

SEVADE: Self-Evolving Multi-Agent Analysis with Decoupled Evaluation for Hallucination-Resistant Sarcasm Detection

Ziqi Liu^{1*}, Ziyang Zhou^{1*}, Yilin Li¹, Mingxuan Hu¹, Yushan Pan¹, Zhijie Xu¹, Yangbin Chen^{1†}

¹School of Advanced Technology, Xi'an Jiaotong-Liverpool University
 {Ziqi.Liu, Ziyang.Zhou, Mingxuan.Hu}22@student.xjtlu.edu.cn, Yilin.Li2202@student.xjtlu.edu.cn,
 {Yushan.Pan, Zhijie.Xu, Yangbin.Chen}@xjtlu.edu.cn

Abstract

Sarcasm detection is a crucial yet challenging Natural Language Processing task. Existing Large Language Model methods are often limited by single-perspective analysis, static reasoning pathways, and a susceptibility to hallucination when processing complex ironic rhetoric, which impacts their accuracy and reliability. To address these challenges, we propose **SEVADE**, a novel Self-Evolving multi-agent Analysis framework with Decoupled Evaluation for hallucination-resistant sarcasm detection. The core of our framework is a Dynamic Agentive Reasoning Engine (DARE), which utilizes a team of specialized agents grounded in linguistic theory to perform a multifaceted deconstruction of the text and generate a structured reasoning chain. Subsequently, a separate lightweight rationale adjudicator (RA) performs the final classification based solely on this reasoning chain. This decoupled architecture is designed to mitigate the risk of hallucination by separating complex reasoning from the final judgment. Extensive experiments on four benchmark datasets demonstrate that our framework achieves state-of-the-art performance, with average improvements of **7.01%** in Accuracy and **6.55%** in Macro-F1 score.

Code — <https://github.com/sunbus100/SEVADE>

1 Introduction

Sarcasm is defined as a rhetorical device or mode of expression, where the intended meaning of a statement is often the opposite of its literal meaning, frequently used to achieve a humorous or sarcastic effect (Booth 1974). Sarcasm detection aims to identify ironic expressions in text, which is crucial for improving applications such as sentiment analysis, content moderation, and public opinion mining (Joshi, Bhattacharyya, and Carman 2017; Yi and Xia 2025). However, sarcasm detection is a challenging Natural Language Processing (NLP) task because of its intrinsic dependence on contextual cues, the semantic incongruity between literal and intended meanings, and its multi-faceted nature which manifests through diverse linguistic phenomena like pragmatic contrast and emotional inversion (Grice 1975; Wilson and Sperber 1992).

*These authors contributed equally.

†Corresponding author.

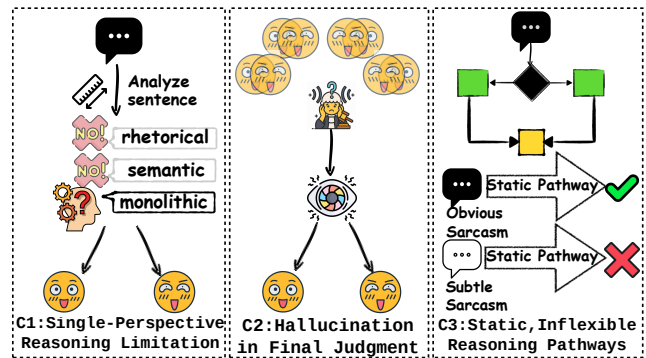


Figure 1: Key limitations of prior LLM-based sarcasm detection: single-perspective analysis, hallucination risk in final judgment, and static, inflexible reasoning pathways.

With the advances of Computer Science and Linguistics, the pursuit of effective sarcasm detection has seen significant evolution. Deep learning models used architectures such as CNNs, LSTMs, and GNNs to learn complex textual representations for training (Joshi, Bhattacharyya, and Carman 2017; Liu, Zhou, and Hu 2025). In particular, Transformer-based models, which used self-attention mechanisms and contextual embeddings, achieved remarkable performance in sarcasm detection by effectively capturing nuanced linguistic patterns and contextual cues within text (González, Hurtado, and Pla 2020). However, these paradigms struggled to grasp the implicit intent driving the intricate ironic rhetoric.

Recently, the emergence of Large Language Models (LLMs) has opened a promising new frontier in sarcasm detection, through both prompt engineering and instruction tuning (Liu, Zhou, and Hu 2025; Epron et al. 2024). But there still remains three critical challenges which limit LLMs' efficacy and reliability in this nuanced domain:

C1: Single-Perspective Reasoning Limitation. Standard LLM-based methods function as monolithic predictors, lacking the capacity to systematically deconstruct and analyze complex sarcasm from multiple linguistic dimensions (Guo et al. 2024).

C2: Hallucination Risk in Final Judgment. LLMs are susceptible to hallucination, leading to unreliable or unfaithful

judgments when synthesizing diverse and conflicting analytical signals into a single conclusion (Paul et al. 2024).

C3: Static and Inflexible Reasoning Pathways. Current LLM-based sarcasm detection models often rely on fixed prompts or architectures, which limits their ability to dynamically adapt their analytical strategies to the specific complexities of the input (Zhang et al. 2024; Yao et al. 2025).

To address these limitations, we introduce a novel framework featuring a decoupled, multi-agent reasoning architecture. The process begins with our **Dynamic Agentive Reasoning Engine (DARE)**, which coordinates six specialized Core Analysis Agents. Each of the six specialized agents is grounded in a distinct linguistic perspective, selected based on linguistic theory and empirical studies (Atardo 2001; Reyes, Rosso, and Veale 2013), to capture the multi-dimensional cues crucial for robust sarcasm detection. The Controller Agent adaptively manages these agents through iterative refinement and targeted team expansion, allowing the reasoning depth and focus to dynamically adjust in response to the complexity of each input. Rather than producing a single prediction, DARE generates a structured reasoning chain, providing a transparent account of the underlying analytical process. This chain is subsequently evaluated by a lightweight Rationale Adjudicator, which delivers the final sarcasm judgment. By explicitly decoupling reasoning from decision-making, our framework mitigates hallucination risk and advances interpretability, surpassing conventional single-model paradigms.

