

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2024.0429000

Fact or Fake? Assessing the Role of Deepfake Detectors in Multimodal Misinformation Detection

A S M SHARIFUZZAMAN SAGAR¹, MOHAMMED BENNAMOUN¹ (Senior Member, IEEE), FARID BOUSSAID¹ (Senior Member, IEEE), NAEHA SHARIF¹, LIAN XU¹, SHAABAN SAHMOUD², and ALI KISHK³

¹Department of Computer Science and Software Engineering, The University of Western Australia, Perth, WA 6009, Australia

²Engineering Faculty, Fatih Sultan Mehmet Vakif University, Istanbul 34664, Türkiye

³Aljazeera Media Network Investigative Department, Doha 23134, Qatar.

Corresponding author: A S M Sharifuzzaman Sagar (e-mail: a.sagar@research.uwa.edu.au).

This research was supported by an Australian Government Research Training Program (RTP) Scholarship doi.org/10.82133/C42F-K220.

ABSTRACT In multimodal misinformation, deception usually arises not just from pixel-level manipulations in an image, but from the *semantic and contextual claim* jointly expressed by the image–text pair. Yet most deepfake detectors, engineered to detect pixel-level forgeries, do not account for claim-level meaning, despite their growing integration in automated fact-checking (AFC) pipelines. This raises a central scientific and practical question: *Do pixel-level detectors contribute useful signal for verifying image–text claims, or do they instead introduce misleading authenticity priors that undermine evidence-based reasoning?* We provide the first systematic analysis of deepfake detectors in the context of multimodal misinformation detection. Using two complementary benchmarks such as MMFakeBench and DGM4, we evaluate: (i) state-of-the-art image-only deepfake detectors, (ii) an evidence-driven fact-checking system that performs tool-guided retrieval via Monte Carlo Tree Search (MCTS) and engages in deliberative inference through Multi-Agent Debate (MAD), and (iii) a hybrid fact-checking system that injects detector outputs as auxiliary evidence. Results across both benchmark datasets show that deepfake detectors offer limited standalone value, achieving F1 scores in the range of 0.26–0.53 on MMFakeBench and 0.33–0.49 on DGM4, and that incorporating their predictions into fact-checking pipelines consistently *reduces* performance by 0.04–0.08 F1 due to non-causal authenticity assumptions. In contrast, the evidence-centric fact-checking system achieves the highest performance, reaching F1 scores of approximately 0.81 on MMFakeBench and 0.55 on DGM4. Overall, our findings demonstrate that multimodal claim verification is driven primarily by semantic understanding and external evidence, and that pixel-level artifact signals do not reliably enhance reasoning over real-world image–text misinformation.

INDEX TERMS Deepfake Detection, Fact-Checking, Large Language Models, Multimodal Misinformation

I. INTRODUCTION

The rapid progress of generative AI has increasingly blurred the perceptual boundary between synthetic and authentic imagery, making it possible to generate, alter, and repurpose photorealistic visuals at an unprecedented scale [1]. However, growing empirical evidence shows that the principal harms of online misinformation stem not merely from the pixels of an image, but from the *semantic narrative* that the image is mobilised to assert [2]–[4]. Modern misinformation is fundamentally multimodal: images are paired with captions, headlines, or short textual claims that jointly construct persuasive

narratives, regardless of whether the underlying visual is synthetic or real. In practice, such deception most often manifests as semantic and contextual misalignment, where *contextual misalignment* involves authentic images reused with incorrect who/when/where framing, and *semantic misalignment* involves a mismatch between what the image depicts and what the text asserts [3, 5] including AI-generated images crafted to support fabricated accounts or benign visuals repurposed to reinforce misleading interpretations. During the COVID-19 pandemic, for example, miscaptioned images such as “doctors collapsing after vaccination” circulated widely despite

originating from unrelated incidents [5], [6]. Similar patterns appear in geopolitical misinformation, where AI-generated images are repeatedly shared as supposed evidence of troop movements in particular cities [7]. *Across all scenarios, a clear pattern emerges: in practical settings, the damaging unit of misinformation is not the image alone, but the image–text claim that the image is leveraged to reinforce.*

Current multimodal misinformation frequently mixes genuine photographs, recycled imagery, and AI-generated content, creating situations where pixel-level authenticity intersects with semantic distortion. Datasets such as MM-FakeBench and DGM⁴ reflect this mix, containing both claim-level manipulations and subsets of synthetic or edited images. As deepfake imagery becomes increasingly common, it remains important to determine whether image-only detectors provide a useful signal for downstream verification.

At the same time, traditional image-based deepfake detectors were developed to capture visual artifacts such as frequency noise, blending seams, or generator fingerprints, achieving strong results on curated forensics benchmarks [8], [9]. These detectors generalize poorly in real-world multimodal settings where images may be authentic but contextually misleading, or synthetic yet semantically plausible [10]. Insights from cognitive science reinforce this limitation: authentic but miscaptioned images substantially increase perceived credibility [11], narrative framing overrides visual fidelity [3], [12], and humans rely heavily on textual context even when images visibly contradict it [13]. These findings highlight that *the epistemic unit of multimodal misinformation is the image–text claim, not the image alone.* Nevertheless, deepfake detectors are widely deployed in moderation pipelines and increasingly used as auxiliary signals in automated fact-checking systems, raising an important open question: whether such artifact-centric cues meaningfully support claim verification or instead introduce misleading priors. This motivates our empirical evaluation of detectors both in isolation and integrated within an evidence-driven fact-checking pipeline.

Building on these insights, we conduct the first systematic study evaluating whether traditional pixel-level deepfake detectors offer any measurable utility for multimodal misinformation detection. As shown in Figure 1, we examine three system families: (i) image-only deepfake detectors, (ii) an evidence-driven multimodal fact-checking system, and (iii) a hybrid variant that incorporates detector outputs as auxiliary cues. This evaluation is guided by the following research question:

RQ: Do image-based deepfake detectors provide any meaningful value for multimodal misinformation verification either as standalone predictors or as auxiliary signals within an evidence-centric fact-checking system, or do their artifact-based predictions introduce misleading priors that ultimately degrade claim-level reasoning?

To answer this question, we used two multimodal misinformation benchmarks datasets such as **MMFakeBench** and **DGM⁴**, which together span diverse narrative themes, manip-

ulation styles, and real-world misinformation patterns. These datasets allow us to assess both semantic/contextual distortion and visually manipulated content, providing a comprehensive testbed for understanding whether pixel-level authenticity signals transfer to claim-level verification.

The remainder of this paper is organized as follows. Section 2 reviews related work in multimodal misinformation detection, deepfake forensics, and automated fact-checking. Section 3 introduces our evidence-centric fact-checking framework and details the integration of deepfake detector outputs in the hybrid setting. Section 4 presents the experimental analysis. Section 6 discuss the case study of evidence-centric and hybrid fact-checking system. Finally, Section 5 concludes the paper.

II. RELATED WORK

A. IMAGE-BASED DEEPAKE DETECTION

Image-based deepfake detection has long been the dominant paradigm for assessing visual authenticity. Early approaches relied on hand-crafted physiological or statistical cues such as eye-blink irregularities, color filter array inconsistencies, and illumination artifacts, while more recent deep learning methods—including transformer-based and frequency-domain models—learn latent generator fingerprints directly from data [8], [9]. Large-scale benchmarks such as FaceForensics++ [14], Celeb-DF [15], and DFDC [16] have standardized evaluation protocols and driven steady progress. Nevertheless, extensive cross-dataset studies consistently show that many detectors overfit to generator-specific artifacts and suffer substantial performance degradation under unseen synthesis methods, heavy compression, and platform-specific distortions [10]. To address these limitations, Hasanaath et al. proposed frequency-enhanced self-blended images, combining self-blending augmentation with discrete wavelet transform features to improve cross-dataset generalization [17]. Huang et al. introduced a generalized detection framework that adapts frozen foundation models such as CLIP via lightweight side-network adapters and released the DiGEN dataset to better capture diverse generative sources and unseen forgeries [18]. Extending beyond binary detection, Huang et al. further proposed SIDA, a large multimodal framework that jointly performs social-media image deepfake detection, tamper localization, and natural language explanation [19]. Complementarily, Kroiß et al. show that transfer-learned CNN backbones can remain strong baselines for face deepfake detection when trained on diverse forgeries such as DFFD [20]. In parallel, Arshed et al. reformulated deepfake detection as a multiclass problem and demonstrated that patch-wise Vision Transformer models can effectively distinguish between GAN, diffusion, and hybrid-generated faces by leveraging global contextual representations, outperforming conventional CNN baselines [21].

Crucially, these detectors answer the question “*Was this image visually manipulated?*” rather than “*Is the associated claim true?*”. They reliably flag synthetic faces but remain insensitive to contextual falsity, such as authentic photos

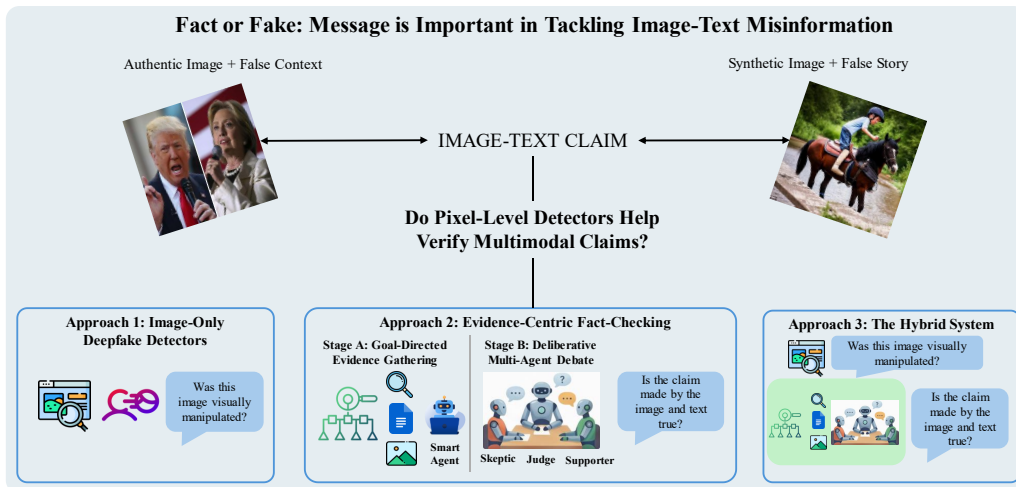


FIGURE 1. Multimodal misinformation often uses authentic or synthetic images paired with false narratives. We compare (1) image-only deepfake detectors, (2) an evidence-centric fact-checking system, and (3) a hybrid that incorporates detector outputs, to ask whether pixel-level authenticity cues help verify image-text claims.

paired with misleading captions. Because most real-world misinformation arises from contextual and semantic distortion rather than pixel-level tampering, image-only detectors are misaligned with claim-level verification. This motivates part of our research question: do artifact-level signals from deepfake detectors provide any value for multimodal truth assessment, or do they introduce noise when authenticity and veracity diverge?

