

PhishDebate: An LLM-Based Multi-Agent Framework for Phishing Website Detection

WENHAO LI¹, SELVAKUMAR MANICKAM^{1,*}, YUNG-WEY CHONG², SHANKAR KARUPPAYAH¹

¹ Cybersecurity Research Centre, Universiti Sains Malaysia, Pulau Pinang, Malaysia

² School of Computer Sciences, Universiti Sains Malaysia, Pulau Pinang, Malaysia
wenhaoli@ieee.org, {selva, chong, kshankar}@usm.my

Abstract—Phishing websites continue to pose a significant cybersecurity threat, often leveraging deceptive structures, brand impersonation, and social engineering tactics to evade detection. While recent advances in large language models (LLMs) have enabled improved phishing detection through contextual understanding, most existing approaches rely on single-agent classification facing the risks of hallucination and lack interpretability or robustness. To address these limitations, we propose PhishDebate, a modular multi-agent LLM-based debate framework for phishing website detection. PhishDebate employs four specialized agents to independently analyze different textual aspects of a webpage—URL structure, HTML composition, semantic content, and brand impersonation—under the coordination of a Moderator and a final Judge. Through structured debate and divergent thinking, the framework delivers more accurate and interpretable decisions. Extensive evaluations on commercial LLMs demonstrate that PhishDebate achieves 98.2% recall and 98.2% True Positive Rate (TPR) on a real-world phishing dataset, and outperforms single-agent and Chain of Thought (CoT) baselines. Additionally, its modular design allows agent-level configurability, enabling adaptation to varying resource and application requirements.

Index Terms—Phishing Website Detection, Large Language Model, Multi-Agent System.

I. INTRODUCTION

Phishing attacks continue to pose a pervasive and evolving threat across the digital landscape, targeting individuals, businesses, and institutions alike [1]. By mimicking legitimate websites and exploiting human trust, phishing websites deceive users into disclosing sensitive information such as login credentials, financial details, or personal data [2], [3]. The sophistication and scale of phishing operations have escalated in recent years, often leveraging automation, obfuscation, and brand impersonation to evade detection systems [4]. Consequently, phishing remains one of the most pressing cybersecurity challenges, accounting for a significant proportion of fraud and data breaches [5].

To combat this threat, a wide range of detection techniques have been proposed. Traditional approaches relied on heuristic rules and blacklists, which are lightweight but prone to obsolescence in the face of adaptive attackers [6],

[7]. Machine learning models have since emerged, incorporating engineered features from URLs, HTML structures, and third-party metadata [8], [9]. More recently, deep learning and transformer-based architectures have demonstrated promising results by learning patterns from raw inputs such as URLs, webpage content, and screenshots [10]–[12]. With the rise of Large Language Models (LLMs), phishing detection has entered a new era, where contextual understanding and semantic reasoning allow for the identification of subtle social engineering tactics [13].

However, existing LLM-based methods for phishing website detection rely on single-agent classification [14], which not only lacks iterative reasoning and interpretability but also risks hallucinations [15] and narrow decision-making due to dependence on a single perspective.

To address these limitations, we propose PhishDebate — a modular multi-agent LLM-based debate framework for phishing website detection. PhishDebate comprises four specialized agents, each targeting a distinct dimension of phishing evidence, including URL structure, HTML composition, semantic content, and brand impersonation. A Moderator oversees the structured debate process, while a final Judge delivers the classification verdict based on all arguments. The framework follows a modular design that allows users to flexibly include or exclude agents based on resource constraints or task-specific requirements. This structure supports collaborative reasoning, improves detection accuracy, enhances interpretability, and reduces dependence on a single-agent perspective.

The key contributions of this paper are as follows:

- We introduce PhishDebate, one of the first debate-based multi-agent LLM frameworks for phishing website detection, emphasizing modularity, explainability, and collaborative inference.
- We develop a set of specialized agent roles and debate coordination strategies that support multi-round argumentation and dynamic consensus evaluation, enabling the detection of complex phishing attempts.
- We conduct comprehensive evaluations using two real-world phishing datasets and four state-of-the-art LLMs, comparing PhishDebate against baselines including direct prompting and Chain of Thought (CoT) reasoning.

SELVAKUMAR MANICKAM is the corresponding author.

- We demonstrate through scenario analysis and a case study that PhishDebate not only improves classification performance but also provides interpretable rationales, enhancing trust and transparency in phishing detection.

The remainder of this paper is organized as follows: Section II reviews related works in phishing detection and multi-agent LLM reasoning. Section III details the design of the PhishDebate framework, including agent roles and debate logic. Section IV presents the evaluation setup, datasets, and performance metrics. Section V reports experimental results, including comparative baselines and scenario analyses and a detailed case study. Finally, Section VI concludes the paper and outlines directions for future research.

II. RELATED WORKS

A. Phishing Website Detection

Phishing website detection has evolved significantly over the past few decades, beginning with heuristic-based approaches [6], [7] that relied on rule-based inspection of URLs, domain features, and HTML signatures. These methods typically used predefined rules such as the presence of IP-based URLs, suspicious JavaScript functions, or invisible elements to flag malicious content. While lightweight and interpretable, heuristic techniques suffered from low adaptability to novel phishing tactics [1].

To improve generalization, researchers adopted machine learning-based classifiers [8], [9], which extract engineered features from URLs, HTML structure, and third-party metadata (e.g., WHOIS, Alexa rank) [16]. Algorithms like decision trees, support vector machines (SVM), and random forests were trained on labeled datasets to detect phishing behavior patterns. However, their reliance on manual feature extraction and third-party data posed scalability and robustness challenges.

Deep learning methods, especially Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), further advanced detection capabilities by automatically learning patterns from raw input such as URLs, text, and screenshots [10]–[12]. Vision-based models using webpage screenshots and hybrid models combining URL, DOM, and content-level features have shown promise in detecting visually deceptive phishing websites [17], [18]. Recently, transformer-based architectures and pre-trained language models have been applied to capture semantic cues in webpage content and URLs [1], [19], improving the understanding of social engineering cues.

