744	0.754	0.747	0.637	0.664	0.646
	MDAM ³	<u>0.853</u>	0.912	0.908	<u>0.886</u>	0.924	<u>0.919</u>	<u>0.891</u>	<u>0.895</u>	<u>0.894</u>	<u>0.729</u>	0.784	0.776

and AP are selected to best reflect the primary evaluation goals of the corresponding benchmarks, and unifying them would risk obscuring task-specific performance characteristics. Accordingly, this survey does not aim to conduct direct numerical comparisons across different tables. Instead, cross-dataset analysis is performed at the level of modeling paradigms and reasoning mechanisms, while numerical comparisons are restricted to within-dataset settings under consistent protocols.

6.1. Results on MM-FakeBench

Table 3 provides a comprehensive evaluation of binary detection performance for various large vision-language models (LVLMs) on the MM-FakeBench benchmark, comparing both validation and large scale test settings. The results highlight the significant role of model size and inference strategies in detecting different types of misinformation distortions.

Smaller models with 7B parameters generally struggle with performance, as reflected by macro-F1 scores below 50. These models exhibit limitations in addressing cross modal inconsistencies, which are critical for distinguishing between authentic and fabricated content. For instance, Otter-Image, MiniGPT4, and InstructBLIP show poor precision and recall, indicating their difficulty in recognizing relevant evidence across text and image modalities. This can be attributed to their inability to handle complex textual veracity distortions, which often involve rumor laden text that is challenging to separate from supporting images. When the model size increases to 13B parameters, performance improves significantly. Models like VILA and InstructBLIP achieve macro-F1 scores above 50 in validation, although they still fall short of human-level performance (F1 = 54.9). The introduction of the MMD-Agent framework notably enhances the models' ability to handle visual veracity distortions, such as fact conflicting images. The MMD-Agent framework employs a structured reasoning process that combines

Table 6

Performance comparisons for different types of misinformation on the Pheme and Twitter16 datasets are conducted. Results are presented for various methods, including SLM, zero-shot and few-shot LLM, and a hybrid LLM + SLM approach. The best performance value for each metric is denoted by **bolded**.

Category	Method	Pheme				Twitter16			
		Acc	Precision	Recall	F1	Acc	Precision	Recall	F1
SLM	RoBERTa [219]	0.714	0.777	0.695	0.734	0.644	0.649	0.641	0.645
	EANN [47]	0.744	0.738	0.745	0.741	0.641	0.621	0.741	0.676
	M ³ FEND [220]	0.746	0.747	0.746	0.746	0.642	0.608	0.816	0.697
	FTT [116]	0.754	0.748	0.764	0.756	0.651	0.649	0.720	0.683
LLM (zero-shot)	Llama2-7B [19]	0.505	0.498	0.697	0.581	0.496	0.501	0.494	0.498
	Llama3-8B [116]	0.535	0.518	0.770	0.620	0.562	0.574	0.531	0.552
	GPT3.5 [221]	0.503	0.500	0.714	0.586	0.583	0.585	0.571	0.578
LLM (few-shot)	Llama2-7B [19]	0.528	0.511	0.894	0.650	0.590	0.598	0.584	0.591
	Llama3-8B [222]	0.549	0.524	0.961	0.679	0.622	0.607	0.717	0.658
	GPT3.5 [221]	0.520	0.507	0.850	0.635	0.621	0.609	0.705	0.653
LLM + SLM	ARG [75]	0.743	0.741	0.779	0.760	0.705	0.698	0.710	0.704
MRCD	Llama2 + RoBERTa [116]	0.772	0.765	0.775	0.770	0.732	0.717	0.619	0.664
	GPT3.5 + RoBERTa [116]	0.781	0.735	0.821	0.778	0.768	0.752	0.734	0.743
	Llama3 + RoBERTa [116]	0.788	0.700	0.900	0.786	0.772	0.765	0.775	0.770
	Llama3 + FTT [116]	0.814	0.788	0.841	0.814	0.794	0.768	0.782	0.774
Improvements	<i>Impr. RoBERTa</i> [219]	+7.4%	/	+20.5%	+5.2%	+12.8%	+11.6%	+13.4%	+12.5%
	<i>Impr. FTT</i> [116]	+6.0%	+4.0%	+7.7%	+5.8%	+14.3%	+11.9%	+13.2%	+12.1%

Table 7

Performance comparison of LVLM-based multimodal misinformation detection methods under the *DLOT* setting within Controlled News Diversity. Since *DLOT* applies to both real and fake instances, we report results (in percentages) for this category as a representative example [224,225].

Infer Type	Data Type	Accuracy	Real			Fake		
			Precision	Recall	F1	Precision	Recall	F1
GPT-4o-mini [223]	Realistic	83.3	89.4	75.6	81.9	78.9	91.1	84.5
	Diversified	64.4 (-18.9)	83.7 (-6.0)	35.7 (-39.9)	50.0 (-31.9)	59.1 (-19.8)	93.1 (+2.0)	72.3 (-12.2)
Claude-3.7-Sonnet [224]	Realistic	88.3	87.7	89.2	88.4	89.0	87.5	88.2
	Diversified	72.4 (-15.9)	86.1 (-1.6)	53.5 (-35.7)	66.0 (-22.4)	66.3 (-22.7)	91.4 (+3.9)	76.8 (-11.4)
Qwen-VL [214]	Realistic	73.7	67.6	91.1	77.6	86.3	56.3	68.2
	Diversified	63.7 (-10.0)	63.9 (-3.7)	62.7 (-28.4)	63.3 (-14.3)	63.5 (-22.8)	64.7 (+8.4)	64.1 (-4.1)
CMIE [225]	Realistic	90.9	88.8	93.6	91.1	93.2	88.2	90.6
	Diversified	72.4 (-18.5)	72.2 (-16.6)	72.8 (-20.8)	72.5 (-18.6)	72.6 (-20.6)	72.1 (-16.1)	72.3 (-18.3)
SNIFFER [34]	Realistic	84.3	78.4	93.2	85.1	92.3	76.2	83.5
	Diversified	73.5 (-10.8)	71.2 (-7.2)	76.0 (-17.2)	73.5 (-11.6)	75.8 (-16.5)	71.1 (-5.1)	73.4 (-10.1)
LEMMA [122]	Realistic	79.6	73.4	92.8	82.0	90.2	66.5	76.5
	Diversified	68.9 (-10.7)	64.7 (-8.7)	83.1 (-9.7)	72.7 (-9.3)	76.3 (-13.9)	54.7 (-11.8)	63.7 (-12.8)

multi step reasoning and confidence calibration, helping the models to better differentiate genuine content from misleading visual cues. For example, VILA's F1 score increases from 51.1 to 56.5, and InstructBLIP's score rises from 51.3 to 56.1, showing the substantial impact of these inference enhancements. At the 34B scale, Nous-Hermes-2 demonstrates strong performance with a macro-F1 of 62.9 in validation, which increases to 67.2 when paired with MMD-Agent. This shows that while model size contributes to performance, it is the targeted application of inference strategies like those in MMD-Agent that enables models to effectively address cross modal consistency distortion. MMD-Agent refines the model's reasoning by enhancing the alignment between text and images, particularly when there are discrepancies such as image/text repurposing or AI-generated content. These refinements lead to more robust detection performance across various distortion types, ensuring the model better handles inconsistencies between the modalities. Finally, GPT-4V, a proprietary LVLM, sets a high benchmark with a macro-F1 of 72.3 in validation and 74.2 on the test set under standard prompting. When enhanced with the MMD-Agent framework, it reaches a macro-F1 of 74.0 in validation and 72.8 in testing, narrowly outperforming human evaluation. The improvements in both precision and recall reflect how MMD-Agent fine tunes the reasoning process to address specific distortions, particularly in visual textual misalignments, making

the model more reliable and accurate in multimodal misinformation detection.

The results underscore that simply increasing model capacity is not enough. The integration of the MMD-Agent framework, with its advanced inference strategies such as multi step reasoning, confidence calibration, and adaptive visual amplification, consistently enhances model performance. This framework allows the models to effectively address different types of distortions, highlighting the importance of structured reasoning and adaptive refinement in achieving high accuracy in multimodal misinformation detection.

6.2. Results on twitter and fakeddit

Table 4 compares the performance of LEMMA with several baseline methods on the Twitter and Fakeddit datasets, both of which are widely used benchmarks for multimodal misinformation detection. The results underscore the importance of structured reasoning and modular architectures in improving performance across these datasets.

On the Twitter dataset, LEMMA achieves an accuracy of 0.824 and an F1 score of 0.816, significantly outperforming the best performing baseline, GPT-4V with Chain-of-Thought (CoT), which attains an accuracy of 0.757 and an F1 score of 0.758. This improvement is

observed consistently across both rumor and non-rumor classification tasks. Specifically, in the rumor category, LEMMA achieves a precision of 0.943 and a recall of 0.741, compared to GPT-4V's 0.866 and 0.670, respectively. In the non-rumor class, LEMMA maintains a precision of 0.721 and a recall of 0.937, demonstrating its superior balance and robustness. The results also highlight a common issue with baseline methods, such as FacTool, which show high recall but suffer from low precision. This indicates a tendency to overfit to superficial cues, such as sensational language or common visual patterns, leading to a high rate of false positives.

On the Fakeddit dataset, LEMMA maintains strong generalization performance with an accuracy of 0.828 and an F1 score of 0.857, once again outperforming all baseline methods. Even the powerful GPT-4V model, under direct prompting, achieves only an accuracy of 0.734 and an F1 score of 0.740. The CoT variant of GPT-4V improves slightly but still lags behind LEMMA by a significant margin. Notably, LEMMA's performance remains stable across both datasets, with minimal variance in precision and recall, suggesting its effectiveness in capturing cross modal evidence without over-relying on domain-specific artifacts or superficial cues.

The ablation studies further validate LEMMA's design choices. Removing the initial stage inference module reduces accuracy to 0.781 on Twitter and 0.803 on Fakeddit, underscoring the critical role of early filtering in improving efficiency and reducing noise. Similarly, excluding visual retrieval results in a performance drop, highlighting the importance of integrating external visual knowledge to support decision-making. These results confirm that LEMMA's strength lies not only in its underlying model capacity but also in its modular architecture, which enables systematic evidence assessment and calibrated decision-making. Compared to prior work, LEMMA avoids the instability observed in models like FacTool, which often exhibit sharp fluctuations in precision recall trade offs due to sensitivity to input phrasing or image composition. LEMMA, on the other hand, consistently maintains high scores across multiple metrics, demonstrating its ability to balance confidence and completeness in explanation generation. This balanced performance ensures that it minimizes both false positives and false negatives, making it particularly well suited for real-world applications where reliability and fairness are paramount.

6.3. Results on MDAM³-DB

Table 5 presents a comparison of various models on the MDAM³-DB dataset, highlighting the performance improvements achieved by integrating the MDAM³ framework across different misinformation types, including fact conflicting content, AI-generated content, offensive content, and out-of-context (OOC) information.

When analyzing the direct application of large visual-language models (LVLMs) without the MDAM³ framework, we observe that models such as BLIP2 and InstructBLIP perform poorly on certain types of misinformation, particularly in detecting factual conflicts and AI-generated content. For instance, BLIP2 achieves an accuracy of 0.467 in detecting fact conflicting content and 0.355 for AI-generated content. Similarly, InstructBLIP shows limited performance, with an accuracy of 0.481 for fact conflicting content and 0.360 for AI-generated content. These results suggest that while these models perform reasonably well in identifying offensive content and cross modal inconsistencies, they struggle to verify factual authenticity reliably, often misclassifying synthetic content as real. This can be attributed to their lack of structured reasoning capabilities, which hinders their ability to make accurate distinctions in complex multimodal contexts. In contrast, GPT-4V demonstrates the strongest baseline performance across all tasks, particularly in handling complex semantics and multimodal information. However, the introduction of the MDAM³ framework leads to substantial improvements. For example, in AI-generated content detection, GPT-4V's accuracy rises from 0.483 to 0.886, and its AUC improves from 0.491 to 0.924. These results underscore the critical role of MDAM³ in correcting model

biases, enhancing fact checking capabilities, and improving classification accuracy. The framework's ability to integrate external knowledge and apply structured reasoning is particularly effective in overcoming the challenges faced by models like BLIP2 and InstructBLIP, leading to notable improvements in their detection capabilities. For example, InstructBLIP's accuracy for detecting fact conflicting content improves from 0.481 to 0.580, and BLIP2 sees a similar improvement, with its accuracy rising from 0.467 to 0.589. Additionally, MDAM³ demonstrates remarkable versatility across multiple types of misinformation. Models like LLaVA and VILA show considerable improvements in both accuracy and AUC for detecting offensive and fact conflicting content after integrating MDAM³. Notably, VILA's performance jumps from an accuracy of 0.484 to 0.730 in fact conflicting content detection, and LLaVA's accuracy increases from 0.503 to 0.712 for the same task. These improvements highlight MDAM³'s ability to enhance model performance across a range of challenging misinformation types, reinforcing its adaptability and robustness. The ablation studies also provide insight into the specific contributions of MDAM³. For instance, the removal of the MDAM³ framework leads to significant drops in performance across all models, particularly in complex tasks like AI-generated content detection and fact conflicting content. These findings emphasize the importance of structured reasoning and external knowledge integration in improving detection accuracy and reducing model bias.