B. MULTIMODAL MISINFORMATION DETECTION

Multimodal misinformation detection systems have traditionally approached the task by assessing how semantically consistent an image-text pair appears, using fusion models that learn joint embeddings through CNN- or Transformer-based encoders. Representative systems such as ViLBERT, CLIP, and subsequent cross-modal attention models classify misinformation by measuring feature alignment across modalities [3], [7]. Francis et al. proposed TamilFacts, the first comprehensive trimodal (text-image-speech) fact-checked dataset for Tamil fake news detection, comprising 7,934 multimodal samples curated from verified fact-checking websites [22]. Similarly, Ojo et al. proposed SmoothDetector, a smoothed Dirichlet-based multimodal framework that integrates probabilistic modeling with deep neural networks to learn continuous latent representations for fake news detection [23]. By replacing discrete or Gaussian latent spaces with a smoothed Dirichlet prior, their approach captures uncertainty and shared cross-modal patterns more effectively while eliminating the need for auxiliary reconstruction subtasks. While such approaches outperform unimodal baselines, they remain fundamentally correlation-driven: an authentic image paired with a persuasive but false caption may still appear semantically coherent in the learned embedding space. To move beyond surface-level coherence, out-of-context (OOC)

detection methods such as COSMOS [24], CLIP-based entailment models [25], and semantic inconsistency approaches [26] explicitly target mismatches between visual content and textual claims. Building on this direction, Qi et al. introduced SNIFFER, a multimodal large language model for explainable OOC detection that combines instruction tuning with retrieval-augmented reasoning over internal and external evidence [27]. Complementary efforts focus on improving fusion itself: CAF-ODNN replaces naive concatenation with complementary bidirectional attention between post text and image captions, explicitly modeling subtle cross-modal discrepancies and improving detection performance across benchmarks [28].

Recent advances increasingly shift toward evidence-grounded and large vision-language model-based verification. Tahmasebi et al. proposed an LVLM-based multimodal misinformation detection pipeline that performs unsupervised multimodal evidence retrieval followed by zero-shot fact verification, achieving strong cross-dataset generalization without task-specific fine-tuning [29]. Yan et al. presented TRUST-VL, a unified and explainable vision-language model trained with structured fact-checking chains to jointly address textual, visual, and cross-modal distortions, improving robustness and interpretability through evidence-grounded reasoning [30]. Extending toward holistic analysis, Xu et al. introduced MDAM3, a comprehensive framework for multitype multimodal misinformation that integrates internal detectors for AI-generated content and cross-modal inconsistencies with external web-based evidence retrieval and LVLM-driven explanation [31]. Gedara et al. proposed an explainable multimodal fake news detection framework that integrates prompt-based Large Language Models with an AutoML-driven image pipeline, fusing modality outputs via a Functional Network [32]. Despite these advances,

existing methods largely frame misinformation detection as a single-stage verdict prediction problem and do not explicitly analyze when visual authenticity cues are informative or misleading for claim-level truth assessment. Collectively, this literature reveals that semantic coherence, contradiction awareness, improved fusion, or even evidence grounding alone does not guarantee factual correctness. Our work builds on these insights by replacing correlation-driven fusion with structured, evidence-seeking verification and by systematically examining the role of visual manipulation signals within a claim-centric multimodal fact-checking framework.

C. AUTOMATED FACT-CHECKING FOR MULTIMODAL MISINFORMATION

Automated fact-checking (AFC) reconceptualizes misinformation detection as a structured reasoning pipeline comprising claim parsing, evidence retrieval, and veracity estimation [33]. Recent multimodal AFC systems explicitly operationalize this framework. DEFAME [34] coordinates specialized retrieval agents for text and vision; LRQ-Fact [35] decomposes complex claims into sub-queries; CRAVE [36] clusters retrieved documents into narrative units; and Holmes [37] integrates retrieval-augmented LLM scoring. Recent work further explores tightly coupled multimodal AFC pipelines, such as Ahmed et al. combine CNN-based image validation with Transformer-based textual reasoning and cross-modal fusion to improve robustness over unimodal baselines [38]. Along similar lines, Kakizaki et al. propose MAFT, which textualizes multimodal inputs (image/video captioning and speech recognition), extracts claims, retrieves external evidence via web APIs, and generates interpretable fact-checking reports using LLMs, while also converting deepfake detector outputs into textual evidence for downstream reasoning [39]. More recently, Yang et al. introduce RAMA, a retrieval-augmented multi-agent framework that refines web queries for multimodal claims, aggregates cross-verified evidence, and ensembles multiple LLM judges to improve robustness and interpretability [40]. Krishnamurthy and Balaji introduced a credibility-aware LLM-powered fact-checking framework that structures claim verification through source reliability modeling and evidence credibility scoring, enabling automated yet interpretable truth assessment grounded in external references [41]. Cui et al. introduced a large-scale multi-task dataset for multi-phase image-based fact verification, explicitly modeling evidence retrieval, visual verification, and claim consistency within a unified benchmark [42]. In addition, Kangur et al. propose MultiReflect, a multimodal self-reflective RAG-based AFC pipeline that iteratively retrieves, filters, ranks, and verifies image–text evidence from the internet, improving multimodal fact-checking accuracy [43]. Complementing retrieval-centric AFC, Chen et al. propose a causal-intervention and counterfactual multimodal fact-checking method that identifies causally relevant regions in retrieved visual evidence and suppresses shortcut unimodal reasoning via counterfactual decomposition, improving generalization on OOC benchmarks [44]. Finally, Zhang et al.

introduce ESCNet and the large-scale CMFC dataset, enhancing multimodal fact-checking by jointly modeling stance reasoning between claims and retrieved evidence and incorporating knowledge-graph-grounded entity relations for stronger semantic and knowledge-level verification [45]. Other work emphasizes evidence alignment and explainability, showing that grounding both image and text in external knowledge improves multimodal verification [46], [47].

Despite these advances, two key gaps remain. **First**, no prior work has systematically compared AFC pipelines with state-of-the-art image-based detectors on identical multimodal benchmarks, leaving open whether artifact-level cues contribute any measurable benefit. **Second**, the interaction between deepfake detectors and reasoning agents is poorly understood: detector outputs might enhance robustness or, alternatively, introduce misleading priors when images are authentic but miscaptioned. Addressing these open questions forms the empirical foundation of our study.

III. METHOD

This section introduces our evidence-centric fact-checking method, a two-stage pipeline combining goal-directed evidence acquisition with deliberative adjudication, and then describes the image-based deepfake detectors used as baselines for comparison.

A. FACT-CHECKING METHOD

Our goal is to determine the truthfulness of an image–text claim using targeted tool use and structured debate-driven reasoning. As shown in Figure 2, a multimodal LLM serves as the central controller: it plans actions, invokes tools, interprets observations, and maintains memory throughout verification. The LLM selects from an extensible toolset depending on the active modality and subtask requirements, enabling unified reasoning across heterogeneous evidence sources.

To efficiently explore potential tool-use trajectories, we adopt a MCTS strategy inspired by T2Agent [48]. T2Agent formulates multimodal misinformation detection using an MCTS framework in which each node is evaluated by a dual reward function, where both the reasoning trajectory score and the confidence score are LLM-estimated and directly combined to guide node selection, backpropagation, and final decision aggregation across forgery-source subtasks. In contrast, our method explicitly disentangles trajectory utility from evidential confidence. The trajectory score S_T is defined as a bounded structural measure that quantifies the quality of evidence acquisition, including progress, tool diversity, non-redundancy, and coherence of the action observation sequence, independent of claim veracity, while the confidence score is reserved exclusively for estimating the strength of factual support or refutation provided by the accumulated evidence for a given subtask. Rather than fusing S_T and S_C into a single decision signal, we use their product only as a search termination criterion, ensuring that high-confidence decisions are reached through sufficiently rich and non-degenerate evidence trajectories. Final claim verification is then performed

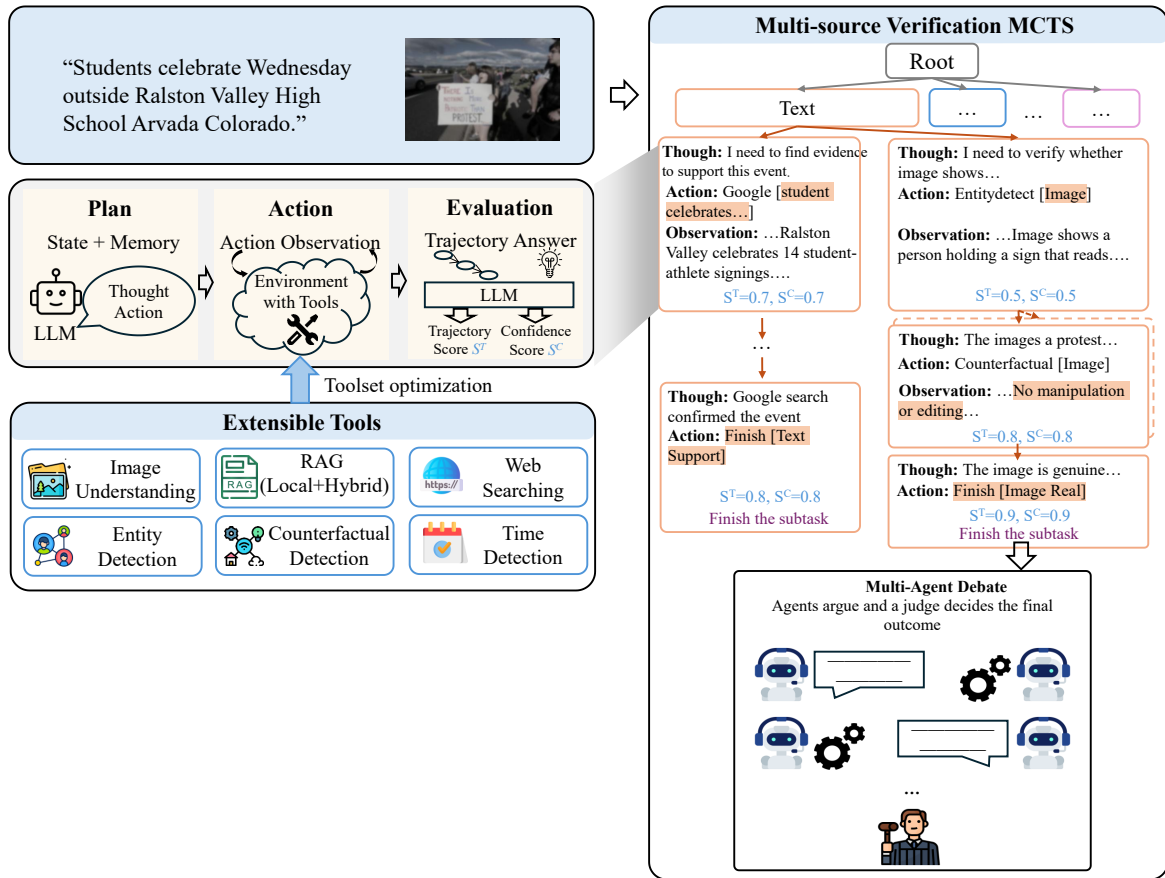


FIGURE 2. Overview of the proposed evidence-centric fact-checking framework. Given an image–text claim, the system performs verification in two stages. (1) **Multi-source verification via MCTS:** a multimodal LLM plans tool calls across an extensible toolkit (e.g., web/RAG search, image understanding, counterfactual/VQA, entity and time detection), executes these actions within an environment, and evaluates each trajectory using utility and confidence scores to generate text-side and image-side stances with their corresponding evidence packs. (2) **Multi-Agent Debate (MAD):** skeptic and supporter agents deliberate over the collected evidence, while a neutral Judge LLM aggregates their arguments to issue the final verdict.

by a MAD module, in which competing verifier agents reason adversarially over the structured evidence sets produced by MCTS, and a neutral judge aggregates their arguments into a final verdict. This separation of (i) trajectory optimization, (ii) confidence estimation, and (iii) deliberative decision making yields a more controlled and interpretable verification process.