The main contributions of this work are summarized as follows:

- Propose a decoupled dynamic multi-agent sarcasm detection framework with adaptive, multi-perspective reasoning which emulates sophisticated human cognition.
- Enhance interpretability and mitigate hallucination via step-wise reasoning chains and separated adjudication.
- Achieve new SOTA on four sarcasm benchmarks, with superior robustness and generalization.

2 Related Work

2.1 Multi-Agent Cooperation

Multi-Agent Systems are a powerful paradigm for solving complex problems through collaborative interactions. Existing research has explored various cooperation patterns, such as debates, which enhance factual accuracy (Du et al. 2023), and structured dialogues for intricate problem-solving (Wu et al. 2023b). Prominent frameworks like CAMEL employ role-playing to simulate nuanced behaviors (Li et al. 2023), while AutoGen offers high versatility through customizable conversation patterns (Wu et al. 2023a). While effective, current agent frameworks suffer from poor adaptability due to their fixed composition. The new research focus is on developing flexible agent systems capable of dynamic reasoning and self-adaptation.

2.2 Self-Evolving Agents

To overcome the limitations of static agent architectures, research has explored self-evolving systems. Initial ap-

proaches focused on dynamically adapting workflow or network topologies based on execution feedback (Niu et al. 2025; Yang et al. 2025). More advanced paradigms involve agents that can rewrite their own source code for enhancement, as demonstrated in systems like ADAS (Hu, Lu, and Clune 2024) and the Darwin-Gödel Machine (Zhang et al. 2025a), which leverage meta agents and empirical validation to drive open ended evolution. The core advantages of these self-evolving architectures such as adaptability are particularly valuable for pragmatic, context dependent language understanding tasks like sarcasm and sarcasm detection, where flexible and situation-aware analysis is crucial.

2.3 Sarcasm Detection

Sarcasm detection is a challenging NLP task due to its inherent contextual and contradictory nature. Early machine learning approaches relied on hand-crafted features, such as lexical cues (Davidov, Tsur, and Rappoport 2010) and sentiment lexicons (Reyes, Rosso, and Veale 2013), but struggled to capture the nuanced semantics of sarcasm. Subsequently, deep learning models leveraged word embeddings (Mikolov et al. 2013; Pennington, Socher, and Manning 2014) and advanced architectures like CNNs (Poria et al. 2016), LSTMs (Zhang, Zhang, and Fu 2016), and GNNs (Liang et al. 2022) to learn hierarchical, sequential, and structural text features, yet often failed to capture the implicit intent behind complex ironic rhetoric. More recently, LLMs have been employed for sarcasm detection through prompting and zero-shot learning (Yao et al. 2025). While promising, these methods typically function as single-model predictors, limiting their capacity for the multi-perspective reasoning necessary to fully comprehend intricate ironic expressions. Recent work has also explored multi-modal sarcasm detection by combining visual and textual cues (Ramakrishnan, Ramakrishnan, and Lee 2025; Zhang et al. 2025b). In parallel, CAF-I (Liu, Zhou, and Hu 2025) investigates collaborative multi-agent LLM frameworks for text-based sarcasm detection using a fixed agent committee. However, achieving adaptive reasoning in single-modality sarcasm detection remains a challenging problem. Therefore, we propose SEVADE, a dynamic and self-evolving multi-agent framework for robust and interpretable sarcasm detection.

3 Methodology

Problem Definition. Let the dataset be denoted as $\mathcal{D} = \{(c_i, y_i)\}_{i=1}^N$, where each instance consists of an input text c_i and its corresponding ground truth label $y_i \in \{0, 1\}$, with 1 indicating sarcasm. The primary objective of sarcasm detection is to learn a predictive function \mathcal{G} that maps an input text c to its corresponding label y , such that $\mathcal{G} : c \mapsto y$.

Overall Framework. First, we design **DARE**, an iterative mechanism that generates a structured reasoning chain. The process unfolds as follows: (1) **Instantiation.** A Controller Agent assembles an initial team of agents with analytical perspectives relevant to the input text. (2) **Targeted Refinement.** In each cycle, the Controller identifies the single most ambivalent agent and prompts it to refine its analysis based on peer conclusions. (3) **Adaptive Expansion.** If

