

Beyond Detection: Exploring Evidence-based Multi-Agent Debate for Misinformation Intervention and Persuasion

Chen Han^{1,2}, Yijia Ma², Jin Tan³, Wenzhen Zheng^{2,4}, Xijin Tang^{1,2}

¹School of Advanced Interdisciplinary Sciences, UCAS

²State Key Laboratory of Mathematical Sciences, Academy of Mathematics and Systems Science, CAS

³School of Economics and Management, UCAS

⁴StepFun

hanchen23@mails.ucas.ac.cn, mayijia25@mails.ucas.ac.cn, tanjin23@mails.ucas.ac.cn,
zhengwenzhen@stepfun.com, xjtang@iss.ac.cn

Abstract

Multi-agent debate (MAD) frameworks have emerged as promising approaches for misinformation detection by simulating adversarial reasoning. While prior work has focused on detection accuracy, the importance of helping users understand the reasoning behind factual judgments has been overlooked. The debate transcripts generated during MAD offer a rich but underutilized resource for transparent reasoning. In this study, we introduce ED2D, an evidence-based MAD framework that extends previous approach by incorporating factual evidence retrieval. More importantly, ED2D is designed not only as a detection framework but also as a persuasive multi-agent system aimed at correcting user beliefs and discouraging misinformation sharing. We compare the persuasive effects of ED2D-generated debunking transcripts with those authored by human experts. Results demonstrate that ED2D outperforms existing baselines across three misinformation detection benchmarks. When ED2D generates correct predictions, its debunking transcripts exhibit persuasive effects comparable to those of human experts; However, when ED2D misclassifies, its accompanying explanations may inadvertently reinforce users' misconceptions, even when presented alongside accurate human explanations. Our findings highlight both the promise and the potential risks of deploying MAD systems for misinformation intervention. We further develop a public community website to help users explore ED2D, fostering transparency, critical thinking, and collaborative fact-checking.

Code —

<https://github.com/hanshenmesen/Debate-to-Detect>

Introduction

Misinformation presents a persistent threat to online discourse, public trust, and democratic institutions (Ansari and Alam 2025; Schmitt et al. 2025). In response, researchers have developed a range of computational approaches to automatically detect misleading content. Among these, multi-agent debate (MAD) frameworks have recently gained attention for their ability to simulate adversarial reasoning (Liang et al. 2024), where large language model (LLM) agents engage in structured argumentation to expose factual incon-

sistencies. These multi-agent systems leverage the complementary strengths of competing perspectives and often yield more robust judgments than a single classifier (Li et al. 2024; Zhang et al. 2025a).

Previous studies employing MAD frameworks for post-verification tasks that focus on identifying the falsehood of claims. However, in the real world, users are active reasoners rather than passive recipients, requiring persuasive explanations and the capacity to resist future misinformation (Danry et al. 2025). Simply labeling a claim as false is insufficient to foster resilience against misinformation (Lyu et al. 2025). Drawing on the adage that **the Truth Becomes Clearer Through Debate** (Han, Zheng, and Tang 2025; Liu et al. 2025), we argue that effective misinformation interventions should prioritize transparent reasoning, evidence grounding, and persuasive debunking.

In this study, we introduce ED2D, an evidence-based MAD framework that extends the Debate-to-Detect (D2D) framework (Han, Zheng, and Tang 2025) by incorporating an evidence retrieval module. It automatically identifies key entities and concepts within a claim, retrieves factual information from external sources, and assesses the stance of the retrieved evidence to the claim. By integrating verifiable facts into the argumentative process, ED2D mitigates hallucinations and achieves superior performance on three standard misinformation detection datasets.

We further evaluate ED2D's persuasive effect on human beliefs using Snopes25, a real-world benchmark consisting of fact-checks authored by professional editors from January to June 2025. It is collected after the training cut-off of GPT-4o, thereby reducing the risk of data leakage and better reflecting current misinformation trends. Using Snopes25, we compare the effects of ED2D-generated debunks with human-expert reports on belief correction and sharing intention. We also analyze failure cases in which ED2D misjudges claims yet still successfully persuades users, highlighting potential risks associated with the deployment of MAD systems.

In summary, our study makes the following contributions:

- We construct **Snopes25**, the first real-world benchmark explicitly designed to compare the persuasive impact of LLMs and human experts. Comprising 448 claims and corresponding fact-check reports authored by profes-

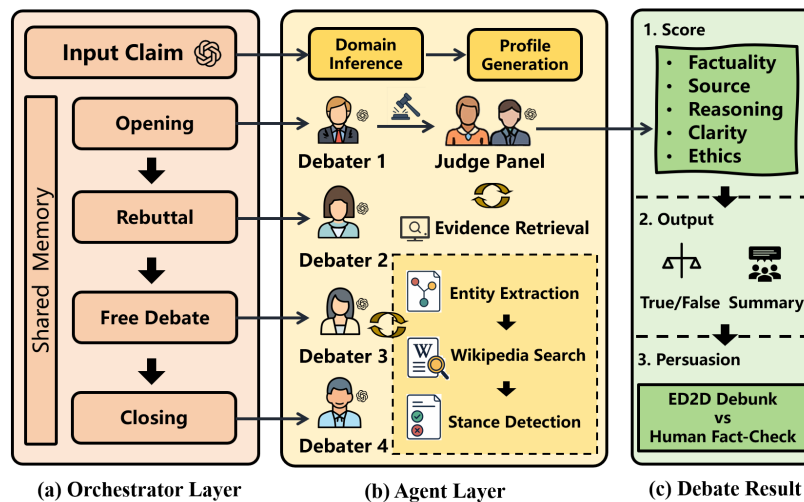


Figure 1: Architecture of the ED2D framework. Given a news claim, LLM agents with domain-specific profiles engage in a structured debate comprising five stages, including Opening, Rebuttal, Free Debate, Closing, and Judgment. During the Free Debate and Judgment, an evidence retrieval module actively retrieves relevant factual information from external sources to support or challenge arguments. All agents share the compressed history memory, enabling coherent multi-turn interactions.