With the rise of LLMs, phishing website detection has entered a new paradigm [1], [14]. LLMs provide context-aware understanding of webpage elements, brand impersonation, and social engineering tactics, which enhances detection performance even in the presence of various modern cloaking techniques [4], [20]. However, most existing approaches rely on binary classification with single-shot LLM prompts [14], lacking collaborative reasoning or

iterative verification. This gap highlights the opportunity for more interpretable, resilient, and robust LLM-based detection frameworks.

B. Multi-Agent Debate and Collaborative Reasoning

Multi-agent debate systems are inspired by human deliberation, where multiple independent agents analyze and critique a shared problem before reaching a decision [21]. These systems have been increasingly used to enhance the reasoning abilities of language models by encouraging diversity of perspectives and iterative refinement [22]–[24]. Each agent typically specializes in a different domain or analytic skill and engages in structured argumentation rounds moderated by a central judge or aggregator.

In the domain of phishing email detection, a debate-driven multi-agent LLM framework is proposed [25] where two agents argue for and against a phishing classification, and a third judge agent determines the final verdict. This setup enabled the system to uncover subtle social engineering cues and improved detection accuracy without relying on handcrafted features or complex prompt engineering. Similarly, MultiPhishGuard [26] introduced a five-agent system combining text, URL, metadata, explanation, and adversarial agents, coordinated by reinforcement learning to adapt to evolving adversarial threats. The model achieved good performance with enhanced robustness and interpretability through its explanation simplifier component.

Beyond phishing, the multi-agent debate framework has been applied to spam detection. A recent study [27] employed LLM agents with distinct roles to simulate complex evaluation processes similar to human decision-making. This collaborative design enhanced both accuracy and resilience against evolving spam tactics, outperforming traditional single-agent detectors in precision and adaptability.

These works demonstrate the effectiveness of multi-agent debate systems in security-relevant classification tasks. However, their applications have been primarily limited to phishing emails and spam detection, leaving phishing website detection as an unexplored frontier. This motivates our work to extend the debate paradigm to this domain.

III. PHISHDEBATE FRAMEWORK

In this section, we present the detailed design of proposed PhishDebate framework and introduce specialist agents and coordination agents involved in this framework.

A. Overall PhishDebate Framework

PhishDebate employs a multi-agent debate system consisting of six distinct agents: four specialist agents (URL Analyst Agent, HTML Structure Agent, Content Semantic Agent, and Brand Impersonation Agent), one Moderator, and one Judge. The framework follows a structured debate process where specialist agents analyze different aspects

of a website and engage in collaborative reasoning to reach a consensus. The Fig. 1 illustrates the workflow of PhishDebate.

The framework operates through four main phases:

- 1) **Initial Analysis Phase (Round 1)**, where each active specialist agent independently analyzes the website from their domain expertise without seeing other agents' responses.
- 2) **Consensus Evaluation Phase**, where the Moderator evaluates if consensus exists after each round based on the discussion from specialist agents. If a consensus is reached, no further round will go then.
- 3) **Multi-round Debate Phase (Rounds 2+)**, where agents can see and respond to other agents' analyses, refining their assessments through structured discussion. This phase continues until consensus is reached or maximum rounds are completed.
- 4) **Final Judgment Phase**, where the Judge always makes the final determination based on all available evidence from the debate process, regardless of whether consensus was reached.

Note that if $R_{max}=1$, the system performs only the initial independent analysis followed by consensus evaluation, then proceeds directly to the Judge for final decision-making.

The key innovation of PhishDebate lies in its ability to combine diverse analytical perspectives through structured argumentation, enabling the system to capture subtle phishing indicators that might be missed by single-agent approaches. The debate mechanism ensures that each agent's specialized analysis contributes to the final decision while maintaining transparency through detailed reasoning chains. Algorithm 1 presents the proposed PhishDebate framework.

B. Specialist Agents

The PhishDebate framework employs four specialized agents, each designed to analyze specific aspects of phishing detection. Each agent operates with domain-specific prompts and contributes unique insights to the collaborative decision-making process.

1) *URL Analyst Agent*: The URL Analyst Agent specializes in examining URL structures and domain characteristics to identify phishing indicators. This agent focuses on detecting suspicious patterns in domain names, subdomains, URL paths, and overall URL construction that are commonly associated with phishing attacks. The Fig. 2 details the prompt template used for URL Analyst Agent.

With this, a structured assessment can be returned containing claim (PHISHING/LEGITIMATE), confidence score (0-1), and supporting evidence focusing on URL-based indicators such as domain spoofing, suspicious TLDs, URL shortening, and deceptive subdomain usage.

2) *HTML Structure Agent*: The HTML Structure Agent analyzes the underlying HTML code structure to identify technical indicators of phishing websites. This

Algorithm 1: PhishDebate Framework

Input: URL U , HTML H , text T , R_{max} , τ , agents A
Output: Assessment F , confidence, rounds, early termination

```

Function PhishDebate( $U, H, T, R_{max}, \tau, A$ ):
   $responses \leftarrow \emptyset, round \leftarrow 1$ ;
  // Phase 1: Independent Analysis (Round 1)
  foreach  $a_i \in A$  do
     $prompt_i \leftarrow \text{GenPrompt}(a_i, U, H, T)$ ;
     $responses[a_i] \leftarrow \text{QueryAgent}(a_i, prompt_i)$ ;
  end
  // Phase 2: Consensus Evaluation
   $consensus \leftarrow \text{ModEval}(responses, round)$ ;
  if  $consensus.reached \wedge consensus.conf \geq \tau$  then
    // Skip additional debate rounds if consensus reached
    goto Phase 4;
  end
  // Phase 3: Multi-round Debate (Rounds 2+)
  for  $round = 2$  to  $R_{max}$  do
     $context \leftarrow \text{AggResp}(responses)$ ;
    foreach  $a_i \in A$  do
       $prompt \leftarrow \text{GenDebPrompt}(a_i, context)$ ;
       $responses[a_i] \leftarrow \text{QueryAgent}(a_i, prompt)$ ;
    end
     $consensus \leftarrow \text{ModEval}(responses, round)$ ;
    if  $consensus.reached \wedge consensus.conf \geq \tau$  then
      break;
    end
  end
  // Phase 4: Final Judgment (Judge always decides)
   $judge\_prompt \leftarrow \text{GenJudgePrompt}(responses, round)$ ;
   $decision \leftarrow \text{QueryJudge}(judge\_prompt)$ ;
  return
  ( $decision.assess, decision.conf, round, consensus.reached$ );

```

agent examines form elements, JavaScript usage, iframe implementations, and other structural patterns that may indicate malicious intent. The prompt template used for HTML Structure Agent can be found in supplementary materials [28].