Thus, multimodal claim verification is decomposed into two coordinated stages:

- **Stage 1 – Goal-Directed Evidence Acquisition.** An MCTS-guided multimodal LLM interacts with a suite of tools to obtain diverse, relevant, and reliable evidence for both text- and image-level subtasks.
- **Stage 2 – Deliberative Decision Making.** The outputs of Stage 1 are adjudicated through MAD, producing a final, evidence-grounded decision.

1) Goal-Directed Evidence Acquisition

Evidence acquisition is framed as an LLM-controlled, goal-directed tree search over an extensible toolset. The system constructs a search tree using the classic *selection* \rightarrow *expansion* \rightarrow *simulation* \rightarrow *backpropagation* cycle, where subtask-specific roots correspond to text and image verification.

At each node, the LLM proposes the next action based on the current state and trajectory, selects an appropriate tool, and processes the resulting observation as evidence for subtask decision making.

a: Search Structure

The search tree captures successive reasoning trajectories guided by the LLM. The root corresponds to the overall verification goal for a claim (I, T) ; its children represent the initial subtask roots $k \in \{\text{text}, \text{image}\}$.

Each node stores a state

$$s = \langle I, T, E, M, k \rangle, \quad (1)$$

where E is the accumulated evidence, M is the short-term memory of recent actions and observations, and k is the active subtask. A node n represents the pair (s, τ) , where the trajectory

$$\tau = (a_{1:t}, o_{1:t})$$

records all actions a_i and observations o_i along the path.

From state s , the planner proposes an action $a = (\text{tool}, \text{args})$, executed in the environment to produce an observation o . This observation is normalized into an evidence atom

$$e = \langle \text{modality} \in \{\text{text}, \text{image}\}, \text{content}, \text{source}, \text{timestamp} \rangle. \quad (2)$$

Each atom is appended to E , incrementally forming the evidence set used for stance estimation and coverage tracking.

b: Trajectory Evaluation

Every partial trajectory $\tau = (a_{1:t}, o_{1:t})$ is evaluated using two complementary bounded scores:

- **Trajectory Utility** $S^T(\tau) \in [0, 1]$: assesses structural quality, coherence, and exploration diversity.
- **Confidence** $S^C(\tau) \in [0, 1]$: measures how strongly the collected observations support the current subtask label.

The final node value used during search is

$$V(n) = \lambda S^T(\tau_n) + (1 - \lambda) S^C(\tau_n). \quad (3)$$

(a) Trajectory Utility S^T . This score rewards progressive, diverse, and coherent evidence acquisition. For a trajectory of length L , we compute:

- **Progress:**

$$\text{prog}(L) = \frac{1}{1 + \exp(-(L - 2))}.$$

- **Tool Diversity:**

$$\text{unique_ratio} = \frac{|\{\text{tools}\}|}{\max(1, L)}.$$

- **Non-redundancy:**

$$\text{rep_ratio} = \frac{\#\{i > 1 : \text{tools}[i] = \text{tools}[i - 1]\}}{\max(1, L - 1)},$$

$$\text{non_repeat} = 1 - \text{rep_ratio}.$$

These are averaged into a structural base score:

$$\text{base } S^T(\tau) = \frac{1}{3} (\text{prog} + \text{unique_ratio} + \text{non_repeat}). \quad (4)$$

A compact LLM grader then assigns a coherence score $\text{llm } S^T(\tau) \in [0, 1]$. The final utility is:

$$S^T(\tau) = \text{clip}\left(\frac{1}{2} \text{base } S^T(\tau) + \frac{1}{2} \text{llm } S^T(\tau), 0, 1\right). \quad (5)$$

(b) Confidence S^C . A grader LLM scores the reliability of the current evidence on a 1–10 scale:

$$S^C(\tau) = \frac{s}{10}. \quad (6)$$

c: Search Algorithm

The MCTS process follows:

- **Selection:**

$$\text{UCT}(n, a) = \frac{\hat{Q}(n, a)}{N(n, a) + 1} + C \sqrt{\frac{\ln(N(n) + 1)}{N(n, a) + 1}}.$$

- **Expansion:** the best (n, a) pair is executed; a child node is created.
- **Simulation:** a lightweight rollout collects provisional scores (S^T, S^C) .
- **Backpropagation:** the combined value V is propagated upward.

Each subtask runs under a fixed step budget B_k and terminates early when its best trajectory τ_k^* satisfies:

$$S^T(\tau_k^*) \cdot S^C(\tau_k^*) \geq \theta_k. \quad (7)$$

Branches yielding high-confidence falsity (e.g., explicit refutation or verified out-of-context provenance) are pruned, and their unused budget is reallocated. This reflects the monotonicity that a claim cannot be Real if any component is confidently Fake.

Stage 1 returns structured summaries:

$$(z_T, E^{\text{text}}, \tau_T^*), \quad (z_I, E^{\text{image}}, \tau_I^*), \quad (8)$$

where

$$z_T \in \{\text{TEXT_REAL}, \text{TEXT_FAKE}\},$$

$$z_I \in \{\text{IMAGE_REAL}, \text{IMAGE_FAKE}\}.$$

Algorithm 1 summarizes our goal-directed evidence acquisition via MCTS. The algorithm takes as input a claim c , a set of subtasks $K = \{\text{Text}, \text{Image}\}$, an extensible tool set \mathcal{T} , subtask-specific step budgets $\{B_k\}_{k \in K}$, termination thresholds $\{\tau_k\}_{k \in K}$, and a mixing weight $\lambda \in [0, 1]$ for combining structural utility and confidence scores. For each subtask $k \in K$, the algorithm returns a discrete label y_k , an evidence set \mathcal{E}_k , and the best reasoning trajectory π_k^* .

For each subtask k , we initialize an empty evidence buffer $\mathcal{E} \leftarrow \emptyset$ and a short-term memory \mathcal{M} that tracks recent actions and observations. Together with the claim c and the active subtask k , these define the initial state s_0 . We initialize an empty trajectory $\pi \leftarrow []$, create the root node $n_0 \leftarrow \text{Node}(s_0)$, and start the search tree with this root.

Each MCTS iteration (up to budget B_k) consists of the four standard phases: selection, expansion, simulation, and backpropagation. In the selection phase, we start from the root and repeatedly select actions according to a UCT-style exploration–exploitation rule. At each internal node, we choose the best action among its outgoing edges, transition to the corresponding child node, and record visited nodes along this path in a list *path*. This continues until reaching a node that is terminal or not yet fully expanded.

In the expansion phase, if the current node is non-terminal, we invoke an LLM-based planner $\text{LLMPlanner}(n.\text{state}, k, \mathcal{T})$ to propose the next action (tool, args). The environment step $\text{StepEnvironment}(n.\text{state}, a)$ executes this tool call and returns an updated state with new evidence and memory, together with an extended trajectory that appends the new action–observation pair. We then create a new child node, attach it to its parent, insert it into the tree, and extend the current path.

In the simulation phase, starting from the expanded node, we perform a lightweight rollout via $\text{RolloutAndScore}(\cdot)$.

Algorithm 1 Goal-Directed Evidence Acquisition

Require: Claim (I, T) ; subtasks $k \in \{\text{text}, \text{image}\}$;
1: tool set \mathcal{A} ; step budgets $\{B_k\}$; thresholds $\{\theta_k\}$; mixing weight $\lambda \in [0, 1]$
Ensure: For each subtask k : label z_k , evidence E_k , best trajectory τ_k^*
2: **for all** $k \in \{\text{text}, \text{image}\}$ **do**
3: Initialize evidence $E \leftarrow \emptyset$, memory $M \leftarrow \emptyset$
4: Initialize state $s_0 \leftarrow \langle I, T, E, M, k \rangle$
5: Initialize trajectory $\tau_0 \leftarrow ()$
6: Create root node $n_0 \leftarrow (s_0, \tau_0)$
7: Initialize search tree $\mathcal{T} \leftarrow \{n_0\}$
8: **for** $b = 1$ to B_k **do** ▷ MCTS iteration
9: $n \leftarrow n_0$; $path \leftarrow [n_0]$
10: **while** n is fully expanded **and** n is not terminal **do**
11: Select action $a^* \leftarrow \arg \max_a \text{UCT}(n, a)$
12: $n \leftarrow \text{Child}(n, a^*)$
13: Append n to $path$
14: **end while**
15: **if** n is not terminal **then**
16: $a \leftarrow \text{LLMPlanner}(s(n), \tau(n))$ ▷
17: $(s', \tau') \leftarrow \text{StepEnvironment}(s(n), \tau(n), a)$
18: Create child node $n' \leftarrow (s', \tau')$ and attach to n
19: $\mathcal{T} \leftarrow \mathcal{T} \cup \{n'\}$
20: $n \leftarrow n'$; append n to $path$
21: **end if**
22: $(S_T, S_C) \leftarrow \text{RolloutAndScore}(n)$ ▷ lightweight rollout
23: $V \leftarrow \lambda S_T + (1 - \lambda) S_C$
24: **for all** $m \in path$ **do**
25: UpdateStats(m, V) ▷ update Q, N for actions leading to m
26: **end for**
27: $n^* \leftarrow \arg \max_{m \in \mathcal{T}} V(m)$
28: $(S_T^*, S_C^*) \leftarrow \text{RecomputeScores}(n^*)$
29: **if** $S_T^* \cdot S_C^* \geq \theta_k$ **then**
30: **break**
31: **end if**
32: **end for**
33: $n_k^* \leftarrow \arg \max_{m \in \mathcal{T}} V(m)$
34: $(s_k^*, \tau_k^*) \leftarrow n_k^*$
35: $E_k \leftarrow s_k^*.E$
36: $z_k \leftarrow \text{SubtaskLabel}(E_k, k)$
37: **end for**
38: **return** $(z_{\text{text}}, E_{\text{text}}, \tau_{\text{text}}^*), (z_{\text{image}}, E_{\text{image}}, \tau_{\text{image}}^*)$

This routine stochastically continues the trajectory for a small number of steps using inexpensive heuristics, then evaluates the resulting trajectory with two bounded scores: (i) a structural utility $U \in [0, 1]$ capturing progress, diversity, and coherence of evidence acquisition; and (ii) a confidence score $S \in [0, 1]$ capturing how strongly the accumulated evidence supports a subtask decision. The overall node value is computed as a convex combination $V = \lambda U + (1 - \lambda)S$.

