

# MIND Your Reasoning: A Meta-Cognitive Intuitive-Reflective Network for Dual-Reasoning in Multimodal Stance Detection

Bingbing Wang<sup>1,2\*</sup>, Zhengda Jin<sup>1\*</sup>, Bin Liang<sup>3</sup>, Wenjie Li<sup>2</sup>, Jing Li<sup>2†</sup>, Ruifeng Xu<sup>1,4†</sup>, Min Zhang<sup>1</sup>

<sup>1</sup>Harbin Institute of Technology, Shenzhen, China

<sup>2</sup>The Hong Kong Polytechnic University, Hong Kong, China

<sup>3</sup>The Chinese University of Hong Kong, Hong Kong, China

<sup>4</sup>Shenzhen Loop Area Institute, Shenzhen, China

{bingbing.wang,220110515}@stu.hit.edu.cn, jing-amelia.li@polyu.edu.hk, xuruifeng@hit.edu.cn

## Abstract

Multimodal Stance Detection (MSD) is a crucial task for understanding public opinion on social media. Existing methods predominantly operate by learning to fuse modalities. They lack an explicit reasoning process to discern how inter-modal dynamics, such as irony or conflict, collectively shape the user's final stance, leading to frequent misjudgments. To address this, we advocate for a paradigm shift from *learning to fuse* to *learning to reason*. We introduce **MIND**, a **Meta-cognitive Intuitive-reflective Network for Dual-reasoning**. Inspired by the dual-process theory of human cognition, MIND operationalizes a self-improving loop. It first generates a rapid, intuitive hypothesis by querying evolving Modality and Semantic Experience Pools. Subsequently, a meta-cognitive reflective stage uses Modality-CoT and Semantic-CoT to scrutinize this initial judgment, distill superior adaptive strategies, and evolve the experience pools themselves. These dual experience structures are continuously refined during training and recalled at inference to guide robust and context-aware stance decisions. Extensive experiments on the MMSD benchmark demonstrate that our MIND significantly outperforms most baseline models and exhibits strong generalization.

## 1 Introduction

The proliferation of social platforms supporting multimodal content, such as text or images (Dar et al., 2025), has made Multimodal Stance Detection (MSD) a critical task for understanding public opinion (Liang et al., 2024; Niu et al., 2024). MSD aims to identify the expressed attitude toward a specific target by integrating both textual and visual information, enabling a more comprehensive analysis (Liang et al., 2022). This task applies widely to political discourse analysis, misinformation detection, and public health communication, making

\*Equal contribution.

†Corresponding authors.

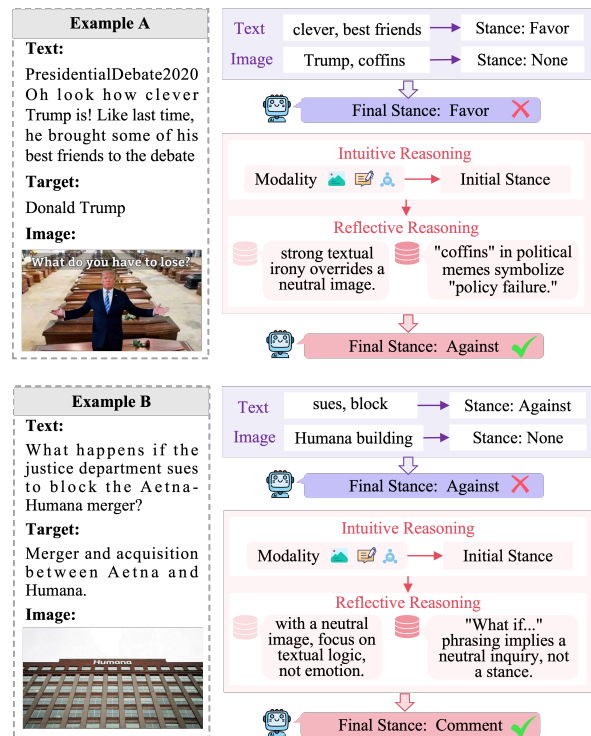


Figure 1: Two examples contrasting existing methods (learning to fuse) in purple with our method (learning to reason) in pink.

it a crucial tool for decoding complex online interactions (Niu et al., 2024; Pangtey et al., 2025). Complementary work studies public opinion and social dynamics through large-scale social network simulation, offering another lens on online behavior (Zhang et al., 2025b, 2026b).

While recent methods have advanced from static fusion to dynamic mechanisms like cross-attention, they remain fundamentally constrained by a paradigm of learning to fuse (Liang et al., 2024; Zhang et al., 2025a). This paradigm excels at weighing feature importance but lacks the explicit, high-level reasoning process needed to interpret inter-modal stance expressions, especially in **semantic traps** where seemingly consistent modalities convey a contradictory stance through nuanced

cues like irony or cultural context. These traps manifest in at least two critical ways. First, models are easily deceived by cross-modal irony, as seen in Figure 1. A fusion-based model would likely interpret the positive text and the image objects as evidence for a FAVOR stance. It fails because it cannot perform the crucial reasoning step: recognizing that the text’s “best friends” sarcastically refers to the image’s “coffins,” a grim symbol requiring world knowledge to interpret as a biting critique. Second, the challenge extends to intra-modal logical fallacies. A text can deceive a model focused on keywords like “sues” and “block” to conclude an AGAINST stance, completely missing that the “What if...” framing makes it a neutral COMMENT. In both scenarios, the failure stems not from an inability to weigh features, but from the absence of a high-level reasoning process to interpret non-literal intent and logical structure.

To bridge this gap, we advocate for a paradigm shift from learning to fuse to learning to reason. We introduce and pioneer **MIND**, a **Meta-cognitive Intuitive-reflective Network for Dual-reasoning** that realizes a more human-like, adaptive reasoning process for a complex multimodal stance detection. Inspired by the dual-process theory, MIND synergizes experience-driven intuitive reasoning with meta-cognitive reflective reasoning. First, an intuitive stage queries the Modality (MEP) and Semantic (SEP) Experience Pools to form an initial hypothesis from past reasoning patterns. This hypothesis is then scrutinized by a meta-cognitive reflective stage, where a Modality Chain of Thought (Modality-CoT) resolves inter-modal dynamics and formulates adaptive fusion strategies, while a complementary Semantic-CoT ensures a robust final decision. Extensive experiments on the public MMSD benchmark demonstrate that MIND significantly outperforms baseline models and exhibits strong generalization, validating the effectiveness of our meta-cognitive approach. Our main contributions are as follows:

- We pioneer a paradigm shift for MSD from learning to fuse to learning to reason. Our dual-reasoning framework is the first to explicitly model a meta-cognitive process for resolving complex stance expressions.
- We operationalize this paradigm with a novel self-improving meta-cognitive loop, where dual Experience Pools store and evolve ex-

PLICIT reasoning strategies that are then refined by a reflective Chain-of-Thought process.