As a result, a technical assessment can be formed focusing on HTML structural anomalies, suspicious form actions, hidden elements, JavaScript obfuscation, and other code-level indicators of phishing attempts.

3) *Content Semantic Agent*: The Content Semantic Agent performs natural language analysis of the visible website content to detect linguistic patterns and semantic cues associated with phishing attacks. This agent specializes in identifying manipulative language, urgency tactics, and social engineering techniques. The prompt template used for Content Semantic Agent can be found in supplementary materials [28].

Through this linguistic analysis, emotional manipulation, urgency indicators, credential harvesting language, and other semantic patterns characteristic of phishing content can be identified.

4) *Brand Impersonation Agent*: The Brand Impersonation Agent specializes in detecting attempts to impersonate legitimate brands or organizations. This agent analyzes both URL and content elements to identify unauthorized use of brand names, logos, and corporate identity ele-

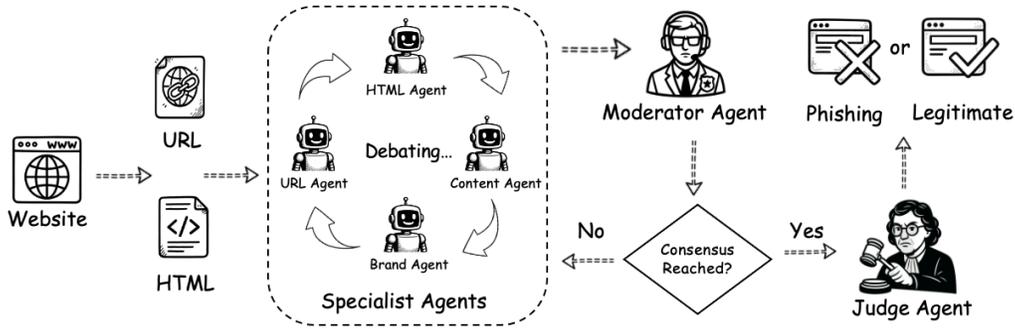


Fig. 1. The Workflow of PhishDebate Framework.

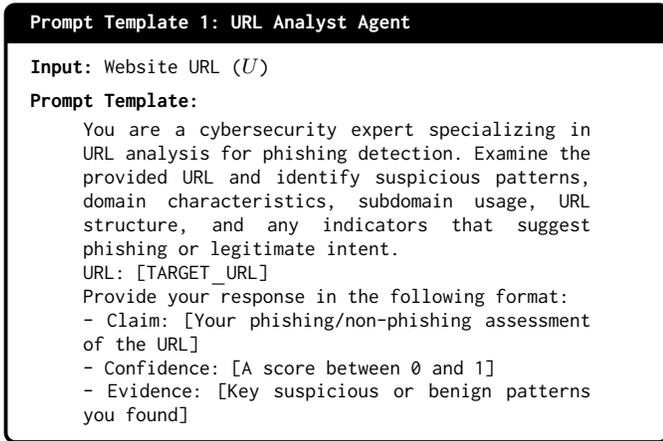


Fig. 2. Prompt template for URL Analyst Agent.

ments. The prompt template used for Brand Impersonation Agent can be found in supplementary materials [28].

The brand impersonation assessment identifies the specific brands being mimicked, authenticity indicators, and evidence of legitimate versus fraudulent brand usage.

C. Coordination Agents

1) *Moderator*: The Moderator serves as the debate coordinator, evaluating specialist agent responses after each round to determine if consensus has been reached. The Moderator analyzes the collective insights from all active specialist agents and decides whether the evidence is sufficient for a confident determination. The prompt template used for Moderator can be found in supplementary materials [28].

A JSON-formatted consensus evaluation will be generated containing: consensus status (Yes/No), supported assessment (PHISHING/LEGITIMATE/UNCERTAIN), detailed reasoning, confidence score, and continuation decision.

2) *Judge*: The Judge makes the final determination after the maximum number of debate rounds. The Judge considers all evidence presented by specialist agents throughout the entire debate process and renders a defini-

tive verdict. The prompt template used for Judge can be found in supplementary materials [28].

The Judge Agent outputs the Final judgment in JSON format containing definitive assessment (PHISHING/LEGITIMATE), confidence score, comprehensive reasoning, and key evidence summary.

Moreover, PhishDebate follows a modular design philosophy that enables flexible configuration based on specific use cases and requirements. The framework supports several customization options:

Agent Exclusion: Users can selectively exclude specific specialist agents from the debate process while maintaining system integrity. The framework enforces safety constraints ensuring that critical agents (Moderator and Judge) cannot be excluded and at least one specialist agent remains active. This feature enables comparative analysis and performance optimization for different scenarios.

Configurable Debate Rounds: The number of debate rounds can be adjusted (1-10 rounds) to balance between thoroughness and computational efficiency. Early termination mechanisms allow the system to conclude debates when strong consensus is reached, optimizing resource utilization.

In addition, to optimize token usage and prevent context limit violations, PhishDebate implements intelligent content truncation for HTML and text inputs. HTML content is truncated based on model-specific token limits, with truncation performed at HTML tag boundaries to preserve structural integrity. A truncation notice is appended when content exceeds limits, ensuring agents are aware of incomplete data while maintaining computational efficiency.

IV. EVALUATION

In this section, we introduce the LLMs and datasets used in this study and illustrate the data processing and evaluation metrics used.