The MDAM³ framework not only addresses the limitations of individual models in specific tasks but also incorporates external information through a structured reasoning process, resulting in significant improvements in the performance of multimodal misinformation detection systems. The observed gains across a variety of metrics, including accuracy, AUC, and AP, highlight that MDAM³ offers a comprehensive and reliable solution for detecting and mitigating different types of misleading information.

6.4. Results on pheme and twitter16

Table 6 presents a performance comparison across different misinformation types on the Pheme and Twitter16 datasets, using various models including small language models (SLMs), large language models (LLMs) in both zero shot and few shot settings, and the hybrid LLM + SLM framework of MRCD.

SLM-based approaches, such as RoBERTa and FTT, demonstrate moderate performance on both datasets, with accuracy values of 0.754 on Pheme and 0.651 on Twitter16. These models struggle to generalize across newly emerging events due to their limited contextual understanding, particularly when facing dynamic misinformation that requires up-to-date knowledge. This is especially evident in their suboptimal performance on complex misinformation types, such as emerging rumors or evolving fake narratives, where deeper reasoning and contextual adaptation are essential. In contrast, LLM based methods, like Llama2-7B and GPT-3.5, even under few shot settings, exhibit limited adaptability, with accuracies remaining below 0.63 across both datasets. Their performance underscores the challenges LLMs face in rapidly evolving domains without explicit domain-specific supervision. Although these models show some improvement over SLMs, they remain less effective at handling nuanced and domain-specific misinformation compared to hybrid approaches that combine the strengths of both model types. The MRCD framework, by integrating LLMs and SLMs through multi round collaboration, achieves a substantial performance boost, particularly in the detection of dynamic and heterogeneous misinformation. For instance, integrating Llama3 with FTT results in the best performance, achieving an accuracy of 0.814 on Pheme and 0.794 on Twitter16, surpassing the strongest standalone SLM baseline by 7.4% and 12.8%, respectively. This improvement is not only observed in accuracy but also across precision, recall, and F1 scores, demonstrating MRCD's balanced enhancement in both detection reliability and robustness. The synergy between LLMs' broad generalization and reasoning capacity and SLMs' domain specific precision allows MRCD

to excel in detecting both emerging and context specific misinformation types, which single model paradigms struggle to address. Further analysis of the results shows that MRCD's multi round collaborative learning mechanism enables dynamic retrieval and application of relevant knowledge, refining the detection process iteratively. This enables the system to continuously adapt to evolving misinformation, a crucial factor in real-world fake news detection scenarios. The improvement of over 20% in performance metrics, such as precision and recall, especially in the context of domain specific fake news detection, underscores the importance of such an approach. Additionally, the hybrid approach mitigates the shortcomings of single model methods, such as overfitting to superficial cues or failing to capture emerging trends in misinformation.

These results confirm that MRCD's collaborative architecture, combining LLMs with SLMs, offers a powerful solution for tackling the complexities of dynamic and evolving misinformation. By dynamically retrieving and integrating the most relevant knowledge, MRCD not only improves detection performance but also ensures that fake news detection systems remain adaptable to new and emerging threats in rapidly changing information environments.

6.5. Results on DriftBench

Table 7 presents a comparison of various LVLM based misinformation detection methods under the DI_OT setting within Controlled News Diversity. The results underscore the vulnerability of current systems to GenAI-driven content variation, which can significantly affect model performance, particularly when models are exposed to diversified news content with stylistic and semantic shifts.

Models like GPT-4o-mini and Claude-3.7-Sonnet, which perform well under realistic conditions, suffer notable declines in performance when exposed to diversified content. For example, GPT-4o-mini's accuracy drops from 83.3% to 64.4%, and its recall for real news falls sharply from 75.6% to 35.7%, illustrating how stylistic and semantic changes can severely impair factual verification. Similarly, Claude-3.7-Sonnet experiences a decline in accuracy from 88.3% to 72.4%, with a substantial drop in recall for fake news. These results emphasize a common challenge in current LVLMs, which is their inability to maintain reliable reasoning across diverse multimodal content, especially when the content deviates from the distribution encountered during training. In contrast, models like SNIFFER and LEMMA exhibit more stability in the face of content diversification, though their F1 scores still decline by more than 10% across both real and fake categories. This indicates that while these models are somewhat resilient to content variation, they still struggle with reasoning consistency, particularly when handling complex distortions such as those induced by GenAI driven content heterogeneity. The integration of multilevel drift scenarios in DriftBench reveals a critical

gap in current detection systems, which is their inability to effectively generalize across evolving distributions. Models that excel in controlled conditions such as GPT-4o-mini and Claude-3.7-Sonnet fail to maintain their performance when confronted with content that deviates stylistically or semantically. This underscores the need for new frameworks capable of adapting to such shifts, stabilizing cross-modal reasoning, and ensuring consistent misinformation detection in rapidly changing information environments. Notably, LEMMA and SNIFFER demonstrate relatively stable performance in these challenging scenarios. This resilience can be attributed to their architectural focus on robust evidence retrieval and adaptive reasoning processes. However, even these models face limitations when dealing with the broader variations introduced by GenAI, underscoring the necessity for frameworks that integrate multi modal, multi round reasoning and continuous adaptation. These findings point to the need for more robust multimodal misinformation detection systems capable of addressing the challenges posed by GenAI driven content diversification.

The results from DriftBench underscore the need for developing frameworks capable of mitigating multi level drift, stabilizing reasoning across diverse contexts, and enhancing the reliability of fake news detection in the face of dynamic, GenAI induced misinformation. While models like LEMMA outperform most others, they still highlight the gaps that must be addressed to achieve robust and generalizable misinformation detection across a wide range of content types.

6.6. Results on FakeSV and FakeTT

Table 8 presents a performance comparison of various models, including both traditional unimodal and state-of-the-art multimodal frameworks, on the FakeSV and FakeTT datasets. The results emphasize the importance of robust multimodal reasoning and the challenges posed by different types of distortions, such as visual manipulation, cross modal inconsistency, and AI generated content.

Models such as BLIP2 and InstructBLIP, which are general purpose vision-language models (VLMs), demonstrate limited capability in handling complex cross modal reasoning. These models achieve accuracies of 70.85% and 78.41% on FakeSV, respectively, which reflects their difficulty in capturing subtle semantic and visual inconsistencies between text and images. The root cause of these limitations lies in their architecture, which lacks the specialized modules required for fine grained multimodal alignment. Specifically, these models struggle with detecting visual distortions, such as manipulated images, or reconciling discrepancies between textual claims and visual evidence, making them less effective at handling nuanced misinformation types. In contrast, more sophisticated multimodal models like TikTec and FANVM introduce explicit feature alignment mechanisms, improving their ability to

Table 8
Performance comparison on two datasets. Best results are shown in **bold**.

Model	FakeSV				FakeTT			
	ACC	M-F1	M-P	M-R	ACC	M-F1	M-P	M-R
GPT-4o-mini [223]	68.08	68.05	69.88	69.49	61.54	61.20	64.41	65.89
GPT-4.1-mini [223]	70.30	70.25	70.61	70.87	49.16	48.54	62.50	59.70
Owen2.5-VL [226]	64.21	60.79	64.55	61.52	45.82	45.31	56.69	55.42
InternVL2.5 [110]	64.39	57.89	68.52	60.50	46.82	45.29	64.92	59.23
InternVL2.5-MPO [110]	65.13	61.07	66.46	62.12	43.14	40.84	61.90	56.23
ViT [61]	70.85	70.66	70.64	70.91	64.88	62.59	62.54	63.80
BERT [227]	78.41	78.25	78.17	78.52	70.90	69.00	68.71	70.60
TikTec [228]	73.06	72.79	72.73	72.93	66.56	65.55	66.50	68.62
FANVM [229]	79.88	78.91	80.98	78.42	71.91	70.85	71.21	73.90
SV-FEND [207]	80.81	80.19	81.08	79.84	77.26	75.55	74.94	77.13
FakingRecipe [210]	84.69	84.39	84.57	84.25	79.26	77.53	76.86	78.89
CA-FVD [230]	85.79	85.28	86.57	84.78	81.61	80.26	79.50	82.17
ExMRD [231]	86.90	86.52	87.31	86.13	84.28	83.13	82.27	85.19
FakeSV-VLM [38]	90.22	89.97	90.55	89.64	89.30	87.98	87.80	88.17

detect cross modal inconsistencies. These models achieve better performance with accuracy scores of 73.06% and 79.88%, respectively, on FakeSV. However, they still fall short in addressing the challenges of fine grained factual verification. Their performance plateaus at around 86.90% accuracy on FakeSV and 84.28% on FakeTT, indicating that while they can capture basic cross modal correlations, they still struggle with more complex distortions like those introduced by generative adversarial techniques or cross modal contradictions. The FakeSV-VLM framework introduces a significant advancement with its Progressive Mixture of Experts (PMOE) architecture, which enables adaptive expert collaboration for multimodal authenticity verification. By leveraging specialized experts for different tasks, FakeSV-VLM demonstrates superior performance, achieving 90.22% accuracy and 89.97% macro-F1 on FakeSV, and 89.30% accuracy with 87.98% macro-F1 on FakeTT. These gains can be attributed to the system's ability to dynamically select the most relevant experts for each specific misinformation task. The framework also incorporates an Alignment-driven Event Checking (ADEC) module, which captures subtle semantic and visual inconsistencies that are crucial for detecting fabricated content, particularly in short videos. This combination of expert reasoning and cross modal alignment allows FakeSV-VLM to outperform both traditional and LVLM based models by a large margin. The performance improvements in FakeSV-VLM can be explained by the integration of adaptive expert collaboration and the ability to handle multimodal distortions across different types of misinformation. The PMOE architecture allows the system to select specialized experts based on the task at hand, providing a tailored approach to each misinformation type, whether it involves visual manipulation, cross modal inconsistency, or semantic distortions. The ADEC module further enhances the system's robustness by identifying and correcting subtle inconsistencies that traditional models fail to capture.

The results from FakeSV-VLM indicate that the combination of structured reasoning, dynamic expert collaboration, and deep multimodal alignment is crucial for achieving state-of-the-art performance in multimodal misinformation detection. While general-purpose models like BLIP2 and InstructBLIP perform well in simpler scenarios, they struggle with the complexities introduced by evolving misinformation tactics. The success of FakeSV-VLM highlights the importance of integrating task-specific adaptability and fine-grained reasoning capabilities in future misinformation detection systems, ensuring that they can handle a wide range of distortion types with high accuracy and interpretability.

6.7. Results on multiple in-domain and out-of-domain datasets

Table 9 compares the performance of TRUST-VL and other baseline VLMs across both in-domain and out-of-domain datasets, highlighting the ability of various models to handle different types of misinformation

distortions, including textual, visual, and cross-modal inconsistencies. The results clearly demonstrate that while general purpose VLMs exhibit certain strengths, they face significant limitations in effectively addressing complex multimodal misinformation, particularly in the presence of advanced distortions like cross-modal inconsistencies or AI-generated manipulations.