- Experiments conducted on the public dataset demonstrate the superiority of our MIND, which significantly outperforms most baselines and exhibits strong generalization.

## 2 Related Works

### 2.1 LLM-based Stance Detection

The design of modern stance detection frameworks has been significantly influenced by LLMs. Early approaches focused on harnessing their zero-shot and few-shot capabilities, using direct prompting and Chain-of-Thought (CoT) to elicit stance predictions without labeled data or fine-tuning (Zhang et al., 2023, 2022). Other works have pursued fine-tuning by fusing retrieved background knowledge with input pairs to improve learning through back-propagation (Li et al., 2023). Subsequent research has sought to further unleash LLM reasoning. Innovations include integrating multi-perspective insights via CoT and prompt-tuning (Ding et al., 2024), formalizing reasoning with first-order logic rules for enhanced interpretability (Zhang et al., 2024), and architecting multi-agent systems where LLMs collaborate for detection (Lan et al., 2024; Wang et al., 2024). Related efforts on LLM logical reasoning propose semantic-aware logical inference via a semiotic framework (Zhang et al., 2026c) and analyze logical phase transitions to characterize reasoning collapse (Zhang et al., 2026a). However, these methods are primarily developed for pure textual stance detection and have not yet explored the integration of multimodal information such as images.

### 2.2 Multimodal Stance Detection

The proliferation of social media has driven the rise of multimodal content (text, image) in forms like memes and reels (Dar et al., 2025; Bansal et al., 2024), which convey implicit opinions through diverse signals, posing challenges for stance detection. Modeling semantically heterogeneous cross-modal interactions remains difficult for traditional models, whereas transformer architectures and LLMs excel here via self-attention and transfer learning. Recent approaches improved stance classification through multimodal fusion. For example, Niu et al. (2024) proposed MmMtCSD and

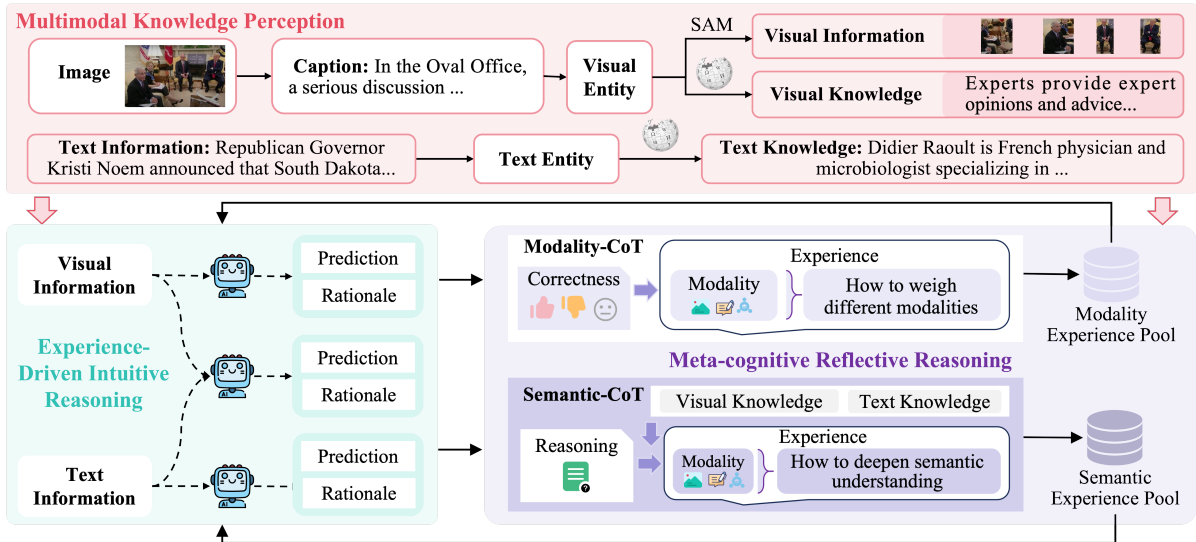


Figure 2: Overview of our MIND framework for multimodal stance detection.

MLLM-SD for text-image pairs in multi-turn dialogues. Kuo et al. (2024) fused textual, visual, and social signals in political discourse using S-BERT, BEiT, and GraphSAGE. Xie et al. (2021) enhanced fake news detection via sentence-guided visual attention. Khiabani and Zubiaga (2023) and Barel et al. (2025) addressed few-shot or noisy settings with graph-based cues. Yao et al. (2023) extended MSD to fact-checking via web evidence integration. However, these approaches typically employ rigid fusion strategies, overlooking the dynamic interplay and context-dependent reliability of modalities, which makes them vulnerable to noisy or contradictory information.

### 3 Method

This section introduces our novel MIND framework for MSD in detail. Drawing inspiration from the dual-process theory of human cognition (Kahneman, 2011), a psychological model that posits two distinct modes of thought. One is an intuitive, fast ‘System 1’ that relies on heuristics and prior experiences to make rapid judgments. The other is a deliberate, slow ‘System 2’ that engages in analytical reasoning to refine and validate these initial assessments. Our MIND framework operationalizes this cognitive paradigm for MSD.

Given a text-image pair  $(u, v)$ , the aim of MIND is to identify the stance label  $y$ , towards a specific target  $t$ . As depicted in Figure 2, our MIND consists of three primary components: 1) **Multimodal Knowledge Perception**, which extracts comprehensive multimodal information and captures rel-

evant external knowledge. 2) **Experience-Driven Intuitive Reasoning** (System 1), which queries the experience pools to form an initial stance hypothesis; and 3) **Meta-cognitive Reflective Reasoning** (System 2), which employs complementary reasoning chains to scrutinize the initial hypothesis and continuously refine the dual experience pools.