During backpropagation, the value V is propagated back along all nodes in $path$. For each node on the path, UpdateStats(\cdot) updates visit counts and value estimates associated with the actions that led to the next node. These statistics are subsequently used by the UCT rule during future selection steps.

After each iteration, we identify the current best node in the tree according to its value estimate and recompute its scores via RecomputeScores(\cdot) to obtain refined structural utility and confidence (U, S) . If their product exceeds the subtask-specific threshold, i.e., $U \cdot S \geq \tau_k$, the search terminates early, reflecting that the system has acquired sufficiently strong and coherent evidence for a stable decision.

At the end of the search (either by budget exhaustion or early stopping), we select the globally best node for subtask k from the tree. Its state contains the final evidence buffer \mathcal{E}_k , and its trajectory encodes the full reasoning path π_k^* used to collect this evidence. A subtask-specific decision function SubtaskLabel(\mathcal{E}_k, k) maps the evidence to the final label y_k (e.g., {Text-Real, Text-Fake} for text and {Image-Real, Image-Fake} for image).

2) Multi-Agent Debate (MAD) for Decision Making

MAD produces the final decision using Stage 1 evidence. For each input (I, T) , the MCTS stage provides modality-specific labels, calibrated confidence scores, and compact evidence packs. MAD transforms these summaries into a consensus through controlled adversarial dialogue.

a: Label Space and Initialization

The overall label space is:

$$Y_4 = \{\text{REAL}, \text{TEXT_FAKE}, \text{IMAGE_FAKE}, \text{BOTH_FAKE}\}.$$

Two agents assume complementary roles: Agent A acts as a skeptic (falsification-first), and Agent B as a supporter (verification-first).

b: Round-Based Debate Process

For R rounds, each agent $j \in \{A, B\}$ outputs:

$$y_r^{(j)} \in Y, \quad c_r^{(j)} \in [0, 1], \quad \rho_r^{(j)},$$

where $y_r^{(j)}$ is the proposed label, $c_r^{(j)}$ is its confidence, and $\rho_r^{(j)}$ is a rationale citing evidence from Stage 1.

Define the novelty factor:

$$\kappa_r^{(j)} = \begin{cases} 1, & \text{if } C_r^{(j)} \setminus C_{<r}^{(j)} \neq \emptyset, \\ \pi, & \text{otherwise,} \end{cases}$$

where $\pi \in (0, 1)$ penalizes repeated evidence use.

c: Judgment and Final Verdict

After the final round, a neutral Judge LLM examines the full transcript and Stage 1 summaries and outputs:

$$y_J \in Y, \quad c_J \in [0, 1].$$

This constitutes the system's final claim-level decision.

Algorithm 2 applies a multi-agent debate (MAD) layer on top of the Stage-1 evidence bundle, which aggregates modality-specific subtasks $k \in K$ with labels $y_k \in Y_4$, confidences $s_k \in [0, 1]$, and cited evidence sets C_k . Two LLM agents debate the claim: a skeptic A (falsification-first) and a supporter B (verification-first). Across up to R rounds, they alternately propose labels $\ell_A, \ell_B \in Y_4$ with confidences $c_A, c_B \in [0, 1]$, rationales ρ_A, ρ_B , and citation sets $C_r^{(A)}, C_r^{(B)}$. Each agent maintains a citation history $C_{<r}^{(j)}$ and receives a novelty factor (1 for new evidence, $\pi < 1$ for fully reused citations) which down-weights its effective confidence $\tilde{c}_r^{(j)} = \kappa_r^{(j)} c_r^{(j)}$. All turns are logged in a global transcript \mathcal{T} , and the debate may terminate early if both agents agree on the same label.

In the judge phase, a separate LLM judge is provided with the full transcript \mathcal{T} and the final turns. It outputs a candidate label $\ell_J \in Y_4$, confidence $c_J \in [0, 1]$, and an override flag $o_J \in \{0, 1\}$. If c_J exceeds a threshold τ_J , the MAD outcome is adopted as the final decision $y^* = \ell_J$; otherwise, the system falls back to a simple fusion of Stage-1 labels and confidences $\{(y_k, s_k)\}_{k \in K}$. The procedure also returns a diagnostic record summarizing whether MAD was trusted, the judge's choice, the number of rounds used, the final agent states, and the full transcript, enabling transparent analysis of how the debate influenced the final claim decision.

B. IMAGE-ONLY DEEPFAKE DETECTORS

To contrast artifact-centric perception with evidence-centric reasoning, we evaluate five representative image-only detectors such as GRAMNet [49], BNext-M [50], NPR [51], FatFormer [52], and ProDet [53], where each exemplifying a major design paradigm in visual forensics. All models operate solely on the image, without access to text or external evidence, and are evaluated under identical preprocessing, thresholding, and scoring protocols to ensure fairness and reproducibility.

- **GRAMNet.** GRAMNet models cross-layer Gram-matrix statistics to capture generator-specific texture and frequency correlations. It is representative of frequency-fingerprint approaches that achieve strong performance on curated benchmarks but degrade under unseen generators or heavy compression. *Justification:* Represents classical frequency-domain detectors.
- **BNext-M.** BNext-M emphasizes physiological and structural consistency, learning illumination and facial micro-geometry cues via convolutional backbones. It performs well on portrait-based forgeries but struggles on contextual or non-face manipulations. *Justification:* Covers structure-aware models common in early deep-fake detection.

Algorithm 2 Multi-Agent Debate (MAD) for Final Decision

Require: Stage-1 evidence bundle with modality-specific subtasks $k \in K$ (e.g., $K = \{\text{text}, \text{image}\}$).

Ensure: Final label $y^* \in Y_4$.

- 1: **Initialize debate:** $\text{transcript} \leftarrow []$; $a_last \leftarrow \perp$;
 $b_last \leftarrow \perp$.
- 2: For each agent $j \in \{A, B\}$, set citation history $C_{<1}^{(j)} \leftarrow \emptyset$.
- 3: Agent roles: $A = \text{skeptic}$, $B = \text{supporter}$.
- 4: **for** $r = 1$ **to** R **do** ▷ Round-based debate
- 5: **Agent A turn (skeptic):**
 Provide inputs: B , transcript , r , R ,
 $\text{role} = \text{skeptic}$.
 Receive $(\ell_A, c_A, \rho_A, C_r^{(A)})$ where $\ell_A \in Y_4$, $c_A \in [0, 1]$.
- 6: Compute novelty factor:
 $\kappa_r^{(A)} \leftarrow \begin{cases} 1, & \text{if } C_r^{(A)} \setminus C_{<r}^{(A)} \neq \emptyset \\ \pi, & \text{otherwise} \end{cases}$
- 7: Effective confidence $\tilde{c}_A \leftarrow \kappa_r^{(A)} c_A$.
- 8: Update history $C_{<r+1}^{(A)} \leftarrow C_{<r}^{(A)} \cup C_r^{(A)}$.
- 9: Append agent-A record to transcript; set $a_last \leftarrow (\ell_A, \tilde{c}_A, \rho_A, C_r^{(A)})$.
- 10: **Agent B turn (supporter):**
 Provide inputs: B , transcript , r , R ,
 $\text{role} = \text{supporter}$.
 Receive $(\ell_B, c_B, \rho_B, C_r^{(B)})$.
- 11: Compute novelty factor:
 $\kappa_r^{(B)} \leftarrow \begin{cases} 1, & \text{if } C_r^{(B)} \setminus C_{<r}^{(B)} \neq \emptyset \\ \pi, & \text{otherwise} \end{cases}$
- 12: Effective confidence $\tilde{c}_B \leftarrow \kappa_r^{(B)} c_B$.
- 13: Update history $C_{<r+1}^{(B)} \leftarrow C_{<r}^{(B)} \cup C_r^{(B)}$.
- 14: Append agent-B record to transcript; set $b_last \leftarrow (\ell_B, \tilde{c}_B, \rho_B, C_r^{(B)})$.
- 15: **if** stop_on_consensus **and** $\ell_A = \ell_B$ **then**
- 16: **break**
- 17: **end if**
- 18: **end for**
- Judge phase:**
 Provide inputs: B , transcript , a_last ,
 b_last .
 Receive judge output (ℓ_J, c_J, o_J) with $\ell_J \in Y_4$.
- 19: **if** $c_J \geq \theta_{\text{judge}}$ **then**
- 20: $y^* \leftarrow \ell_J$
- 21: **else**
- 22: Obtain y^* by simple fusion of Stage-1.
- 23: **end if**

- **NPR** NPR leverages sensor-noise priors and patch-residual statistics to flag local inconsistencies between natural and synthesized regions. It generalizes moderately across generators but is easily disrupted by benign image transformations. *Justification*: Captures residual-based anomaly detectors.
- **FatFormer**. FatFormer employs transformer encoders with artifact-aware objectives, modeling long-range dependencies and achieving state-of-the-art accuracy on several forensics benchmarks. However, it tends to over-flag stylistically atypical yet authentic images. *Justification*: Represents high-capacity transformer-based detectors common in recent leaderboards.
- **ProDet**. ProDet focuses on provenance-agnostic cues such as camera-pipeline and compression irregularities, aiming to generalize across unknown generators but often trading precision for recall. *Justification*: Provides a manipulation agnostic baseline for general-purpose detection.

C. HYBRID FACT-CHECKING SYSTEM WITH DEEPFAKE-DETECTOR

For the hybrid model used to address RQ2, we extend the fact-checking pipeline by incorporating traditional image-based deepfake detector predictions as additional image-side evidence. The detector's binary authenticity verdict (*real* or *fake*) and its associated confidence score are first converted into a structured evidence unit, which is then injected into both stages of the fact-checking process.

During the MCTS-guided search, this detector-derived evidence unit is added to the initial evidence set E available to the planner. The multimodal LLM may query or reference this unit when forming intermediate hypotheses about the image's authenticity or when deciding whether further provenance, temporal, or semantic evidence is required. In this way, the detector's output can influence the search trajectory and is treated as one evidence source among others, alongside retrieved items from web/RAG, image analysis, and other tools.

The same evidence unit is passed to the Multi-Agent Debate module as part of the shared evidence pack provided to the supporter and skeptic agents. Both agents may cite the authenticity cue when constructing arguments, and the Judge model considers the detector's output when synthesizing the debate transcript and issuing the final decision. This integration strategy ensures that deepfake detectors function strictly as auxiliary evidence within the decision process, mirroring real-world conditions in which artifact-level authenticity cues shape users' prior beliefs and judgments. The hybrid model thus provides a controlled setting for evaluating whether pixel-level forgery signals genuinely contribute to or instead interfere with evidence-centric multimodal reasoning.