sional editors between January and June 2025, Snopes25 provides high-quality fact-checking explanations and reflects contemporary misinformation patterns.

- We propose **ED2D**, an evidence-based MAD framework that integrates factual evidence retrieval to generate verifiable and coherent argumentation. ED2D outperforms existing baselines across three benchmarks including Snopes25, demonstrating both high detection accuracy and strong interpretability.
- We conduct the first controlled human-subject experiment to evaluate a **persuasive multi-agent system**, directly comparing the debunking effectiveness of ED2D-generated transcripts with expert-written fact-checks. Results indicate that ED2D achieves expert-level persuasive efficacy when correct, but reinforces misinformation when incorrect, revealing a fundamental tradeoff in the design of persuasive AI systems.
- We deploy a **public platform** that allows users to interactively engage with ED2D’s debate process. The system enhances transparency, fosters epistemic vigilance, and promotes collective fact-checking through discourse generated by LLM-agents.

Related Work

LLM-based Misinformation Detection

Early approaches focused on fine-tuning pretrained transformers such as BERT for binary veracity classification tasks (Li et al. 2021; Zhu et al. 2025). While these approaches perform well on benchmark datasets, they often lack interpretability and generalize poorly to novel claims (Pelrine et al. 2023; Han and Tang 2026). To overcome these limitations, recent work has shifted toward prompting-based paradigms using LLMs. Zero-shot and few-shot prompting enable flexible adaptation to misinformation detection tasks

without the need for explicit fine-tuning (Wu et al. 2023). Moreover, techniques such as Chain-of-Thought (Wei et al. 2022) and Self-Reflection (Madaan et al. 2023) enhance reasoning transparency by enabling models to decompose complex judgments into intermediate steps or engage in iterative self-critique. However, as these approaches depend entirely on internal model knowledge, they are prone to hallucinations and factual errors (Huang et al. 2025; Yu, Huang, and Liu 2025). To address these reliability concerns, an emerging line of research incorporates external evidence into the reasoning process. For example, Retrieval-Augmented Generation (RAG) has been widely adopted to ground LLM outputs in verifiable content. Nevertheless, most existing RAG systems still rely on a fixed knowledge base, which limits their effectiveness in open-domain misinformation detection scenarios (Wang et al. 2025).

Multi-Agent Debate for Adversarial Reasoning

MAD frameworks have gained traction as an approach to improve the deliberative quality of LLM outputs by simulating adversarial interactions between agents with opposing views (Du et al. 2024; Liang et al. 2024; Li et al. 2024). Drawing inspiration from human deliberation and dialectical reasoning, such frameworks allow agents to challenge assumptions, correct errors, and elicit deeper justifications (Zhang et al. 2025b).

Existing MAD systems vary in structure, ranging from simple two-agent dialogues to more complex setups involving multiple rounds, and role-specific prompting. The D2D framework introduces a five-stage debate process, including Opening, Rebuttal, Free Debate, Closing, and Judgment. Agents are assigned fixed stances and judged across multiple dimensions. While D2D improves factual accuracy and interpretability, the exclusive reliance on internal model knowledge introduces vulnerability to hallucinations

and limits robustness when handling emerging or unfamiliar claims. In response, ED2D incorporates an evidence retrieval component that enables agents to support or challenge claims with verifiable external information. This integration strengthens both adversarial reasoning and fact-grounded argumentation, enhancing the reliability and generalizability of the debate process.

Persuasive Interventions and Debunking

Beyond accurate classification, effective misinformation mitigation requires persuasive interventions that influence user beliefs and behaviors (Ecker et al. 2022). Research in cognitive psychology has explored the design of corrective messages, demonstrating that informative and clear explanations can reduce belief in falsehoods and discourage the sharing intention (Jahanbakhsh et al. 2021). Factors such as emotion, epistemic trust, and narrative coherence have been shown to moderate the effectiveness of fact-checking efforts (Pennycook et al. 2021; Scherer et al. 2021).

Professional fact-checking organizations, such as Snopes and PolitiFact, consistently produce structured and evidence-based rebuttals that represent the current gold standard in combating online misinformation. However, expert-written fact-checks are often costly and inherently lack scalability. Consequently, there is growing interest in automating persuasive debunking using LLMs (Salvi et al. 2025; Schoenegger et al. 2025). However, few studies have directly compared the persuasive efficacy of AI-generated versus expert-written explanations.

Our work contributes to this field by evaluating ED2D not only as a detection system but also as a persuasive agent. Using the Snopes25 benchmark, we compare ED2D’s impact on belief revision, sharing intention, and emotional alignment with that of expert fact-checking. Furthermore, we examine the failure cases in which ED2D outputs persuasive but factually inaccurate explanations, underscoring the dual-edged potential of LLM-generated content and the need for appropriate safeguards in real-world deployment.