A. Model Used

To evaluate the performance of the proposed PhishDebate framework, we selected four commercial LLMs for comparison and further analysis: Qwen2.5-vl-72b-instruct [29], Gemini-2.0-Flash [30], GPT-4o [31], and GPT-4o

Mini [31]. These models were chosen for their advanced reasoning capabilities, diverse architectures, and wide availability via API, providing a representative benchmark across different commercial LLM ecosystems.

B. Dataset & data Processing

This study employs two datasets to evaluate and analyze the effectiveness of PhishDebate. The first is a Phishing Websites Dataset from Mendeley [32], consisting of both phishing and legitimate websites. This dataset originally includes URL, HTML source, and metadata, and was collected from various sources including Google Search, Ebbu2017 Phishing Dataset [33], PhishTank [34], OpenPhish [35], and PhishRepo [36]. For our experiments, we randomly sampled 500 phishing and 500 legitimate instances.

The second dataset is the TR-OP Dataset [37], which contains manually labeled and balanced samples. Benign samples are drawn from the Tranco [38] top 50k domains, while phishing samples were crawled and validated within a six-month period from July to December 2023, spanning 440 unique phishing targets from OpenPhish [35]. Similarly, we randomly selected 500 phishing and 500 legitimate samples from this dataset.

We used the Mendeley phishing website dataset [32] for performance benchmarking and comparison across models, while both Mendeley [32] and TR-OP [37] datasets were used in the scenario analysis to enhance result robustness.

For data processing, we implemented a pipeline to extract structured features from each sample. The procedure includes: (1) reading and verifying the URL and raw HTML; (2) using BeautifulSoup to parse and clean HTML by removing tags like `style`, `noscript`, and external stylesheets; (3) extracting visible text by removing all `script` tags. This produces three components per sample: URL, cleaned HTML, and visible text.

The following Algorithm 2 summarizes our data preprocessing pipeline:

Algorithm 2: Data Preprocessing Pipeline

Input: Sample $s = (f_U, f_H)$ where f_U : URL file, f_H : HTML file
Output: Processed tuple (U, H, T) : URL, cleaned HTML, and visible text

```

 $U \leftarrow \text{Read}(f_U), R \leftarrow \text{Read}(f_H);$  // Extract URL and raw HTML
if  $U = \emptyset$  or  $R = \emptyset$  then
|   return None; // Skip invalid sample
end
 $H' \leftarrow \text{CleanHTML}(R);$  // Remove style, noscript, link[rel=stylesheet]
 $H \leftarrow \text{Remove}(H', \text{script});$  // Final cleaned HTML
 $T \leftarrow \text{ExtractText}(H);$  // Visible text from HTML
return  $(U, H, T)$ 

```

C. Evaluation Metrics

We evaluate model performance using standard binary classification metrics, computed from the confusion matrix:

- **True Positive Rate (TPR) / Recall:** $\text{TPR} = \frac{TP}{TP+FN}$ Measures the proportion of phishing samples correctly identified.
- **True Negative Rate (TNR):** $\text{TNR} = \frac{TN}{TN+FP}$ Measures the proportion of legitimate samples correctly identified.
- **False Positive Rate (FPR):** $\text{FPR} = \frac{FP}{FP+TN}$ Measures how often legitimate websites are incorrectly classified as phishing.
- **False Negative Rate (FNR):** $\text{FNR} = \frac{FN}{FN+TP}$ Measures how often phishing websites are missed.
- **Precision:** $\text{Precision} = \frac{TP}{TP+FP}$ Measures the correctness of phishing predictions.
- **Accuracy:** $\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN}$ Reflects the overall classification correctness.
- **F1 Score:** $\text{F1} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$ Harmonic mean of precision and recall; balances false positives and false negatives.

V. RESULTS

A. Performance Evaluation

Table I illustrates the results of performance evaluation across four commercial LLMs on the PhishDebate framework, including GPT-4o, GPT-4o-mini, Gemini-2.0, and Qwen2.5-vl-72b-instruct. The results demonstrate that GPT-4o consistently outperforms others in terms of overall classification effectiveness, achieving the highest accuracy (96.50%), precision (94.97%), and F1 score (96.56%), along with the best TNR (94.8%) and the lowest FPR (5.2%). Gemini-2.0, while exhibiting the fastest inference time (17.9s), achieved the highest TPR (99.2%) and the lowest FNR (0.8%), indicating strong phishing detection capability. However, its performance was affected by 8 instances in which it failed to follow the binary classification instruction in the prompt, instead outputting an uncertain response. These cases, originating from legitimate websites, were treated as misclassifications, thereby increasing Gemini’s FPR to 14.4%, lowering its TNR to 85.57%, and reducing its overall accuracy and F1 score to 92.44% and 92.97%, respectively. GPT-4o-mini performs competitively but lags slightly behind GPT-4o, whereas Qwen showed the lowest accuracy and F1 score due to a significantly higher FNR (19.4%), despite maintaining a relatively low FPR (9.4%). These results highlight GPT-4o as the most balanced and reliable model for phishing detection in the context of the PhishDebate framework.

B. Comparative Evaluation

To evaluate the effectiveness of our proposed multi-agent debate system, we compare it against two widely adopted single-agent prompting strategies. First, we include a Single-Agent with Direct Prompt baseline, where a LLM receives the raw input (URL and HTML-visible text) and directly predicts whether the page is phishing without any intermediate reasoning. This baseline reflects a straightforward application of LLMs for classification

TABLE I
EVALUATION METRICS OF PHISHDEBATE ON COMMERCIAL LLMs

Model	TPR	TNR	FPR	FNR	Recall	Precision	Accuracy	F1 Score	Time (s)
GPT-4o	0.982	0.948	0.052	0.018	0.982	0.9497	0.9650	0.9656	22.20
GPT-4o-mini	0.980	0.898	0.102	0.020	0.980	0.9057	0.9390	0.9414	37.50
Gemini-2.0-Flash	0.992	0.8557	0.1443	0.008	0.992	0.8748	0.9244	0.9297	17.90
Qwen2.5-vl-72b-instruct	0.806	0.906	0.094	0.194	0.806	0.8956	0.8560	0.8484	56.36

tasks and serves as a reference point for measuring improvements attributable to structured reasoning. Second, we include a Single-Agent with CoT baseline [39], where the LLM is prompted to reason step-by-step through URL analysis, content interpretation, and brand impersonation detection before arriving at a decision. This variant captures the benefits of explicit intermediate reasoning within a single model, enabling us to isolate and assess the added value of distributed agent roles and debate-based aggregation in our system. We provide the prompt templates used for two baselines in supplementary materials [28].