IV. RESULTS AND ANALYSIS

A. BENCHMARKS

We evaluate our systems on two multimodal misinformation benchmarks: **MMFakeBench** and **DGM4**. Following estab-

lished practice, we use the official *test* split of MMFakeBench and the official *validation* split of DGM4.

The datasets span diverse narrative domains (politics, crisis, health, entertainment), mixtures of authentic and synthetic imagery, and frequent contextual misalignment (authentic images paired with false captions). MMFakeBench emphasizes caption-image consistency with broad topical coverage, whereas DGM4 focuses on manipulation diversity and generator heterogeneity. Together, they capture both (i) narrative/context falsity and (ii) artifact diversity.

We evaluate all methods using **accuracy**, **precision**, **recall**, and **F1-score**. These metrics jointly measure overall correctness and the trade-off between false alarms and missed detections—both crucial for multimodal misinformation detection. Accuracy reflects global reliability, while precision, recall, and F1-score capture sensitivity to false claims and robustness against misclassifying true ones, providing a fair and interpretable basis for comparison.

B. QUANTITATIVE ANALYSIS

To ensure clarity and direct comparability, we report results using a single unified table per dataset. Each table includes (i) five representative deepfake detectors, (ii) the hybrid detector-augmented fact-checking model, and (iii) the evidence-centric fact-checking system. This consolidated reporting eliminates ambiguity and allows strict side-by-side comparison under identical evaluation settings.

1) MMFakeBench

Table 1 presents the unified comparison on MMFakeBench. Among the detectors, FatFormer achieves the highest performance (Acc = 0.79, F1 = 0.53), followed by BNext-M and NPR (F1 \approx 0.53). GRAMNet attains reasonable accuracy (0.75) but weaker balanced metrics (F1 = 0.48), while ProDet performs poorly under claim-level supervision (F1 = 0.26). Overall, artifact-only cues plateau around F1 \approx 0.53, confirming that pixel-based signals do not transfer well to a benchmark dominated by caption-driven and out-of-context misinformation.

Integrating detector predictions into the fact-checking system *reduces* performance (Acc = 0.76, F1 = 0.74). Although the hybrid model benefits from evidence-gathering modules, the injected pixel-level authenticity prior introduces non-causal noise particularly harmful when authentic images accompany deceptive captions. Our evidence-centric system achieves the strongest results (Acc = 0.82, F1 = 0.81), outperforming all baselines across all metrics.

Table 1 highlights a clear trend: *artifact-centric detectors plateau, hybridization degrades results, and evidence-driven reasoning performs best*.

A deeper analysis explains why detector integration consistently harms performance which can be seen from Figure 3. Most MMFakeBench “fake” instances stem from textual distortions, repurposing, semantic contradiction, or temporal mismatch; only a minority contain visual manipulations detectable by pixel-level forensic models. Consequently, deep-

fake detectors tend to predict *real* for many false claims. When these misleading priors are injected into the fact-checker, they conflict with retrieved provenance or semantic evidence. Confusion matrices confirm that the hybrid variant exhibits higher false-negative *and* false-positive rates, indicating that detector signals suppress the fact-checker's ability to flag false claims and simultaneously inflate errors on true claims. These results verify that on MMFakeBench, where misinformation is overwhelmingly semantic, pixel-level authenticity signals undermine, rather than enhance, multimodal verification.

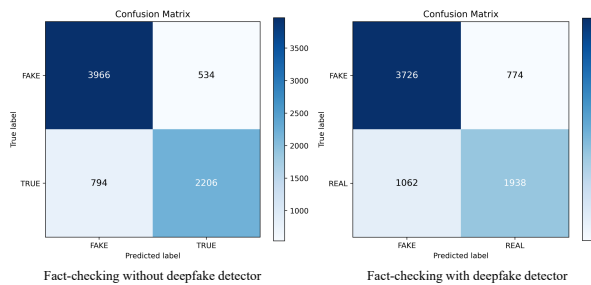


FIGURE 3. The confusion matrix of the fact-checking with and without image-based deepfake detector on MMFakebench dataset.

2) DGM4

Unified results for DGM4 are shown in Table 2. DGM4 is more heterogeneous, with substantial generator diversity and mixed manipulation styles. Among the detectors, ProDet achieves the best F1-score (0.49), outperforming FatFormer, GRAMNet, and NPR ($F1 \approx 0.33-0.35$). Yet all detectors fall within a narrow accuracy band (0.46–0.52), again highlighting the limited transferability of artifact cues to claim-level tasks.

The hybrid fact-checking + detector model yields $F1 = 0.51$, slightly below our evidence-centric system. While it benefits from structured evidence acquisition, the detector's artifact priors introduce ambiguity for cases involving authentic images paired with false captions or temporal inconsistencies. Our evidence-centric system achieves the highest performance overall ($Acc = 0.55, F1 = 0.55$). Its behavior reflects DGM4's difficulty: the model expresses high precision for FAKE and high recall for REAL, consistent with the challenge of reasoning across unseen generators and diverse narrative structures.

A more granular analysis reveals why detector integration again fails to help. Although DGM4 includes genuine visual manipulations, such as face swaps (FS) and face attribute edits (FA), the manipulated regions are often tiny, subtle, or sentiment-driven. As a result, many fake claims remain visually plausible, and detectors frequently predict *real* even for manipulated cases, which is evident from Figure 4. Furthermore, hybrid (face+text) manipulations constitute only a minority, weakening the correlation between pixel artifacts and claim-level falsity. Injecting these weak priors into

the fact-checker introduces systematic bias, diminishing the influence of provenance, entailment, or contradiction cues. Confusion-matrix analyses confirm that the hybrid model exhibits elevated misclassification rates across classes. Thus, even in a dataset with more visible manipulations, artifact cues remain unreliable for claim verification.

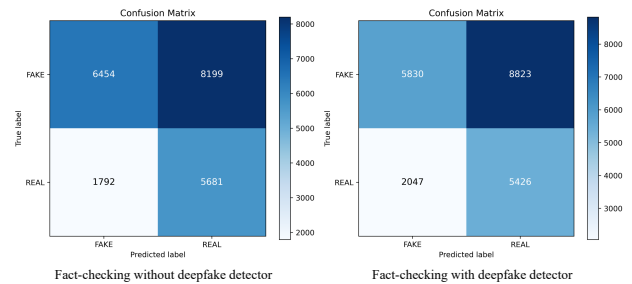


FIGURE 4. Confusion matrices on the DGM⁴ dataset for the evidence-centric fact-checking system and the hybrid system integrating a deepfake detector.

C. QUALITATIVE ANALYSIS

Figure 5 shows qualitative examples comparing our evidence-based fact-checking system with the hybrid variant that also uses a deepfake detector signal. The top examples highlight how evidence-centric fact-checking can succeed even when the image itself is visually plausible. In the REAL case, the system finds supporting sources that match the key details in the caption (e.g., event context and timing), and the image content appears consistent with that context. In the FAKE case, the image looks ordinary and could fit many narratives, but external sources and provenance checks do not support the specific story asserted by the text. In both cases, the model treats the image as helpful context, while letting the final decision be driven by whether the claim can be substantiated.

The bottom row illustrates the contrasting behavior of the hybrid fact-checking system, which integrates deepfake detector outputs as auxiliary cues. In the left example, the hybrid model correctly labels a clearly synthetic image as FAKE because the forgery detector produces a confident IMAGE_FAKE signal that aligns with the claim's falsity. However, the rightmost example exposes a critical failure mode: the detector labels a subtly manipulated or synthetic image as REAL with high confidence, despite the claim being false. This erroneous authenticity cue biases the downstream reasoning process, causing the system to underweight semantic and contextual inconsistencies and ultimately misclassify the claim as REAL. This case concretely illustrates how artifact-centric priors can override evidence-based reasoning when pixel-level manipulations are subtle or fall outside the detector's learned artifact space. This suggests deepfake detector outputs should be treated carefully as an auxiliary hint rather than a deciding factor—and balanced against evidence consistency and contextual checks during claim verification.

TABLE 1. Unified performance comparison on MMFakeBench across image-only deepfake detectors, the evidence-centric fact-checking system, and the hybrid (detector-integrated) variant.

Method	Accuracy	Precision	Recall	F1-score
GRAMNet	0.74	0.48	0.49	0.47
BNext-M	0.74	0.58	0.51	0.53
NPR	0.64	0.52	0.53	0.52
FatFormer	0.79	0.61	0.53	0.53
ProDet	0.26	0.54	0.52	0.25
Fact-checking + deepfake detector	0.76	0.74	0.74	0.74
Fact-checking (ours)	0.82	0.82	0.81	0.81

TABLE 2. Unified performance comparison on DGM4 across image-only deepfake detectors, the evidence-centric fact-checking system, and the hybrid (detector-integrated) variant.

Method	Accuracy	Precision	Recall	F1-score
GRAMNet	0.46	0.51	0.50	0.35
BNext-M	0.46	0.49	0.49	0.37
NPR	0.45	0.46	0.49	0.33
FatFormer	0.46	0.60	0.50	0.33
ProDet	0.51	0.50	0.50	0.49
Fact-checking + deepfake detector	0.51	0.56	0.55	0.51
Fact-checking (ours)	0.55	0.60	0.61	0.55

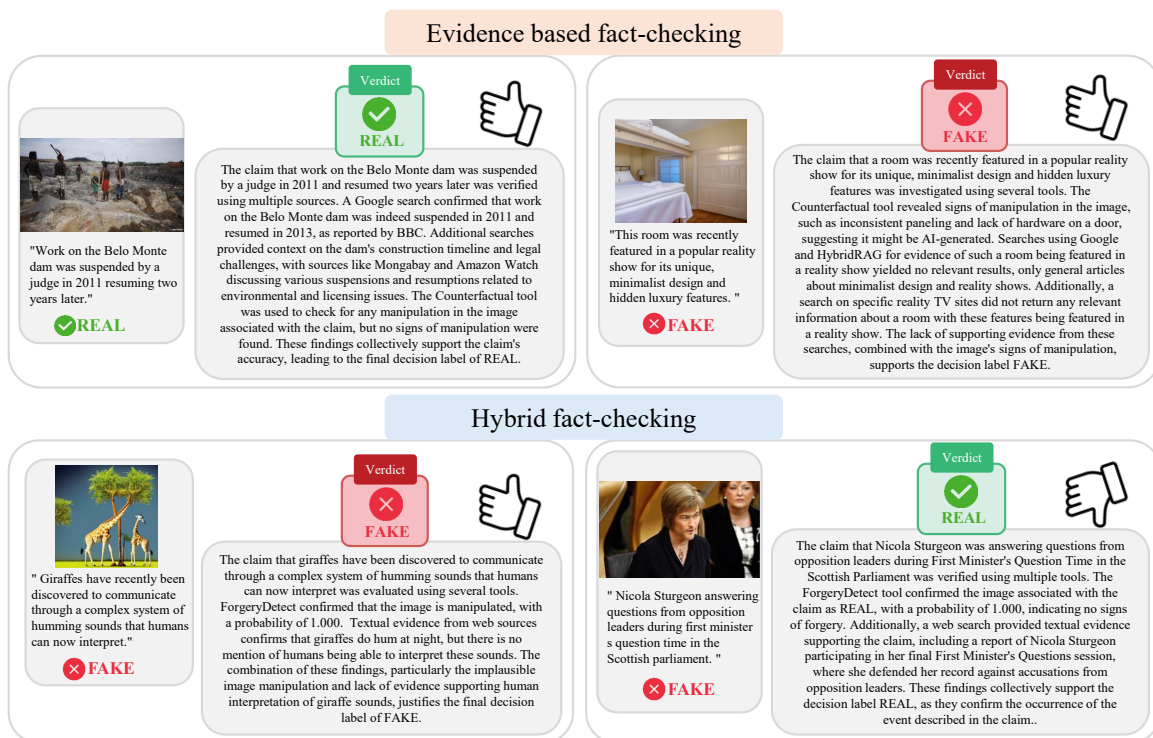


FIGURE 5. Qualitative cases showing robust evidence-centric verification and a failure mode of hybrid fact-checking caused by misleading deepfake detector outputs.