Models such as BLIP2 and InstructBLIP, despite showing relatively high performance in some cases, struggle significantly with specific distortion types. These models often misclassify cross modal or AI-generated distortions as genuine content, leading to substantial drops in accuracy. For example, BLIP2's performance on datasets like MOCHEG and VERITE is markedly lower, as it fails to capture the nuanced inconsistencies between text and supporting visual content, especially in cases where visual veracity is distorted. These results suggest that the standard training and prompting strategies used by these models are insufficient for tackling complex visual and cross modal distortions, highlighting a clear gap in their ability to adapt across varied distortion types. In contrast, TRUST-VL, which introduces the Question Aware Visual Amplifier (QAVA) module for adaptive visual representation enhancement, shows substantial improvements across all datasets, including both in-domain and out-of-domain settings. By dynamically adjusting the visual representations based on contextual textual cues, TRUST-VL excels in addressing both textual veracity distortions (e.g., rumor-laden text) and visual veracity distortions (e.g., fact conflicting images). Its performance on Factify2, achieving near perfect accuracy, underscores its ability to handle specific types of distortion, particularly in scenarios involving high level visual manipulations. The model's average accuracy of 86.16%, surpassing the second best model by +8.42%, reflects the strength of its unified reasoning approach, which effectively integrates structured textual and visual reasoning. Moreover, TRUST-VL's robustness is evident in out-of-domain datasets like MOCHEG and Fakeddit-M, where traditional models like GPT-4o show severe performance degradation. These models fail to generalize effectively when faced with data that deviate from their training distribution, particularly in domains with cross-modal inconsistencies. TRUST-VL, however, maintains strong performance in these settings, demonstrating that its training paradigm, which combines reasoning based instruction with multimodal alignment, allows it to generalize better to unseen data. This performance is further amplified by its integration of the QAVA module, which optimally balances textual and visual cues for more accurate cross modal reasoning, ensuring that the model adapts well across a wide range of distortions.

The stark contrast in performance between TRUST-VL and other models reveals the limitations of traditional approaches, particularly in their inability to maintain robustness across distortion types. The success

Table 9

Performance (%) comparison between TRUST-VL and other baseline VLMs across in-domain and out-of-domain datasets. The best score is bolded, and the second-best score is underlined.

Methods	Avg. Acc.	In-domain								Out-of-domain							
		MMFakeBench		Factify2		DGM ⁴ -Face		NewsCLIPpings		MOCHEG		Fakeddit-M		VERITE			
		Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1	Acc.	F1		
General-purpose VLMs																	
BLIP2 [94]	53.36	37.40	34.45	54.30	42.38	47.70	34.35	50.14	34.28	62.50	57.16	70.75	70.19	50.75	37.35		
InstructBLIP [95]	58.41	57.30	56.38	66.83	66.48	50.40	48.66	53.85	50.71	63.25	60.85	64.75	62.83	52.50	49.60		
LLaVA [24]	60.25	62.60	61.72	79.59	79.10	46.41	38.14	45.87	48.54	66.50	64.71	68.00	66.67	52.75	49.80		
xGen-MM	62.20	65.40	62.77	86.03	86.04	50.10	49.68	59.87	59.18	59.50	56.32	60.00	53.45	54.50	54.41		
LLaVA-NeXT [87]	62.35	71.60	65.99	79.60	79.09	53.40	52.21	59.86	59.37	58.25	52.52	59.00	52.36	54.75	54.57		
Qwen2-VL [109]	69.85	67.00	66.28	89.40	89.37	48.10	41.63	70.94	69.91	66.25	64.57	77.25	76.96	70.00	68.94		
GPT-4o [223]	76.16	83.10	80.88	88.37	88.21	57.14	49.24	86.51	86.51	77.00	76.81	73.50	73.12	67.50	67.57		
o1 [232]	<u>77.74</u>	<u>83.90</u>	<u>82.41</u>	<u>96.90</u>	<u>96.90</u>	50.06	38.06	86.80	86.54	<u>81.50</u>	<u>81.38</u>	73.25	73.07	71.75	71.66		
Misinformation detectors																	
MMD-Agent [211]	56.11	69.10	48.68	71.03	69.35	48.30	48.29	53.06	41.12	54.25	43.72	42.25	42.24	54.75	47.00		
SNIFFER [34]	61.17	51.40	51.33	61.00	55.97	47.20	37.96	<u>88.85</u>	<u>88.85</u>	53.75	50.73	53.50	51.13	<u>72.50</u>	<u>72.02</u>		
LRQ-FACT [233]	66.60	71.30	74.00	86.63	89.79	41.80	44.14	68.19	73.45	66.25	69.25	67.25	71.77	64.75	68.32		
TRUST-VL [138]	86.16	87.30	85.42	99.50	99.50	88.50	88.39	90.35	90.35	82.75	82.58	82.50	82.20	73.75	73.61		

of TRUST-VL highlights the necessity of integrated, adaptive models that not only scale across model sizes but also incorporate more sophisticated reasoning strategies. This work establishes a new benchmark for multimodal misinformation detection, setting a precedent for future models that aim to achieve both high accuracy and generalization capabilities across diverse distortion types.

7. Future directions

While Large Vision-Language Models (LVLMs) have demonstrated remarkable capabilities in understanding and reasoning over multimodal content, their application to the nuanced and adversarial domain of fake news detection remains nascent and fraught with unresolved challenges. Drawing from the limitations and gaps identified in this survey, we outline several high-impact research directions that are critical to advancing the field toward robust, trustworthy, and deployable systems.

7.1. Causal and counterfactual reasoning for explainable detection

Current large vision-language models (LVLMs) detect fake news by learning statistical patterns from massive image–text datasets. While effective on benchmark tasks, these models often rely on superficial correlations rather than genuine evidential reasoning. For example, an LVLM may incorrectly associate certain visual styles, such as low-resolution images, particular color tones, or recurring backgrounds, with “fakeness” merely because these features frequently appear in the false examples within its training data. Similarly, it may link emotionally charged language (e.g., “shocking,” “you won’t believe”) or particular news sources to misinformation, even when the underlying claim is factually accurate. This reliance on non-causal cues leads to two critical failures. First, the model’s decisions become brittle: when confronted with novel manipulation tactics or content from underrepresented domains, performance degrades sharply. Second, and more importantly, the model lacks the ability to explain why a piece of content is classified as false based on concrete inconsistencies between visual and textual evidence, such as mismatched timestamps, manipulated objects, or claims that contradict the visible context. Without such reasoning, the system remains a black box, unsuitable for high-stakes applications like journalism, legal review, or public policy. Standard LVLMs are optimized to maximize predictive accuracy on observed data without distinguishing between causal features that make a claim false and correlated features that simply co-occur with falsity in the training distribution. This limitation arises from inherent biases in real-world fake news datasets, which mirror historical disinformation patterns, platform-specific moderation rules, and cultural contexts. As a result, the model often learns superficial associations, such as treating content from certain regions or containing protest imagery as deceptive, rather than reasoning about whether the visual evidence genuinely contradicts the accompanying text. In essence, today’s LVLMs answer the question “What usually goes with fake news?” instead of “What proves this specific claim is false?”

To address this limitation, future LVLMs should be explicitly designed to perform causal and counterfactual reasoning, enabling them to identify which elements of the multimodal input contribute to the veracity judgment and to evaluate how modifying those elements would affect the final decision. A crucial step is to disentangle causally relevant signals, such as temporal inconsistencies, object mismatches, or semantic contradictions, from confounding factors like source bias, stylistic patterns, or demographic correlations. This can be accomplished by training LVLMs across diverse environments, including news from different countries, platforms, or time periods, and ensuring that their internal representations of evidence remain consistent even when confounding factors vary. Methods such as invariant risk minimization and domain-adversarial training can be adapted to reduce the model’s dependence on environment-specific cues while maintaining its ability to detect genuine cross-modal factual inconsistencies. Another effective strategy is to

enhance training with counterfactual examples, which are modified versions of real multimodal samples where only non-essential attributes are changed. For example, an image in a fake news instance can be replaced with a visually similar but factually neutral photo while keeping the misleading text intact. A causally aware model should retain its “fake” prediction only if the textual claim remains unsupported by the new image. Conversely, if a true claim is paired with a manipulated image implying falsity, the model should detect the inconsistency and revise its judgment accordingly. Creating such counterfactual examples through image editing, text paraphrasing, or retrieval from verified media archives offers direct supervision for teaching the model to focus on the truly causal factors underlying veracity assessment. Finally, advancing in this direction requires the development of evaluation metrics that go beyond traditional accuracy. A causally competent model should demonstrate counterfactual consistency, meaning its predictions change only when causally relevant evidence is altered and remain stable when irrelevant features such as font style or background scenery are modified. Furthermore, the model’s explanations, including the image regions it highlights and the textual phrases it emphasizes, should align with human judgments of what constitutes decisive evidence. Collaborations with fact-checking organizations can support the creation of such evaluation benchmarks, ensuring that models are evaluated on genuine reasoning ability rather than statistical pattern recognition.

7.2. Adversarially robust and out-of-distribution generalization

Fake news is not a static phenomenon but an adversarial and adaptive process. Malicious actors continuously refine their strategies to evade detection, exploiting the very patterns that current LVLMs depend on. Early misinformation often involved simple image–text mismatches, whereas modern disinformation campaigns employ more sophisticated techniques such as context swapping, where real images are reused with fabricated captions; deepfakes with semantic alignment, which generate synthetic media that coherently supports false narratives; and multimodal prompt injection, which crafts inputs that subtly mislead LVLMs toward incorrect conclusions. These tactics are intentionally designed to maintain surface-level plausibility while eroding factual integrity, revealing the inherent fragility of existing models. Compounding this issue is the distributional shift inherent in real-world deployment. LVLMs trained on historical datasets, which are often dominated by English-language content from a few major platforms, tend to struggle when confronted with news from underrepresented regions, emerging social media formats (e.g., short videos, memes), or novel event types (e.g., pandemics, geopolitical crises). In such out-of-distribution (OOD) scenarios, models frequently default to spurious heuristics or fail silently, producing high-confidence but incorrect predictions. The root cause lies in the passive learning paradigm of most LVLMs: they are optimized for average-case performance on fixed, curated datasets, not for worst-case robustness or adaptability. Their representations are highly sensitive to input perturbations that preserve semantics to humans but alter model-internal features, such as minor color shifts, object repositioning, or synonym substitutions. Moreover, because training data rarely includes examples of how fake news evolves over time, models lack mechanisms to recognize or adapt to new manipulation strategies. Crucially, standard evaluation protocols mask this vulnerability. Benchmarks like FakeNewsNet or Weibo21 consist of static snapshots of past misinformation, offering no test of a model’s ability to generalize to future or unseen attack vectors. As a result, reported performance often overestimates real-world effectiveness.

To address these challenges, future work must move beyond static training and embrace proactive robustness by design. Below are several interlinked strategies that can significantly improve generalization under adversarial pressure and distribution shifts. First, LVLMs should be trained not only on clean examples but also on realistic adversarial variants that mimic actual disinformation tactics. This requires developing perturbation models tailored to multimodal content. Specifically, for

images, apply context-preserving modifications such as object removal or replacement using diffusion-based inpainting, adjustments to lighting or color tone, and the generation of subtle deepfakes that maintain overall scene coherence. For text, apply meaning-preserving manipulations such as paraphrasing that introduces subtle false implications, substituting key entities with misleading alternatives, or inserting plausible yet fabricated details to distort the narrative while maintaining fluency. For cross-modal alignment, introduce semantic drift in which the image and text remain individually coherent but become jointly deceptive, such as pairing a real photo of a flood with a claim about a different location. By training on such augmented data, ideally produced through an iterative red-team and blue-team process, LVLMs can gradually learn invariant features that remain robust against manipulation. A more principled strategy is to ensure that the model's internal representation of "fakeness" relies only on features that remain consistent across diverse environments, such as different news sources, languages, or event types. This can be accomplished by partitioning the training data into multiple environments, for example based on platform, country, or time period, and applying algorithms such as invariant risk minimization or domain-conflict-aware optimization. These methods encourage the model to build decision boundaries that depend on signals consistently associated with veracity across all environments rather than on environment-specific artifacts. The resulting representation captures the fundamental essence of multimodal inconsistency rather than the peculiarities of any single dataset. Finally, robust systems must know when they don't know. LVLMs should be equipped with calibrated uncertainty estimation, for instance, via ensemble methods, Monte Carlo dropout, or density-based out-of-distribution scoring. When encountering inputs that deviate significantly from the training manifold (e.g., a novel meme format or a deepfake with unusual artifacts), the model should flag high uncertainty and defer to human reviewers or external verification tools, rather than outputting a confident but erroneous verdict.

