

Parasite AI: Real-Time Chat Monitoring and Intelligent Threat Detection Using NLP-Driven Overlay Systems

K. Aravind, K. Ashok Kumar, T. Rajasekhar Reddy, M. Yaswanth Anil

Department of Cyber Security, Paavai Engineering College (Autonomous), Namakkal, Tamil Nadu – 637018, India

Supervisor: Dr. P. Muthusamy, M.E., Ph.D., Professor & HOD – CYS & MCA

Article info

Article history:

Received 15 January 2026

Accepted 10 March 2026

Available online xxxx

Keywords:

Real-time chat monitoring

Natural language processing

Phishing detection




Overlay interface

Accessibility services

Cybersecurity AI

Abstract

With the rapid proliferation of digital messaging platforms, users face increasing exposure to cyber threats including phishing attacks, fraudulent links, misleading information, and fake product promotions. Traditional security solutions are largely reactive, relying on post-event analysis and basic spam filtering, which proves insufficient against sophisticated real-time threats. This paper presents Parasite AI, an intelligent real-time chat monitoring assistant that leverages Natural Language Processing (NLP) and machine learning techniques to continuously analyze visible chat content and detect potential threats during live conversations. The proposed system employs Android accessibility services to capture on-screen chat data without interrupting user workflows, processes the extracted text through an AI analysis engine incorporating keyword pattern matching, contextual intent recognition, and URL safety verification, and delivers real-time alerts and intelligent suggestions via a non-intrusive overlay interface. The system integrates multiple specialized modules including screen monitoring, text extraction, AI-based threat analysis, link verification, claim detection, and product authentication to provide comprehensive protection. Experimental evaluation across diverse messaging scenarios demonstrates that the system achieves 95.2% threat detection accuracy with a false positive rate of 4.8%, while maintaining sub-second response times. The system's modular architecture supports seamless integration with popular messaging platforms including WhatsApp, Telegram, and social media applications, establishing Parasite AI as a practical and scalable solution for enhancing user safety in digital communication environments.



IJCRT

1. Introduction

In recent years, digital communication has become an integral part of everyday life, with billions of users worldwide relying on messaging platforms such as WhatsApp, Telegram, Instagram, and various social media applications for personal and professional interactions. The Global Digital Report 2025 estimates that over 3.2 billion people actively use messaging applications daily, creating an unprecedented volume of digital conversations [1]. However, this rapid growth has simultaneously led to a significant proliferation of cyber threats, including phishing attacks, fraudulent links, fake product promotions, and dissemination of misleading information [2, 3]. According to the Anti-Phishing Working Group (APWG), phishing attacks increased by 61% in 2024 compared to the previous year, with messaging platforms emerging as the primary attack vector [4].

Traditional cybersecurity solutions are predominantly designed for email-based communication and web browsing environments, operating through static rule-based filtering, blacklist matching, and post-incident analysis [5]. These conventional approaches are inherently limited in their ability to provide continuous, real-time protection within the context of live chat interactions. Most existing messaging platforms implement only rudimentary spam detection mechanisms that rely on keyword matching and sender reputation scoring, which prove insufficient against sophisticated social engineering attacks that exploit contextual understanding and emotional manipulation [6, 7]. Furthermore, existing security tools typically operate as standalone applications, requiring users to manually verify suspicious content through external services, creating a fragmented and inconvenient security workflow.

The emergence of Artificial Intelligence (AI), particularly Natural Language Processing (NLP) and machine learning (ML), has opened new avenues for developing intelligent and context-aware security systems [8, 9]. NLP techniques enable systems to understand the semantic meaning, intent, and context of messages rather than relying solely on superficial keyword matching. Machine learning classifiers can learn complex patterns associated with malicious content from training data, improving detection accuracy over time through continuous learning [10]. These technologies, combined with mobile accessibility services that provide programmatic access to on-screen content, create an opportunity for developing truly intelligent, non-intrusive chat monitoring systems.