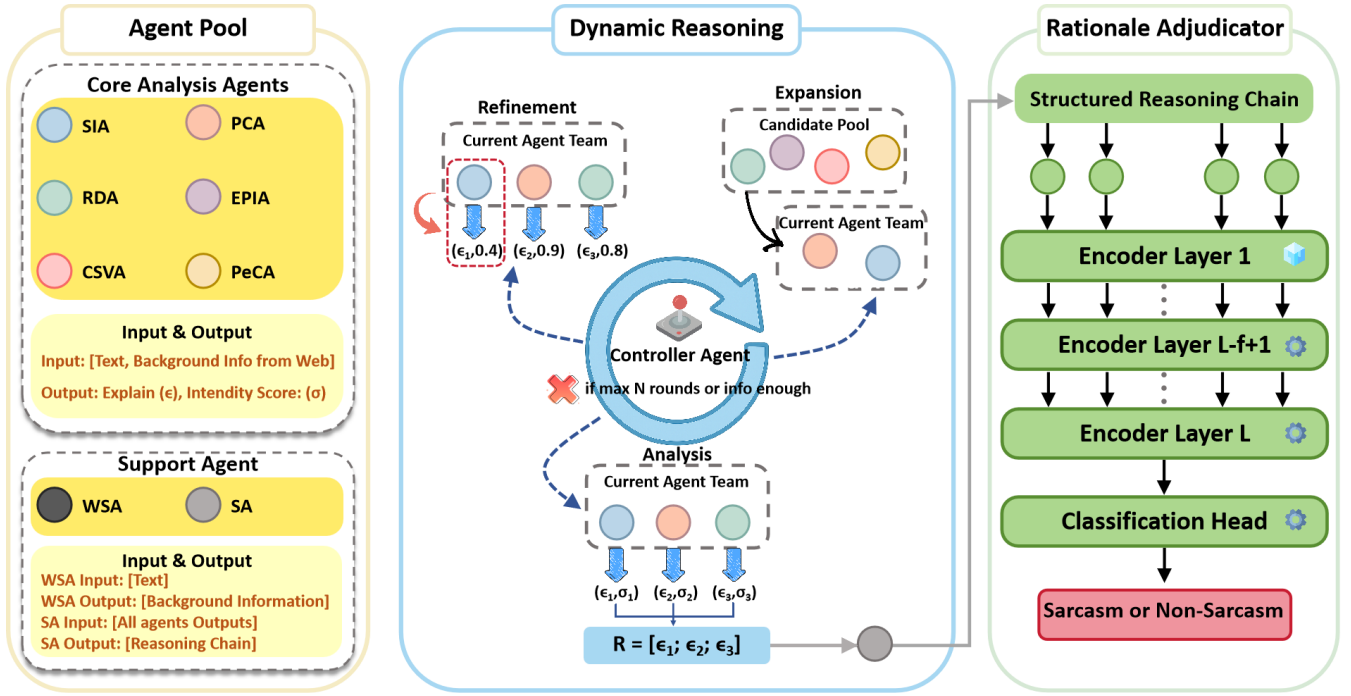


Figure 2: The overall framework of SEVADE. On the left, a pool of Core Analysis Agents, each with a distinct analytical perspective, and Support Agents with auxiliary functions are maintained. In the middle, the Controller Agent dynamically selects, refines, and expands the agent team to generate a structured reasoning chain. On the right, the Rationale Adjudicator produces the final sarcasm prediction based solely on the reasoning chain, decoupling reasoning from judgment.

the reasoning stagnates, the Controller recruits a new agent from an inactive pool to introduce a complementary viewpoint. (4) **Summarization.** Upon conclusion, a designated agent synthesizes the collective findings into a final reasoning chain R . Second, we introduce the Rationale Adjudicator, a lightweight classifier. It receives the reasoning chain R as its sole input to perform the final classification, ensuring a robust prediction grounded in the preceding logical analysis. An overview of our proposed framework is shown in Figure 2.

3.1 The Dynamic Agentive Reasoning Engine

Controller Agent and the Reasoning Process The reasoning process is governed by a central Controller Agent, which presides over the operational workflow of the DARE module to transform an input text into a final reasoning chain R . A foundational concept to this process is the structured output of each analytical agent, a_i . This output is a tuple, $O_i = (\sigma_i, \epsilon_i)$, comprising a scalar sarcasm intensity score, $\sigma_i \in [0, 1]$, representing perceived magnitude of the sarcastic signal, and a corresponding textual explanation, ϵ_i .

The process is initiated with **instantiation**, where the Controller Agent adaptively selects the initial agent team, $\mathcal{A}_{\text{active}}^{(0)}$, from the full pool $\mathcal{A}_{\text{pool}}$, by prompting a language model to identify the analytical roles most relevant to the input text.

Following instantiation, the system enters an iterative cycle of **targeted refinement**. At each iteration t , the Con-

troller identifies the single most ambivalent agent, $a_{\text{amb}}^{(t)}$, whose sarcasm intensity score exhibits maximum uncertainty.

$$a_{\text{amb}}^{(t)} = \underset{a_i \in \mathcal{A}_{\text{active}}^{(t)}}{\operatorname{argmin}} \left| \sigma_i^{(t)} - 0.5 \right| \quad (1)$$

This designated agent is subsequently tasked with a refinement of its analysis. This refinement is executed via the function $\mathcal{F}_{\text{refine}}$, which operates upon the agent’s prior output, $O_{\text{amb}}^{(t)}$, and is conditioned on the set of all peer outputs, yielding an updated analysis $O_{\text{amb}}^{(t+1)}$.

$$O_{\text{amb}}^{(t+1)} = \mathcal{F}_{\text{refine}}(O_{\text{amb}}^{(t)}, \{O_j^{(t)} \mid a_j \in \mathcal{A}_{\text{active}}^{(t)} \setminus \{a_{\text{amb}}^{(t)}\}\}) \quad (2)$$

Concurrently, the Controller performs a meta-cognitive check to determine if **adaptive expansion** is warranted. This check prompts the Controller to first render a judgment on whether the collective analysis is deemed *incomplete*, *contradictory*, or *logically stuck*, and if affirmative, to then execute a meta-reasoning task to select the agent, a_{new} , from the inactive pool $\mathcal{A}_{\text{inactive}}^{(t)}$, that can most effectively address the identified analytical gaps. The active team is subsequently expanded:

$$\mathcal{A}_{\text{active}}^{(t+1)} = \mathcal{A}_{\text{active}}^{(t)} \cup \{a_{\text{new}}\} \quad (3)$$

This iterative loop of refinement and expansion persists until the Controller Agent either deems the collective analysis to have achieved sufficient consistency. or when the pool of inactive agents is exhausted. The entire workflow is summarized in Algorithm 1.