7.3. Efficient and modular architectures for real-time deployment

While recent large vision-language models (LVLMs) have shown impressive performance on curated fake news detection benchmarks, their practical usefulness in real-world content moderation remains severely limited. The main challenge lies not only in accuracy but also in computational efficiency, latency, and adaptability under operational constraints. Social media platforms must process millions of multimodal posts every minute, including images, short videos, memes, and text captions, often on edge devices or within strict response-time limits of less than 500 milliseconds per post. Current monolithic LVLMs, which contain billions of parameters and rely on complex cross-modal fusion mechanisms, are poorly suited for such environments. They demand costly GPU clusters, consume large amounts of energy, and lack the flexibility and transparency required for timely updates or inspection when new manipulation tactics appear. Moreover, the "one-size-fits-all" architecture of most LVLMs ignores the heterogeneity of real-world content. A breaking news photo, a satirical meme, and a deepfake video demand fundamentally different verification strategies, yet today's models apply the same computationally intensive pipeline to all inputs, wasting resources on low-risk or obviously benign content. This mismatch between research prototypes and deployment realities creates a critical gap: the most accurate models are too slow to be useful, while fast heuristics lack the reasoning depth needed for nuanced disinformation.

To bridge this gap, future systems must abandon the paradigm of "bigger is better" in favor of intelligent, modular, and adaptive architectures that allocate computational resources only when and where they are needed. This requires rethinking both the structure of the model and the flow of inference. Instead of applying a full LVLM to every post, a more scalable solution is to adopt a multi-stage verification cascade. In the first stage, lightweight screening models such as distilled transformers or vision-text co-occurrence filters quickly identify obviously benign content like personal photos and product advertisements,

as well as clear violations such as known deepfake signatures or black-listed sources. These models operate within milliseconds and can filter out more than 80% of total traffic. The second stage involves specialized verifiers that are activated only for ambiguous or high-risk cases. For example, a temporal inconsistency detector can be used when a claim references a recent event, while a semantic alignment analyzer focuses on narrative-heavy posts. Each module is trained for a narrow and well-defined subtask and can be independently updated or replaced. The third stage is a full LVLMLM fallback, reserved for rare cases that require comprehensive multimodal reasoning, ensuring that high-cost processing is applied only when necessary. This hierarchical design mirrors the workflow of human fact-checkers, where rapid preliminary screening is followed by more detailed analysis for uncertain or disputed claims. Furthermore, not all inputs demand the same level of processing. Early-exit mechanisms enable the model to produce confident predictions at intermediate layers when sufficient evidence has already been gathered, avoiding the need to engage deeper and more computationally expensive modules. For example, if a meme's text includes a clearly false statistic and the accompanying image offers no conflicting visual evidence, the system can terminate the analysis early without performing resource-intensive cross-attention over visual features. Similarly, input-adaptive routing can direct different modalities through tailored subnetworks. A text-dominant post (e.g., a quote screenshot) might skip heavy visual encoding, while a video-based claim could activate a dedicated temporal consistency module. Such routing can be learned jointly with the main task or guided by lightweight meta-classifiers. Finally, efficiency must be co-designed with hardware constraints. Techniques such as model quantization (reducing numerical precision from 32-bit to 8-bit), structured pruning (removing redundant attention heads or feed-forward neurons), and knowledge distillation (training small student models to mimic large teachers) can drastically reduce model size and latency without significant accuracy loss. Crucially, these optimizations should preserve the model's sensitivity to subtle multimodal inconsistencies, avoiding aggressive compression that erases fine-grained evidential signals.

7.4. Inference-time mitigation of multimodal hallucination

Multimodal hallucination represents a critical yet underexplored challenge in LVLMLM based fake news detection. Unlike general reasoning errors, hallucinated visual entities, events, or cross modal causal relations can artificially increase the plausibility of false claims, misleading both automated detectors and downstream users by providing high confidence yet unfounded evidence. This issue is especially concerning in misinformation scenarios, where models must reason beyond directly observable content. Hallucinated visual or textual content can lead to the generation of false narratives that appear credible, thereby undermining the reliability of detection systems. The problem is compounded by the seamless reinforcement of misleading textual claims through misinterpreted or manipulated visual content, creating more convincing but ultimately false narratives.

While most existing solutions focus on training time alignment or data curation, emerging research indicates that inference time mitigation offers a lightweight, model agnostic alternative to reduce hallucinated reasoning without the need for retraining large scale LVLMLMs. This approach is particularly advantageous for real-world applications, where retraining large models on massive datasets can be computationally expensive and impractical. A prominent example is Instruction Contrastive Decoding (ICD) [45], which suppresses generation paths dominated by language priors by explicitly contrasting faithful and unfaithful decoding trajectories. This approach ensures that the model's output remains grounded in both textual and visual modalities, leading to more accurate visual grounding and reducing the risk of hallucinations. ICD is especially relevant for fake news detection, where even subtle hallucinations such as fabricated visual content or misrepresented events can distort veracity assessments by introducing fictitious evidence. Beyond ICD, innovative methods such as Cogsteer [234] have

shown considerable promise in mitigating hallucinations during inference. Cogsteer introduces a selective layer intervention mechanism that draws inspiration from cognitive processes of attention and filtering. This method guides the model to prioritize relevant, contextually accurate information while downplaying less reliable or hallucinated details. By intervening at specific layers of the model during inference, Cogsteer effectively steers large models toward more grounded, factually accurate outputs, significantly reducing the likelihood of generating hallucinated cross modal inferences without requiring large scale retraining or additional data curation. The combination of selective layer interventions and contrastive decoding offers a robust strategy for mitigating hallucinations across both visual and textual modalities. While ICD addresses hallucinated textual content, Cogsteer ensures that the model's attention aligns with relevant visual evidence, thus creating a more holistic solution to hallucination mitigation. These techniques, when combined, form a powerful framework for enhancing the reliability of LVLMs in multimodal misinformation detection tasks. Furthermore, additional inference time mechanisms, such as evidence-conditioned generation and self-consistency verification, provide complementary controls to enhance cross modal faithfulness. Evidence-conditioned generation verifies the consistency of generated content with available, trusted evidence during the generation process. Self-consistency verification enables models to cross check their outputs against multiple sources or versions of the input, ensuring that the generated content aligns with verified data. While these techniques have yet to be fully explored in the context of multimodal misinformation detection, they offer promising avenues for further reducing hallucination rates and improving model trustworthiness.

Future research should focus on integrating hallucination aware inference strategies such as ICD, Cogsteer, and self-consistency mechanisms into LVLM based fake news detectors. Evaluation metrics should prioritize faithfulness oriented measures that explicitly quantify evidence grounding and hallucination rates, rather than relying solely on end task detection accuracy. These faithfulness metrics would offer a more nuanced evaluation of model performance, particularly in complex real-world scenarios where identifying hallucinations is just as critical as detecting false claims. Incorporating such metrics will significantly advance the development of more reliable, robust, and trustworthy LVLM based systems for multimodal fake news detection.

7.5. Knowledge-enhanced LVLM adaptation for veracity reasoning

The discussion of knowledge enhanced LVLM adaptation remains underdeveloped, despite mounting evidence that integrating contextual world knowledge can substantially enhance multimodal reasoning, particularly in tasks requiring complex, cross modal understanding. Existing frameworks often frame knowledge enhancement narrowly, focusing primarily on external fact retrieval or static knowledge graph augmentation. However, recent advancements suggest that LVLMs can not only leverage structured external knowledge but also dynamically generate and contextualize world knowledge to improve decision-making processes during inference.

A prime example of this is WisdoM [235], which demonstrates the capacity of LVLMs to generate, contextualize, and fuse commonsense and background knowledge with both visual and textual modalities. By actively eliciting relevant contextual knowledge during inference, WisdoM improves multimodal sentiment analysis, especially under ambiguous conditions or when context is sparse. This approach goes beyond static knowledge integration by enabling models to adjust their knowledge based on the immediate reasoning context, leading to more accurate multimodal predictions. This dynamic knowledge generation and fusion approach has strong implications for veracity reasoning in fake news detection. Fake news often presents claims that, while visually plausible, may be contextually implausible. For instance, claims may involve geopolitical contexts or events that require implicit world knowledge, such as an understanding of regional political dynamics, typical event

timelines, or physical plausibility. These forms of knowledge are often not directly observable in the visual modality, which makes their integration crucial for accurate veracity assessments. Building on these insights, future research should focus on adapting knowledge enhanced paradigms to veracity reasoning tasks. LVLMs must be able to dynamically incorporate both external and internal knowledge contextually relevant background information that aids in the identification of false claims while maintaining cross modal consistency. Unlike traditional approaches that treat knowledge as an isolated post hoc component, this vision advocates for the seamless integration of knowledge generation and grounding mechanisms throughout the model's reasoning process. Such a holistic approach would enable LVLMs to not only detect inconsistencies between visual and textual cues but also assess whether a claim, while visually credible, aligns with contextual knowledge, leading to more robust and reliable fake news detection systems.

Knowledge enhanced LVLMs, particularly those capable of dynamic knowledge generation and fusion, present significant potential for advancing multimodal reasoning in fake news detection. By enabling LVLMs to distinguish between visually plausible yet contextually false claims, this approach offers a promising path toward improving veracity assessment, providing a more comprehensive and accurate solution to the challenges posed by multimodal misinformation.

8. Conclusion

As one of the fastest-growing areas in artificial intelligence, large visual-language models (LVLMs) for multimodal fake news detection have made significant progress in recent years. Therefore, we provide a comprehensive review of this research area. First, we introduce its background and motivation and highlight the unique challenges posed by LVLMs. Second, we present some preliminary research results, including the definition of multimodal fake news, the evolution of traditional multimodal methods, and LVLM-based methods. Third, we propose a taxonomy of current methods, categorizing them into parameter-frozen paradigm and parameter-tuning paradigm. Each paradigm provides a complementary perspective on the development of this field. We then review representative models within each paradigm, tracing their development history and methodological innovations. Fourth, we introduce commonly used evaluation metrics in multimodal fake news detection. Fifth, we compare the performance of different methods on various multimodal fake news detection datasets. Sixth, we summarize several commonly used multimodal fake news detection datasets. Finally, we outline unresolved challenges and future research directions.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

- [1] B. Wang, J. Ma, H. Lin, Z. Yang, R. Yang, Y. Tian, Y. Chang, Explainable fake news detection with large language model via defense among competing wisdom, in: Proceedings of the ACM Web Conference, 2024, pp. 2452–2463.
- [2] S. Mohseni, F. Yang, S. Pentyala, M. Du, Y. Liu, N. Lupfer, X. Hu, S. Ji, E. Ragan, Machine learning explanations to prevent overtrust in fake news detection, in: Proceedings of the International AAAI Conference on Web and Social Media, vol. 15, 2021, pp. 421–431.
- [3] F. Qian, C. Gong, K. Sharma, Y. Liu, Neural user response generator: fake news detection with collective user intelligence, in: International Joint Conference on Artificial Intelligence, vol. 18, 2018, pp. 3834–3840.
- [4] U.S.S. Varshini, R.P. Sree, M. Srinivas, R.B.V. Subramanyam, Rdgt-gan: robust distribution generalization of transformers for Covid-19 fake news detection, IEEE Trans. Comput. Soc. Syst. 11 (2) (2023) 2418–2432.
- [5] A. Rossler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, M. Nießner, Faceforensics++: Learning to detect manipulated facial images, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 1–11.

[6] S. Abdali, S. Shaham, B. Krishnamachari, Multi-modal misinformation detection: Approaches, challenges and opportunities, *ACM Comput. Surv.* 57 (3) (2024) 1–29.

[7] S. Aneja, C. Bregler, M. Nießner, Cosmos: Catching out-of-context misinformation with self-supervised learning, *arXiv preprint arXiv:2101.06278*, 2021.

[8] Y. Shou, T. Meng, W. Ai, K. Li, Dynamic graph neural ODE network for multi-modal emotion recognition in conversation, in: *Proceedings of the 31st International Conference on Computational Linguistics*, 2025, pp. 256–268.

[9] H. Liu, W. Wang, H. Sun, A. Rocha, H. Li, Robust domain misinformation detection via multi-modal feature alignment, *IEEE Trans. Inf. Forensics Secur.* 19 (2023) 793–806.

[10] M. Hosseini, A.J. Sabet, S. He, D. Aguiar, Interpretable fake news detection with topic and deep variational models, *Online Soc. Netw. Media* 36 (2023) 100249.

[11] K. Shu, S. Wang, H. Liu, Beyond news contents: the role of social context for fake news detection, in: *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 2019, pp. 312–320.

[12] W. Shang, K. Song, J. Ji, T. Yi, J. Cai, X. Li, Semantic space aligned multimodal fake news detection, *Inf. Fusion* (2025) 103469.

[13] L. Wu, Y. Long, C. Gao, Z. Wang, Y. Zhang, Mfir: multimodal fusion and inconsistency reasoning for explainable fake news detection, *Inf. Fusion* 100 (2023) 101944.