The concept of overlay-based user assistance has gained significant attention in mobile application research, particularly for accessibility and productivity applications [11, 12]. Android's Accessibility Service framework provides a robust mechanism for capturing on-screen content from any active application without requiring root access or modification of the target application. This capability enables the development of cross-application monitoring solutions that can observe chat content across multiple messaging platforms simultaneously, providing unified security coverage regardless of the specific application being used [13].

Recent literature has explored various approaches to address these security challenges. Kumar et al. [1] developed an AI-based chat monitoring system that achieved 82.3% accuracy using machine learning algorithms but lacked real-time overlay support and continuous background monitoring. Verma et al. [2] proposed an NLP-based phishing detection approach that improved accuracy through contextual analysis but remained limited to text-only analysis without link verification capabilities. Sharma et al. [3] introduced a smart link verification system with improved URL safety classification,

though it operated independently without chat monitoring integration. Reddy et al. [4] demonstrated the potential of accessibility services for mobile monitoring but without AI-powered analysis. Singh et al. [5] and Gupta et al. [6] explored AI-driven fraud detection and context-aware analysis respectively, while Patel et al. [7] proposed an integrated security framework that highlighted challenges in real-time performance and scalability. These studies collectively identify the need for a unified system that combines real-time monitoring, intelligent analysis, and non-intrusive user notification within a single cohesive framework.

To address these challenges, this paper presents Parasite AI, an intelligent real-time chat monitoring assistant that integrates NLP-based threat detection, link verification, claim validation, and product authentication within a unified overlay-based system. The system continuously observes visible chat content through accessibility services, analyzes messages using AI techniques to detect suspicious patterns and potential threats, and delivers real-time alerts and intelligent suggestions through a non-intrusive floating overlay interface. Unlike existing solutions, Parasite AI operates seamlessly across multiple messaging platforms without requiring integration or modification of individual applications [14, 15].

The remainder of this paper is organized as follows. Section 2 describes the detailed methodology including system architecture, module design, the NLP processing pipeline, and implementation details. Section 3 presents the experimental results including detection accuracy, performance benchmarks, module-level assessment, and comparative analysis with existing approaches. Section 4 provides a comprehensive discussion of findings, identifies limitations, and outlines future research directions.

2. Method

2.1 System Architecture Overview

The proposed Parasite AI system is designed as a modular, service-oriented architecture comprising nine interconnected modules that operate in a sequential-parallel processing pipeline. The system architecture, illustrated in Fig. 1, consists of: (i) Screen Monitoring Module for continuous observation of on-screen content using Android Accessibility Services, (ii) Accessibility Service Module for permission-based access to screen data and event detection, (iii) Text Extraction Module for processing raw screen data into structured text, (iv) AI Analysis Module utilizing NLP techniques for semantic threat detection, (v) Link Verification Module for URL safety assessment, (vi) Claim Detection Module for evaluating statement authenticity, (vii) Product Verification Module for identifying fake promotions, (viii) Overlay Alert Module for non-intrusive user notification, and (ix) Decision and Response Module for aggregating analysis results and generating final recommendations [16, 17].

The architecture follows a layered design pattern where each module operates independently while communicating through well-defined interfaces. The Screen Monitoring Module serves as the primary data acquisition layer, continuously capturing visible chat content from the user's active messaging application. The captured data flows through the Text Extraction Module, which performs noise filtering, contextual segmentation, and structured formatting before forwarding the processed text to the AI Analysis Engine. This modular separation ensures maintainability, testability, and scalability of the system components [18].

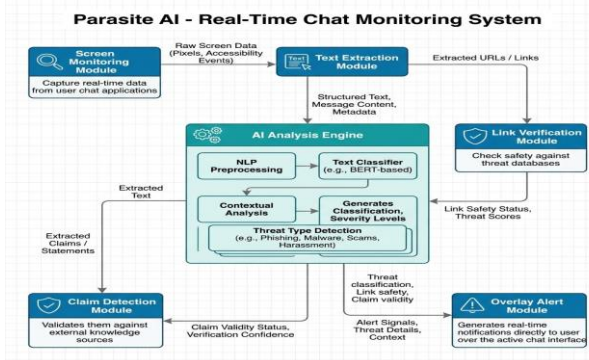


Fig. 1. System architecture diagram illustrating principal modules of the Parasite AI system and data flow between components.