#### 3.1 Multimodal Knowledge Perception

This initial stage constructs a rich semantic foundation by extracting and refining knowledge from both textual and visual modalities. The goal is to capture not only high-level conceptual information but also grounded, fine-grained evidence, which is crucial for nuanced reasoning.

##### 3.1.1 Textual Knowledge Attainment

For the text input  $u$ , we perform a three-step knowledge extraction process. First, we employ a Multimodal Large Language Model (MLLM) with strategically designed prompts to identify a set of key semantic entities, denoted as  $\mathcal{E}_u$ . To ground these entities in external world knowledge, we retrieve corresponding background information for each entity  $e \in \mathcal{E}_u$  from the Wikipedia. As the raw retrieved knowledge is often verbose and contains noisy information, we leverage the MLLM again to distill it into a concise, semantically-rich summary. This final summary, denoted as the text knowledge  $\mathcal{K}_u$ , serves as the knowledge for the text modality.

##### 3.1.2 Visual Information Grounding

For the image input  $v$ , we capture both high-level semantic concepts and low-level visual informa-

tion. First, to extract semantic concepts, an MLLM generates an image caption, identifies a set of visual entities  $\mathcal{E}_v$  from the caption. These entities are then used to query Wikipedia, and the resulting information is summarized into a condensed visual knowledge  $\mathcal{K}_v$ . However, knowledge alone is insufficient to capture the visual manifestation of these concepts. To address this, we ground the semantic entities in the image’s pixel space. We leverage the identified entity set  $\mathcal{E}_v$  to guide the open-vocabulary segmentation model SAM (Ravi et al., 2024) in precisely segmenting the image regions corresponding to each entity, yielding a set of entity-aligned visual information  $\mathcal{V}$  where each element is a segmented image directly showcasing the visual appearance of a specific semantic entity.

$$\mathcal{V} = \{\bar{v}_i\}_{i=1}^N = \text{SegModel}(v, \mathcal{E}_v) \quad (1)$$

where  $N$  indicates the number of sub-images.

### 3.2 Experience-Driven Intuitive Reasoning

This stage simulates rapid, intuitive judgment by leveraging experience pools accumulated during training to quickly formulate an initial stance hypothesis. The process consists of two main steps: modal synergistic experience retrieval and experience-driven stance prediction.

#### 3.2.1 Modal Synergistic Experience Retrieval

To retrieve relevant past experiences, we introduce a modal synergistic retrieval mechanism. Unlike conventional text-only systems, our approach leverages both text and visual information to query the modality and semantic experience pool. Each pool stores a collection of experiences  $\{E_j\}$ , where each experience is a key-value pair. The key is a pre-computed bimodal embedding  $(k_u^j, k_v^j)$ , representing the experience’s core features, while the value contains the associated experiential data.

The retrieval process commences by encoding the current text input  $u$  and sub-images  $\mathcal{V}$  into their respective query vectors,  $q_u$  and  $q_v$ :

$$q_u = \text{Enc}(u), \quad q_v = \frac{1}{N} \sum_{i=1}^N \text{Enc}(\bar{v}_i) \quad (2)$$

where  $q_u$  and  $q_v$  are produced by the BGE-VL-CLIP Enc (Zhou et al., 2025). The stored key vectors  $k_u^j$  and  $k_v^j$  for each Experience  $E_j$  are computed analogously as  $q_u$  and  $q_v$ . With the query vectors established, we measure their similarity to the key vectors of each stored experience  $E_j$  by

computing unimodal relevance scores for text  $S_u$  and vision  $S_v$  using cosine similarity:

$$S_u(q_u, k_u^j) = \frac{q_u \cdot k_u^j}{\|q_u\| \|k_u^j\|}, \quad S_v(q_v, k_v^j) = \frac{q_v \cdot k_v^j}{\|q_v\| \|k_v^j\|} \quad (3)$$

To obtain a holistic measure of relevance, we fuse these scores into a unified Modality Relevance Score (MRS), denoted as  $S(E_j)$ , via a weighted linear combination:

$$S(E_j) = \alpha \cdot S_u(q_u, E_j) + (1 - \alpha) \cdot S_v(q_v, E_j) \quad (4)$$

where  $\alpha \in [0, 1]$  balances the weight of each modality for retrieval. Ultimately, experiences are ranked in descending order based on their MRS. The top- $k$  experiences with scores exceeding a relevance threshold  $\tau$  are then selected, yielding the modality experience set  $\{\mathcal{S}_i^{\text{ME}}\}_{i=1}^k$  and the semantic experience set  $\{\mathcal{S}_i^{\text{SE}}\}_{i=1}^k$ .

#### 3.2.2 Experience-driven Stance Prediction

Following the retrieval of relevant past experiences, our MIND agent  $A_R$  proceeds to generate initial, intuitive stance hypotheses. We formulate a comprehensive prompt  $\mathcal{P}$  that integrates two key components: (1) text information  $u$  and visual information  $\mathcal{V}$ ; (2) the retrieved experience sets  $\{\mathcal{S}_i^{\text{ME}}\}_{i=1}^k$  and  $\{\mathcal{S}_i^{\text{SE}}\}_{i=1}^k$ , which serve as in-context exemplars. This prompt instructs  $A_R$  to perform a preliminary analysis from three distinct perspectives: text-only, vision-only, and combined multimodal. For each perspective, the agent generates an initial stance prediction  $\hat{y}$  and a corresponding rationale  $r$ . This process is formally expressed as:

$$(\hat{y}_c, r_c) = A_R(\mathcal{P}(c, \{\mathcal{S}_i^{\text{ME}}\}_{i=1}^k, \{\mathcal{S}_i^{\text{SE}}\}_{i=1}^k)) \quad (5)$$

where  $c \in \{u, \mathcal{V}, (u, \mathcal{V})\}$  denotes the context of evaluation: the text information  $u$ , the visual information  $\mathcal{V}$  or the full multimodal pair  $(u, \mathcal{V})$ . These generated pairs  $(\hat{y}_c, r_c)$  serve as essential intermediate reasoning traces that are scrutinized and refined during the subsequent reflective reasoning stage.