Our Framework

Architectural Overview

Figure 1 presents the architecture of ED2D. The framework builds upon the five-stage debate structure: Opening Statement, Rebuttal, Free Debate, Closing Statement, and Judgment. At the core of the system lies the Agent Layer, comprising two debating teams—the Affirmative and the Negative—each consisting of four agents. These agents are assigned domain-specific profiles relevant to the input and are fixed to either a “True” or “Fake” stance. Within each team, agents collaborate to construct coherent arguments that support or refute the veracity of the claim.

The debate is evaluated by a panel of judge agents, who observe the full dialogue and score it across five dimensions: Factuality, Source Reliability, Reasoning Quality, Clarity, and Ethical Considerations. Each dimension is evaluated using a complementary scoring scheme in which paired scores sum to seven, thereby precluding any possibility of a tie. The

aggregated scores result in a definitive classification of the claim as either REAL or FAKE.

The Orchestrator Layer governs the overall debate process, assigning roles, scheduling turns across the five structured stages, and maintaining a compressed shared dialogue memory. To mitigate the context-length limitations of LLMs, the system performs stage-wise context compression, distilling salient information into concise summaries that guide subsequent reasoning stages. This mechanism ensures continuity and coherence throughout the multi-turn debate.

Evidence Retrieval and Integration

The main extension introduced in ED2D is an evidence retrieval module integrated into the Free Debate stage. Unlike prior frameworks that rely solely on language model internal knowledge, ED2D dynamically extracts key entities and relations from the input claim, retrieves external information from sources such as Wikipedia, and incorporates factual evidence into the ongoing debate. This module operates through four steps:

1. **Entity and Relation Extraction:** Using in-context prompting with LLMs, the system identifies up to five salient entities or concepts from the input claim.
2. **Evidence Retrieval:** The extracted elements are used to formulate structured queries to a Wikipedia-based API, returning a ranked list of relevant content segments.
3. **Stance Classification:** Retrieved content is evaluated using LLMs to determine its stance toward the original claim. The model classifies each evidence segment as supporting, refuting, or neutral, enabling targeted use of evidence in the subsequent debate.
4. **Evidence Integration:** During the Free Debate stage, debater agents incorporate supporting or refuting evidence into their responses, using it to support or challenge the veracity of the claim. Neutral evidence is preserved to support objective assessment by judge agents.

By grounding debates in retrieved evidence, ED2D enhances factual accuracy, mitigates hallucination risks, and supports more persuasive and transparent reasoning.

Judgment and Output

In the final stage, ED2D generates two outputs: a binary veracity label and a structured summary of the debate. The summary highlights key arguments, evidence-based rebuttals, and controversial points raised by both sides. The output enhances transparency and interpretability, supporting both machine-side evaluation and human-side engagement. In contrast to D2D, which emphasizes multi-dimensional scoring, ED2D prioritizes concise and interpretable decision-making. Such a design enables downstream integration into user-facing debunking interfaces and facilitates empirical analyses of persuasive effects on users’ beliefs and behaviors.

Method	Weibo21				FakeNewsDataset				Snopes25			
	Acc	Prec	Rec	F1	Acc	Prec	Rec	F1	Acc	Prec	Rec	F1
BERT	75.64	78.50	77.06	77.77	78.30	78.60	81.33	79.94	73.55	74.02	75.41	74.71
RoBERTa	79.82	80.42	81.75	81.08	81.17	81.03	83.39	82.19	75.26	78.12	76.08	77.09
ZS	67.11	65.74	68.90	67.28	66.31	65.57	68.67	67.09	60.04	64.78	63.49	64.13
w/ evidence	74.26	72.59	74.90	73.73	73.34	72.05	74.41	73.21	69.41	68.81	70.90	69.84
CoT	74.04	72.74	75.35	74.02	72.32	71.14	75.11	73.07	66.74	70.20	71.03	70.61
w/ evidence	78.21	77.12	79.80	78.44	75.39	75.22	77.09	76.14	68.02	71.50	69.40	70.43
SR	76.33	75.68	76.32	76.00	73.71	74.29	72.53	73.40	64.96	70.74	64.29	67.36
w/ evidence	78.25	79.20	78.79	78.99	75.08	76.29	76.52	76.40	69.32	72.35	71.40	71.87
SMAD	77.02	76.76	76.27	76.52	74.79	74.42	75.54	74.97	68.53	72.84	70.24	71.52
w/ evidence	78.51	77.73	80.20	78.95	77.24	75.32	79.96	77.57	70.72	73.51	72.33	72.92
D2D	82.17	81.39	82.55	81.97	81.65	80.67	83.26	81.94	74.11	77.20	76.59	76.89
ED2D	83.59	82.64	83.74	83.18	84.45	82.22	84.65	83.41	77.90	80.24	80.56	80.40

Table 1: Performance (%) of all methods on Weibo21, FakeNewsDataset, and Snopes25. Rows with “w/ evidence” indicate the use of external factual evidence as additional input to the LLM. ED2D consistently achieves the best results across all datasets.

Misinformation Detection

We first evaluate ED2D as a misinformation detection system. This section addresses the following question:

- **RQ1:** Does evidence-based method improve the accuracy for misinformation detection?