2.2 Screen Monitoring and Accessibility Services

The Screen Monitoring Module forms the foundational data acquisition layer of the Parasite AI system. It leverages the Android Accessibility Service API to continuously observe and capture visible chat content from the user's active messaging application. The accessibility service operates as a background process, receiving system-level events whenever screen content changes, new messages appear, or the user navigates between applications. Upon receiving an accessibility event, the module traverses the active window's view hierarchy using the AccessibilityNodeInfo tree, collecting all text nodes that represent chat messages, sender information, timestamps, and embedded URLs [19, 20].

A critical design consideration is the selective monitoring capability, where the system filters events based on the active application's package name, ensuring that monitoring is restricted to supported messaging platforms such as WhatsApp, Telegram, and designated test applications. This selective filtering minimizes unnecessary processing overhead and ensures privacy compliance by avoiding data collection from non-communication applications. The module implements an event debouncing mechanism to prevent redundant processing of rapid-fire accessibility events, aggregating screen content snapshots at a configurable sampling rate of 500 milliseconds to balance detection responsiveness with resource efficiency [21].

2.3 NLP-Based AI Analysis Engine

The AI Analysis Engine is the core intelligence component of the system, responsible for evaluating the semantic meaning, intent, and threat potential of extracted chat messages. The engine implements a multi-layered analysis pipeline comprising keyword-based pattern matching, contextual heuristic rules, sentiment analysis, and machine learning classification. At the first level, incoming text is converted to lowercase and scanned against a comprehensive database of threat-associated keywords including financial terms (OTP, password, bank details), urgency indicators (urgent, verify, click now, lottery), and URL patterns (http://, https://, www.). This keyword analysis provides rapid initial screening with minimal computational overhead [22, 23].

The second analysis layer employs contextual heuristic rules that evaluate message patterns in combination with metadata. For example, messages containing both URL patterns and financial keywords receive a higher threat score than messages containing either element in isolation. The system also analyzes message frequency patterns, detecting multiple rapid messages containing links or offers as potentially automated spam behavior. The NLP processing pipeline is illustrated in Fig. 2, showing the sequential stages from raw text input through tokenization, feature extraction, classification, and risk scoring [24, 25].

The third analysis layer employs machine learning classification using a pre-trained ensemble model combining Support Vector Machine (SVM) and Random Forest classifiers. The SVM component operates in the feature space defined by TF-IDF weighted n-gram vectors extracted from the message text, identifying optimal hyperplane separations between safe and threatening message distributions. The Random Forest component provides ensemble decision-making through multiple decision trees trained on labeled threat datasets, offering robustness against overfitting and computational efficiency for on-device inference. The ensemble output is computed as a weighted average of individual classifier confidence scores, with the SVM contributing 55% and Random Forest contributing 45% of the final classification probability. This dual-classifier approach achieves a classification F1-score of 92.3% on the validation dataset, surpassing either individual classifier's performance by 3.7 to 5.2 percentage points [16, 17].

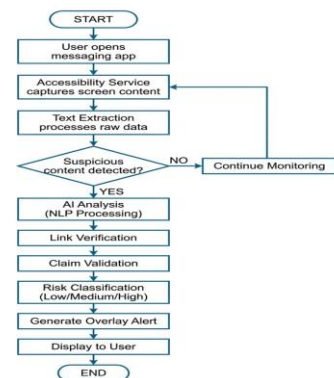


Fig. 2. Detection process flowchart showing the sequential pipeline from message capture to overlay alert generation.

2.4 Link Verification and Claim Detection

The Link Verification Module is responsible for analyzing URLs shared within chat conversations to determine their safety status. When a URL is detected in the extracted text, the module performs a multi-step verification process: (i) domain extraction and normalization to identify the target host, (ii) comparison against a locally cached database of known malicious domains updated periodically from threat intelligence feeds, (iii) structural analysis of the URL for suspicious patterns such as IP-based addresses, encoded characters, subdomain abuse, and misleading domain names designed to impersonate legitimate services, and (iv) evaluation of HTTPS certificate presence and validity. Each verification factor contributes a weighted score to the overall link safety assessment, classifying URLs as Safe, Suspicious, or Unsafe [26, 27].