#### 3.3 Meta-cognitive Reflective Reasoning

Meta-cognitive Reflective Reasoning is analogous to the ‘System 2’ slow thinking process in dual-process cognitive theory and serves as the refinement module of our framework. This process is governed by experience formulation, where novel insights are generated through structured reflection, and experience evolution, where the experience pools are refined by integrating these new insights.

### 3.3.1 Experiences Formulation

The primary objective of experience formulation is to facilitate meta-cognitive reflection, enabling the agent to distill actionable experience from its performance outcomes. To this end, the MIND agent employs a structured, multi-step reasoning process, analogous to a CoT, to generate new insights.

To generate the modality experience, the Modality-CoT initiates a self-diagnostic process. It evaluates the accuracy of the initial predictions for text  $\hat{y}_u$ , vision  $\hat{y}_v$ , and combined multimodal input  $\hat{y}_{(u,v)}$ , along with their corresponding rationales  $r_u$ ,  $r_v$ , and  $r_{(u,v)}$ . By analyzing the failures and successes of these initial judgments, it distills a high-level *modality insight*  $\mathcal{I}_m$ . This insight is not merely a summary, but an adaptive fusion strategy that dictates how to dynamically balance unimodal and multimodal cues in future reasoning.

In parallel, the Semantic-CoT is tasked with forming the **semantic experience**. This process involves a deep interrogation of the initial reasoning’s semantic grounding. It synthesizes the agent’s internal rationales with the externally-grounded textual knowledge  $\mathcal{K}_u$  and visual knowledge  $\mathcal{K}_v$ . This synthesis allows the agent to identify and rectify content-level biases or shallow interpretations. The outcome is a semantic insight  $\mathcal{I}_s$ , an actionable directive that refines the model’s strategies for achieving deeper semantic and contextual understanding.

### 3.3.2 Experiences Evolution

Upon deriving new insights  $\mathcal{I}_m$  and  $\mathcal{I}_s$ , MIND updates its experience pools to consolidate and refine existing experiences by first employing the modal synergistic retrieval mechanism to retrieve the top- $k$  most relevant prior experiences from each pool. These experiences, denoted as  $\{\bar{\mathcal{S}}_i^{\text{ME}}\}_{i=1}^k$  from MEP and  $\{\bar{\mathcal{S}}_i^{\text{SE}}\}_{i=1}^k$  from SEP, have scores exceeding a relevance threshold  $\tau$  and are selected based on the current text information  $u$  and visual information  $v$ . If no experiences meet this threshold, the new insights are encapsulated as new key-value entries and appended to their respective pools. Conversely, MIND performs targeted updates on these entries to prevent redundant memories.

This evolutionary update hinges on a fusion operation. For each retrieved modality experience, MIND initiates a re-reasoning process, using a prompt to synthesize the prior experience  $\{\bar{\mathcal{S}}_i^{\text{ME}}\}_{i=1}^k$  with the new modality insight  $\mathcal{I}_m$ . This generates a more comprehensive, updated modality experience pool  $\hat{\mathcal{S}}^{\text{ME}}$ , which then overwrites the

original entry. A parallel process is executed for the semantic experiences, updating  $\{\bar{\mathcal{S}}_i^{\text{SE}}\}_{i=1}^k$  with  $\mathcal{I}_s$  to produce new semantic experience pool  $\hat{\mathcal{S}}^{\text{SE}}$ . This dual-update ensures both experience pools remain synchronized and continuously refined. In essence, this mechanism enables MIND to build a compact and coherent knowledge base through continuous refinement, supporting efficient learning.

### 3.3.3 Stance Inference

During the inference phase, MIND first leverages the SAM to extract visual information  $v$  from the input image, capturing fine-grained visual manifestations of semantic entities. This visual information is then paired with the text input  $u$  to form a multimodal context. It then employs its modality-synergistic retrieval mechanism to fetch the relevant prior experiences from both the modality and semantic experience pools. Finally, MIND integrates these experiences with multimodal context and the target to derive the final stance prediction, ensuring decisions are based on immediate data and guided by its refined long-term understanding.

## 4 Experiments

### 4.1 Experiment Settings

**Dataset and Metrics.** To evaluate our MIND, we benchmark its performance on the MMSD dataset (Liang et al., 2024), which comprises five distinct domains: Multi-modal COVID-CQ (MCCQ), Multi-modal Russo-Ukrainian Conflict (MRUC), Multi-modal Taiwan Question (MTWQ), Multi-modal Twitter Stance Election 2020 (MTSE), and Multi-modal Will-They-Won’t-They (MWTWT). Following the MMSD dataset, we adopt the Macro-F1 score as the primary metric for evaluation.

**Implementation Details.** Our MIND method is implemented upon the Qwen-2.5-VL-32B model as its backbone. For the experience pool retrieval, we set the coefficient  $\alpha$  to 0.7, a relevance threshold  $\tau$  to 0.7, and retrieve the top  $k = 3$  experiences to balance context with noise. Entity-related knowledge is sourced from Wikipedia using the Python Wikipedia library.

### 4.2 Baseline Methods.

We evaluate our method against three categories of baselines: **Text-only baselines:** 1) BERT (Devlin et al., 2019), using uncased BERT-base model; 2) RoBERTa (Liu et al., 2019), using RoBERTa-base model; 3) KEBERT (Kawintiranon and Singh,