[14] Y. Shou, H. Lan, X. Cao, Contrastive graph representation learning with adversarial cross-view reconstruction and information bottleneck, *Neural Networks* 184 (2025) 107094.

[15] J. Xie, J. Liu, Z.-J. Zha, Towards effective and transferable detection for multimodal fake news in the social media stream, *IEEE Trans. Knowl. Data Eng.* (2025).

[16] F. Wang, L. Ding, J. Rao, Y. Liu, L. Shen, C. Ding, Can linguistic knowledge improve multimodal alignment in vision-language pretraining? *ACM Trans. Multimed. Comput. Commun. Appl.* 20 (12) (2024) 1–22.

[17] S. Kuntur, A. Wróblewska, M. Paprzycki, M. Ganzha, Under the influence: A survey of large language models in fake news detection, *IEEE Trans. Artif. Intell.* (2024).

[18] Y. Liu, Y. Liu, Z. Li, R. Yao, Y. Zhang, D. Wang, Modality interactive mixture-of-experts for fake news detection, in: *Proceedings of the ACM on Web Conference*, 2025, pp. 5139–5150.

[19] J. Wu, J. Guo, B. Hooi, Fake news in sheep's clothing: robust fake news detection against llm-empowered style attacks, in: *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024, pp. 3367–3378.

[20] Y. Zhou, Y. Yang, Q. Ying, Z. Qian, X. Zhang, Multimodal fake news detection via clip-guided learning, in: *IEEE International Conference on Multimedia and Expo (ICME)*, IEEE, 2023, pp. 2825–2830.

[21] C. Wang, T. Zhang, Multimodal fake news detection based on contrastive learning and data augmentation, in: *International Conference on Electronic Communication and Artificial Intelligence (ICECAI)*, IEEE, 2024, pp. 343–346.

[22] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds, et al., Flamingo: a visual language model for few-shot learning, *Adv. Neural Inf. Process. Syst.* 35 (2022) 23716–23736.

[23] S. Huang, L. Dong, W. Wang, Y. Hao, S. Singh, S. Ma, T. Lv, L. Cui, O.K. Mohammed, B. Patra, et al., Language is not all you need: Aligning perception with language models, *Adv. Neural Inf. Process. Syst.* 36 (2023) 72096–72109.

[24] H. Liu, C. Li, Y. Li, Y.J. Lee, Improved baselines with visual instruction tuning, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 26296–26306.

[25] H. Lyu, J. Huang, D. Zhang, Y. Yu, X. Mou, J. Pan, Z. Yang, Z. Wei, J. Luo, Gpt-4v (ision) as a social media analysis engine, *ACM Trans. Intell. Syst. Technol.* 16 (3) (2025) 1–54.

[26] Y. Chen, D. Li, P. Zhang, J. Sui, Q. Lv, L. Tun, L. Shang, Cross-modal ambiguity learning for multimodal fake news detection, in: *Proceedings of the ACM Web Conference*, 2022, pp. 2897–2905.

[27] L. Hu, Z. Chen, Z. Zhao, J. Yin, L. Nie, Causal inference for leveraging image-text matching bias in multi-modal fake news detection, *IEEE Trans. Knowl. Data Eng.* 35 (11) (2022) 11141–11152.

[28] X. Gao, X. Wang, Z. Chen, W. Zhou, S.C.H. Hoi, Knowledge enhanced vision and language model for multi-modal fake news detection, *IEEE Trans. Multimed.* 26 (2024) 8312–8322.

[29] L. Peng, S. Jian, Z. Kan, L. Qiao, D. Li, Not all fake news is semantically similar: contextual semantic representation learning for multimodal fake news detection, *Inf. Process. Manag.* 61 (1) (2024) 103564.

[30] M. Li, Y. Zhang, H. Xu, X. Li, C. Gao, Z. Wang, Learning complex heterogeneous multimodal fake news via social latent network inference, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, 2025, pp. 433–441.

[31] L. Zhang, X. Zhang, Z. Zhou, F. Huang, C. Li, Reinforced adaptive knowledge learning for multimodal fake news detection, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, 2024, pp. 16777–16785.

[32] L. Zhang, X. Zhang, Z. Zhou, X. Zhang, P.S. Yu, C. Li, Knowledge-aware multimodal pre-training for fake news detection, *Inf. Fusion* 114 (2025) 102715.

[33] Q. Zhang, J. Liu, F. Zhang, J. Xie, Z.-J. Zha, Natural language-centered inference network for multi-modal fake news detection, in: *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 2024, pp. 2542–2550.

[34] P. Qi, Z. Yan, W. Hsu, M.L. Lee, Sniffer: multimodal large language model for explainable out-of-context misinformation detection, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 13052–13062.

[35] W. Hu, Y. Wang, Y. Jia, Q. Liao, B. Zhou, A multi-modal prompt learning framework for early detection of fake news, in: *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 18, 2024, pp. 651–662.

[36] Y. Tong, W. Lu, Z. Zhao, S. Lai, T. Shi, Mmdfnd: Multi-modal multi-domain fake news detection, in: *Proceedings of the 32nd ACM International Conference on Multimedia*, 2024, pp. 1178–1186.

[37] J. Yin, M. Gao, K. Shu, W. Li, Y. Huang, Z. Wang, Graph with sequence: Broad-range semantic modeling for fake news detection, in: *Proceedings of the ACM on Web Conference*, 2025, pp. 2838–2849.

[38] J. Wang, Y. Wang, L. Cheng, Z. Zhong, Fakesv-vlm: Taming vlm for detecting fake short-video news via progressive mixture-of-experts adapter, *arXiv preprint arXiv:2508.19639*, 2025.

[39] X. Xu, X. Li, T. Wang, Y. Jiang, Ample: Emotion-aware multimodal fusion prompt learning for fake news detection, in: *International Conference on Multimedia Modeling*, Springer, 2025, pp. 86–100.

[40] E.J. Hu, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al., Lora: Low-rank adaptation of large language models, in: *International Conference on Learning Representations*, 2021.

[41] A. Silva, L. Luo, S. Karunasekera, C. Leckie, Embracing domain differences in fake news: Cross-domain fake news detection using multi-modal data, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, 2021, pp. 557–565.

[42] H. Niu, L. Xie, J. Lin, S. Zhang, Exploring semantic consistency and style diversity for domain generalized semantic segmentation, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, 2025, pp. 6245–6253.

[43] M. Abulaish, M. Fazil, M.J. Zaki, Domain-specific keyword extraction using joint modeling of local and global contextual semantics, *ACM Trans. Knowl. Discov. Data* 16 (4) (2022) 1–30.

[44] X. Li, Y. Zhang, E.C. Malthouse, Large language model agent for fake news detection, *arXiv preprint arXiv:2405.01593*, 2024.

[45] X. Wang, J. Pan, L. Ding, C. Biemann, Mitigating hallucinations in large vision-language models with instruction contrastive decoding, in: *Findings of the Association for Computational Linguistics: ACL*, 2024.

[46] J. He, H. Zhang, Y. Xiao, W. Guo, S. Yao, R. Liu, Factguard: Event-centric and commonsense-guided fake news detection, *arXiv preprint arXiv:2511.10281*, 2025.

[47] Y. Wang, F. Ma, Z. Jin, Y. Yuan, G. Xun, K. Jha, L. Su, J. Gao, Eann: event adversarial neural networks for multi-modal fake news detection, in: *Proceedings of the 24th ACM Sigkdd International Conference on Knowledge Discovery and Data Mining*, 2018, pp. 849–857.

[48] L. Cui, S. Wang, D. Lee, Same: sentiment-aware multi-modal embedding for detecting fake news, in: *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2019, pp. 41–48.

[49] D. Khattar, J.S. Goud, M. Gupta, V. Varma, Mvae: multimodal variational autoencoder for fake news detection, in: *The World Wide Web Conference*, 2019, pp. 2915–2921.

[50] S. Singhal, R.R. Shah, T. Chakraborty, P. Kumaraguru, S. Satoh, Spotfake: A multi-modal framework for fake news detection, in: *IEEE Fifth International Conference on Multimedia Big Data*, IEEE, 2019, pp. 39–47.

[51] R.K. Kaliyar, A. Goswami, P. Narang, S. Sinha, Fndnet—a deep convolutional neural network for fake news detection, *Cogn. Syst. Res.* 61 (2020) 32–44.

[52] J. Jing, H. Wu, J. Sun, X. Fang, H. Zhang, Multimodal fake news detection via progressive fusion networks, *Inf. Process. Manag.* 60 (1) (2023) 103120.

[53] S. Qian, J. Hu, Q. Fang, C. Xu, Knowledge-aware multi-modal adaptive graph convolutional networks for fake news detection, *ACM Trans. Multimed. Comput. Commun. Appl.* 17 (3) (2021) 1–23.

[54] V. Pérez-Rosas, B. Kleinberg, A. Lefevre, R. Mihalcea, Automatic detection of fake news, in: *Proceedings of the 27th International Conference on Computational Linguistics*, 2018, pp. 3391–3401.

[55] P.K. Verma, P. Agrawal, I. Amorim, R. Prodan, Welfake: word embedding over linguistic features for fake news detection, *IEEE Trans. Comput. Soc. Syst.* 8 (4) (2021) 881–893.

[56] K. Shu, L. Cui, S. Wang, D. Lee, H. Liu, Defend: explainable fake news detection, in: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019, pp. 395–405.

[57] J. Wang, H. Zhang, C. Liu, X. Yang, Fake news detection via multi-scale semantic alignment and cross-modal attention, in: *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2024, pp. 2406–2410.

[58] N. Ruchansky, S. Seo, Y. Liu, CSI: A hybrid deep model for fake news detection, in: *Proceedings of the ACM on Conference on Information and Knowledge Management*, 2017, pp. 797–806.

[59] J. Xue, Y. Wang, Y. Tian, Y. Li, L. Shi, L. Wei, Detecting fake news by exploring the consistency of multimodal data, *Inf. Process. Manag.* 58 (5) (2021) 102610.

[60] H. Wang, Y. Dou, C. Chen, L. Sun, P.S. Yu, K. Shu, Attacking fake news detectors via manipulating news social engagement, in: *Proceedings of the ACM Web Conference*, 2023, pp. 3978–3986.

[61] L. Yuan, Y. Chen, T. Wang, W. Yu, Y. Shi, Z.-H. Jiang, F.E.H. Tay, J. Feng, S. Yan, Tokens-to-token VIT: training vision transformers from scratch on imagenet, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 558–567.

[62] Z. Liu, X. Qi, P.H.S. Torr, Global texture enhancement for fake face detection in the wild, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8060–8069.

[63] B. Guo, Y. Ding, L. Yao, Y. Liang, Z. Yu, The future of false information detection on social media: new perspectives and trends, *ACM Comput. Surv.* 53 (4) (2020) 1–36.

[64] J. Ma, W. Gao, K.-F. Wong, Detect rumors on Twitter by promoting information campaigns with generative adversarial learning, in: *The World Wide Web Conference*, 2019, pp. 3049–3055.

[65] X. Zhou, R. Zafarani, A survey of fake news: fundamental theories, detection methods, and opportunities, *ACM Comput. Surv.* 53 (5) (2020) 1–40.

[66] W.Y. Wang, “Liar, liar pants on fire”: A new benchmark dataset for fake news detection, in: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), Association for Computational Linguistics, 2017.

[67] F. Kou, B. Wang, H. Li, C. Zhu, L. Shi, J. Zhang, L. Qi, Potential features fusion network for multimodal fake news detection, *ACM Trans. Multimed. Comput. Commun. Appl.* 21 (3) (2025) 1–24.

[68] Y. Shou, X. Cao, D. Meng, Speglc: Self-supervised graph spectrum contrastive learning without positive samples, arXiv preprint [arXiv:2410.10365](https://arxiv.org/abs/2410.10365), 2024.

[69] S. Hiriyannaiah, A.M.D. Srinivas, G.K. Shetty, G.M. Siddes, K.G. Srinivasa, A computationally intelligent agent for detecting fake news using generative adversarial networks, in: Hybrid Computational Intelligence, Elsevier, 2020, pp. 69–96.

[70] F. Yu, Q. Liu, S. Wu, L. Wang, T. Tan, et al., A convolutional approach for misinformation identification, in: International Joint Conference on Artificial Intelligence, vol. 2017, 2017, pp. 3901–3907.

[71] R. Yang, X. Wang, Y. Jin, C. Li, J. Lian, X. Xie, Reinforcement subgraph reasoning for fake news detection, in: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2022, pp. 2253–2262.