Experimental Setup

Dataset. We conduct experiments on three datasets: two publicly available benchmarks and one newly constructed resource. The public datasets include Weibo21 (Nan et al. 2021) and FakeNewsDataset (Pérez-Rosas et al. 2018). In addition, we collect Snopes25, a new benchmark compiled from fact-checked real-world claims by professional editors on Snopes, covering the period from January to June 2025. This period is selected to follow the GPT-4o training cutoff to reduce the risk of data leakage from memorized knowledge. Each claim in Snopes25 is annotated as either True or False and paired with an expert-written fact-checking article. Table 2 summarizes the key statistics of the datasets.

Dataset	Fake	Real	Total
Weibo21	2,373	2,461	4,843
FakeNewsDataset	466	466	932
Snopes25	252	196	448

Table 2: Statistics of the three datasets

Baselines. We compare ED2D with following baselines:

- **BERT** (Devlin et al. 2019): A fine-tuned BERT-base model for binary classification.
- **RoBERTa** (Liu et al. 2019): A fine-tuned RoBERTa-base model with the same setup as BERT, serving as a stronger discriminative baseline.
- **Zero-Shot (ZS):** A single LLM directly predicts the veracity of each news item without any intermediate reasoning or task-specific adaptation.

- **Chain-of-Thought (CoT)** (Wei et al. 2022): The LLM is prompted to generate a step-by-step reasoning trace before producing a final prediction.
- **Self-Reflect (SR)** (Madaan et al. 2023): The model iteratively critiques and revises its own outputs until a convergence criterion is met, typically based on self-evaluation of quality or confidence.
- **Standard Multi-Agent Debate (SMAD):** Two LLM agents engage in a four-turn debate, and a judge agent provides a binary prediction based on the dialogue.
- **Debate-to-Detect (D2D)** (Han, Zheng, and Tang 2025): A MAD framework with the same five structured stages. Agents are assigned domain-specific profiles and fixed stances but do not have access to external evidence.

In addition, each LLM-based baseline is further extended with a comparable evidence retrieval module, enabling evaluation of the impact of external factual evidence across prompting strategies.

Implementation. All experiments use GPT-4o as the base model. LLM-agents are initialized with predefined prompts provided in Appendix A. Agent response lengths are capped at 1024 tokens. Domain inference and final judgment are conducted with a temperature of 0.0 to ensure stability. To encourage diversity, profile generation and debate responses across all stages use a temperature of 0.7. The Free Debate stage defaults to a single round, but the number of rounds is configurable for more complex tasks.

Main Results

Table 1 summarizes model performance across the three benchmark datasets using standard metrics—accuracy (Acc), precision (Prec), recall (Rec), and F1-score (F1). ED2D achieves the strongest results on every dataset and metric, substantially outperforming both fine-tuned transformers (BERT, RoBERTa) and prompting-based methods such as CoT, SR, and SMAD. Although several deep learning baselines yield competitive accuracy, their limited interpretability constrains practical deployment. These results

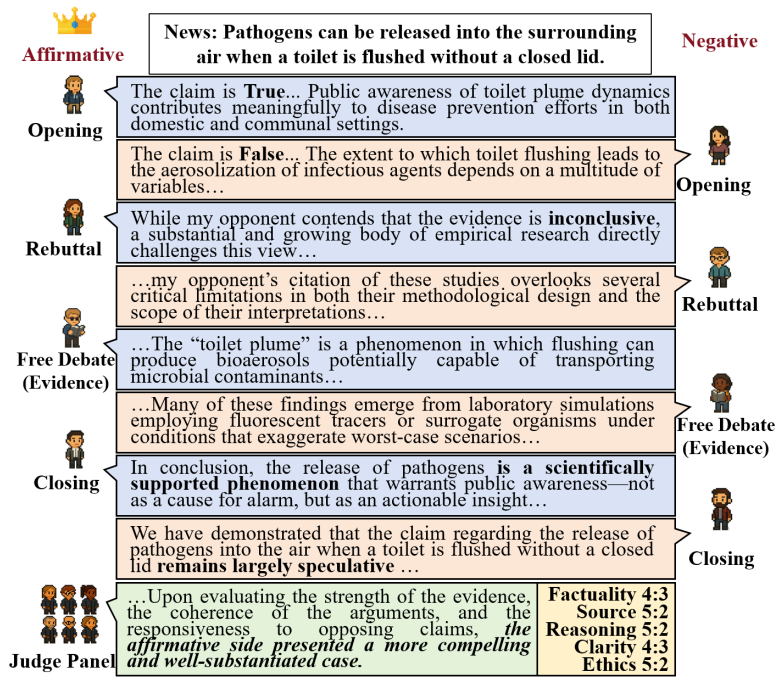


Figure 2: A debate example on the claim that toilet flushing releases airborne pathogens. The case demonstrates the ED2D reasoning process, with the affirmative side prevailing based on evidence-grounded argumentation.

demonstrate that structured debate, when augmented with external factual evidence, produces more accurate and reliable misinformation detection. ED2D’s consistent superiority on Snopes25 further indicates strong generalization to real-world, post-training claims.

Evidence-based grounding improves performance. For each LLM-based method, we evaluate a variant that incorporates retrieved evidence as contextual input. Grounding yields uniform improvements across all methods, with the largest gains observed for simpler prompting approaches such as ZS, highlighting the central role of factual support in LLM reasoning.