MODALITY	METHOD	MTSE		MCCQ	MWTWT				MRUC		MTWQ		
		DT	JB	CQ	CA	CE	AC	AH	DF	RUS	UKR	MOC	TOC
Text-only	BERT	48.25	52.04	66.57	75.62	60.85	63.05	59.24	81.53	41.25	46.80	57.77	45.91
	RoBERTa	58.39	60.79	66.57	69.56	65.03	69.74	67.99	79.21	39.52	57.66	55.22	48.88
	KEBERT	64.50	69.81	66.84	71.67	67.56	69.29	69.74	80.57	41.55	59.01	58.15	47.75
	LLaMA2	53.23	52.67	47.40	34.89	41.95	49.09	44.32	30.21	38.84	38.54	55.31	46.51
	GPT4	68.74	66.39	65.84	63.14	65.12	<u>69.93</u>	<b>71.62</b>	52.69	41.64	53.76	58.05	49.81
Vision-only	ResNet	37.89	38.59	47.16	39.89	42.20	43.52	37.05	50.34	35.10	40.00	42.02	33.94
	ViT	40.48	40.42	46.64	46.63	50.00	40.16	46.32	50.86	33.31	39.87	38.63	35.53
	SwinT	39.89	40.43	48.80	46.30	46.99	41.02	47.39	51.32	35.01	40.89	35.03	35.47
Multimodal	BERT+ViT	41.86	45.82	61.32	63.20	44.71	56.45	46.85	73.71	39.28	48.41	47.47	40.86
	ViLT	35.32	48.24	47.85	62.70	56.44	58.06	60.22	73.66	34.62	42.41	44.43	59.51
	CLIP	53.22	65.83	63.65	70.93	67.17	67.43	70.86	79.06	44.99	59.86	55.29	40.98
	TMPT	55.41	61.61	67.67	<u>76.60</u>	63.19	67.25	62.92	81.19	43.56	59.24	55.68	46.82
	TMPT+CoT	66.61	68.75	71.79	74.40	<u>69.96</u>	68.43	63.00	<b>82.71</b>	<u>45.04</u>	<u>60.52</u>	<u>68.95</u>	<u>59.87</u>
	Qwen-VL	43.31	45.13	50.51	43.06	45.49	49.79	46.04	27.73	36.50	40.78	42.14	39.34
	MiMo	66.21	65.89	61.54	45.32	42.19	50.69	42.36	40.26	39.30	58.95	60.07	52.39
	LLaVa	49.48	45.80	60.30	20.48	28.48	33.83	38.94	31.57	36.53	47.47	50.45	47.07
	Qwen2.5-VL	68.07	70.82	63.33	75.27	68.44	67.15	68.91	63.21	41.04	51.34	65.62	56.35
GPT4-Vision	<u>70.46</u>	<u>72.82</u>	61.63	44.59	47.07	57.47	57.90	37.61	44.83	56.40	66.72	56.90	
Backbones	<b>MIND</b>	<b>70.87</b> ↑	<b>73.31</b> ↑	<b>72.90</b> *↑	<b>78.49</b> ↑	<b>71.85</b> ↑	<b>71.89</b> *↑	<u>70.62</u> ↑	65.54↑	<b>50.81</b> *↑	<b>62.42</b> *↑	<b>69.38</b> ↑	<b>61.13</b> ↑
	- w/ Qwen-VL	65.14↑	63.84↑	63.51↑	55.76↑	53.46↑	50.67↑	52.23↑	50.06↑	39.96↑	49.30↑	57.36↑	52.75↑
	- w/ MiMo	66.43↑	69.91↑	70.84↑	59.25↑	53.34↑	61.26↑	56.35↑	58.18↑	42.52↑	60.22↑	61.29↑	54.82↑
	- w/ LLaVa	58.26↑	51.61↑	62.38↑	24.86↑	28.88↑	40.67↑	43.14↑	35.77↑	36.79↑	50.15↑	51.81↑	48.04↑

Table 1: In-target results (%). Best and second-best scores are in **bold** and underline, respectively. \* indicates our MIND model significantly outperforms baselines ( $p < 0.05$ ). The dashed line separates fine-tuned from non-fine-tuned methods. w/ denotes the backbone, and ↑ indicates improvement over the original MLLM.

2022); 4) LLaMA2 (Touvron et al., 2023), using LLaMA2-70B-chat; 5) GPT4. **Vision-only baselines:** 1) ResNet (He et al., 2016), using ResNet-50 v1.5; 2) ViT (Dosovitskiy et al., 2020), using ViT-base-patch16-224 model; 3) Swin Transformer (SwinT) (Liu et al., 2021), using SwinV2-base-patch4-window12-192-22k model. **Multimodal baselines:** 1) ViLT (Kim et al., 2021), using vilt-b32-mlm model; 2) CLIP (Radford et al., 2021), using CLIP-ViT-base-patch32 model; 3) BERT+ViT, a model that uses BERT and ViT as the text and image encoders; 4) Qwen-VL (Bai et al., 2023), using the Qwen-VL-Chat7B. 5) Qwen2.5-VL (Bai et al., 2025), using the Qwen2.5-VL-32B-Instruct; 6) GPT4-Vision ; 7) MiMo (Xiaomi et al., 2025), using MiMo-7B-RL, 8) LLaVa (Li et al., 2024), using LLaVa-v1.6-Mistral-7B; 9) TMPT and TMPT+CoT (Liang et al., 2024).

### 4.3 Main Results

#### 4.3.1 In-target Multimodal Stance Detection.

As presented in Table 1, our MIND achieves competitive performance across most baselines, even powerful LLMs such as GPT4 and GPT4-Vision in multiple scenarios. While TMPT+CoT shows stronger results on some targets, it relies on training with GPT4-generated CoT examples. Instead, MIND highlights its superiority by leveraging ac-

cumulated experience for adaptive modality contribution rethinking.

A deeper analysis of baseline performance reveals the intricate challenges in multimodal fusion. The consistent outperformance of text-only models over vision-only counterparts confirms that text serves as the primary modality for stance determination. Yet, the superior performance of advanced multimodal models like GPT4-Vision compared to their text-only variants validates the supplementary value of visual cues, underscoring the necessity of cross-modality integration. Notably, the inconsistent performance gains across fusion approaches, combined with the struggles of simpler frameworks like CLIP and TMPT, underscore a critical deficiency, as these models fail to handle the varying contribution of each modality on a per-sample basis, allowing noisy or irrelevant visual data to degrade overall performance. In contrast, MIND leverages its dual-reasoning paradigm to dynamically rethink modality contributions, thereby enhancing its final predictive accuracy.