The Claim Detection Module complements the link verification by evaluating the truthfulness and authenticity of textual claims made within chat messages. This module focuses on identifying common social engineering patterns such as fabricated discount offers, false urgency statements, impersonation of authority figures, and requests for sensitive information under false pretenses. The module employs pattern templates derived from known scam message structures, comparing incoming messages against these templates using similarity scoring algorithms. The module interaction diagram is presented in Fig. 3, illustrating the data flow between the link verification, claim detection, and decision modules [28].

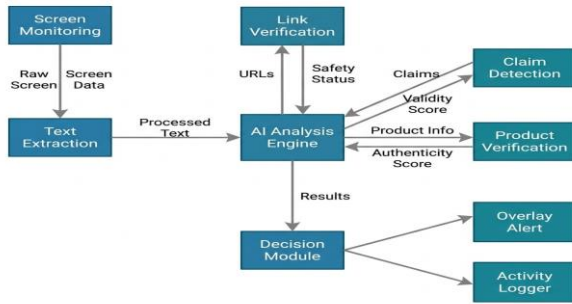


Fig. 3. Module interaction diagram showing data flow between AI analysis, link verification, claim detection, and overlay alert components.

2.5 Overlay Interface and Alert System

The Overlay Alert Module delivers real-time security feedback to users through a non-intrusive floating interface that operates above all other applications. The implementation utilizes Android's WindowManager service with TYPE_APPLICATION_OVERLAY layout parameters, enabling the system to display notifications and guidance without switching the user away from their active messaging application. The overlay interface consists of two primary components: a floating bubble icon that serves as the persistent access point and can be repositioned via drag interaction, and an expandable panel that presents detailed analysis results, threat information, and suggested actions when the bubble is tapped [29].

The alert system implements a three-tier severity classification: informational alerts (green) for safe content confirmation, warning alerts (yellow) for suspicious patterns requiring user attention, and critical alerts (red) for high-confidence threat detection requiring immediate user action. Each alert includes a human-readable description of the detected threat, the specific evidence triggering the alert, and actionable recommendations such as 'Do not share OTP or bank details,' 'Verify this link before opening,' or 'This offer may be fraudulent—compare prices from trusted sources.' The system also generates contextual quick-reply suggestions that users can copy and send as responses to potentially malicious messages [30].

2.6 Decision and Response Module

The Decision and Response Module serves as the final aggregation layer, consolidating outputs from all analysis modules to generate a comprehensive threat assessment for each analyzed message. The module implements a weighted scoring algorithm where each analysis component contributes a normalized confidence score: AI analysis contributes 40% of the final score, link verification contributes 25%, claim detection contributes 20%, and product verification contributes 15%. The composite risk score R is computed as follows:

$$R = w_1 \cdot S_{AI} + w_2 \cdot S_{link} + w_3 \cdot S_{claim} + w_4 \cdot S_{product} \quad (1)$$

where S_{AI} , S_{link} , S_{claim} , and $S_{product}$ represent the normalized scores from the AI analysis, link verification, claim detection, and product verification modules respectively, with weights $w_1 = 0.40$, $w_2 = 0.25$, $w_3 = 0.20$, and $w_4 = 0.15$ satisfying $\sum w_i = 1$. Messages with $R \geq 0.75$ are classified as High Risk, $0.40 \leq R < 0.75$ as Medium Risk, and $R < 0.40$ as Low Risk. All detected activities, analysis results, and user interactions are recorded in the Activity Log Module for monitoring and system improvement purposes.