Algorithm 1: Dynamic and Iterative Reasoning Process

```
1: Input: Text  $T$ 
2: Output: Reasoning Chain  $R$ 
3:  $\mathcal{A}_{\text{pool}} \leftarrow$  Initialize all Core & Support Agents
4:  $\mathcal{A}_{\text{active}} \leftarrow$  Controller.InitialAgents( $T, \mathcal{A}_{\text{pool}}$ )
5: loop
6:    $a_{\text{amb}} \leftarrow$  Controller.FindAmbivalentAgent( $\mathcal{A}_{\text{active}}$ )
7:    $a_{\text{amb}}.\text{refine\_analysis}(\text{context} = \mathcal{A}_{\text{active}})$ 
8:   if not Controller.ExpansionNeeded( $\mathcal{A}_{\text{active}}$ ) then
9:     break
10:  end if
11:   $\mathcal{A}_{\text{inactive}} \leftarrow \mathcal{A}_{\text{pool}} \setminus \mathcal{A}_{\text{active}}$ 
12:  if  $\mathcal{A}_{\text{inactive}}$  is empty then
13:    break
14:  end if
15:   $a_{\text{new}} \leftarrow$  Controller.ComplementAgent( $T, \mathcal{A}_{\text{inactive}}$ )
16:   $\mathcal{A}_{\text{active}}.\text{add}(a_{\text{new}})$ 
17: end loop
18:  $R \leftarrow$  SummarizationAgent.summarize( $\mathcal{A}_{\text{active}}$ )
19: return  $R$ 
```

Agent Roles The DARE module is populated by a diverse set of agents engineered for specialized functions, which are categorized as either Core Analysis Agents or Support Agents.

Core Analysis Agents. This agent pool consists of six agents grounded in linguistic and rhetorical theory, ensuring a multi-faceted deconstruction of the text.

- **Semantic Incongruity Agent (SIA):** Identifies and quantifies conflicts between the text’s literal meaning and established world knowledge.
- **Pragmatic Contrast Agent (PCA):** Analyzes the discordance between an utterance’s formulation and its pragmatic context.
- **Rhetorical Device Agent (RDA):** Detects key figures of speech indicative of sarcasm, such as hyperbole and understatement.
- **Emotion Polarity Inverter Agent (EPIA):** Measures the contradiction between the text’s overtly expressed emotion and the sentiment that would be objectively inferred from the situation.
- **Common Sense Violation Agent (CSVA):** Evaluates if the text’s content violates widely held principles of common sense.
- **Persona Conflict Agent (PeCA):** Examines and reports on inconsistencies between the speaker’s projected persona and the content of their statement.

Support Agents. This category includes agents that provide essential auxiliary functions to the reasoning process. The **Web Search Agent (WSA)** is selectively invoked before the main analysis stage, based on a decision by the Controller Agent. Given the input text, if external context is deemed necessary, the Controller uses an extraction function $\mathcal{F}_{\text{extract}}$ to generate one or two concise keywords:

$$\mathcal{K} = \mathcal{F}_{\text{extract}}(T) \quad (4)$$

These keywords are used by the WSA to query external sources and retrieve the top three most relevant background knowledge snippets. The retrieved evidence, together with the original text, is then provided to each analysis agent, supporting more informed and context aware reasoning. In contrast, the **Summarization Agent (SA)** executes the final step of the DARE workflow, synthesizing the structured output of all active agents into a single, coherent reasoning chain R .

3.2 Rationale Adjudicator

The Rationale Adjudicator (RA) is our framework’s final component, designed to robustly adjudicate the reasoning chain generated by the DARE module. To mitigate the hallucination risks inherent in a large model’s synthesis of diverse analyses, we employ a lightweight, fine-tuned model for this task, ensuring reliable and specialized final judgment. A critical design choice is that the Rationale Adjudicator’s sole input is the reasoning chain, R , generated by the DARE module. This forces the model to base its judgment exclusively on the logical coherence and semantic patterns of the provided rationale.

We implement the Rationale Adjudicator with BERT (Devlin et al. 2019), fine-tuning only the parameters of the last f layers while freezing all preceding layers. In all experiments, we set $f = 2$:

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

This strategy allows the model to adapt its high-level representations for the specific task of interpreting reasoning chains. The entire model is then optimized by minimizing the Binary Cross-Entropy loss, \mathcal{L}_{BCE} , between its prediction \hat{y} and the ground-truth label y .

$$\mathcal{L}_{\text{BCE}} = -\frac{1}{N} \sum_{i=1}^N [y_i \cdot \log(\hat{y}_i) + (1 - y_i) \cdot \log(1 - \hat{y}_i)] \quad (6)$$

4 Experiments

4.1 Experiment Setup

An overview of the experimental setup is outlined below.

Datasets We conduct evaluations on four sarcasm detection benchmarks. Includes **IAC-V1** (Lukin and Walker 2013), **IAC-V2** (Oraby et al. 2017), **MuSTARD** and **SemEval-2018 Task 3** (Van Hee, Lefever, and Hoste 2018). Statistics of datasets are summarized in Table 1.

Dataset	Train	Val	Test	Avg. Len
IAC-V1	1595	80	320	68
IAC-V2	5216	262	1042	43
SemEval-2018	3634	200	784	14
MuSTARD	552	–	138	14

Table 1: Overview of the benchmark datasets used for evaluating sarcasm detection.

Baselines We conduct a comparative analysis against a diverse suite of baselines organized into three main categories. **LLM-Based:** We include approaches from SarcasmCue (Yao et al. 2025), such as GPT-4o zero-shot and three prompting strategies, Chain of Contradiction, Graph of Cues, and Bagging of Cues, with performance reported from their original results. We also add GPT-5 as a stronger reasoning baseline for completeness. **Fine-tuned PLMs:** We evaluate BERT-base (Devlin et al. 2019) and RoBERTa-base (Liu et al. 2019), both fine-tuned on each target dataset. **Deep Learning Methods:** We compare with representative methods including MIARN (Tay et al. 2018), SAWS (Pan et al. 2020), and DC-Net (Liu et al. 2021), with results taken from previous studies (Qiu et al. 2024; Zhang et al. 2024; Tay et al. 2018; Xue et al. 2024; Pan et al. 2020).