From D2D to ED2D. While D2D benefits from structured multi-agent deliberation, ED2D’s integration of explicit evidence retrieval during the free-debate phase produces consistent 2–3 point improvements across all metrics. By combining multi-agent argumentation with evidence-based reasoning, ED2D offers a robust and interpretable architecture well suited to real-world fact-checking and misinformation mitigation workflows.

Case Study

Figure 2 illustrates a representative Snopes25 case concerning whether flushing a toilet with the lid open releases pathogens into the air. The affirmative team anchored its argument in scientific summaries retrieved from Wikipedia, emphasizing the formation of bioaerosols and the relevance of preventative hygiene measures. The negative team challenged the practical significance of these findings, citing limited causal evidence and cautioning against distraction from more established transmission routes. The judges fa-

vored the affirmative team on the basis of stronger evidence synthesis, clearer reasoning, and appropriate reliance on the precautionary principle.

To support real-world use, we additionally developed a publicly accessible ED2D community platform that allows users to generate and inspect structured debates for custom claims, as illustrated in Figure 3. By exposing the full deliberation process, the system aims to improve user resilience to misinformation. A demonstration video is included in the supplementary materials.

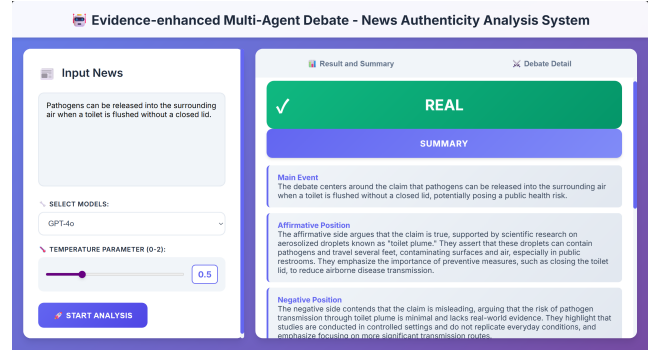


Figure 3: A demonstration of the ED2D community website.

Persuasive Debunking Evaluation

Based on Snopes25, ED2D yields 203 correct and 49 incorrect classifications for false claims, and 146 correct and 50 incorrect for true claims. Building on these results, we eval-

Evaluation	Claim Type	Condition	Accuracy	Belief	Share	Emotion
RQ2	False	Control	63.60%	3.46	3.15	3.66
		ED2D (Correct)	80.40%	2.85	2.84	3.45
		Snopes	85.60%	2.77	2.55	3.15
		Combined	88.00%	2.40	2.53	3.20
	True	Control	67.40%	3.56	3.06	3.40
		ED2D (Correct)	81.20%	4.70	4.32	4.42
		Snopes	88.80%	5.05	4.55	4.77
		Combined	92.40%	5.13	4.59	4.90
RQ3	False	Control	59.60%	3.80	3.09	3.47
		ED2D (Incorrect)	42.00%	4.44	3.55	3.90
		Snopes	81.20%	2.97	2.86	3.55
		Combined	68.00%	3.35	2.95	3.73
	True	Control	68.80%	4.28	3.95	3.42
		ED2D (Incorrect)	52.00%	2.95	3.07	3.15
		Snopes	82.80%	4.44	4.20	4.09
		Combined	75.60%	4.05	4.12	3.59

Table 3: User outcomes following exposure to different explanation conditions, stratified by claim veracity and ED2D judgment accuracy. Higher accuracy and stronger alignment between user responses and claim veracity indicate more effective persuasion.

uate ED2D’s effectiveness in shaping user beliefs under different conditions. Specifically, we investigate the following research questions:

- **RQ2:** How persuasive are ED2D-generated debunks compared to those written by human experts?
- **RQ3:** Can ED2D mistakenly persuade users of false claims when its judgment is incorrect?

Experimental Setup

Study Design. We recruit 200 native English speakers and divide them into two independent cohorts of 100 participants. The first cohort addresses **RQ2**, assessing persuasive effectiveness when ED2D outputs the same result with Snopes. The second cohort addresses **RQ3**, evaluating the risk of misleading persuasion in cases where ED2D misclassifies the claim. Within each cohort, participants are randomly assigned to one of four conditions, with 25 individuals allocated to each condition:

- **Control:** Participants view the claim and judge its truthfulness based solely on their prior knowledge.
- **ED2D:** Participants view the claim along with ED2D’s judgment and explanation, including the full debate transcript and retrieved evidence.
- **Snopes:** Participants view the claim along with corresponding expert-written explanation from Snopes.
- **Combined:** Participants are shown both ED2D and Snopes explanations. For **RQ2**, both sources provide the same correct label, enabling assessment of reinforcement effects. For **RQ3**, the sources provide conflicting labels, allowing assessment of persuasive influence when ED2D is incorrect and potentially misleading.

Each participant evaluates 10 true and 10 false claims. Prior to the task, they are informed that some explanations

may be AI-generated and potentially unreliable, and that performance-based bonuses are awarded for accurate responses. In addition to binary veracity judgments, participants rate each claim on following three subjective dimensions using a 7-point Likert scale:

- **Belief in the claim:** Perceived truthfulness of the claim (1 = certainly false, 7 = certainly true), used as the primary measure of belief change.
- **Willingness to share:** Likelihood of sharing the claim with others (1 = not at all, 7 = very likely), reflecting behavioral diffusion risk.
- **Emotional agreement:** Perceived alignment between the claim and one’s personal values (1 = not at all, 7 = strongly), capturing affective resonance.