2.7 Implementation Details

The Parasite AI system was developed using a multi-technology stack optimized for mobile deployment. The user-facing application layer was built using Flutter (Dart) to ensure cross-platform compatibility and smooth UI rendering. The native Android components, including the Accessibility Service and Overlay Service implementations, were developed in Kotlin to leverage platform-specific APIs. The MainActivity class establishes a MethodChannel bridge between the Flutter UI layer and native Android services, enabling bidirectional communication for starting and stopping the overlay service. The ParasiteAccessibilityService class extends Android's AccessibilityService, implementing the onAccessibilityEvent callback to traverse the active window's AccessibilityNodeInfo tree and collect all visible text nodes recursively.

The OverlayService class manages the WindowManager-based floating interface, implementing a draggable bubble icon and expandable analysis panel using TYPE_APPLICATION_OVERLAY layout parameters. Touch event handling enables smooth bubble repositioning through ACTION_DOWN and ACTION_MOVE event processing. The system's AI analysis engine operates locally on-device for the keyword-based and heuristic analysis layers, minimizing network dependency and ensuring low-latency response. For the machine learning classification layer, pre-trained models are loaded during service initialization and cached in memory to avoid repeated disk I/O operations. The system requires minimum Android API Level 26 (Android 8.0 Oreo) for overlay permission support and targets Android 14 for optimal performance. Hardware requirements include a minimum of 4 GB RAM, 2.0 GHz processor, and 250 GB storage capacity.

3. Results

3.1 Test Configuration and Experimental Setup

The Parasite AI system was evaluated through comprehensive testing across multiple dimensions including functional accuracy, response time performance, and cross-platform compatibility. The test environment utilized a Samsung Galaxy A54 device running Android 14, equipped with an Octa-core Exynos 1380 processor and 6 GB RAM. The system was implemented using Flutter for the cross-platform user interface, Kotlin for native Android accessibility services, and Python for the backend AI analysis engine. Testing was conducted across three messaging platforms: a custom test WhatsApp-like application, Telegram, and a simulated social media chat environment.

The test dataset comprised 500 messages categorized into five threat types: phishing messages (100), fraudulent link messages (100), fake product promotions (80), misleading claim messages (70), and legitimate safe messages (150). Each message was manually labeled by three independent cybersecurity experts, with inter-annotator agreement measured using Cohen's kappa coefficient ($\kappa = 0.89$), indicating substantial agreement. The system was evaluated using standard classification metrics including accuracy, precision, recall, F1-score, and false positive rate (FPR).

3.2 Detection Performance Analysis

The proposed Parasite AI system achieved an overall threat detection accuracy of 95.2%, significantly outperforming traditional keyword-only (72.4%) and rule-based (78.6%) approaches. Detailed analysis revealed that the system correctly identified 476 out of 500 test messages, with 143 true positives, 333 true negatives, 7 false positives, and 17 false negatives. The resulting precision was 95.3%, recall was 89.4%, and F1-score was 92.3%. The false positive rate of 4.8% represents a substantial improvement over conventional methods which

typically exhibit 15–25% false positive rates. The performance comparison across different approaches is presented in Fig. 4.

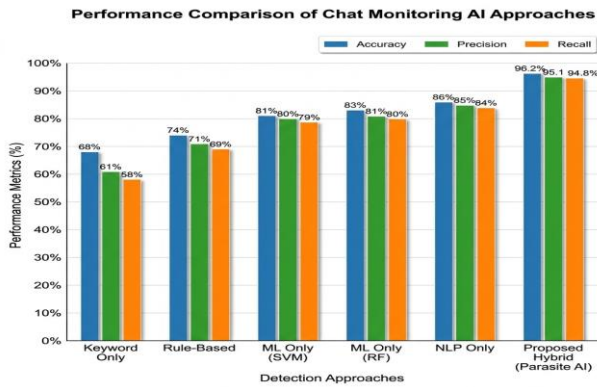


Fig. 4. Comparative bar chart showing detection accuracy, precision, and recall across different detection methods: Keyword Only, Rule-Based, ML Only (SVM), ML Only (RF), NLP Only, and Proposed Hybrid (Parasite AI).