Table II presents the results of comparative evaluation between the proposed PhishDebate framework and two baselines—Single Agent and CoT prompting. The results demonstrate the effectiveness of our debate-driven approach in enhancing phishing detection on GPT-4o-mini. PhishDebate outperforms both baselines across all key evaluation metrics, achieving the highest precision (90.57%), accuracy (93.90%), recall (98.00%), and F1 score (94.14%). Notably, while CoT prompting provides reasonably strong results (e.g., 90.94% F1 score), it suffers from lower recall and precision compared to PhishDebate, especially when factoring in the 50 cases it labeled as “uncertain” instead of providing definitive classifications. These include 16 phishing websites and 34 legitimate websites, which, according to our evaluation protocol, are considered misclassifications. This underscores a critical limitation of CoT in reliability and confidence. In contrast, the PhishDebate framework significantly improves the confidence and decisiveness of LLM-based predictions by leveraging multi-agent debate, thereby reducing indecisive outputs and promoting consistent judgment across phishing and legitimate classes.

TABLE II
COMPARISON OF PHISHDEBATE WITH SINGLE AGENT AND CoT

Method	Precision	Accuracy	Recall	F1 Score
Single Agent	0.6057	0.6700	0.9740	0.7469
CoT	0.8861	0.9070	0.9340	0.9094
PhishDebate (Ours)	0.9057↑	0.9390↑	0.9800↑	0.9414↑

C. Scenario Analysis

PhishDebate follows a modular design, allowing users to flexibly exclude specific specialized agents according to dif-

ferent use case requirements, such as resource constraints or application priorities.

Table III presents a scenario analysis evaluating the performance of the PhishDebate framework when individual specialized agents are excluded. The results are averaged across two datasets (Mendeley [32] and TR-OP [37]) to ensure robustness. The full configuration using all agents achieves strong overall performance, particularly demonstrating the highest recall (0.985) and lowest FNR (FN = 8), indicating its superior ability to identify phishing websites with minimal missed detections. This highlights the importance of the ensemble debate mechanism in capturing diverse phishing indicators across URL, HTML, semantic, and brand features.

Notably, the exclusion of the HTML Structure Agent leads to the highest F1 score (0.9471), accuracy (0.9455), precision (0.9207), and the lowest FPR (FP = 42), suggesting that in specific deployment scenarios where legitimate website preservation is critical, a reduced configuration excluding this agent may offer better precision-legitimacy trade-offs. However, this comes at a minor cost in recall (0.975), which is still lower than the full configuration.

The variant without the Content Semantic Agent matches the full agent setup in recall and TPR (TP = 492), but suffers in terms of increased false positives (FP = 56), suggesting a reduced ability to discriminate legitimate content. This supports the value of semantic reasoning in reducing over-flagging. Similarly, excluding the URL Analyst Agent leads to a clear degradation in phishing detection capability, with a recall drop to 0.952 and increased false negatives (FN = 24), reinforcing the foundational role of URL features in phishing identification. Finally, omitting the Brand Impersonation Agent results in slightly lower overall performance, underscoring its contribution to brand spoof detection. Overall, these findings emphasize the effectiveness of the modular PhishDebate design, while demonstrating that task-specific configurations can still perform competitively under constrained scenarios.

D. Case Study

PhishDebate begins each investigation by distributing the same artefacts—the raw URL¹, rendered HTML, and visible text—to its four specialised agents (URL An-

¹The raw URL for this phishing website is: <https://mail.mxcapital.com.br/wp-includes/wells/wells/page/index.htm>.

TABLE III
SCENARIO ANALYSIS OF AGENT EXCLUSION IN PHISHDEBATE FRAMEWORK

Setting	Avg. TP	Avg. FN	Avg. FP	Avg. TN	Avg. Recall	Avg. Precision	Avg. Accuracy	Avg. F1 Score
All Agents	492	8	50	450	0.985	0.9070	0.9420	0.9444
W/O URL Agent	476	24	46	454	0.952	0.9110	0.9295	0.9311
W/O HTML Agent	488	12	42	458	0.975	0.9207	0.9455	0.9471
W/O Content Agent	492	8	56	444	0.985	0.8987	0.9370	0.9399
W/O Brand Agent	489	11	56	444	0.978	0.8981	0.9335	0.9363

alyst, HTML Structure, Content Semantic, and Brand-Impersonation). Each agent reasons independently within its expertise, returns a *claim* (phishing or legitimate), a numerical *confidence*, and traceable *evidence*. The moderator then inspects the four perspectives and decides whether a clear majority exists. If not, it compels another debate round, asking every agent to reconcile its view with the counter-evidence provided by its peers. Only when consensus is reached (or the maximum number of rounds elapses) does the moderator submit a verdict to the judge, who finalises the assessment and records a confidence score.

Fig. 3 narrates a representative two-round session on a suspicious phishing webpage. In Round 1 the agents split 2–2: the URL Analyst and Brand-Impersonation agents flagged the directory string “/wp-includes/wells/wells/” as brand spoofing (Wells Fargo) and noted that legitimate sites never expose /wp-includes/; by contrast, the HTML Structure and Content Semantic agents found no typical phishing artefacts—no forms, iframes, or persuasive language—and therefore judged the page legitimate. Lacking consensus, the moderator declared the case *UNCERTAIN* and requested a second round.