Persuasion Results

Table 3 presents group-level means for factual accuracy and user ratings of belief, willingness to share, and emotional alignment on a 7-point Likert scale. We assess effect reliability using mixed-effects regression with random intercepts for participants and claims to account for repeated observations. Binary accuracy is modeled with logistic mixed-effects regression, and Likert-scale outcomes with linear mixed-effects regression. Tukey-adjusted post-hoc comparisons indicate significant differences across conditions (all $p < 0.05$) in both objective accuracy and subjective belief-related measures.

RQ2: Persuasiveness When ED2D Is Correct. When ED2D produces correct labels, its explanations are strongly persuasive and comparable to expert fact-checking. Relative to the no-explanation baseline, ED2D and Snopes both improve participants’ truthfulness judgments for true and false claims, lower belief in misinformation, and strengthen belief

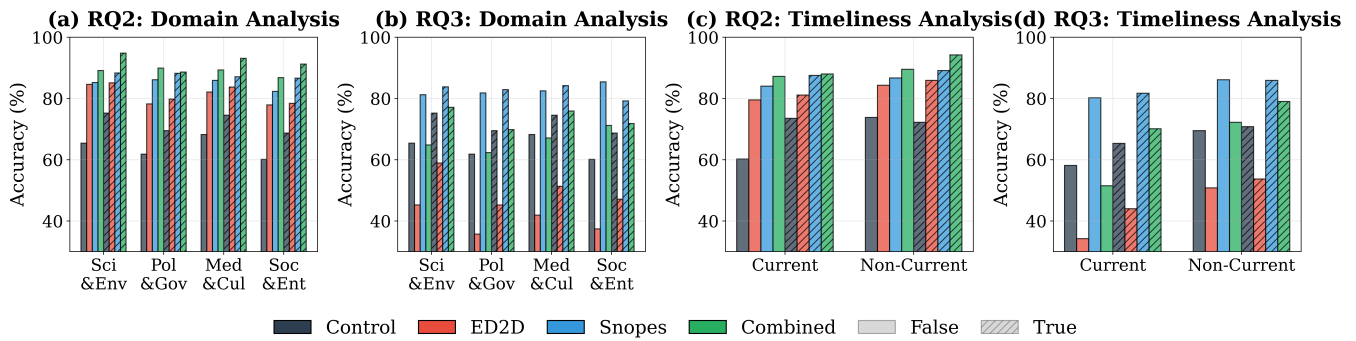


Figure 4: Accuracy comparison by topical domain and claim timeliness. Subfigures (a) and (b) show domain-level accuracy under RQ2 and RQ3 across four domains: Science&Environment, Politics&Government, Medicine&Culture, and Society&Entertainment. Subfigures (c) and (d) analyze accuracy by timeliness, contrasting Current versus Non-Current claims.

in accurate content. These explanations also reduce willingness to share false claims and weaken emotional alignment with them, while increasing alignment and sharing intent for true claims.

RQ3: Misleading Persuasion When ED2D Is Incorrect. When ED2D provides incorrect labels, its explanations systematically distort user judgment. For false claims misclassified as true, ED2D increases belief and sharing intention, whereas Snopes mitigates belief in misinformation and reduces sharing. When both explanations are shown, ED2D’s misleading influence partially counteracts Snopes’ corrective effect. For true claims misclassified as false, ED2D suppresses belief in accurate information and lowers judgment accuracy, while Snopes helps preserve more appropriate evaluations. Overall, these findings show that persuasive but erroneous AI explanations can undermine human fact-checking—even when presented alongside authoritative guidance.

Figure 4 reports participants’ judgment accuracy across claim domains and timeliness conditions. Under **RQ2**, ED2D explanations consistently increase accuracy in every domain. Relative to the Control group, ED2D yields marked gains in politics, science, and culture, achieving accuracy levels close to those of human experts. Comparable improvements appear for both current-event and static knowledge claims, indicating that ED2D’s persuasive effectiveness is largely insensitive to temporal context.

Under **RQ3**, when ED2D produces incorrect labels, its explanations uniformly depress accuracy across domains. The reduction is most pronounced for politicized and entertainment-related content, where participants are more easily influenced by misleading reasoning. Timeliness analysis further shows that current-event claims are especially susceptible to misclassification-driven misinformation effects, whereas static, knowledge-based claims exhibit slightly greater robustness.

Post-exposure Comparison

To evaluate whether exposure to ED2D explanations improves participants’ ability to independently detect misinformation, we conduct a post-test within the **RQ2** condition.

After completing the main task, participants from the Control, ED2D, and Snopes groups assess a new set of ten claims (five true and five false) without any explanatory material, basing their veracity judgments solely on the claim text.

Group	True	False	Overall
Control	60.6	73.2	66.9
ED2D	70.2	81.6	75.9
Snopes	78.4	84.8	78.6

Table 4: Post-test accuracy on new claims under RQ2.

As shown in Table 4, all groups exhibit higher post-test accuracy compared with baseline performance in the original Control condition. Participants previously exposed to ED2D explanations show substantial improvements, especially in detecting false claims, indicating that engagement with evidence-based AI explanations can enhance transferable reasoning skills. Participants in the Snopes condition attain the highest overall accuracy, reaffirming the lasting benefit of expert fact-checks in promoting epistemic vigilance and effective misinformation detection. Taken together, these findings suggest that MAD-style debunking can strengthen users’ ability to evaluate online information even in the absence of explicit explanatory support.