#### 4.3.2 Zero-shot Multimodal Stance Detection.

As shown in Table 2, the results expose the inherent fragility of models reliant on in-domain training. While smaller architectures like BERT and RoBERTa suffer a precipitous performance

MODALITY	METHOD	MTSE		MWTWT				MRUC		MTWQ	
		DT	JB	CA	CE	AC	AH	RUS	UKR	MOC	TOC
Text-only	BERT	32.52	29.97	63.55	61.30	59.18	52.89	22.01	15.45	28.04	9.57
	RoBERTa	26.60	32.21	59.22	59.22	64.86	57.46	27.10	19.98	30.62	15.84
	KEBERT	26.17	31.81	59.70	62.56	63.92	55.53	24.68	28.18	29.17	19.80
	LLaMA2	53.57	53.92	32.47	38.37	48.08	46.13	31.86	36.34	51.46	44.10
	GPT4	70.78	68.83	57.19	60.56	65.63	<b>69.01</b>	40.22	49.18	62.10	52.12
Vision-only	ResNet	25.52	29.70	23.01	24.11	25.21	25.27	23.88	25.57	27.59	24.88
	ViT	28.63	29.70	24.59	28.18	34.06	33.40	27.26	28.51	29.37	23.69
	SwinT	28.54	30.85	28.53	28.50	35.87	34.33	25.44	24.54	27.90	19.69
Multimodal	BERT+ViT	26.70	31.57	59.21	59.30	65.04	59.28	23.33	15.21	24.76	11.70
	ViLT	28.08	29.74	38.33	46.00	55.01	48.55	21.56	23.96	23.54	19.18
	CLIP	28.21	28.99	61.08	55.67	63.80	60.06	25.62	27.40	27.21	15.69
	TMPT	31.69	32.65	66.36	66.39	66.32	61.56	23.87	24.71	32.18	26.48
	TMPT+CoT	54.30	58.46	<u>67.28</u>	63.73	64.87	54.26	<u>48.99</u>	<u>51.75</u>	45.32	43.70
	Qwen-VL	47.62	46.14	<u>38.57</u>	43.36	47.82	41.01	<u>36.95</u>	41.39	44.32	44.08
	MiMo	66.58	66.01	38.11	40.10	44.20	39.05	44.85	50.24	64.65	46.25
	LLaVa	50.26	46.74	26.53	30.77	34.96	36.57	40.67	38.69	48.97	39.80
	Qwen2.5-VL	68.11	67.02	62.59	60.21	62.52	65.19	48.22	45.82	<u>67.08</u>	51.23
	GPT4-Vision	<u>72.68</u>	<u>71.28</u>	42.23	45.92	54.59	53.19	42.09	47.00	65.00	<u>52.36</u>
Backbones	<b>MIND</b>	<b>73.78*</b> ↑	<b>72.69*</b> ↑	<b>70.83</b> ↑	<b>66.90</b> ↑	<b>71.82*</b> ↑	67.25↑	<b>50.07*</b> ↑	<b>52.57*</b> ↑	<b>69.77</b> ↑	<b>64.39*</b> ↑
	- w/ QwenVL	59.48↑	56.47↑	59.71↑	52.10↑	58.09↑	50.94↑	44.12↑	42.42↑	60.16↑	45.79↑
	- w/ MiMo	67.44↑	66.09↑	44.70↑	40.12↑	45.59↑	41.27↑	46.62↑	52.23↑	65.57↑	46.47↑
	- w/ LLaVa	53.10↑	49.52↑	27.21↑	30.42	44.13↑	38.25↑	42.43↑	40.64↑	50.63↑	45.76↑

Table 2: Zero-shot results (%). Best and second-best scores are in **bold** and underline, respectively. \* indicates our MIND model significantly outperforms baselines ( $p < 0.05$ ). The dashed line separates fine-tuned from non-fine-tuned methods. w/ denotes the backbone, and ↑ indicates improvement over the original MLLM.

METHOD	MTSE	MCCQ	MWTWT	MRUC	MTWQ
<i>in-target setting</i>					
MIND	<b>70.88</b>	<b>72.90</b>	<b>72.18</b>	<b>71.41</b>	<b>69.76</b>
-w/o MEP	69.04	68.38	69.42	70.33	68.45
-w/o SEP	70.65	69.56	67.48	67.05	69.11
-w/o CoT	70.78	67.92	69.13	65.45	69.17
<i>zero-shot setting</i>					
MIND	<b>71.19</b>	-	<b>70.59</b>	<b>70.74</b>	<b>71.39</b>
-w/o MEP	66.95	-	63.26	65.71	68.55
-w/o SEP	64.47	-	64.35	68.74	70.41
-w/o CoT	70.85	-	64.77	69.11	71.22

Table 3: Experimental results (%) of the ablation study on in-target and zero-shot multimodal stance detection.

collapse, even powerful LLMs such as GPT4 and GPT4-Vision demonstrate significant performance volatility when confronted with unseen targets. This underscores a fundamental limitation in current paradigms: true zero-shot generalization remains a formidable and largely unsolved challenge. In this demanding scenario, MIND establishes a new standard for robustness and adaptability. It achieves state-of-the-art results on a majority of the evaluated datasets. This stability, particularly when contrasted with the erratic outputs of leading LLM-based baselines, provides compelling evidence of a superior generalization capability.

#### 4.4 Ablation Study

To evaluate the impact of each component within our MIND, we conduct an ablation study, with results reported in Table 3. The findings unequivocally demonstrate the synergistic importance of all modules, as the full MIND model consistently outperforms any ablated version across both in-target and zero-shot settings. The most critical components are the experience pools. Without using modality experience (w/o MEP) results in the most significant performance degradation, particularly in the zero-shot scenario, highlighting the indispensable role of assessing cross-sample modality contribution for robust generalization. Similarly, the exclusion of using semantic experience (w/o SEP) causes a severe decline, confirming the necessity of leveraging semantic context for accurate stance inference. Finally, while removing the CoT (w/o CoT) module also impairs performance, the impact is less pronounced, suggesting that CoT functions as a higher-level reasoning synthesizer, which organizes and refines the foundational evidence provided by the MEP and SEP to formulate a coherent final stance.

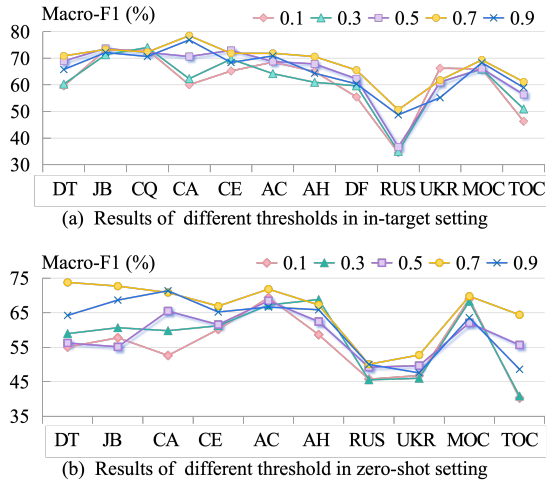


Figure 3: Results on different relevance thresholds.