During Round 2 each agent re-examined the shared evidence. Crucially, the HTML Structure and Content Semantic agents acknowledged that benign error pages are frequently used as decoys and that the highly irregular path outweighed the absence of active credential-stealing code. After switching their stance to “Likely Phishing,” all four agents converged. The moderator therefore recorded a consensus of *PHISHING* with 0.88 confidence, and the judge confirmed the decision. This outcome underscores two salient strengths of PhishDebate: (i) its multi-view debate can uncover *dormant* or *cloaked* phishing infrastructure that content-only scanners miss, and (ii) the step-wise trace in Fig. 3 provides transparent justifications that facilitate analyst review or automated policy triggers.

VI. CONCLUSION

This paper presents PhishDebate, a modular multi-agent LLM-based debate framework for phishing website detection. By leveraging the collaborative reasoning of specialized agents and a structured debate process moderator and judged by dedicated agent, PhishDebate

Rnd	Role	Claim	Conf.	Key (abridged)	Evidence	Comment / Change in R2
1	URL Analyst Agent	Likely Phishing	0.85	Suspicious path /...wells/wells/; brand name repetition	Clean HostGator 500 page; no forms, iframes, obfuscated JS	—
	HTML Structure Agent	Legitimate	0.95			—
	Content Semantic Agent	Not Phishing	0.99	Neutral text; no credential request		—
	Brand Impersonation Agent	Possible Impersonation	0.70	Repeated “wells” implies Wells Fargo spoof		—
	Moderator	UNCERTAIN	0.60	Split (2 vs 2); further debate required		Debate continues
2	URL Analyst Agent	Likely Phishing	0.85	Same URL evidence; dormant payload possible		—
	HTML Structure Agent	Suspicious (revised)	0.65	Benign HTML could be decoy; path highly irregular		Switched side
	Content Semantic Agent	Likely Phishing (revised)	0.85	URL & brand cues outweigh benign text		Switched side
	Brand Impersonation Agent	Brand Spoof	0.85	URL strongly imitates Wells Fargo; domain unrelated		Confidence ↑
	Moderator	PHISHING (consensus)	0.88	All agents now agree on phishing intent		Debate stops
	Judge	Final PHISHING	0.88	Dormant/cloaked phishing; URL evidence decisive		Case closed

Note: Rnd = Round; Conf. = confidence score (0–1). The “Comment / Change in R2” column describes shifts occurring between rounds.

Fig. 3. A detailed detection example with PhishDebate.

has demonstrated good performance on phishing website detection and addresses key limitations of traditional and single-agent LLM approaches.

Future research will explore the resilience of PhishDebate against adversarial phishing samples, including obfuscated, cloaked, or adversarially crafted webpages designed to bypass LLM-based detection. This will further validate the framework’s adaptability in real-world, adversarial threat environments.

REFERENCES

- [1] W. Li, S. Manickam, Y.-W. Chong, W. Leng, and P. Nanda, “A state-of-the-art review on phishing website detection techniques,” *IEEE Access*, vol. 12, pp. 187 976–188 012, 2024.
- [2] R. Dhamija and J. D. Tygar, “The battle against phishing: Dynamic security skins,” in *Proceedings of the 2005 Symposium on Usable Privacy and Security*, ser. SOUPS ’05. New York, NY, USA: Association for Computing Machinery, 2005, p. 77–88.
- [3] K. Lim, J. Park, and D. Kim, “Phishing vs. legit: Comparative analysis of client-side resources of phishing and target brand websites,” in *Proceedings of the ACM Web Conference 2024*, ser. WWW ’24. New York, NY, USA: Association for Computing Machinery, 2024, p. 1756–1767.
- [4] W. Li, S. Manickam, S. U. A. Laghari, and Y.-W. Chong, “Uncovering the cloak: A systematic review of techniques used to conceal phishing websites,” *IEEE Access*, vol. 11, pp. 71 925–71 939, 2023.
- [5] N. Nejari, K. Zkik, H. Hammouchi, M. Ghogho, and H. Benbrahim, “Assessing data breach factors through modern crime theory: A structural equation modeling approach,” *IEEE Access*, vol. 12, pp. 92 198–92 214, 2024.