V. CASE STUDY

In the first example shown in Figure 6, the fact-checking system correctly classifies a REAL claim describing Texas Governor George W. Bush watching election returns with his father on November 7, 2000. As illustrated in the figure, the

image depicts a plausible historical setting with recognizable individuals and a natural indoor context, which is consistent with the claim. Rather than relying on visual appearance alone, the system retrieves multiple independent textual sources, including archival news descriptions and stock image

Evidence-centric system



News Text: Texas Gov George W Bush left the Republican presidential candidate watches elections returns with his father former president George HW Bush on Nov 7 2000 in Austin Texas.
GT:REAL
Label: REAL

Claim Report

The claim that Texas Governor George W. Bush watched election returns with his father, former President George H.W. Bush, on November 7, 2000, in Austin, Texas, was verified through multiple checks. The Counterfactual tool confirmed the authenticity of an image depicting George W. Bush, George H.W. Bush, and Barbara Bush in a natural setting, with no signs of manipulation. Additionally, a Google search provided supporting evidence, including a stock image description and news articles confirming that George W. Bush was indeed with his father during the election night in Austin [1][2]. These findings collectively support the claim's accuracy, justifying the label REAL.

References:

1. <https://www.bridgemanimages.com/en/naartistknown/bush-with-father-during-election-night-in-austin-2000-11-07-photo/photograph/asset/7799527>
2. <http://georgebush.chez.com/en/news/2000/2000-11-07/index.html>

Decision Trace

[depth 0](image)

Thought: We need to verify the authenticity of the image. A direct question about potential manipulation could reveal decisive evidence.

Action: Counterfactual[is there any sign of manipulation?]

Observation: Upon analyzing the provided image and its crops, I have examined various aspects for signs of manipulation, including face swap, face attribute manipulation, implausible elements, contradictions, or anomalies.

Here are the observations:

1. ****Overall Coherence:**** The image depicts three individuals (identifiable as George W. Bush, George H.W. Bush, and Barbara Bush) in what appears to be a domestic or office setting. The composition, lighting, and depth of field appear natural and consistent.

2. ****Facial Analysis:****

* ****George W. Bush (foreground):**** His face appears authentic. The skin texture, lighting, and expression (listening intently on the phone) are consistent with a genuine photograph. There are no signs of unnatural blending, sharp edges, or inconsistent features that wo...

[depth 2](image)

Thought: The Counterfactual analysis indicates no signs of manipulation or anomalies in the image. The image appears authentic based on the comprehensive analysis provided.

Action: Finish [REAL]

Observation: [FINISH]

[depth 0](Text)

Thought: The claim involves a historical event where George W. Bush watched election returns with his father on November 7, 2000. I need to verify the accuracy of this event using text evidence.

Action: Google["George W. Bush election night November 7 2000"]

Observation:

- 2000 United States presidential election recount in Florida Bush and Al Gore. The Florida vote was ultimately settled in Bush's favor by a margin of 537 votes out of 5,963,110 cast when the U.S. Supreme Court, in Bush v. https://en.wikipedia.org/wiki/2000_United_States_presidential_election_recount_in_Florida

- Election Night Rally Watch | Video Election night rallies for Governor George W. Bush and Vice President Al Gore were shown in a split screen format so that viewers could see ...

https://www.c-span.org/program/public-affairs-event/election-night-rally-watch/102231

- Bush Election Stirs Political and Legal Controversy

The 2000 U.S. presidential election, particularly the voting results in Florida, ignited significant political and legal controversy that has been the ...

https://www...

[depth 1](Text)

Thought: The previous search did not provide specific evidence about George W. Bush watching election returns with his father on November 7, 2000. I need to find a more targeted source that confirms this specific event.

Action: Wikipedia[George W. Bush 2000 election night]

Observation: George W. Bush — President of the United States from 2001 to 2009: George Walker Bush is an American politician, businessman, and former U.S. Air Force officer who served as the 43rd president of the United States from 2001 to 2009. A member of the Republican Party and the eldest son of the 41st president, George H. W. Bush, he served as the 46th governor of Texas from 1995 to 2000.

[depth 0](Text)

Thought: The claim states that George W. Bush, then Governor of Texas, watched election returns with his father on November 7, 2000. This is a historical event, and I need to verify if this specific detail is supported by reputable sources. Previous searches have not provided decisive evidence.

Action: Google["George W. Bush election night November 7 2000 with father"]

Observation: - BUSH WITH FATHER DURING ELECTION NIGHT IN ...

Download stock image by Rick Wilking - BUSH WITH FATHER DURING ELECTION NIGHT IN AUSTIN, 2000-11-07 - High quality fine art images, ...

https://www.bridgemanimages.com/en/naartistknown/bush-with-father-during-election-night-in-austin-2000-11-07-photo/photograph/asset/7799527

- George Bush > News > 2000 > Election Night

Austin (Texas), November 7, 2000: Texas Governor George W. Bush waits for election results with his father, former President George Bush, in the Governor's ...

http://georgebush.chez.com/en/news/2000/2000-11-07/index.html

- The Florida Recount Of 2000: A Nightmare That Goes On ...

The weeks-long battle over "hanging chads" that ultimately landed the fate of the presidency in the U.S. Supreme Court, continues to cast a ...

https...

[depth 2](Text)

Thought: The previous Google search provided strong support that George W. Bush watched election returns with his father on November 7, 2000, in Austin, Texas. Multiple sources confirm this event, and there are no refutations. This is sufficient to conclude the textual claim is true.

Action: Finish[REAL]

Observation: [FINISH]

===== Multi-Agent Debate Transcript =====

[Round 1] Agent A: REAL (conf=0.70, cites=[T1], [I1])

rationale: The text is supported by multiple sources confirming the event of George W. Bush watching election returns with his father on November 7, 2000 [T1]. The image analysis shows no signs of manipulation, with all elements appearing natural and consistent [I1]. Both subtasks indicate authenticity, making the claim and image likely real.

[Round 2] Agent B: REAL (conf=0.70, cites=[T1], [I1])

rationale: The text claim is corroborated by multiple sources confirming the event of George W. Bush watching election returns with his father on November 7, 2000, in Austin, Texas [T1]. The image analysis reveals no signs of manipulation, with all elements appearing natural and consistent [I1]. Both the textual and visual evidence strongly support the authenticity of the claim and image, justifying the label 'REAL'.

===== Agent Finals =====

Agent A → REAL (conf=0.70)

rationale: The text is supported by multiple sources confirming the event of George W. Bush watching election returns with his father on November 7, 2000 [T1]. The image analysis shows no signs of manipulation, with all elements appearing natural and consistent [I1]. Both subtasks indicate authenticity, making the claim and image likely real.

cites:[T1], [I1]

Agent B → REAL (conf=0.70)

rationale: The text claim is corroborated by multiple sources confirming the event of George W. Bush watching election returns with his father on November 7, 2000, in Austin, Texas [T1]. The image analysis reveals no signs of manipulation, with all elements appearing natural and consistent [I1]. Both the textual and visual evidence strongly support the authenticity of the claim and image, justifying the label 'REAL'.

cites:[T1], [I1]

===== Judge =====

Judge label: REAL

FIGURE 6. The case study of correct REAL classification through evidence-centric fact-checking, combining external source verification with image consistency checks.

records, which explicitly confirm the event, its date, and location. Image analysis tools further find no indications of manipulation, reinforcing but not determining the assessment. During the Multi-Agent Debate, both verifier agents cite corroborating textual evidence alongside the image consistency checks and identify no temporal or contextual contradictions. The Judge therefore issues a REAL verdict. This example demonstrates how the proposed evidence-centric framework uses the image as contextual support while grounding the final decision primarily in verifiable external evidence and historical records.

Figure 7 illustrates a representative failure case of the hybrid system, where the ground-truth label is FAKE due to image manipulation performed using an InfoSwap-based face-swapping method, but the system ultimately predicts REAL. As shown in the figure, the manipulated image appears visually natural, with consistent lighting, facial structure, and background context, making the swap difficult to detect through surface-level inspection. As InfoSwap produces highly realistic results with minimal pixel-level artifacts, the deepfake detector assigns a confident REAL label to the image. When this detector output is incorporated into the fact-checking pipeline, it acts as an authenticity cue that subtly biases the reasoning process. In the decision trace, the system terminates visual analysis early based on the detector's high-confidence prediction, while subsequent textual evidence, although factually correct in isolation does not trigger further scrutiny of image authenticity. During the Multi-Agent Debate, both verifier agents reference the detector output alongside corroborating text sources, and the Judge issues a REAL verdict. This case highlights how artifact-centric signals, when faced with realistic face-swapping techniques, may inadvertently reinforce incorrect assumptions about image authenticity, underscoring the importance of treating deepfake detector outputs as auxiliary signals rather than decisive evidence within multimodal claim verification

VI. CONCLUSION

This work examined a central question in multimodal misinformation detection: whether image-based deepfake detectors provide any meaningful value for verifying image-text claims, either as standalone predictors or as auxiliary signals within an evidence-centric fact-checking system or whether their artifact-centric predictions instead introduce misleading priors that degrade claim-level reasoning. To answer this question, we evaluated five state-of-the-art image-only detectors alongside our MCTS-guided, tool-augmented fact-checking system with Multi-Agent Debate across MMFakeBench and DGM⁴. Across both benchmarks, the evidence-centric fact-checking framework consistently outperforms all detector baselines, achieving F1 \approx 0.81 on MMFakeBench and F1 \approx 0.55 on DGM⁴, compared with F1 values of only 0.33–0.53 for image-based detectors. Moreover, integrating detector outputs into the fact-checking pipeline systematically *reduces* performance by 0.06–0.08 F1 on MMFakeBench and by roughly 0.04 F1 on DGM⁴

and increases both false negatives and false positives. These findings demonstrate that multimodal misinformation verification is inherently evidence-centric: the truth of an image-text claim is determined primarily by its semantic content and contextual grounding, not by pixel-level artifact signals. Our findings suggest that although image-based deepfake detectors are commonly deployed as upstream authenticity filters in real-world moderation systems and occasionally incorporated into multimodal fact-checking pipelines, they should not be treated as primary evidential signals for determining the truth of an image-text claim. Effective verification systems must instead prioritize targeted retrieval, provenance, and temporal checks, and deliberative adjudication. Artifact-level cues, when present, should be treated only as optional and non-binding auxiliary signals rather than decisive indicators of claim authenticity.