Metrics and Implementation Details Our framework employs the **GPT-4o** model as backbone for all agents, accessed via the official OpenAI API with a temperature of 0 to ensure reproducibility. Following standard practices (Yao et al. 2025), performance is evaluated using **Accuracy** and primarily the **Macro-F1 score**.

4.2 Main Result

Table 2 presents a comprehensive comparison of our framework against various baselines. From the results, we have the following observations.

First, our framework establishes a new state-of-the-art across all datasets, achieving an average Accuracy of 78.14 and Macro-F1 of 77.90, improving by 7.01% and 6.55% over the strongest baseline, DC-Net. This performance stems from our synergistic two-stage architecture, the DARE module generates structured rationales through multi-agent analysis, and the lightweight Rationale Adjudicator classifies solely based on them. This design decouples reasoning from judgment, ensuring the decision is grounded in explicit logic and mitigating hallucination risks. Even against strong reasoning model like GPT-5, SEVADE remains clearly superior, confirming its design advantage.

Second, our findings reveal a core weakness in generalist LLMs, they lack the specific task adaptation required for nuanced sarcasm detection, while fine-tuning is computationally prohibitive and impractical. Our framework is designed to resolve this exact trade-off through the specific mechanism of the DARE module. By structuring the reasoning process around a dynamic, multi-perspective analysis, DARE provides an effective form of task adaptation. This targeted approach explains its ability to robustly handle complex cases where monolithic LLM architectures falter.

Finally, SEVADE’s advantage is most evident on the complex MuSTARD and SemEval datasets, with accuracy gains of 7.75% and 10.61% over the strongest baseline. This suggests these benchmarks require reasoning that transcends surface-level pattern matching, often hinging on common sense or external world knowledge. Traditional models lack mechanisms for such explicit reasoning. Our framework addresses this gap directly; agents like the Common Sense Violation Agent and the on-demand Web Search Agent explicitly integrate external validation, enabling our model to re-

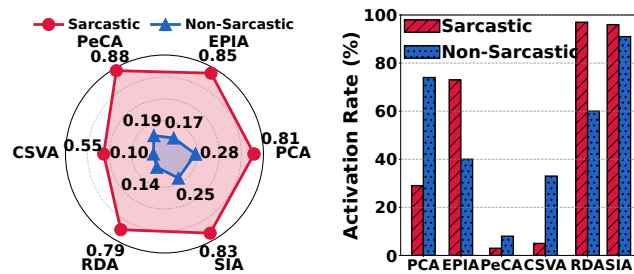
solve intricate sarcastic expressions that confound other approaches.

4.3 Ablation Study

We comprehensively validate our framework with the ablation study detailed in Table 3. We first assess the contribution of six core analysis agent by creating variants in which they are individually removed. Furthermore, we evaluate the core architecture through two mechanism-focused variants: **w/o Evolving**, which forces a static analysis by disabling iterative refinement and expansion; and **w/o RA**, which replaces our specialized Rationale Adjudicator with the base LLM to classify the final reasoning chain R .

The experiment results highlight the following conclusions: **(1) Each Core Analysis Agent provides essential insights.** The removal of any single agent leads to a noticeable drop in performance across all datasets. This demonstrates that our chosen set of agents offers a comprehensive, non-redundant analytical foundation, and a multi-faceted approach is critical for robust sarcasm detection; **(2) The impact of agents is dataset dependent.** For IAC-V1, performance is most affected by the absence of the Semantic Incongruity and Pragmatic Contrast agents. For Semeval-2018, the Pragmatic Contrast and Emotion Polarity Inverter agents are the most influential, highlighting the need for context and sentiment analysis for this dataset; **(3) The dynamic evolution process is crucial.** The **w/o Evolving** variant, which uses a static analysis, shows a substantial performance drop on both datasets. This result confirms that our iterative process of refinement and expansion is superior to a static agent committee approach; **(4) The superior performance of our full model over the w/o RA variant.** This finding validates our specialized Rationale Adjudicator, unlike the general-purpose base LLM, our lightweight classifier’s constrained nature mitigates hallucination risks, ensuring a more reliable judgment that is faithfully grounded in the preceding analysis.

4.4 Model Explainability



(a) Mean intensity of Agents. (b) Activation Rate of Agents.

Figure 3: Visualization of agent dynamics in processing sarcastic and non-sarcastic samples. (a) shows the mean intensity score. (b) shows the activation rate, indicating agent participation frequency.

We also investigate the internal dynamics of the DARE framework to validate the adaptability and interpretability