#### 4.5 Effect of Different Backbones

We assess the performance of MIND with different MLLM backbones in both in-target and zero-shot settings, as detailed in Table 1 and Table 2. Smaller backbones generally yield weaker performance, likely due to their limited visual perception capabilities. However, integrating any MLLM as MIND’s backbone still brings significant improvements marked by  $\uparrow$  compared to standalone operation, with Qwen-VL showing particularly notable gains. This consistent enhancement across diverse architectures highlights that our framework effectively handles varying modality contributions by dynamically weighting them based on their actual expressive power for target stances. It thus compensates for limitations in smaller backbones, ensuring robust performance regardless of the backbone’s inherent biases toward text or vision.

#### 4.6 Different Relevance Thresholds

We analyzed the effect of varying relevance thresholds (0.1, 0.3, 0.5, 0.7, and 0.9) on retrieval performance for both in-target and zero-shot settings, as shown in Figure 3. The results indicate that in both settings, performance tends to be weaker at lower thresholds and gradually degrades once the threshold exceeds 0.7. This trend highlights a critical trade-off between the quality and quantity of retrieved experiences. A low threshold risks retrieving irrelevant or poor-quality experiences, which can introduce noise and impair the reasoning process. Conversely, an excessively high threshold becomes too restrictive, retrieving an insufficient number of experiences and thus limiting the model’s ability to leverage its accumulated knowl-

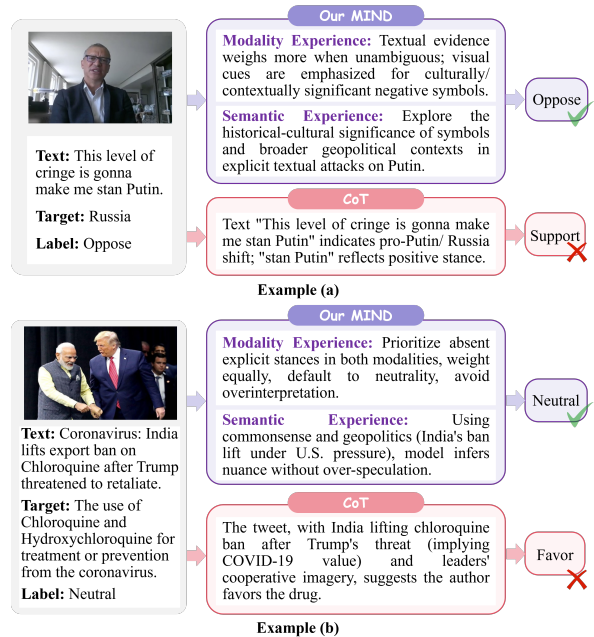


Figure 4: Case study between our MIND and directly CoT method based on the same LLM backbone.

edge for robust inference.

#### 4.7 Case Study

To further underscore the strengths of MIND’s dual reasoning, we analyze two examples in Figure 4 by comparing the results of direct CoT with MIND. In Example (a), direct CoT literally interprets ‘stan Putin’ in the text as support but ignores the negative context of ‘cringe’, which renders ‘stan’ sarcastic. MIND’s modality experience emphasizes weighting textual evidence via contextual judgment, enabling recognition of implicit sarcasm through the contradiction between *cringe* and *stan* to avoid being misled by isolated explicit terms. Its semantic experience clarifies that such sarcasm criticizes Russia-related phenomena, thus linking the text to the target stance of opposing Russia.

For Example (b), direct CoT infers support for the drug from surface information like India lifting the export ban in text and leaders’ collaboration in images. In contrast, MIND’s modality experience stresses prioritizing explicit stance information across modalities, equal weighting, and default neutrality, while it deepens understanding via common sense and geopolitical context, revealing that the text describes objective policy adjustments amid international games rather than the author’s subjective endorsement of the drug’s efficacy. By curbing surface information misinterpretation and calibrating reasoning with background knowledge,

MIND accurately captures the neutral stance.

## 5 Conclusion

This paper proposes MIND, a framework for multimodal stance detection that employs a dual-reasoning paradigm inspired by human cognition. By integrating experience-driven intuitive reasoning with deliberate reflective reasoning, it dynamically captures the variable contributions of different modalities, overcoming the limitations of fixed fusion strategies. The dual-pool mechanism enables the accumulation and refinement of modality and semantic insights, enhancing robustness and context awareness. Experiments on MMSD show MIND outperforms most baselines in both in-target and zero-shot settings, with strong generalization.

## Limitations

A primary limitation of this work is its computational cost and inference latency. To emulate the complex dual-reasoning process of human cognition, our MIND framework involves multiple calls to the large multimodal model at various stages, including knowledge perception, intuitive reasoning, and reflective reasoning. While this deliberate design is crucial for achieving dynamic and adaptive reasoning, it inevitably leads to significant computational overhead, making it less suitable for real-time applications. Future work could explore techniques such as knowledge distillation or model compression to create a more lightweight version of the model while preserving its core reasoning capabilities.

## Acknowledgments

This work was partially supported by the National Natural Science Foundation of China 62576120, the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. T41-517/25-N and PolyU/25200821), and the Innovation and Technology Fund (Project No. PRP/047/22FX).

## References

- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 1(2):3.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie

Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.

- Shubhi Bansal, Kushaan Gowda, Mohammad Zia Ur Rehman, Chandravardhan Singh Raghaw, and Nagendra Kumar. 2024. A hybrid filtering for micro-video hashtag recommendation using graph-based deep neural network. *Engineering Applications of Artificial Intelligence*, 138:109417.
- Guy Barel, Oren Tsur, and Dan Vilenchik. 2025. Acquired taste: Multimodal stance detection with textual and structural embeddings. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6492–6504.
- Shahid Shafi Dar, Mohammad Zia Ur Rehman, Karan Bais, Mohammed Abdul Haseeb, and Nagendra Kumar. 2025. A social context-aware graph-based multimodal attentive learning framework for disaster content classification during emergencies. *Expert Systems with Applications*, 259:125337.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Daijun Ding, Rong Chen, Liwen Jing, Bowen Zhang, Xu Huang, Li Dong, Xiaowen Zhao, and Ge Song. 2024. Cross-target stance detection by exploiting target analytical perspectives. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 10651–10655. IEEE.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, and 1 others. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Daniel Kahneman. 2011. Fast and slow thinking. *Allen Lane and Penguin Books, New York*.
- Kornraphop Kawintiranon and Lisa Singh. 2022. Polibertweet: a pre-trained language model for analyzing political content on twitter. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 7360–7367.
- Parisa Jamadi Khiabani and Arkaitz Zubiaga. 2023. Few-shot learning for cross-target stance detection by aggregating multimodal embeddings. *IEEE Transactions on Computational Social Systems*, 11(2):2081–2090.