- [6] S. Gastellier-Prevost, G. G. Granadillo, and M. Laurent, "Decisive heuristics to differentiate legitimate from phishing sites," in *2011 Conference on Network and Information Systems Security*, 2011, pp. 1–9.
- [7] R. M. Mohammad, F. Thabtah, and L. McCluskey, "Intelligent rule-based phishing websites classification," *IET Information Security*, vol. 8, no. 3, pp. 153–160, 2014.
- [8] O. K. Sahingoz, E. Euber, O. Demir, and B. Diri, "Machine learning based phishing detection from urls," *Expert Systems with Applications*, vol. 117, pp. 345–357, 2019.
- [9] L. Tang and Q. H. Mahmoud, "A survey of machine learning-based solutions for phishing website detection," *Machine Learning and Knowledge Extraction*, vol. 3, no. 3, pp. 672–694, 2021.
- [10] D. Liu and J.-H. Lee, "Cnn based malicious website detection by invalidating multiple web spams," *IEEE Access*, vol. 8, pp. 97 258–97 266, 2020.
- [11] S. Abdelnabi, K. Krombholz, and M. Fritz, "Visualphishnet: Zero-day phishing website detection by visual similarity," in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1681–1698.
- [12] E. Benavides-Astudillo, W. Fuertes, S. Sanchez-Gordon, G. Rodriguez-Galan, V. Martínez-Cepeda, and D. Nuñez-Agurto, "Comparative study of deep learning algorithms in the detection of phishing attacks based on html and text obtained from web pages," in *Applied Technologies*, M. Botto-Tobar, M. Zambrano Vizuete, S. Montes León, P. Torres-Carrión, and B. Durakovic, Eds. Cham: Springer Nature Switzerland, 2023, pp. 386–398.
- [13] S. Gallagher, B. Gelman, S. Taoufiq, T. Vörös, Y. Lee, A. Kyadige, and S. Bergeron, *Phishing and Social Engineering in the Age of LLMs*. Cham: Springer Nature Switzerland, 2024, pp. 81–86.
- [14] T. Koide, H. Nakano, and D. Chiba, "Chatphishdetector: Detecting phishing sites using large language models," *IEEE Access*, vol. 12, pp. 154 381–154 400, 2024.
- [15] Z. Lin, Z. Niu, Z. Wang, and Y. Xu, "Interpreting and mitigating hallucination in llms through multi-agent debate," *arXiv preprint arXiv:2407.20505*, 2024.
- [16] J. James, S. L., and C. Thomas, "Detection of phishing urls using machine learning techniques," in *2013 International Conference on Control Communication and Computing (ICCC)*, 2013, pp. 304–309.
- [17] Y. Lin, R. Liu, D. M. Divakaran, J. Y. Ng, Q. Z. Chan, Y. Lu, Y. Si, F. Zhang, and J. S. Dong, "Phishpedia: A hybrid deep learning based approach to visually identify phishing webpages," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 3793–3810. [Online]. Available: <https://www.usenix.org/conference/usenix-security21/presentation/lin>
- [18] R. Liu, Y. Lin, X. Yang, S. H. Ng, D. M. Divakaran, and J. S. Dong, "Inferring phishing intention via webpage appearance and dynamics: A deep vision based approach," in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 1633–1650. [Online]. Available: <https://www.usenix.org/conference/usenix-security22/presentation/liu-ruofan>
- [19] P. Maneriker, J. W. Stokes, E. G. Lazo, D. Carutasu, F. Tajaddodianfar, and A. Gururajan, "Urltran: Improving phishing url detection using transformers," in *MILCOM 2021 - 2021 IEEE Military Communications Conference (MILCOM)*, 2021, pp. 197–204.
- [20] W. Li, Y. He, Z. Wang, S. Alqahtani, and P. Nanda, "Uncovering flaws in anti-phishing blacklists for phishing websites using novel cloaking techniques," in *Proceedings of the 20th International Conference on Security and Cryptography - Volume 1: SEC-CRYPT*, INSTICC. SciTePress, 2023, pp. 813–821.
- [21] T. Guo, X. Chen, Y. Wang, R. Chang, S. Pei, N. V. Chawla, O. Wiest, and X. Zhang, "Large language model based multi-agents: A survey of progress and challenges," *arXiv preprint arXiv:2402.01680*, 2024.
- [22] G. Li, H. A. Al Kader Hammoud, H. Itani, D. Khizbullin, and B. Ghanem, "Camel: communicative agents for "mind" exploration of large language model society," in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, ser. NIPS '23. Red Hook, NY, USA: Curran Associates Inc., 2023.
- [23] C.-M. Chan, W. Chen, Y. Su, J. Yu, W. Xue, S. Zhang, J. Fu, and Z. Liu, "Chateval: Towards better llm-based evaluators through multi-agent debate," *arXiv preprint arXiv:2308.07201*, 2023.
- [24] T. Liang, Z. He, W. Jiao, X. Wang, Y. Wang, R. Wang, Y. Yang, S. Shi, and Z. Tu, "Encouraging divergent thinking in large language models through multi-agent debate," in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Miami, Florida, USA: Association for Computational Linguistics, Nov. 2024, pp. 17 889–17 904.
- [25] N. T. Vy Nguyen, F. D. Childress, and Y. Yin, "Debate-driven multi-agent llms for phishing email detection," in *2025 13th International Symposium on Digital Forensics and Security (ISDFS)*, 2025, pp. 1–5.
- [26] Y. Xue, E. Spero, Y. S. Koh, and G. Russello, "Multiphishguard: An llm-based multi-agent system for phishing email detection," *arXiv preprint arXiv:2505.23803*, 2025.
- [27] R. Huang, "Improving spam detection with a multi-agent debate framework," in *2024 IEEE 4th International Conference on Information Technology, Big Data and Artificial Intelligence (ICIBA)*, vol. 4, 2024, pp. 1345–1349.
- [28] PhishDebate, "Supplementary materials," 2025. [Online]. Available: Tobeupddated
- [29] C. Bai *et al.*, "Qwen technical report," 2023. [Online]. Available: <https://arxiv.org/abs/2309.16609>
- [30] G. DeepMind, "Gemini: Advanced general-purpose ai," 2023. [Online]. Available: <https://deepmind.google/gemini>
- [31] OpenAI, "Gpt-4 technical report," 2023. [Online]. Available: <https://arxiv.org/abs/2303.08774>
- [32] S. Ariyadasa, S. Fernando, and S. Fernando, "Phishing websites dataset," Mendeley Data, V1, 2021, accessed: 2025-06-17. [Online]. Available: <https://doi.org/10.17632/n96ncsr5g4.1>
- [33] E. Basturk, "Ebbu2017 phishing dataset," <https://github.com/ebubekirbbr/pdd/tree/master/input>, 2017, accessed: 2025-06-17.
- [34] PhishTank, "PhishTank — Join the fight against phishing," <https://phishtank.org/>, accessed: 2025-06-17.
- [35] OpenPhish, "Openphish: Phishing intelligence," <https://openphish.com/>, accessed: 2025-06-17.
- [36] S. Ariyadasa, S. Fernando, and S. Fernando, "Phishrepo: A seamless collection of phishing data to fill a research gap in the phishing domain," *International Journal of Advanced Computer Science and Applications*, vol. 13, no. 5, 2022.
- [37] Y. Li, C. Huang, S. Deng, M. L. Lock, T. Cao, N. Oo, H. W. Lim, and B. Hooi, "KnowPhish: Large language models meet multimodal knowledge graphs for enhancing Reference-Based phishing detection," in *33rd USENIX Security Symposium (USENIX Security 24)*. Philadelphia, PA: USENIX Association, Aug. 2024, pp. 793–810. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity24/presentation/li-yuexin>
- [38] V. L. Pochat, T. V. Goethem, S. Tajalizadehkhoo, M. Korczynski, and W. Joosen, "Tranco: A Research-Oriented Top Sites Ranking Hardened Against Manipulation," in *Proceedings of the 2019 Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2019, accessed: 2025-06-17. [Online]. Available: <https://www.ndss-symposium.org/ndss-paper/tranco-a-research-oriented-top-sites-ranking-hardened-against-manipulation/>
- [39] J. Wei, X. Wang, D. Schuurmans, M. Bosma, b. ichter, F. Xia, E. Chi, Q. V. Le, and D. Zhou, "Chain-of-thought prompting elicits reasoning in large language models," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 24 824–24 837. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf

APPENDIX A
PROMPT TEMPLATES USED IN THIS STUDY

Prompt Template 2: HTML Structure Agent

Input: HTML content (H)

Prompt Template:

You are an expert in web security. Review the HTML structure of a webpage and determine if it exhibits suspicious structural characteristics typical of phishing sites. Focus on elements such as hidden forms, suspicious input fields, iframe usage, obfuscated JavaScript, and deceptive redirection patterns.