Baseline	IAC-V1		IAC-V2		MuSTARD		SemEval-2018		Avg.	
	Acc.	M-F1	Acc.	M-F1	Acc.	M-F1	Acc.	M-F1	Acc.	M-F1
MIARN	0.6321	0.6318	0.7275	0.7275	0.6460	0.6390	0.6850	0.6780	0.6726	0.6691
SAWS	0.6613	0.6560	0.7620	0.7620	0.6971	0.7095	0.6990	0.6890	0.7048	0.7041
DC-Net	0.6650	0.6640	0.7800	0.7790	<u>0.7128</u>	0.7143	<u>0.7630</u>	<u>0.7670</u>	<u>0.7302</u>	<u>0.7311</u>
BERT	0.6530	0.6520	0.7640	0.7620	0.6430	0.6430	0.6990	0.6840	0.6897	0.6852
RoBERTa	0.7010	0.6990	0.7660	0.7670	0.6610	0.6600	0.7020	0.6910	0.7075	0.7042
GPT-4o	0.7063	0.7005	0.7303	0.7199	0.6724	0.6579	0.6403	0.6317	0.6873	0.6775
GPT-4o+CoC	<u>0.7219</u>	<u>0.7152</u>	0.7336	0.7231	0.6942	0.6848	0.7079	0.7060	0.7144	0.7073
GPT-4o+GoC	0.6500	0.6291	0.6497	0.6130	0.7069	0.6991	0.7403	0.7402	0.6867	0.6704
GPT-4o+BoC	0.6875	0.6736	0.7135	0.6939	0.6942	0.6845	0.6212	0.6185	0.6791	0.6676
GPT-5	0.7031	0.7084	0.7493	0.7442	0.7076	0.6994	0.7257	0.7183	0.7214	0.7178
Ours	0.7304	0.7298	0.7832	0.7830	0.7681	0.7638	0.8439	0.8394	0.7814	0.7790
Improv.	1.17% \uparrow	2.04% \uparrow	0.51% \uparrow	0.41% \uparrow	7.75% \uparrow	6.92% \uparrow	10.61% \uparrow	9.43% \uparrow	7.01% \uparrow	6.55% \uparrow

Table 2: Overall performance comparison across four benchmark datasets. All LLM strategies are zero-shot. **Acc.** denotes Accuracy and **Ma-F1** signifies Macro-F1. Best results are in **bold**, second-best are underlined. Scores are reported as decimals.

Dataset	IAC-V1		Semeval-2018	
	Acc.	M-F1	Acc.	M-F1
w/o SIA	0.6708	0.6617	0.8091	0.8059
w/o PCA	0.6783	0.6747	0.7176	0.7166
w/o RDA	0.6891	0.6889	0.7734	0.7734
w/o EPIA	0.6825	0.6808	0.7361	0.7352
w/o CSVA	0.7047	0.7033	0.7328	0.7319
w/o PeCA	0.7038	0.7031	0.7349	0.7342
w/o Evolving	0.6928	0.6976	0.7754	0.7754
w/o RA	0.7147	0.7138	0.7558	0.7555
Full Model	0.7304	0.7298	0.8439	0.8394

Table 3: Ablation study on key components of our framework.

of its process by analyzing agent behavior (Figure 3).

From the result we present two key findings: **(1) Agents are highly specialized.** While all agents show baseline sensitivity by assigning appropriate high/low intensity scores when dealing with sarcasm/non sarcasm sample 3a, they exhibit profound functional specialization. The RDA and SIA agents dominate sarcastic detection with both high activation rates and the highest intensity scores, acting as primary incongruity detectors. Conversely, the PCA is more likely to be activated for non-sarcastic samples as a dedicated contextual congruity verifier; **(2) The model follows an evolutionary reasoning process driven by refinement and expansion.** This strategy is evidenced by the differential activation patterns. The model’s reasoning process evolves by shifting its analytical focus based on initial evidence. For sarcastic samples containing strong signals, the framework converges efficiently without needing extensive contextual analysis from PCA. For non-sarcastic samples that lack such obvious signals, the framework adaptively invokes PCA to perform a deeper contextual non-sarcasm validation. In short, this evolution is co-driven by refinement on a single perspective for analytical depth, and expansion with new perspectives for conceptual breadth. In conclusion, DARE’s evolutionary reasoning process culminates in an interpretable reasoning chain, which shows how the decision evolved.

4.5 Model Generalizability




Figure 4: Cross-dataset generalization performance. The left chart shows model that train on the IAC-V1 and test on semeval. The right chart shows the reverse scenario. Our proposed model is compared against BERT and RoBERTa.

To assess our framework’s real world applicability, we conducted a rigorous cross-dataset evaluation, training on one dataset and testing on the other. This setup evaluates the model’s ability to generalize across different data distributions and linguistic styles against BERT and RoBERTa baselines. As detailed in Figure 4, our framework demonstrates superior generalization. For instance, when trained on IAC-V1 and tested on semeval, our model achieves an M-F1 score of 66.97, surpassing the RoBERTa by over 27%. This significant performance advantage is consistent in the reverse setting as well, highlighting our model’s robustness.

This superior robustness stems directly from our framework’s design philosophy. Unlike monolithic models that overfit to superficial textual features, our approach first employs a set of specialized agents. These agents are explicitly designed based on formal linguistic and semantic principles, enabling them to extract robust, generalizable signals of sarcasm rather than dataset-specific artifacts. The pivotal Summarization Agent then synthesizes these high-quality, multi-perspective analyses into a structured reasoning chain R . By operating on this clean and logically coherent representation, the final Rationale Adjudicator makes its prediction based on universal sarcastic patterns, effectively overcoming the challenges of domain shift.

Input Text: “Do you mean someone who was wrong? I asked you before to give an example of where Newton found empirical evidence of God. You failed to do so. I suspect that you can’t.”
Ground Truth: NOT SARCASTIC **Model Prediction:** SARCASTIC

Agent	Intensity Score	Agent’s Rationale
SIA	0.8	”... due to the accusatory tone ... and the mismatch between the literal question and the speaker’s clear skepticism , which implies dismissal ...”
PCA	0.2	”The tone reads as critical but ... straightforward, with no strikingly jarring mismatch between the situation (a debate ...) and the language style . It seems literal rather than sarcastic.”
RDA	0.8	”... Pragmatic Contrast Agent highlights the critical tone ... Refining my explanation, the rhetorical question ‘Do you mean someone who was wrong?’ mock the possibility of evidence ... ”

Summarization: ”The agents predominantly indicate sarcasm conveyed through skepticism and rhetorical dismissal ...”