- Wonjae Kim, Bokyoung Son, and Ildoo Kim. 2021. Vilt: Vision-and-language transformer without convolution or region supervision. In *International conference on machine learning*, pages 5583–5594. PMLR.
- Kuan-Hung Kuo, Ming-Hung Wang, Hung-Yu Kao, and Yu-Chen Dai. 2024. Advancing stance detection of political fan pages: A multimodal approach. In *Companion Proceedings of the ACM Web Conference 2024*, pages 702–705.
- Xiaochong Lan, Chen Gao, Depeng Jin, and Yong Li. 2024. Stance detection with collaborative role-infused llm-based agents. In *Proceedings of the international AAAI conference on web and social media*, volume 18, pages 891–903.
- Ang Li, Bin Liang, Jingqian Zhao, Bowen Zhang, Min Yang, and Ruifeng Xu. 2023. Stance detection on social media with background knowledge. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, pages 15703–15717.
- Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. 2024. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*.
- Bin Liang, Zixiao Chen, Lin Gui, Yulan He, Min Yang, and Ruifeng Xu. 2022. Zero-shot stance detection via contrastive learning. In *Proceedings of the ACM web conference 2022*, pages 2738–2747.
- Bin Liang, Ang Li, Jingqian Zhao, Lin Gui, Min Yang, Yue Yu, Kam-Fai Wong, and Ruifeng Xu. 2024. Multi-modal stance detection: New datasets and model. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 12373–12387.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022.
- Fuqiang Niu, Zebang Cheng, Xianghua Fu, Xiaojiang Peng, Genan Dai, Yin Chen, Hu Huang, and Bowen Zhang. 2024. Multimodal multi-turn conversation stance detection: A challenge dataset and effective model. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 3867–3876.
- Lata Pangtey, Anukriti Bhatnagar, Shubhi Bansal, Shahid Shafi Dar, and Nagendra Kumar. 2025. Large language models meet stance detection: A survey of tasks, methods, applications, challenges and future directions. *arXiv preprint arXiv:2505.08464*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR.
- Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, and 1 others. 2024. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, and 1 others. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Xiaolong Wang, Yile Wang, Sijie Cheng, Peng Li, and Yang Liu. 2024. Deem: Dynamic experienced expert modeling for stance detection. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4530–4541.
- LLM Xiaomi, Bingquan Xia, Bowen Shen, Dawei Zhu, Di Zhang, Gang Wang, Hailin Zhang, Huaqiu Liu, Jiebao Xiao, Jinhao Dong, and 1 others. 2025. Mimo: Unlocking the reasoning potential of language model—from pretraining to posttraining. *arXiv preprint arXiv:2505.07608*.
- Jianhui Xie, Song Liu, Ruixin Liu, Yinghong Zhang, and Yuesheng Zhu. 2021. Sern: Stance extraction and reasoning network for fake news detection. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2520–2524. IEEE.
- Barry Menglong Yao, Aditya Shah, Lichao Sun, Jin-Hee Cho, and Lifu Huang. 2023. End-to-end multimodal fact-checking and explanation generation: A challenging dataset and models. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2733–2743.
- Bowen Zhang, Daijun Ding, Zhichao Huang, Ang Li, Yangyang Li, Baoquan Zhang, and Hu Huang. 2024. Knowledge-augmented interpretable network for zero-shot stance detection on social media. *IEEE Transactions on Computational Social Systems*.
- Bowen Zhang, Daijun Ding, Liwen Jing, Genan Dai, and Nan Yin. 2022. How would stance detection techniques evolve after the launch of chatgpt? *arXiv preprint arXiv:2212.14548*.
- Bowen Zhang, Xianghua Fu, Daijun Ding, Hu Huang, Genan Dai, Nan Yin, Yangyang Li, and Liwen Jing. 2023. Investigating chain-of-thought with chatgpt for stance detection on social media. *arXiv preprint arXiv:2304.03087*.

- Bowen Zhang, Jun Ma, Xianghua Fu, and Genan Dai. 2025a. Logic augmented multi-decision fusion framework for stance detection on social media. *Information Fusion*, page 103214.
- Xinglang Zhang, Yunyao Zhang, ZeLiang Chen, Junqing Yu, Wei Yang, and Zikai Song. 2026a. [Logical phase transitions: Understanding collapse in llm logical reasoning](#). *Preprint*, arXiv:2601.02902.
- Yunyao Zhang, Yihao Ai, Zuocheng Ying, Qirui Mi, Junqing Yu, Wei Yang, and Zikai Song. 2026b. [Coupling macro dynamics and micro states for long-horizon social simulation](#). *Preprint*, arXiv:2604.05516.
- Yunyao Zhang, Zikai Song, Hang Zhou, Wenfeng Ren, Yi-Ping Phoebe Chen, Junqing Yu, and Wei Yang. 2025b. [ga - s<sup>3</sup>: Comprehensive social network simulation with group agents](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 8950–8970, Vienna, Austria. Association for Computational Linguistics.
- Yunyao Zhang, Xinglang Zhang, Junxi Sheng, Wenbing Li, Junqing Yu, Yi-Ping Phoebe Chen, Wei Yang, and Zikai Song. 2026c. [Semantic-aware logical reasoning via a semiotic framework](#). *Preprint*, arXiv:2509.24765.
- Junjie Zhou, Yongping Xiong, Zheng Liu, Ze Liu, Shitao Xiao, Yueze Wang, Bo Zhao, Chen Jason Zhang, and Defu Lian. 2025. Megapairs: Massive data synthesis for universal multimodal retrieval. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 19076–19095.