Conclusion

In this paper, we present ED2D, an evidence-based multi-agent debate (MAD) framework for misinformation detection and persuasive debunking. ED2D augments structured debate with an evidence retrieval module that grounds agent arguments in verifiable facts, improving both accuracy and the interpretability of explanations. Across three real-world datasets, ED2D consistently outperforms strong baselines. A controlled user study shows that its explanations are as persuasive as expert fact-checks. However, the results also reveal a dual-use risk: when ED2D is wrong, it can still shift user beliefs in misleading directions. Future work will focus on more efficient scaling, richer reasoning beyond binary labels, and safeguards against deceptive or adversarial use.

Acknowledgements

This work is supported by the National Social Science Fund of China (No. 23 & ZD331).

References

- Ansari, S.; and Alam, M. Z. 2025. From AI Fact-Checks to User Understanding: Explaining Misinformation Detection to Non-Expert Audiences. In *Proceedings of the 4th ACM International Workshop on Multimedia AI against Disinformation*, MAD' 25, 28–36. New York, NY, USA: Association for Computing Machinery. ISBN 9798400718915.
- Cao, B.; Wu, Q.; Cao, J.; Liu, B.; and Gui, J. 2025. External Reliable Information-enhanced Multimodal Contrastive Learning for Fake News Detection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(1): 31–39.
- Danry, V.; Pataranutaporn, P.; Groh, M.; and Epstein, Z. 2025. Deceptive Explanations by Large Language Models Lead People to Change their Beliefs About Misinformation More Often than Honest Explanations. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25. New York, NY, USA: Association for Computing Machinery.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2024. Improving factuality and reasoning in language models through multiagent debate. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org.
- Ecker, U. K.; Lewandowsky, S.; Cook, J.; et al. 2022. The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology*, 1: 13–29. Published 12 January 2022; Accepted 30 September 2021.
- Feng, Y.; Li, W.; Wang, Y.; Wang, J.; Liu, F.; and Han, Z. 2025. Contradicted in Reliable, Replicated in Unreliable: Dual-Source Reference for Fake News Early Detection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(22): 23896–23904.
- Goforth, M. P.; Boone, S. A.; Clark, J.; Valenzuela, P. B.; McKinney, J.; Ijaz, M. K.; and Gerba, C. P. 2024. Impacts of lid closure during toilet flushing and of toilet bowl cleaning on viral contamination of surfaces in United States restrooms. *American Journal of Infection Control*, 52(2): 141–146.
- Han, C.; Li, Y.; and Tang, X. 2025. DocPolicyKG: A Lightweight LLM-Based Framework for Knowledge Graph Construction from Chinese Policy Documents. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, CIKM '25, 4753–4757. New York, NY, USA: Association for Computing Machinery. ISBN 9798400720406.
- Han, C.; and Tang, X. 2026. An Agent-Based Simulation Framework for Misinformation Susceptibility Test with LLMs: Insights from Psychological Factors. In *Knowledge and Systems Sciences*, 355–368. ISBN 978-981-95-4990-0.
- Han, C.; Zheng, W.; and Tang, X. 2025. Debate-to-Detect: Reformulating Misinformation Detection as a Real-World Debate with Large Language Models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 15125–15140. Suzhou, China: Association for Computational Linguistics.
- Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; and Liu, T. 2025. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Trans. Inf. Syst.*, 43(2).
- Jahanbakhsh, F.; Zhang, A. X.; Berinsky, A. J.; Pennycook, G.; Rand, D. G.; and Karger, D. R. 2021. Exploring Lightweight Interventions at Posting Time to Reduce the Sharing of Misinformation on Social Media. *Proc. ACM Hum.-Comput. Interact.*, 5(CSCW1).
- Li, X.; Xia, Y.; Long, X.; Li, Z.; and Li, S. 2021. Exploring Text-Transformers in AAAI 2021 Shared Task: COVID-19 Fake News Detection in English. In Chakraborty, T.; Shu, K.; Bernard, H. R.; Liu, H.; and Akhtar, M. S., eds., *Combating Online Hostile Posts in Regional Languages during Emergency Situation*, 106–115. Cham: Springer International Publishing.
- Li, Y.; Du, Y.; Zhang, J.; Hou, L.; Grabowski, P.; Li, Y.; and Ie, E. 2024. Improving Multi-Agent Debate with Sparse Communication Topology. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Findings of the Association for Computational Linguistics: EMNLP 2024*, 7281–7294. Miami, Florida, USA: Association for Computational Linguistics.
- Liang, T.; He, Z.; Jiao, W.; Wang, X.; Wang, Y.; Wang, R.; Yang, Y.; Shi, S.; and Tu, Z. 2024. Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 17889–17904. Miami, Florida, USA: Association for Computational Linguistics.
- Liu, Y.; Liu, Y.; Zhang, X.; Chen, X.; and Yan, R. 2025. 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*, SIGIR '25, 504–514. New York, NY, USA: Association for Computing Machinery.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. arXiv:1907.11692.
- Lyu, X.; Sun, M.; Cao, C.; Liang, B.; and Xu, R. 2025. Verifying ambiguous claims with a reasoner-translator framework. *Neurocomputing*, 647: 130536.
- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhunoye, S.; Yang,