[72] Y. Dun, K. Tu, C. Chen, C. Hou, X. Yuan, KAN: Knowledge-aware attention network for fake news detection, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, 2021, pp. 81–89.

[73] X. Li, J. Qiao, S. Yin, L. Wu, C. Gao, Z. Wang, X. Li, A survey of multimodal fake news detection: a cross-modal interaction perspective, *IEEE Trans. Emerg. Top. Comput. Intell.* (2025).

[74] Y. Feng, W. Li, Y. Wang, J. Wang, F. Liu, Z. Han, Contradicted in reliable, replicated in unreliable: dual-source reference for fake news early detection, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 39, 2025, pp. 23896–23904.

[75] B. Hu, Q. Sheng, J. Cao, Y. Shi, Y. Li, D. Wang, P. Qi, Bad actor, good advisor: Exploring the role of large language models in fake news detection, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, 2024, pp. 22105–22113.

[76] S. Yang, K. Shu, S. Wang, R. Gu, F. Wu, H. Liu, Unsupervised fake news detection on social media: A generative approach, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, 2019, pp. 5644–5651.

[77] M.K. Ngueajio, S. Aryal, M. Atemkeng, G. Washington, D. Rawat, Decoding fake news and hate speech: A survey of explainable AI techniques, *ACM Comput. Surv.* 57 (7) (2025) 1–37.

[78] F. Yang, S.K. Pentyala, S. Mohseni, M. Du, H. Yuan, R. Linder, E.D. Ragan, S. Ji, X. Hu, Xfake: Explainable fake news detector with visualizations, in: The World Wide Web Conference, 2019, pp. 3600–3604.

[79] M. Huh, A. Liu, A. Owens, A.A. Efros, Fighting fake news: Image splice detection via learned self-consistency, in: Proceedings of the European Conference on Computer Vision (ECCV), 2018, pp. 101–117.

[80] W. Xu, J. Wu, Q. Liu, S. Wu, L. Wang, Evidence-aware fake news detection with graph neural networks, in: Proceedings of the ACM Web Conference, 2022, pp. 2501–2510.

[81] Y. Jin, X. Wang, R. Yang, Y. Sun, W. Wang, H. Liao, X. Xie, Towards fine-grained reasoning for fake news detection, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, 2022, pp. 5746–5754.

[82] A. Choudhary, A. Arora, Linguistic feature based learning model for fake news detection and classification, *Expert Syst. Appl.* 169 (2021) 114171.

[83] H. Shen, X. Li, X. Wang, Y. Dai, T. Wang, G. Bai, Llm-augmented contrastive learning for misinformation detection in social networks, *IEEE Trans. Comput. Soc. Syst.* (2025).

[84] J. Si, Y. Zhu, D. Zhou, Exploring faithful rationale for multi-hop fact verification via salience-aware graph learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 37, 2023, pp. 13573–13581.

[85] L. Wu, L. Wang, Y. Zhao, Unified evidence enhancement inference framework for fake news detection, in: Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, 2024, pp. 6541–6549.

[86] L. Wu, P. Liu, Y. Zhao, P. Wang, Y. Zhang, Human cognition-based consistency inference networks for multi-modal fake news detection, *IEEE Trans. Knowl. Data Eng.* 36 (1) (2023) 211–225.

[87] X. Zheng, Z. Zeng, H. Wang, Y. Bai, Y. Liu, M. Luo, From predictions to analyses: Rationale-augmented fake news detection with large vision-language models, in: Proceedings of the ACM on Web Conference, 2025, pp. 5364–5375.

[88] R. Shao, T. Wu, Z. Liu, Detecting and grounding multi-modal media manipulation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 6904–6913.

[89] Y. Wang, W. Yang, F. Ma, J. Xu, B. Zhong, Q. Deng, J. Gao, Weak supervision for fake news detection via reinforcement learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, 2020, pp. 516–523.

[90] K. Li, B. Guo, J. Liu, J. Wang, H. Ren, F. Yi, Z. Yu, Dynamic probabilistic graphical model for progressive fake news detection on social media platform, *ACM Trans. Intell. Syst. Technol.* 13 (5) (2022) 1–24.

[91] A. Radford, J.W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al., Learning transferable visual models from natural language supervision, in: International Conference on Machine Learning, PMLR, 2021, pp. 8748–8763.

[92] C. Jia, Y. Yang, Y. Xia, Y.-T. Chen, Z. Parekh, H. Pham, Q. Le, Y.-H. Sung, Z. Li, T. Duerig, Scaling up visual and vision-language representation learning with noisy text supervision, in: International Conference on Machine Learning, PMLR, 2021, pp. 4904–4916.

[93] Q. Xu, H. Chen, H. Du, H. Zhang, S. Łukasik, T. Zhu, X. Yu, M3a: A multimodal misinformation dataset for media authenticity analysis, *Comput. Vis. Image Underst.* 249 (2024) 104205.

[94] J. Li, D. Li, S. Savarese, S. Hoi, Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models, in: International Conference on Machine Learning, PMLR, 2023, pp. 19730–19742.

[95] W. Dai, J. Li, D. Li, A. Tiong, J. Zhao, W. Wang, B. Li, P.N. Fung, S. Hoi, Instructblip: Towards general-purpose vision-language models with instruction tuning, *Adv. Neural Inf. Process. Syst.* 36 (2023) 49250–49267.

[96] Z. Yang, L. Li, K. Lin, J. Wang, C.-C. Lin, Z. Liu, L. Wang, The dawn of lmms: Preliminary explorations with gpt-4v (isison), arXiv preprint [arXiv:2309.17421](https://arxiv.org/abs/2309.17421), 2023.

[97] L. Yao, R. Huang, L. Hou, G. Lu, M. Niu, H. Xu, X. Liang, Z. Li, X. Jiang, C. Xu, Filip: Fine-grained interactive language-image pre-training, in: International Conference on Learning Representations, 2024.

[98] W. Liu, X. Yang, D. Tao, J. Cheng, Y. Tang, Multiview dimension reduction via hessian multiset canonical correlations, *Inf. Fusion* 41 (2018) 119–128.

[99] Z.-Y. Dou, Y. Xu, Z. Gan, J. Wang, S. Wang, L. Wang, C. Zhu, P. Zhang, L. Yuan, N. Peng, et al., An empirical study of training end-to-end vision-and-language transformers, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 18166–18176.

[100] A. Mukherjee, S. Ghosh, Unite-fnd: Reframing multimodal fake news detection through unimodal scene translation, arXiv preprint [arXiv:2502.11132](https://arxiv.org/abs/2502.11132), 2025.

[101] Team, Gemini and Georgiev, Petko and Lei, Ving Ian and Burnell, Ryan and Bai, Libin and Gulati, Anmol and Tanzer, Garrett and Vincent, Damien and Pan, Zhufeng and Wang, Shibo and others, Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context, arXiv preprint [arXiv:2403.05530](https://arxiv.org/abs/2403.05530), 2024.

[102] D. Zhu, J. Chen, X. Shen, X. Li, M. Elhoseiny, Minigpt-4: Enhancing vision-language understanding with advanced large language models, arXiv preprint [arXiv:2304.10592](https://arxiv.org/abs/2304.10592), 2023.

[103] J. Li, D. Li, C. Xiong, S. Hoi, BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation, in: International Conference on Machine Learning, PMLR, 2022, pp. 12888–12900.

[104] B. Zhang, D. Ding, Z. Huang, A. Li, Y. Li, B. Zhang, H. Huang, Knowledge-augmented interpretable network for zero-shot stance detection on social media, *IEEE Trans. Comput. Soc. Syst.* (2024).

[105] A. Hurst, A. Lerer, A.P. Goucher, A. Perelman, A. Ramesh, A. Clark, A.J. Ostrow, A. Welihinda, A. Hayes, A. Radford, et al., Gpt-4o system card, arXiv preprint [arXiv:2410.21276](https://arxiv.org/abs/2410.21276), 2024.

[106] Microsoft, Phi-3-vision-128k-instruct, 2024.

[107] H. Liu, C. Li, Q. Wu, Y.J. Lee, Visual instruction tuning, in: A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, S. Levine (Eds.), Advances in Neural Information Processing Systems, vol. 36, Curran Associates, Inc., 2023, pp. 34892–34916.

[108] Mistral AI Team, Pixtral 12b - the First-Ever Multimodal Mistral Model, Tech. rep., Mistral, 2024.

[109] P. Wang, S. Bai, S. Tan, S. Wang, Z. Fan, J. Bai, K. Chen, X. Liu, J. Wang, W. Ge, Y. Fan, K. Dang, M. Du, X. Ren, R. Men, D. Liu, C. Zhou, J. Zhou, J. Lin, Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution, arXiv preprint [arXiv:2409.12191](https://arxiv.org/abs/2409.12191), 2024.

[110] Z. Chen, J. Wu, W. Wang, W. Su, G. Chen, S. Xing, M. Zhong, Q. Zhang, X. Zhu, L. Lu, et al., Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 24185–24198.

[111] Z. Wu, X. Chen, Z. Pan, X. Liu, W. Liu, D. Dai, H. Gao, Y. Ma, C. Wu, B. Wang, et al., Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding, arXiv preprint [arXiv:2412.10302](https://arxiv.org/abs/2412.10302), 2024.

[112] X. Chen, Z. Wu, X. Liu, Z. Pan, W. Liu, Z. Xie, X. Yu, C. Ruan, Janus-pro: Unified multimodal understanding and generation with data and model scaling, arXiv preprint [arXiv:2501.17811](https://arxiv.org/abs/2501.17811), 2025.

[113] GLM, Team and Zeng, Aohan and Xu, Bin and Wang, Bowen and Zhang, Chenhui and Yin, Da and Zhang, Dan and Rojas, Diego and Feng, Guanyu and Zhao, Hanlin and others, Chatglm: A family of large language models from glm-130b to glm-4 all tools, arXiv preprint [arXiv:2406.12793](https://arxiv.org/abs/2406.12793), 2024.

[114] A. At Meta, Llama 3.2: Revolutionizing Edge Ai and Vision with Open, Customizable Models, Tech. rep., Meta (2024).

[115] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F.L. Aleman, D. Almeida, J. Alten Schmidt, S. Altman, S. Anadkat, et al., Gpt-4 technical report, arXiv preprint [arXiv:2303.08774](https://arxiv.org/abs/2303.08774), 2023.

[116] Z. Zhou, X. Zhang, S. Tan, L. Zhang, C. Li, Collaborative evolution: Multi-round learning between large and small language models for emergent fake news detection, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 39, 2025, pp. 1210–1218.

[117] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E.P. Xing, H. Zhang, J.E. Gonzalez, I. Stoica, Judging llm-as-a-judge with mt-bench and chatbot arena, [arXiv:2306.05685](https://arxiv.org/abs/2306.05685), 2023.

[118] Z. Cai, M. Cao, H. Chen, K. Chen, K. Chen, X. Chen, Z. Chen, Z. Chen, P. Chu, et al., Internlm2 technical report, arXiv preprint [arXiv:2403.17297](https://arxiv.org/abs/2403.17297), 2024.

[119] Mistral AI Team, Mistral 7b the Best 7b Model to Date, Tech. rep., Mistral, 2023.

[120] A. Yang, B. Yang, B. Hui, B. Zheng, B. Yu, C. Zhou, C. Li, C. Li, D. Liu, F. Huang, et al., Qwen2 technical report, arXiv preprint [arXiv:2407.10671](https://arxiv.org/abs/2407.10671), 2024.

[121] R. Huang, L. Dugan, Y. Yang, C. Callison-Burch, Miragenews: Multimodal realistic ai-generated news detection, in: Association for Computational Linguistics, 2024.

[122] K. Xuan, L. Yi, F. Yang, R. Wu, Y.R. Fung, H. Ji, Lemma: towards lvm-enhanced multimodal misinformation detection with external knowledge augmentation, arXiv preprint [arXiv:2402.11943](https://arxiv.org/abs/2402.11943), 2024.

[123] X. Zheng, M. Luo, X. Wang, Unveiling fake news with adversarial arguments generated by multimodal large language models, in: Proceedings of the 31st International Conference on Computational Linguistics, 2025, pp. 7862–7869.

[124] C. Li, Z. Li, C. Jing, S. Liu, W. Shao, Y. Wu, P. Luo, Y. Qiao, K. Zhang, Searchlvm: A plug-and-play framework for augmenting large vision-language models by searching up-to-date internet knowledge, *Adv. Neural Inf. Process. Syst.* 37 (2024) 64582–64603.