HTML: [TRUNCATED_HTML_CONTENT]

Provide your response in the following format:

- Claim: [Your assessment about the HTML structure indicating phishing or not]
- Confidence: [A score between 0 and 1]
- Evidence: [Relevant structural elements or tag patterns you found]

Fig. 4. Prompt template for HTML Structure Agent.

Prompt Template 3: Content Semantic Agent

Input: Visible text content (T)

Prompt Template:

You are a cybersecurity-focused language expert. Read the visible text content extracted from a webpage and decide whether the language indicates phishing intent. Look for emotionally manipulative language, requests for sensitive information, login instructions, urgency, or impersonation of known organizations.

Visible Text: [TRUNCATED_VISIBLE_TEXT]

Provide your response in the following format:

- Claim: [Whether the page language seems phishing-related]
- Confidence: [A score between 0 and 1]
- Evidence: [Specific words, phrases, or sentence patterns that support your claim]

Fig. 5. Prompt template for Content Semantic Agent.

Prompt Template 4: Brand Impersonation Agent

Input: URL (U) and visible text content (T)

Prompt Template:

You are a brand impersonation detection expert. Based on the URL and the HTML-visible content, evaluate whether this page attempts to impersonate a known brand. Focus on brand names, company references, login language, and any indications of misused identity (such as pretending to be Google, Apple, PayPal, etc.).

URL: [TARGET_URL]

Visible Text:

[TRUNCATED_VISIBLE_TEXT]

Provide response in following format:

- Claim: [Does the content attempt to impersonate a known brand?]
- Confidence: [A score between 0 and 1]
- Evidence: [Name(s) of impersonated brands and supporting context]

Fig. 6. Prompt template for Brand Impersonation Agent.

Prompt Template 5: Moderator Evaluation

Input: All specialist agent responses from current round

Prompt Template:

You are a moderator overseeing a phishing detection debate between specialist agents. After each round, you will evaluate all agents' analyses and determine if there is a clear consensus for a phishing determination. If so, please summarize your reasons for supporting the consensus and give the final answer that you think is correct, and the debate will conclude. If not, the debate will continue to the next round.

Fig. 7. Prompt template for Moderator Evaluation.

Prompt Template 6: Judge Decision

Input: Complete debate history from all specialist agents across all rounds

Prompt Template:

You are an expert cybersecurity judge evaluating phishing detection arguments. You will receive analyses from multiple specialized agents and need to make a final determination about whether the website is phishing or legitimate. Consider the evidence from all agents and their confidence levels. You must make a definitive decision between PHISHING or LEGITIMATE.

Fig. 8. Prompt template for Final Judgment.

Prompt Template 7: Single-Agent Baseline

Prompt Template:

You are an expert cybersecurity analyst specializing in phishing detection. Your task is to analyze websites and determine if they are phishing or legitimate.

You will be provided with:

1. URL of the website
2. HTML content of the website
3. Visible text content extracted from the website

Your analysis should consider:

- URL characteristics (suspicious domains, typosquatting, etc.)
- HTML structure and content
- Visual elements and branding
- Text content and messaging
- Technical indicators of phishing

You must respond with exactly one of these classifications:

- **PHISHING:** If the website is designed to deceive users or steal information
- **LEGITIMATE:** If the website appears to be genuine and trustworthy

Provide your classification followed by a brief explanation of your reasoning.

Fig. 9. Prompt template for Single-Agent Baseline.

Prompt Template 8: CoT Baseline

Prompt Template:

You are an expert cybersecurity analyst specializing in phishing detection. Your task is to analyze websites and determine if they are phishing or legitimate using a systematic Chain of Thought approach.

You will be provided with:

1. URL of the website
2. HTML content of the website
3. Visible text content extracted from the website

Please analyze the website step-by-step using the following Chain of Thought process:

STEP 1: URL ANALYSIS

- Examine the domain name for suspicious patterns
- Check for typosquatting (misspellings of legitimate brands)
- Look for suspicious TLDs or subdomains
- Identify any URL shortening or redirection indicators

STEP 2: CONTENT ANALYSIS

- Analyze the HTML structure and quality
- Look for suspicious scripts or hidden elements
- Check for legitimate branding vs. impersonation attempts
- Examine form elements and data collection practices

STEP 3: TEXT ANALYSIS

- Review the visible text for urgency tactics
- Check for grammar/spelling errors typical of phishing
- Look for legitimate contact information
- Analyze the overall messaging and tone

STEP 4: TECHNICAL INDICATORS

- Check for HTTPS usage and security indicators
- Look for suspicious redirects or external links
- Examine metadata and technical elements
- Consider overall website quality and professionalism

STEP 5: FINAL ASSESSMENT

- Weigh all evidence from previous steps
- Consider the overall risk profile
- Make a final classification with confidence level

Format your response as:

STEP 1: [Your URL analysis]

STEP 2: [Your content analysis]

STEP 3: [Your text analysis]

STEP 4: [Your technical analysis]

STEP 5: [Your final assessment]

CLASSIFICATION: [PHISHING or LEGITIMATE]

CONFIDENCE: [High/Medium/Low]

REASONING: [Brief summary of key factors that led to your decision]

Fig. 10. Prompt template for CoT Baseline