Table 4: Case study of a false positive sample on the IAC-V1 dataset.

4.6 Influences of LLM scales

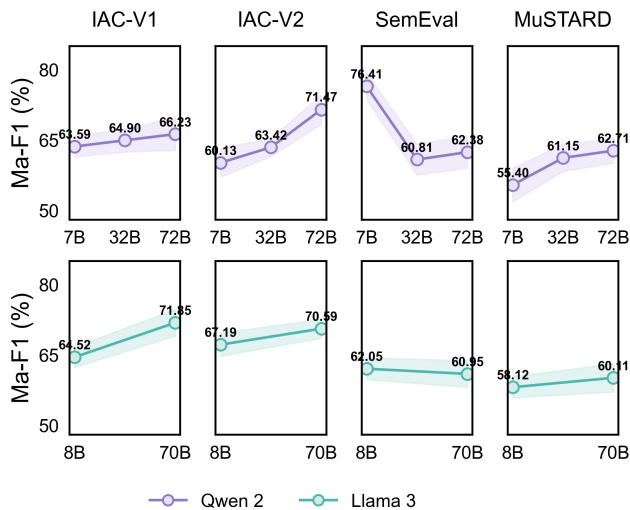


Figure 5: The effect of model scale on performance, using the Qwen 2 and Llama 3 series.

In our examination of model scale’s influence, a clear pattern emerges across most datasets. As illustrated in Figure 5, the performance of our framework generally improves as the underlying LLM’s parameter count increases. This aligns with the understanding that larger models possess more sophisticated reasoning capabilities, and enable DARE to produce more nuanced analyses.

Conversely, the results on the SemEval dataset present an important counter-trend, where smaller models achieve superior performance. This outcome stems from our two-stage architectural design. On datasets characterized by explicit cues like SemEval, smaller models generate concise and logically direct reasoning chains, which our rationale adjudicates reliably. Larger models, in contrast, exhibit a propensity for over-analysis, producing convoluted rationales laden with secondary or spurious details. Our adjudicator evaluates these less coherent chains as lower qual-

ity, leading to diminished performance. This phenomenon illustrates our framework’s architectural strength, its efficacy is determined by the coherence of the generated reasoning rather than the raw parameter scale. This inherent robustness underscores the framework’s adaptability and its utility in resource-constrained scenarios.

4.7 Error Analysis

Our error analysis indicates a predominant failure mode of misclassifying non-sarcastic texts as sarcastic, resulting in a higher false positive rate, a finding that is consistent with observations in similar advanced models (Yao et al. 2025). To investigate this common over-detection challenge, we deconstruct a representative case in Table 4.

From the case we can see that while SIA and RDA identified strong textual cues for sarcasm, such as *mismatch between the literal question and the speaker’s clear skepticism* and a *rhetorical question* used to *mock the possibility of evidence* the PCA provided crucial counterevidence that language is appropriate for debate circumstance. The actual misclassification, which results from this benign internal disagreement, highlights a core direction for future research, developing more advanced, context-sensitive perspective fusion mechanisms that go beyond simple cue identification or weighting, enabling models to flexibly and holistically integrate multiple linguistic perspectives.

5 Conclusion

In this work, we proposed SEVADE, a novel multi-agent framework that overcomes key sarcasm detection challenges like single-perspective, static analysis and LLM hallucination. Its core, the DARE, uses specialized linguistic agents to perform a multi-faceted analysis, generating a structured reasoning chain. A separate, lightweight Rationale Adjudicator then makes the final judgment based solely on this rationale, a decoupled architecture that effectively mitigates hallucination. Our approach achieves new state-of-the-art performance across four benchmarks, demonstrating superior accuracy and generalization.

Acknowledgements

This research is supported by XJTLU RDF-24-02-008, National Natural Science Foundation of China (Grant No. 62502396), and Leadership Talent Program (Science and Education) of SIP, KJL2024104.