[125] F. Li, J. Wu, T. Fu, Y. Dong, B. Song, W. Zhou, Drifting away from truth: Genai-driven news diversity challenges lvm-based misinformation detection, *arXiv preprint arXiv:2508.12711*, 2025.

[126] X. Liu, P. Li, H. Huang, Z. Li, X. Cui, J. Liang, L. Qin, W. Deng, Z. He, Fka-owl: Advancing multimodal fake news detection through knowledge-augmented lvm, in: *Proceedings of the 32nd ACM International Conference on Multimedia*, 2024, pp. 10154–10163.

[127] Y. Jiang, T. Wang, X. Xu, Y. Wang, X. Song, D. Maynard, Cross-modal augmentation for few-shot multimodal fake news detection, *Eng. Appl. Intell.* 142 (2025) 109931.

[128] W. Gao, M. Ni, H. Deng, X. Zhu, P. Zeng, X. Hu, Few-shot fake news detection via prompt-based tuning, *J. Intell. Fuzzy Syst.* 44 (6) (2023) 9933–9942.

[129] J. Wu, S. Li, A. Deng, M. Xiong, B. Hooi, Prompt-and-align: prompt-based social alignment for few-shot fake news detection, in: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2023, pp. 2726–2736.

[130] Y. Ouyang, P. Wu, L. Pan, Cool: Comprehensive knowledge enhanced prompt learning for domain adaptive few-shot fake news detection, *arXiv preprint arXiv:2406.10870*, 2024.

[131] J. Yuan, C. Chen, C. Hou, X. Yuan, Fskd: Detecting fake news with few-shot knowledge distillation, in: *International Conference on Advanced Data Mining and Applications*, Springer, 2023, pp. 421–436.

[132] X. Liu, C. Liu, Z. Zhang, C. Li, L. Wang, Y. Lan, C. Shen, Stablept: Towards stable prompting for few-shot learning via input separation, in: *Findings of the Association for Computational Linguistics: EMNLP*, 2024, pp. 9259–9273.

[133] Y. Jiang, Y. Wang, Imfd: In-context multimodal fake news detection with large visual-language models, *Knowl.-Based Syst.* (2025) 113880.

[134] K. Yan, M. Liu, Y. Liu, R. Fu, Z. Wen, J. Tao, X. Liu, Debunk and infer: Multimodal fake news detection via diffusion-generated evidence and lvm reasoning, *arXiv preprint arXiv:2506.21557*, 2025.

[135] X. Shen, M. Huang, Z. Hu, S. Cai, T. Zhou, Multimodal fake news detection with contrastive learning and optimal transport, *Front. Comput. Sci.* 6 (2024) 1473457.

[136] W.U. Yin, Z. Zhang, W.A.N.G. Fuling, Y. Luo, H. Xiong, N. Tang, Detecting out-of-context misinformation via multi-agent and multi-grained retrieval, 2025.

[137] A. Bhattacharya, D. Brahma, S. Nagaje, A. Asati, V. Verma, S. Biswas, Can out-of-domain data help to learn domain-specific prompts for multimodal misinformation detection? in: *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, IEEE, 2025, pp. 2808–2817.

[138] Z. Yan, P. Qi, W. Hsu, M.L. Lee, Trust-vl: An explainable news assistant for general multimodal misinformation detection, *arXiv preprint arXiv:2509.04448*, 2025.

[139] J. Lai, X. Yang, W. Luo, L. Zhou, L. Li, Y. Wang, X. Shi, Rumorllm: A rumor large language model-based fake-news-detection data-augmentation approach, *Appl. Sci.* 14 (8) (2024) 3532.

[140] J. Wu, Y. Fu, N. Yu, G. Fu, E2lvm: Evidence-enhanced large vision-language model for multimodal out-of-context misinformation detection, *arXiv preprint arXiv:2502.10455*, 2025.

[141] S. Tahmasebi, E. Müller-Budack, R. Ewerth, Multimodal misinformation detection using large vision-language models, in: *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2024, pp. 2189–2199.

[142] R. Jin, R. Fu, Z. Wen, S. Zhang, Y. Liu, J. Tao, Fake news detection and manipulation reasoning via large vision-language models, *arXiv preprint arXiv:2407.02042*, 2024.

[143] F. Zeng, W. Li, W. Gao, Y. Pang, Multimodal misinformation detection by learning from synthetic data with multimodal llms, in: *EMNLP Conference on Empirical Methods in Natural Language Processing, Findings of EMNLP*, Association for Computational Linguistics (ACL), 2024, pp. 10467–10484.

[144] S. Tahmasebi, E. Müller-Budack, R. Ewerth, Verifying cross-modal entity consistency in news using vision-language models, in: *European Conference on Information Retrieval*, Springer, 2025, pp. 339–354.

[145] S. Kumar, A. Kumar, A. Mallik, R.R. Singh, Optnet-fake: fake news detection in socio-cyber platforms using grasshopper optimization and deep neural network, *IEEE Trans. Comput. Soc. Syst.* 11 (4) (2023) 4965–4974.

[146] X. Liu, P. Li, H. Huang, Z. Li, X. Cui, J. Liang, L. Qin, W. Deng, Z. He, Fakenewsgpt4: advancing multimodal fake news detection through knowledge-augmented lvm, *CoRR* (2024).

[147] W. Jin, Y. Gao, T. Tao, X. Wang, N. Wang, B. Wu, B. Zhao, Veracity-oriented context-aware large language models-based prompting optimization for fake news detection, *Int. J. Intell. Syst.* 2025 (1) (2025) 5920142.

[148] C. Zhang, Z. Feng, Z. Zhang, J. Qiang, G. Xu, Y. Li, Is llms hallucination usable? Llm-based negative reasoning for fake news detection, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, 2025, pp. 1031–1039.

[149] Y. Zhang, M. Li, C. Gao, X. Li, Confidence breeds success: Improving fake news video detection via lvm-assisted inference, *IEEE Int. Conf. Multimed. Expo* (2025) 1–6.

[150] H. Wan, J. Wu, M. Luo, X. Kong, Z. Ma, Z. Zeng, Difar: Enhancing multimodal misinformation detection with diverse, factual, and relevant rationales, *arXiv preprint arXiv:2508.10444*, 2025.

[151] S. Wang, H. Lin, Z. Luo, Z. Ye, G. Chen, J. Ma, Mfc-bench: Benchmarking multimodal fact-checking with large vision-language models, in: *The Thirteenth International Conference on Learning Representations*, 2025.

[152] Y. Wang, Z. Gu, S. Zhang, S. Zheng, T. Wang, T. Li, H. Feng, Y. Xiao, Llm-gan: Constructing generative adversarial network through large language models for explainable fake news detection, in: *ICASSP IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2025, pp. 1–5.

[153] C. Wang, M. Gao, Z. Wang, J. Yin, K. Shu, C. Lin, Prompt-induced linguistic fingerprints for lvm-generated fake news detection, *arXiv preprint arXiv:2508.12632*, 2025.

[154] Y. Liu, Y. Liu, X. Zhang, X. Chen, R. Yan, The truth becomes clearer through debate! multi-agent systems with large language models unmask fake news, in: *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2025, pp. 504–514.

[155] S. Yang, Z. Yu, Z. Ying, Y. Dai, G. Wang, J. Lan, J. Xu, J. Li, E.C.H. Ngai, Rama: Retrieval-augmented multi-agent framework for misinformation detection in multimodal fact-checking, *arXiv preprint arXiv:2507.09174*, 2025.

[156] B. Fan, F. Ding, G. Zhu, J. Huang, S. Kwong, P. Atrey, S. Lyu, Generating higher-quality anti-forensics deepfakes with adversarial sharpening mask, *ACM Trans. Multimed. Comput. Commun. Appl.* 21 (6) (2025) 1–18.

[157] M. Mostafa, A.S. Almogren, M. Al-Qurishi, M. Alrubaian, Modality deep-learning frameworks for fake news detection on social networks: A systematic literature review, *ACM Comput. Surv.* 57 (3) (2024) 1–50.

[158] L. Shen, Y. Long, X. Cai, I. Razzak, G. Chen, K. Liu, S. Jameel, Gamed: Knowledge adaptive multi-experts decoupling for multimodal fake news detection, in: *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining*, 2025, pp. 586–595.

[159] H. Lin, P. Yi, J. Ma, H. Jiang, Z. Luo, S. Shi, R. Liu, Zero-shot rumor detection with propagation structure via prompt learning, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, 2023, pp. 5213–5221.

[160] M. Choudhary, S.S. Chouhan, S.S. Rathore, Beyond text: Multimodal credibility assessment approaches for online user-generated content, *ACM Trans. Intell. Syst. Technol.* 15 (5) (2025) 1–33.

[161] C. Raj, A. Mukherjee, Z. Zhu, True and fair: Robust and unbiased fake news detection via interpretable machine learning, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023, pp. 962–963.

[162] K. Sharma, F. Qian, H. Jiang, N. Ruchansky, M. Zhang, Y. Liu, Combating fake news: A survey on identification and mitigation techniques, *ACM Trans. Intell. Syst. Technol.* 10 (3) (2019) 1–42.

[163] M.V. Nezafat, S. Samet, Fake news detection with retrieval augmented generative artificial intelligence, in: *International Conference on Foundation and Large Language Models (FLLM)*, IEEE, 2024, pp. 160–167.

[164] Q. Nan, J. Cao, Y. Zhu, Y. Wang, J. Li, Mdfend: Multi-domain fake news detection, in: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2021, pp. 3343–3347.

[165] H. Guo, W. Zeng, J. Tang, X. Zhao, Interpretable fake news detection with graph evidence, in: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2023, pp. 659–668.

[166] S. Qin, M. Zhang, Boosting generalization of fine-tuning BERT for fake news detection, *Inf. Process. Manag.* 61 (4) (2024) 103745.

[167] K. Shu, X. Zhou, S. Wang, R. Zafarani, H. Liu, The role of user profiles for fake news detection, in: *Proceedings of IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2019, pp. 436–439.

[168] L. Zhang, X. Zhang, Z. Zhou, X. Zhang, S. Wang, P.S. Yu, C. Li, Early detection of multimodal fake news via reinforced propagation path generation, *IEEE Trans. Knowl. Data Eng.* (2024).

[169] A. Mosallanezhad, M. Karami, K. Shu, M.V. Mancenido, H. Liu, Domain adaptive fake news detection via reinforcement learning, in: *Proceedings of the ACM Web Conference*, 2022, pp. 3632–3640.

[170] W.-Y. Kao, A.-Z. Yen, How we refute claims: Automatic fact-checking through flaw identification and explanation, in: *Companion Proceedings of the ACM Web Conference*, 2024, pp. 758–761.

[171] R. Shao, T. Wu, L. Nie, Z. Liu, Deepfake-adapter: Dual-level adapter for deepfake detection, *Int. J. Comput. Vis.* 133 (6) (2025) 3613–3628.

[172] X. Zhang, S. Dadkhah, A.G. Weismann, M.A. Kanaani, A.A. Ghorbani, Multimodal fake news analysis based on image-text similarity, *IEEE Trans. Comput. Soc. Syst.* 11 (1) (2023) 959–972.

[173] L. Zong, W. Lin, J. Zhou, X. Liu, X. Zhang, B. Xu, S. Wu, Text-guided fine-grained counterfactual inference for short video fake news detection, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, 2025, pp. 1237–1245.

[174] P. Li, X. Sun, H. Yu, Y. Tian, F. Yao, G. Xu, Entity-oriented multi-modal alignment and fusion network for fake news detection, *IEEE Trans. Multimed.* 24 (2021) 3455–3468.

[175] O.D. Apuke, B. Omar, User motivation in fake news sharing during the Covid-19 pandemic: an application of the uses and gratification theory, *Online Inf. Rev.* 45 (1) (2021) 220–239.

[176] S.K. Hamed, M.J. Ab Aziz, M.R. Yaakub, Enhanced feature representation for multimodal fake news detection using localized fine-tuning of improved BERT and vgg-19 models, *Arab. J. Sci. Eng.* 50 (10) (2025) 7423–7439.

[177] N. Housley, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. De Laroussilhe, A. Gesmundo, M. Attariyan, S. Gelly, Parameter-efficient transfer learning for NLP, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 2790–2799.

[178] X.L. Li, P. Liang, Prefix-tuning: Optimizing continuous prompts for generation, in: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 4582–4597.

[179] B. Lester, R. Al-Rfou, N. Constant, The power of scale for parameter-efficient prompt tuning, in: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 3045–3059.