REFERENCES

- [1] Q. Peng, Y. Lu, Y. Peng, S. Qian, X. Liu, and C. Shen, "Crafting synthetic realities: Examining visual realism and misinformation potential of photorealistic ai-generated images," in *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2025.
- [2] H. Farid, "Mitigating the harms of manipulated media: Confronting deepfakes and digital deception," *PNAS nexus*, vol. 4, no. 7, p. pgaf194, 2025.
- [3] J. Lee, M. Hameleers, and S. Y. Shin, "The emotional effects of multimodal disinformation: How multimodality, issue relevance, and anxiety affect misperceptions about the flu vaccine," *new media & society*, vol. 26, no. 12, pp. 6838–6860, 2024.
- [4] A. O. Ajibare, A. T. Ajibare, and O. V. Ndubueze, "Disinformation in the digital era: A systematic review of its evolution, impact, and mitigation strategies," 2025.
- [5] A. Fichera, "Video shows doctor collapsing in early 2020, not covid vaccine effect," *AP News*, 2025.
- [6] M. Davey, "'alarming': convincing ai vaccine and vaping disinformation generated by australian researchers," *The Guardian*, 2023.
- [7] F. Shalabi, H. H. Nguyen, H. Felouat, C.-C. Chang, and I. Echizen, "Image-text out-of-context detection using synthetic multimodal misinformation," in *2023 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pp. 605–612, IEEE, 2023.
- [8] A. Yermakov, J. Cech, J. Matas, and M. Fritz, "Deepfake detection that generalizes across benchmarks," *arXiv preprint arXiv:2508.06248*, 2025.
- [9] F.-A. Croitoru, A.-I. Hiji, V. Hondru, N. C. Ristea, P. Irofti, M. Popescu, C. Rusu, R. T. Ionescu, F. S. Khan, and M. Shah, "Deepfake media generation and detection in the generative ai era: a survey and outlook," *arXiv preprint arXiv:2411.19537*, 2024.
- [10] F. Cocchi, L. Baraldi, S. Poppi, M. Cornia, L. Baraldi, and R. Cucchiara, "Unveiling the impact of image transformations on deepfake detection: An experimental analysis," in *International Conference on Image Analysis and Processing*, pp. 345–356, Springer, 2023.
- [11] S. Lewandowsky, U. K. Ecker, C. M. Seifert, N. Schwarz, and J. Cook, "Misinformation and its correction: Continued influence and successful debiasing," *Psychological science in the public interest*, vol. 13, no. 3, pp. 106–131, 2012.
- [12] L. K. Fazio, N. M. Brashier, B. K. Payne, and E. J. Marsh, "Knowledge does not protect against illusory truth.," *Journal of experimental psychology: general*, vol. 144, no. 5, p. 993, 2015.
- [13] G. S. Cheema, S. Hakimov, E. Müller-Budack, C. Otto, J. A. Bateman, and R. Ewerth, "Understanding image-text relations and news values for multimodal news analysis," *Frontiers in artificial intelligence*, vol. 6, p. 1125533, 2023.
- [14] A. Rossler, D. Cozzolino, L. Verdoliva, C. Riess, J. Thies, and M. Nießner, "Faceforensics++: Learning to detect manipulated facial images," in *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1–11, 2019.
- [15] Y. Li, X. Yang, P. Sun, H. Qi, and S. Lyu, "Celeb-df: A large-scale challenging dataset for deepfake forensics," in *Proceedings of the IEEE/CVF*

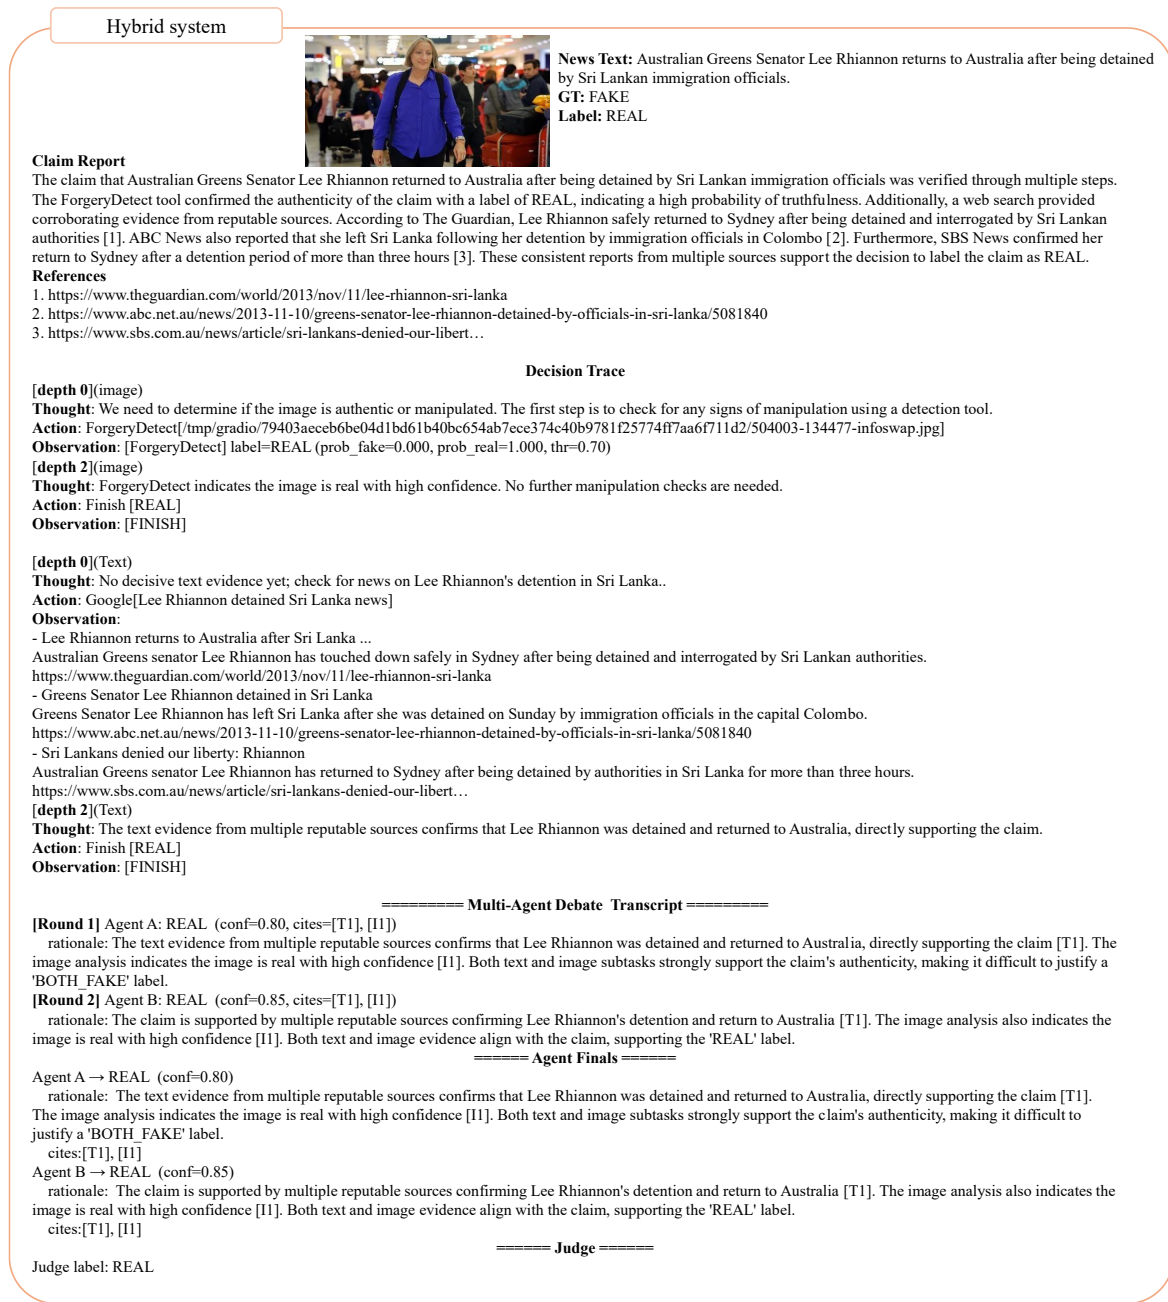


FIGURE 7. Case study of hybrid fact-checking system misclassifies a FAKE claim after the deepfake detector predicts the image as REAL which highlights the need to treat detector signals carefully in the fact-checking system.

conference on computer vision and pattern recognition, pp. 3207–3216, 2020.

[16] B. Dolhansky, J. Bitton, B. Pflaum, J. Lu, R. Howes, M. Wang, and C. C. Ferrer, “The deepfake detection challenge (dfdc) dataset,” *arXiv preprint arXiv:2006.07397*, 2020.

[17] A. A. Hasanaath, H. Luqman, R. Katib, and S. Anwar, “Fsb: Deepfake detection with frequency enhanced self-blended images,” *Image and Vision Computing*, vol. 154, p. 105418, 2025.

[18] T.-M. Huang, Y.-H. Han, E. Chu, S.-T. Lo, K.-L. Hua, and J.-C. Chen, “Generalized image-based deepfake detection through foundation model adaptation,” in *International Conference on Pattern Recognition*, pp. 194–210, Springer, 2024.

[19] Z. Huang, J. Hu, X. Li, Y. He, X. Zhao, B. Peng, B. Wu, X. Huang, and G. Cheng, “Sida: Social media image deepfake detection, localization and explanation with large multimodal model,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 28831–28841, 2025.

[20] L. Kroiß and J. Reschke, “Deepfake detection of face images based on a convolutional neural network,” *arXiv preprint arXiv:2503.11389*, 2025.

[21] M. A. Arshed, S. Mumtaz, M. Ibrahim, C. Dewi, M. Tanveer, and S. Ahmed, “Multiclass ai-generated deepfake face detection using patch-wise deep learning model,” *Computers*, vol. 13, no. 1, p. 31, 2024.

[22] M. A. Francis, A. R. Kurup, B. Premjith, and B. R. Chakravarthi, “Multimodal fake news classification in tamil using fact-checked social media content and cost-sensitive learning,” *IEEE Access*, 2025.