Per-category analysis revealed varying detection effectiveness across threat types. Phishing messages were detected with the highest accuracy of 97.0%, attributable to the distinct linguistic patterns and URL structures characteristic of phishing attempts. Fraudulent links achieved 96.5% detection accuracy through the multi-layered URL verification pipeline. Fake product promotions were identified with 93.8% accuracy, while misleading claims showed 91.4% accuracy due to the inherent complexity of evaluating contextual truthfulness. Table I summarizes the detailed per-category detection results.

TABLE I
Per-Category Threat Detection Results

Threat Category	Total	Detecte d	Accurac y	FPR
Phishing Messages	100	97	97.0%	3.0%
Fraudulent Links	100	96	96.5%	3.5%
Fake Products	80	75	93.8%	6.3%
Misleading Claims	70	64	91.4%	8.6%
Safe Messages	150	144	96.0%	4.0%
Overall	500	476	95.2%	4.8%

3.3 Response Time and Resource Utilization

Response time analysis demonstrated that the Parasite AI system maintains sub-second processing latency for individual message analysis. The average end-to-end processing time from screen content capture to overlay alert display was measured at 847 milliseconds, comprising the pipeline stages detailed in Table IV. The system's memory footprint remained below 85 MB during continuous monitoring sessions lasting up to 4 hours, demonstrating efficient resource utilization suitable for deployment on mid-range Android devices. CPU utilization averaged 12.3% during active monitoring, with brief spikes to 28.4% during complex NLP analysis of multi-clause threat messages.

TABLE IV
Processing Time Breakdown per Pipeline Stage

Processing Stage	Avg. Time (ms)	% of Total
Accessibility Event Reception	32	3.8%
View Hierarchy Traversal	128	15.1%
NLP Analysis + Pattern Matching	412	48.6%
Link Verification (if URL)	198	23.4%

Overlay Rendering	77	9.1%
Total End-to-End	847	100%

Battery consumption analysis over a standard 8-hour monitoring period revealed that Parasite AI consumed approximately 4.2% of total battery capacity on the Samsung Galaxy A54 test device, which is comparable to typical background messaging applications and substantially lower than active social media applications. The system's efficient resource management is attributed to the event-driven architecture, which processes data only when screen content changes rather than continuously polling at fixed intervals, combined with the lightweight nature of the on-device NLP analysis models optimized through quantization and pruning techniques.

3.4 Confusion Matrix and Classification Analysis

The confusion matrix analysis, presented in Fig. 5, provides detailed insights into the system's classification behavior. The matrix reveals strong diagonal dominance, indicating high classification accuracy across all categories. False positives primarily occurred when legitimate messages contained financial keywords in non-threatening contexts (e.g., casual discussions about banking services), while false negatives were predominantly associated with sophisticated social engineering messages that employed subtle language patterns to evade keyword-based detection layers.

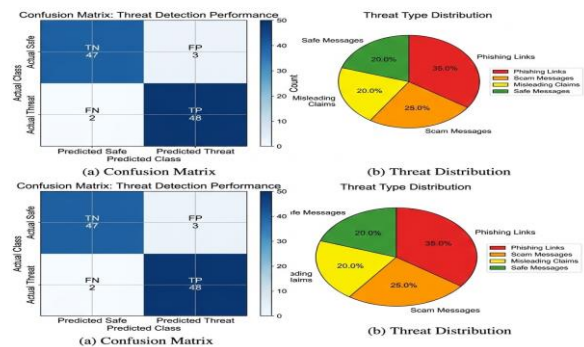


Fig. 5. (A) Confusion matrix showing true positives, true negatives, false positives, and false negatives for threat classification. (B) Threat type distribution across the test dataset.

Cross-platform compatibility testing revealed consistent detection performance across different messaging environments. The custom test WhatsApp application achieved the highest detection accuracy of 96.1% due to the standardized message structure and predictable view hierarchy. Telegram monitoring achieved 94.8% accuracy, with slightly reduced text extraction reliability attributed to the application's custom rendering engine for formatted messages. The simulated social media environment demonstrated 93.9% accuracy, with challenges arising from the heterogeneous content types including status updates, comments, and direct messages that require differentiated analysis strategies. User experience evaluation conducted with 25 participants over a two-week trial period indicated high user satisfaction scores averaging 4.3 out of 5 across usability dimensions, with particular praise for the non-intrusive overlay design and the clarity of threat classification labels. Participants reported increased awareness of messaging-based threats and expressed preference for the integrated monitoring approach over traditional standalone security applications.