[180] H. Liu, D. Tam, M. Muqeeth, J. Mohta, T. Huang, M. Bansal, C.A. Raffel, Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning, *Adv. Neural Inf. Process. Syst.* 35 (2022) 1950–1965.

[181] E. Ben-Zaken, S. Ravfogel, Y. Goldberg, Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models, in: 60th Annual Meeting of the Association for Computational Linguistics, ACL, Association for Computational Linguistics (ACL), 2022, pp. 1–9.

[182] T. Dettmers, A. Pagnoni, A. Holtzman, L. Zettlemoyer, Qlora: Efficient finetuning of quantized llms, *Adv. Neural Inf. Process. Syst.* 36 (2023) 10088–10115.

[183] X. Zhou, R. Zafarani, Network-based fake news detection: A pattern-driven approach, *ACM SIGKDD Explorations Newsletter* 21 (2) (2019) 48–60.

[184] J. Qiao, X. Li, C. Gao, L. Wu, J. Feng, Z. Wang, Improving multimodal fake news detection by leveraging cross-modal content correlation, *Inf. Process. Manag.* 62 (5) (2025) 104120.

[185] G. Shan, B. Zhao, J.R. Clavin, H. Zhang, S. Duan, Poligraph: Intrusion-tolerant and distributed fake news detection system, *IEEE Trans. Inf. Forensics Secur.* 17 (2021) 28–41.

[186] X. Liu, A. Nourbakhsh, Q. Li, R. Fang, S. Shah, Real-time rumor debunking on Twitter, in: Proceedings of the 24th ACM International Conference on Information and Knowledge Management, 2015, pp. 1867–1870.

[187] J. Ma, W. Gao, P. Mitra, S. Kwon, B.J. Jansen, K.-F. Wong, M. Cha, Detecting rumors from microblogs with recurrent neural networks, 2016.

[188] Z. Jin, J. Cao, H. Guo, Y. Zhang, J. Luo, Multimodal fusion with recurrent neural networks for rumor detection on microblogs, in: Proceedings of the 25th ACM International Conference on Multimedia, 2017, pp. 795–816.

[189] A. Zubia, M. Liakata, R. Procter, G. Wong Sak Hoi, P. Tolmie, Analysing how people orient to and spread rumours in social media by looking at conversational threads, *PLoS One* 11 (3) (2016) e0150989.

[190] K. Nakamura, S. Levy, W.Y. Wang, Fakeddit: A new multimodal benchmark dataset for fine-grained fake news detection, in: Proceedings of the Twelfth Language Resources and Evaluation Conference, 2020, pp. 6149–6157.

[191] G. Luo, T. Darrell, A. Rohrbach, Newsclippings: Automatic generation of out-of-context multimodal media, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 2021, pp. 6801–6817.

[192] A.F. Biten, L. Gomez, M. Rusinol, D. Karatzas, Good news, everyone! context driven entity-aware captioning for news images, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 12466–12475.

[193] R. Tan, B. Plummer, K. Saenko, Detecting cross-modal inconsistency to defend against neural fake news, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020, pp. 2081–2106.

[194] Y. Li, B. Jiang, K. Shu, H. Liu, Mm-covid: A multilingual and multimodal data repository for combating Covid-19 disinformation, 2020.

[195] C. Yang, X. Zhou, R. Zafarani, Checked: Chinese Covid-19 fake news dataset, *Soc. Netw. Anal. Min.* 11 (1) (2021) 58.

[196] Q. Xu, H. Du, S. Lukasik, T. Zhu, S. Wang, X. Yu, Mdam3: A misinformation detection and analysis framework for multitype multimodal media, in: The ACM Web Conference, Association for Computing Machinery, 2025.

[197] Y. Zhu, Y. Wang, Z. Yu, Multimodal fake news detection: Mfnd dataset and shallow-deep multitask learning, *arXiv preprint arXiv:2505.06796*, 2025.

[198] K. Shu, A. Sliva, S. Wang, J. Tang, H. Liu, Fake news detection on social media: A data mining perspective, *ACM SIGKDD Explorations Newsletter* 19 (1) (2017) 22–36.

[199] F.K.A. Salem, R. Al Feel, S. Elbassuoni, M. Jaber, M. Farah, Fa-kes: A fake news dataset around the Syrian war, in: Proceedings of the International AAAI Conference on Web and Social Media, vol. 13, 2019, pp. 573–582.

[200] L. Cui, D. Lee, Coaid: Covid-19 healthcare misinformation dataset, *CoRR* (2020).

[201] X. Hu, Z. Guo, G. Wu, A. Liu, L. Wen, P.S. Yu, Chef: A pilot Chinese dataset for evidence-based fact-checking, in: Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2022, pp. 3362–3376.

[202] E. Min, Y. Rong, Y. Bian, T. Xu, P. Zhao, J. Huang, S. Ananiadou, Divide-and-conquer: Post-user interaction network for fake news detection on social media, in: Proceedings of the ACM Web Conference, 2022, pp. 1148–1158.

[203] D.S. Nielsen, R. McConville, Mumin: A large-scale multilingual multimodal fact-checked misinformation social network dataset, in: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2022, pp. 3141–3153.

[204] D.K. Sharma, S. Garg, Ifnd: a benchmark dataset for fake news detection, *Complex Intell. Syst.* 9 (3) (2023) 2843–2863.

[205] X. Hu, Z. Guo, J. Chen, L. Wen, P.S. Yu, Mr2: A benchmark for multimodal retrieval-augmented rumor detection in social media, in: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2023, pp. 2901–2912.

[206] B.M. Yao, A. Shah, L. Sun, J.-H. Cho, L. Huang, End-to-end multimodal fact-checking and explanation generation: A challenging dataset and models, in: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2023, pp. 2733–2743.

[207] P. Qi, Y. Bu, J. Cao, W. Ji, R. Shui, J. Xiao, D. Wang, T.-S. Chua, 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, vol. 37, 2023, pp. 14444–14452.

[208] S. Suryavardan, S. Mishra, P. Patwa, M. Chakraborty, A. Rani, A.N. Reganti, A. Chaddha, A. Das, A.P. Sheth, M. Chinnakotla, et al., Factify 2: A multimodal fake news and satire news dataset, in: Proceedings of the AAAI Conference on Artificial Intelligence, 2023.

[209] S.-I. Papadopoulos, C. Koutlis, S. Papadopoulos, P.C. Petrantonakis, Verite: a robust benchmark for multimodal misinformation detection accounting for unimodal bias, *Int. J. Multimed. Inf. Retr.* 13 (1) (2024) 4.

[210] Y. Bu, Q. Sheng, J. Cao, P. Qi, D. Wang, J. Li, Fakingrecipe: Detecting fake news on short video platforms from the perspective of creative process, in: Proceedings of the 32nd ACM International Conference on Multimedia, 2024, pp. 1351–1360.

[211] X. Liu, Z. Li, P.P. Li, H. Huang, S. Xia, X. Cui, L. Huang, W. Deng, Z. He, Mmfakebench: A mixed-source multimodal misinformation detection benchmark for llvms, in: The Thirteenth International Conference on Learning Representations, 2025.

[212] B. Li, Y. Zhang, L. Chen, J. Wang, F. Pu, J.A. Cahyono, J. Yang, C. Li, Z. Liu, Otter: a multi-modal model with in-context instruction tuning, *IEEE Trans. Pattern Anal. Mach. Intell.* (2025).

[213] D. Zhu, J. Chen, X. Shen, X. Li, M. Elhoseiny, Minigpt-4: Enhancing vision-language understanding with advanced large language models, in: The Thirteenth International Conference on Learning Representations, 2024.

[214] T. Liang, S. Wu, J. Fang, G. Yang, W. Wang, F. Lv, Damage analysis via bidirectional multi-task cascaded multimodal fusion, in: Proceedings of the ACM on Web Conference, 2025, pp. 625–636.

[215] J. Lin, H. Yin, W. Ping, P. Molchanov, M. Shoeybi, S. Han, Vila: On pre-training for visual language models, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 26689–26699.

[216] Y. Su, T. Lan, H. Li, J. Xu, Y. Wang, D. Cai, Pandappt: One model to instruction-follow them all, in: Proceedings of the 1st Workshop on Taming Large Language Models: Controllability in the Era of Interactive Assistants!, 2023, pp. 11–23.

[217] Q. Ye, H. Xu, J. Ye, M. Yan, A. Hu, H. Liu, Q. Qian, J. Zhang, F. Huang, Mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 13040–13051.

[218] H. Liu, C. Li, Q. Wu, Y.J. Lee, Visual instruction tuning, *Adv. Neural Inf. Process. Syst.* 36 (2023) 34892–34916.

[219] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, *arXiv preprint arXiv:1907.11692*, 2019.

[220] Y. Zhu, Q. Sheng, J. Cao, Q. Nan, K. Shu, M. Wu, J. Wang, F. Zhuang, Memory-guided multi-view multi-domain fake news detection, *IEEE Trans. Knowl. Data Eng.* 35 (7) (2023) 7178–7191.

[221] D.N. Hong, Y. Hashimoto, I. Paik, T.C. Thang, Leveraging llms for llm-generated fake news detection insights from Covid-19 misinformation, in: IEEE 16th International Conference on Computational Intelligence and Communication Networks (CICN), IEEE, 2024, pp. 1460–1466.

[222] M. Maggini, P.G. Otero, Leveraging advanced prompting strategies in llama3-8b for enhanced hyperpartisan news detection, in: Proceedings of the 10th Italian Conference on Computational Linguistics, 2024, pp. 531–539.

[223] K.J. Roumeliotis, N.D. Tseliakis, D.K. Nasiopoulos, Fake news detection and classification: a comparative study of convolutional neural networks, large language models, and natural language processing models, *Futur. Internet* 17 (1) (2025).

[224] S. Yang, Y. Dai, G. Wang, X. Zheng, J. Xu, J. Li, Z. Ying, W. Wang, E.C.H. Ngai, Realfactbench: A benchmark for evaluating large language models in real-world fact-checking, *arXiv preprint arXiv:2506.12538*, 2025.

[225] F. Li, J. Wu, C. He, W. Zhou, Cmie: Combining llm insights with external evidence for explainable out-of-context misinformation detection, *arXiv preprint arXiv:2505.23449*, 2025.

[226] S. Bai, K. Chen, X. Liu, J. Wang, W. Ge, S. Song, K. Dang, P. Wang, S. Wang, J. Tang, et al., Qwen2. 5-vl technical report, *arXiv preprint arXiv:2502.13923*, 2025.

[227] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, 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, Volume 1 (Long and Short Papers), 2019, pp. 4171–4186.

[228] L. Shang, Z. Kou, Y. Zhang, D. Wang, A multimodal misinformation detector for Covid-19 short videos on tiktok, in: IEEE International Conference on Big Data (Big Data), IEEE, 2021, pp. 899–908.

[229] H.Y.E.W.O.N. Choi, Y.O.U.N.G.J.O.O.N.G. Ko, Topic-agnostic fake news video detection using adversarial learning and topic modeling, *IEEE Access* 9 (2021) 164846–164853.

[230] J. Wang, J. Liu, N. Zhang, Y. Wang, Consistency-aware fake videos detection on short video platforms, in: International Conference on Intelligent Computing, Springer, 2025, pp. 200–210.

[231] R. Hong, J. Lang, J. Xu, Z. Cheng, T. Zhong, F. Zhou, Following clues, approaching the truth: explainable micro-video rumor detection via chain-of-thought reasoning, in: Proceedings of the ACM on Web Conference, 2025, pp. 4684–4698.

[232] A. Jaech, A. Kalai, A. Lerer, A. Richardson, A. El-Kishky, A. Low, A. Helyar, A. Madry, A. Beutel, A. Carney, et al., Openai o1 system card, *arXiv preprint arXiv:2412.16720*, 2024.

[233] A. Beigi, B. Jiang, D. Li, T. Kumarage, Z. Tan, P. Shaeri, H. Liu, Lrq-fact: Llm-generated relevant questions for multimodal fact-checking, *arXiv e-prints arXiv-2410*, 2024.

[234] X. Wang, J. Pan, L. Ding, L. Wang, L. Jiang, X. Li, C. Biemann, Cogsteer: cognition-inspired selective layer intervention for efficiently steering large language models, in: Findings of the Association for Computational Linguistics: ACL, 2025, pp. 25507–25522.

[235] W. Wang, L. Ding, L. Shen, Y. Luo, H. Hu, D. Tao, Wisdom: improving multimodal sentiment analysis by fusing contextual world knowledge, in: Proceedings of the 32nd ACM International Conference on Multimedia, 2024, pp. 2282–2291.