- Y.; Gupta, S.; Majumder, B. P.; Hermann, K.; Welleck, S.; Yazdanbakhsh, A.; and Clark, P. 2023. Self-Refine: Iterative Refinement with Self-FeedbackRR. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Nan, Q.; Cao, J.; Zhu, Y.; Wang, Y.; and Li, J. 2021. MD-FEND: Multi-domain Fake News Detection. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management, CIKM '21*, 3343–3347. New York, NY, USA: Association for Computing Machinery.
- Pelrine, K.; Imouza, A.; Thibault, C.; Reksoprodjo, M.; Gupta, C.; Christoph, J.; Godbout, J.-F.; and Rabbany, R. 2023. Towards Reliable Misinformation Mitigation: Generalization, Uncertainty, and GPT-4. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 6399–6429. Singapore: Association for Computational Linguistics.
- Pennycook, G.; Epstein, Z.; Mosleh, M.; Arechar, A. A.; Eckles, D.; and Rand, D. G. 2021. Shifting attention to accuracy can reduce misinformation online. *Nature*, 592: 590–595. Published 17 March 2021; Issue Date 22 April 2021; Accepted 8 February 2021.
- Pérez-Rosas, V.; Kleinberg, B.; Lefevre, A.; and Mihalcea, R. 2018. Automatic Detection of Fake News. In Bender, E. M.; Derczynski, L.; and Isabelle, P., eds., *Proceedings of the 27th International Conference on Computational Linguistics*, 3391–3401. Santa Fe, New Mexico, USA: Association for Computational Linguistics.
- Salvi, F.; Horta Ribeiro, M.; Gallotti, R.; Fiorucci, M.; De Domenico, M.; and Lewandowsky, S. 2025. On the conversational persuasiveness of GPT-4. *Nature Human Behaviour*. Published 19 May 2025; Accepted 28 March 2025.
- Scherer, L. D.; McPhetres, J.; Pennycook, G.; Kempe, A.; Allen, L. A.; Knoepke, C. E.; Tate, C. E.; and Matlock, D. D. 2021. Who is susceptible to online health misinformation? A test of four psychosocial hypotheses. *Health Psychology*, 40(4): 274–284. Epub 2021 Mar 1.
- Schmitt, V.; Bezzaoui, I.; Jakob, C.; Sahitaj, P.; Wang, Q.; Hilbert, A.; Upravitelev, M.; Fegert, J.; Möller, S.; and Solopova, V. 2025. Beyond Transparency: Evaluating Explainability in AI-Supported Fact-Checking. In *Proceedings of the 4th ACM International Workshop on Multimedia AI against Disinformation, MAD' 25*, 63–72. New York, NY, USA: Association for Computing Machinery. ISBN 9798400718915.
- Schoenegger, P.; Salvi, F.; Liu, J.; Nan, X.; Debnath, R.; Fasolo, B.; Leivada, E.; Recchia, G.; Günther, F.; Zarifhonarvar, A.; Kwon, J.; Islam, Z. U.; Dehnert, M.; Lee, D. Y. H.; Reinecke, M. G.; Kamper, D. G.; Kobaş, M.; Sandford, A.; Kgombo, J.; Hewitt, L.; Kapoor, S.; Oktar, K.; Kucuk, E. E.; Feng, B.; Jones, C. R.; Gainsburg, I.; Olschewski, S.; Heinzelmann, N.; Cruz, F.; Tappin, B. M.; Ma, T.; Park, P. S.; Onyonka, R.; Hjorth, A.; Slattery, P.; Zeng, Q.; Finke, L.; Grossmann, I.; Salatiello, A.; and Karger, E. 2025. Large Language Models Are More Persuasive Than Incentivized Human Persuaders. arXiv:2505.09662.
- Wang, H.; Prasad, A.; Stengel-Eskin, E.; and Bansal, M. 2025. Retrieval-Augmented Generation with Conflicting Evidence. arXiv:2504.13079.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E. H.; Le, Q. V.; and Zhou, D. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*. Red Hook, NY, USA: Curran Associates Inc.
- Wu, J.; Li, S.; Deng, A.; Xiong, M.; and Hooi, B. 2023. 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, CIKM '23*, 2726–2736. ACM.
- Yu, B.; Huang, W.; and Liu, K. 2025. SR-KI: Scalable and Real-Time Knowledge Integration into LLMs via Supervised Attention. arXiv:2511.06446.
- Zhang, J.; Cai, K.; Fan, Y.; Liu, N.; and Wang, K. 2025a. MAT-Agent: Adaptive Multi-Agent Training Optimization. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Zhang, J.; Huang, Z.; Fan, Y.; Liu, N.; Li, M.; Yang, Z.; Yao, J.; Wang, J.; and Wang, K. 2025b. KABB: Knowledge-Aware Bayesian Bandits for Dynamic Expert Coordination in Multi-Agent Systems. In *Forty-second International Conference on Machine Learning*.
- Zhu, J.; Wang, J.; Yu, B.; Wu, X.; Li, J.; Wang, L.; and Xu, N. 2025. TableEval: A Real-World Benchmark for Complex, Multilingual, and Multi-Structured Table Question Answering. In Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; and Peng, V., eds., *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 7137–7157. Suzhou, China: Association for Computational Linguistics. ISBN 979-8-89176-332-6.