References

- Attardo, S. 2001. *Humorous Texts: A Semantic and Pragmatic Analysis*. Mouton de Gruyter.
- Booth, W. C. 1974. *A rhetoric of irony*. University of Chicago Press.
- Davidov, D.; Tsur, O.; and Rappoport, A. 2010. Semi-supervised recognition of sarcasm in Twitter and Amazon. In *Proceedings of the fourteenth conference on computational natural language learning*, 107–116.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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)*, 4171–4186.
- Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2023. Improving factuality and reasoning in language models through multiagent debate. In *Forty-first International Conference on Machine Learning*.
- Epron, P.; Renard, M.; Zhang, S.; Couciero, M.; and Guibon, G. 2024. Which Approach to Detect Irony?
- González, J. Á.; Hurtado, L.-F.; and Pla, F. 2020. Transformer based contextualization of pre-trained word embeddings for irony detection in Twitter. *Information Processing & Management*, 57(4): 102262.
- Grice, H. P. 1975. Logic and conversation. In *Speech acts*, 41–58. Brill.
- Guo, T.; Chen, X.; Wang, Y.; Chang, R.; Pei, S.; Chawla, N. V.; Wiest, O.; and Zhang, X. 2024. Large language model based multi-agents: A survey of progress and challenges. *arXiv preprint arXiv:2402.01680*.
- Hu, S.; Lu, C.; and Clune, J. 2024. Automated design of agentic systems. *arXiv preprint arXiv:2408.08435*.
- Joshi, A.; Bhattacharyya, P.; and Carman, M. J. 2017. Automatic sarcasm detection: A survey. *ACM Computing Surveys (CSUR)*, 50(5): 1–22.
- Li, G.; Hammoud, H.; Itani, H.; Khizbullin, D.; and Ghanem, B. 2023. Camel: Communicative agents for “mind” exploration of large language model society. *Advances in Neural Information Processing Systems*, 36: 51991–52008.
- Liang, B.; Lou, C.; Li, X.; Yang, M.; Gui, L.; He, Y.; Pei, W.; and Xu, R. 2022. Multi-modal sarcasm detection via cross-modal graph convolutional network. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1767–1777.
- 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 preprint arXiv:1907.11692*.
- Liu, Y.; Wang, Y.; Sun, A.; Meng, X.; Li, J.; and Guo, J. 2021. A dual-channel framework for sarcasm recognition by detecting sentiment conflict. *arXiv preprint arXiv:2109.03587*.
- Liu, Z.; Zhou, Z.; and Hu, M. 2025. CAF-I: A Collaborative Multi-Agent Framework for Enhanced Irony Detection with Large Language Models. *arXiv preprint arXiv:2506.08430*.
- Lukin, S.; and Walker, M. 2013. Really? Well. Apparently Bootstrapping Improves the Performance of Sarcasm and Nastiness Classifiers for Online Dialogue. *NAACL 2013*, 30.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Niu, B.; Song, Y.; Lian, K.; Shen, Y.; Yao, Y.; Zhang, K.; and Liu, T. 2025. Flow: Modularized agentic workflow automation. *arXiv preprint arXiv:2501.07834*.
- Oraby, S.; Harrison, V.; Reed, L.; Hernandez, E.; Riloff, E.; and Walker, M. 2017. Creating and characterizing a diverse corpus of sarcasm in dialogue. *arXiv preprint arXiv:1709.05404*.
- Pan, H.; Lin, Z.; Fu, P.; and Wang, W. 2020. Modeling the incongruity between sentence snippets for sarcasm detection. In *ECAI 2020*, 2132–2139. IOS Press.
- Paul, D.; West, R.; Bosselut, A.; and Faltings, B. 2024. Making reasoning matter: Measuring and improving faithfulness of chain-of-thought reasoning. *arXiv preprint arXiv:2402.13950*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Poria, S.; Cambria, E.; Hazarika, D.; and Vij, P. 2016. A deeper look into sarcastic tweets using deep convolutional neural networks. *arXiv preprint arXiv:1610.08815*.
- Qiu, Z.; Yu, J.; Zhang, Y.; Lai, H.; Rao, Y.; Su, Q.; and Yin, J. 2024. Detecting Emotional Incongruity of Sarcasm by Commonsense Reasoning. *arXiv preprint arXiv:2412.12808*.
- Ramakrishnan, A. A.; Ramakrishnan, A. A.; and Lee, D. 2025. IRONIC: Coherence-Aware Reasoning Chains for Multi-Modal Sarcasm Detection. *arXiv preprint arXiv:2505.16258*.
- Reyes, A.; Rosso, P.; and Veale, T. 2013. A multidimensional approach for detecting irony in twitter. *Language resources and evaluation*, 47: 239–268.
- Tay, Y.; Tuan, L. A.; Hui, S. C.; and Su, J. 2018. Reasoning with sarcasm by reading in-between. *arXiv preprint arXiv:1805.02856*.
- Van Hee, C.; Lefever, E.; and Hoste, V. 2018. Semeval-2018 task 3: Irony detection in english tweets. In *Proceedings of the 12th international workshop on semantic evaluation*, 39–50.
- Wilson, D.; and Sperber, D. 1992. On verbal irony. *Lingua*, 87(1): 53–76.
- Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Li, B.; Zhu, E.; Jiang, L.; Zhang, X.; Zhang, S.; Liu, J.; et al. 2023a. Autogen:

Enabling next-gen llm applications via multi-agent conversation. *arXiv preprint arXiv:2308.08155*.

Wu, Y.; Jia, F.; Zhang, S.; Li, H.; Zhu, E.; Wang, Y.; Lee, Y. T.; Peng, R.; Wu, Q.; and Wang, C. 2023b. Mathchat: Converse to tackle challenging math problems with llm agents. *arXiv preprint arXiv:2306.01337*.

Xue, H.; Xu, L.; Tong, Y.; Li, R.; Lin, J.; and Jiang, D. 2024. Breakthrough from Nuance and Inconsistency: Enhancing Multimodal Sarcasm Detection with Context-Aware Self-Attention Fusion and Word Weight Calculation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 2493–2503.

Yang, Y.; Chai, H.; Shao, S.; Song, Y.; Qi, S.; Rui, R.; and Zhang, W. 2025. Agentnet: Decentralized evolutionary coordination for llm-based multi-agent systems. *arXiv preprint arXiv:2504.00587*.

Yao, B.; Zhang, Y.; Li, Q.; and Qin, J. 2025. Is sarcasm detection a step-by-step reasoning process in large language models? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 25651–25659.

Yi, P.; and Xia, Y. 2025. Irony Detection, Reasoning and Understanding in Zero-shot Learning. *arXiv preprint arXiv:2501.16884*.

Zhang, J.; Hu, S.; Lu, C.; Lange, R.; and Clune, J. 2025a. Darwin Godel Machine: Open-Ended Evolution of Self-Improving Agents. *arXiv preprint arXiv:2505.22954*.

Zhang, M.; Zhang, Y.; and Fu, G. 2016. Tweet sarcasm detection using deep neural network. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: technical papers*, 2449–2460.

Zhang, Y.; Zou, C.; Lian, Z.; Tiwari, P.; and Qin, J. 2024. SarcasmBench: Towards Evaluating Large Language Models on Sarcasm Understanding. *arXiv preprint arXiv:2408.11319*.

Zhang, Y.; Zou, C.; Wang, B.; and Qin, J. 2025b. Commander-GPT: Fully Unleashing the Sarcasm Detection Capability of Multi-Modal Large Language Models. *arXiv preprint arXiv:2503.18681*.