[23] A. O. Ojo, F. Najar, N. Zamzami, H. T. Hindi, and N. Bouguila, “Smoothdetector: A smoothed dirichlet multimodal approach for combat-

- ing fake news on social media," *IEEE Access*, vol. 13, pp. 39289–39305, 2025.
- [24] S. Aneja, C. Bregler, and M. Nießner, "Cosmos: Catching out-of-context misinformation with self-supervised learning," *arXiv preprint arXiv:2101.06278*, 2021.
- [25] Y. Zhou, A. Pang, and G. Yu, "Clip-gcn: an adaptive detection model for multimodal emergent fake news domains," *Complex & Intelligent Systems*, vol. 10, no. 4, pp. 5153–5170, 2024.
- [26] H. Yu, H. Wu, X. Fang, M. Li, and H. Zhang, "Sr-cibn: Semantic relationship-based consistency and inconsistency balancing network for multimodal fake news detection," *Neurocomputing*, vol. 635, p. 129997, 2025.
- [27] P. Qi, Z. Yan, W. Hsu, and 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*, pp. 13052–13062, 2024.
- [28] A. M. Luvembe, W. Li, S. Li, F. Liu, and X. Wu, "Caf-odnn: Complementary attention fusion with optimized deep neural network for multimodal fake news detection," *Information Processing & Management*, vol. 61, no. 3, p. 103653, 2024.
- [29] S. Tahmasebi, E. Müller-Budack, and R. Ewerth, "Multimodal misinformation detection using large vision-language models," in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 2189–2199, 2024.
- [30] Z. Yan, P. Qi, W. Hsu, and M.-L. Lee, "Trust-vl: An explainable news assistant for general multimodal misinformation detection," in *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 5588–5604, 2025.
- [31] Q. Xu, H. Du, S. Lukasik, T. Zhu, S. Wang, and X. Yu, "Mdam3: A misinformation detection and analysis framework for multitype multimodal media," in *Proceedings of the ACM on Web Conference 2025*, pp. 5285–5296, 2025.
- [32] T. M. Gedara, R. Gaeta, S. Tomasiello, and R. El Shawi, "Toward explainable multimodal fake news detection: A fusion of prompt learning and automated machine learning," *IEEE Access*, vol. 13, pp. 188150–188169, 2025.
- [33] M. Akhtar, M. Schlichtkrull, Z. Guo, O. Cocarascu, E. Simperl, and A. Vlachos, "Multimodal automated fact-checking: A survey," *arXiv preprint arXiv:2305.13507*, 2023.
- [34] T. Braun, M. Rothermel, M. Rohrbach, and A. Rohrbach, "Defame: Dynamic evidence-based fact-checking with multimodal experts," *arXiv preprint arXiv:2412.10510*, 2024.
- [35] A. Beigi, B. Jiang, D. Li, Z. Tan, P. Shaeri, T. Kumarage, A. Bhattacharjee, and H. Liu, "Can llms improve multimodal fact-checking by asking relevant questions?," *arXiv preprint arXiv:2410.04616*, 2024.
- [36] A. U. Dey, M. J. Awan, G. Channing, C. S. de Witt, and J. Collomosse, "Fact-checking with contextual narratives: Leveraging retrieval-augmented llms for social media analysis," *arXiv preprint arXiv:2504.10166*, 2025.
- [37] H. Ou, G. Deng, X. Han, J. Zhang, X. He, H. Qiu, S. Guo, and T. Zhang, "Holmes: Automated fact check with large language models," *arXiv preprint arXiv:2505.03135*, 2025.
- [38] K. R. Ahmed, M. S. Khan, M. A. R. Chowdhury, S. T. Nesar, M. A. Hosien, M. R. Karim, and N. Sultana, "Detecting misinformation with multimodal ai: Leveraging vision and nlp for fact-checking," in *2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN)*, pp. 1–6, IEEE, 2025.
- [39] K. Kakizaki, Y. Matsunaga, and R. Furukawa, "Maft: Multimodal automated fact-checking via textualization," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, pp. 29646–29648, 2025.
- [40] S. Yang, Z. Yu, Z. Ying, Y. Dai, G. Wang, J. Lan, J. Xu, J. Li, and E. C. Ngai, "Rama: Retrieval-augmented multi-agent framework for misinformation detection in multimodal fact-checking," *arXiv preprint arXiv:2507.09174*, 2025.
- [41] K. Vallidevi and V. Balaji, "Yours truly: A credibility framework for effortless llm-powered fact checking," *IEEE Access*, 2024.
- [42] S. Cui, H. H. Nguyen, T.-N. Le, C.-S. Lu, and I. Echizen, "Lookupforensics: A large-scale multi-task dataset for multi-phase image-based fact verification," *IEEE Access*, 2025.
- [43] U. Kangur, K. Agrawal, Y. Singh, A. Sabir, and R. Sharma, "Multireflect: Multimodal self-reflective rag-based automated fact-checking," in *Proceedings of the 1st Workshop on Multimodal Augmented Generation via Multimodal Retrieval (MAGMar 2025)*, pp. 1–17, 2025.
- [44] Z. Chen, Q. Zhang, J. Liu, Y. Wang, H. Lv, L. Wang, J. Duan, M. Xv, and H. Wang, "Counterfactual multimodal fact-checking method based on causal intervention," in *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, pp. 582–595, Springer, 2024.
- [45] F. Zhang, J. Liu, J. Xie, Q. Zhang, Y. Xu, and Z.-J. Zha, "Escnet: Entity-enhanced and stance checking network for multi-modal fact-checking," in *Proceedings of the ACM Web Conference 2024*, pp. 2429–2440, 2024.
- [46] M. N. S. Shopnil, S. Duwal, A. Tyagi, and A. M. Prama, "Mirage: Agentic framework for multimodal misinformation detection with web-grounded reasoning," *arXiv preprint arXiv:2510.17590*, 2025.
- [47] S.-I. Papadopoulos, C. Koutlis, S. Papadopoulos, and P. C. Petrantonakis, "Red-dot: Multimodal fact-checking via relevant evidence detection," *IEEE Transactions on Computational Social Systems*, 2025.
- [48] X. Cui, Y. Zou, Z. Li, P. Li, X. Xu, X. Liu, and H. Huang, "T² agent: a tool-augmented multimodal misinformation detection agent with monte carlo tree search," *arXiv preprint arXiv:2505.19768*, 2025.
- [49] Z. Liu, X. Qi, and P. H. Torr, "Global texture enhancement for fake face detection in the wild," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8060–8069, 2020.
- [50] R. Lanzino, F. Fontana, A. Diko, M. R. Marini, and L. Cinque, "Faster than lies: Real-time deepfake detection using binary neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3771–3780, 2024.
- [51] C. Tan, Y. Zhao, S. Wei, G. Gu, P. Liu, and Y. Wei, "Rethinking the up-sampling operations in cnn-based generative network for generalizable deepfake detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28130–28139, 2024.
- [52] H. Liu, Z. Tan, C. Tan, Y. Wei, J. Wang, and Y. Zhao, "Forgery-aware adaptive transformer for generalizable synthetic image detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10770–10780, 2024.
- [53] J. Cheng, Z. Yan, Y. Zhang, Y. Luo, Z. Wang, and C. Li, "Can we leave deepfake data behind in training deepfake detector?," *Advances in Neural Information Processing Systems*, vol. 37, pp. 21979–21998, 2024.



human-computer interaction.

A S M SHARIFUZZAMAN SAGAR received the bachelor's degree from China University of Mining and Technology, China, in 2020, and the master's degree from Sejong University, Seoul, Republic of Korea, in 2022. He is currently a PhD student at the Department of Computer Science and Software Engineering, The University of Western Australia. His current research interests include deep learning, fact-checking, agentic systems, embedded AI systems, the Internet of Things, and

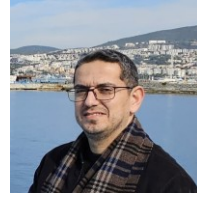


MOHAMMED BENNAMOUN is a Winthrop Professor at The University of Western Australia, Perth. He has led breakthroughs in 3D vision, deep learning, multimodal video understanding, human-robot interaction, biometrics, and large-scale AI systems. His innovations support critical capabilities in defense, crime prevention, transportation, agriculture, clinical diagnostics, and geospatial intelligence. With over 500 peer-reviewed publications and four books, his leadership in fundamental algorithms and industry translation sets new benchmarks for AI deployment across academia, government, and industry.



FARID BOUSSAID (Senior Member, IEEE) received the M.S. and Ph.D. degrees in microelectronics from the National Institute of Applied Science (INSA), Toulouse, France, in 1996 and 1999, respectively. He joined Edith Cowan University, Perth, WA, Australia, as a Postdoctoral Research Fellow and a member of the Visual Information Processing Research Group, in 2000. In 2005, he joined The University of Western Australia, Perth, WA, where he is currently a Professor. His current

research interests include neuromorphic engineering, smart sensors, and machine learning



SHAABAN SAHMOUD is an Associate Professor of Computer Engineering at Fatih Sultan Mehmet Vakıf University, Istanbul, Türkiye. He received his Ph.D. in Computer Engineering from Marmara University in 2019. He has more than 15 years of experience in software development and over 10 years of experience in Artificial Intelligence research and applications. His expertise centers on applied AI, including machine learning, computer vision, evolutionary algorithms, dynamic multi-

objective optimization, and intelligent data-driven systems. He has supervised numerous AI projects across academia and industry, translating advanced research into practical, scalable solutions. Dr. Sahnoud has authored more than 30 research papers in the fields of Artificial Intelligence and optimization, published in reputable international journals and conferences. In addition to his academic career, he worked as a Data Scientist at Al Jazeera Media Network and served for two years as an AI Consultant at the Presidency Office of the Republic of Türkiye, contributing to strategic AI initiatives and advanced intelligent system development. He actively bridges academia and industry by designing and deploying real-world AI applications that address complex challenges in media, government, and engineering sectors.



NAEHA SHARIF is a Lecturer at the Department of Computer Science and Software Engineering, UWA, specializing in Computer Vision (CV) and Natural Language Processing (NLP). She completed her PhD in 2021, focusing on a contemporary problem of Artificial Intelligence (AI)- Natural Language description of images. During her doctoral research, Dr. Sharif developed state-of-the-art captioning metrics and models. Her work emphasized the crucial role of linguistic knowl-

edge in enhancing machines' visual understanding. In recognition of her exceptional achievements, she was honoured with the prestigious "ACS 1962 Medal" for the most outstanding PhD in Western Australia in the field of Information Technology and Computer Science. Dr. Sharif's research is published in renowned venues such as the European Conference on Computer Vision (ECCV), the International Journal of Computer Vision (IJCV), the Association of Computational Linguistics (ACL), and the Winter Conference on Applications of Computer Vision (WACV). Her research interests include Vision-Language Models, Multimodal Learning, Automatic Image/Video Description, Captioning/Machine Translation metrics, Medical Image Analysis, and Automatic Medical Report Generation

ALI KISHK is a documentary filmmaker and investigative journalist who works with AlJazeera's long-form programmes. He co-produced several Al-Jazeera World documentaries on human-rights abuse and political struggles in conflict zones.

...



LIAN XU received the Ph.D. degree from the Department of Computer Science and Software Engineering, The University of Western Australia. She is currently a Research Fellow with The University of Western Australia. Her research interests include computer vision and machine learning.