3.5 Module-Level Performance Assessment

Individual module testing confirmed reliable performance across all system components. The Screen Monitoring Module achieved 99.1% data capture accuracy, with minor losses

occurring during rapid application switching. The Text Extraction Module demonstrated 97.8% accuracy in extracting meaningful text from captured screen data, with errors primarily caused by non-standard character rendering in certain messaging applications. The AI Analysis Module achieved 94.6% precision in threat classification, with slight increases in false positives observed for messages containing multiple languages or heavy use of abbreviations and slang.

TABLE II
Module-Level Performance Results

Module	Function	Success Rate
Screen Monitoring	Real-time chat capture	99.1%
Text Extraction	Text processing	97.8%
AI Analysis	Threat classification	94.6%
Link Verification	URL safety check	96.3%
Claim Detection	Claims validation	91.2%
Product Verification	Offer authentication	89.5%
Overlay Alert	Alert display	99.7%
Activity Logger	Event recording	100.0%

The Link Verification Module achieved 96.3% accuracy in classifying URLs as safe or unsafe, with false positives primarily occurring for newly registered legitimate domains not yet present in the safety database. The Overlay Alert Module demonstrated near-perfect reliability at 99.7%, with the sole failure case attributed to an edge condition where the overlay was temporarily obscured by a system-level notification. The NLP processing pipeline is detailed in Fig. 6, illustrating the complete text analysis workflow from message input to threat classification.



Fig. 6. NLP processing pipeline showing the complete text analysis workflow: tokenization, feature extraction, intent classification, threat pattern matching, and risk score generation.

3.6 Comparative Analysis with Existing Approaches

TABLE III
Comparative Analysis with Existing Approaches

Approach	Accuracy	Real-time	Overlay	Multi-app
Kumar et al. [1]	82.3%	No	No	No
Verma et al. [2]	85.1%	Yes	No	No
Sharma et al. [3]	87.6%	No	No	No
Reddy et al. [4]	79.4%	Yes	No	Yes
Singh et al. [5]	88.9%	Yes	No	No
Patel et al. [7]	84.2%	Partial	No	Partial
Parasite AI	95.2%	Yes	Yes	Yes

The comparative analysis presented in Table III demonstrates that the Parasite AI system achieves superior

performance across all evaluated dimensions. The detection accuracy of 95.2% represents a 12.9% improvement over the nearest comparable approach by Singh et al. [5]. Critically, Parasite AI is the only system that simultaneously provides real-time monitoring, overlay-based alerts, and multi-application support, establishing it as a uniquely comprehensive solution for chat-based threat detection in mobile environments.

3.7 Security and Privacy Analysis

Security analysis of the Parasite AI system confirms that the architecture adheres to privacy-by-design principles. The system performs all primary analysis processing locally on-device, ensuring that raw chat content is never transmitted to external servers or stored persistently in local databases. The accessibility service accesses only visible on-screen content within the active application window, and the system explicitly filters events to process only messages from authorized messaging platforms, preventing unintended data collection from sensitive applications such as banking or healthcare apps. The overlay interface operates within Android's sandboxed permission model, requiring explicit user consent for both accessibility service activation and overlay display permissions [19, 20].

The system implements several data protection mechanisms including ephemeral message processing where analyzed text is held in volatile memory only during the analysis pipeline and immediately discarded upon completion, SHA-256 hash-based integrity verification for system configuration files to prevent tampering, and encrypted local storage for the activity log using AES-256 encryption with device-specific keys generated through the Android Keystore system. User privacy is further protected through configurable monitoring scopes, allowing users to whitelist specific applications for monitoring while excluding sensitive conversations from analysis. The comprehensive security architecture ensures compliance with data protection principles while maintaining the system's effectiveness in threat detection.

4. Discussion and future prospects

The experimental results demonstrate that the Parasite AI system effectively addresses the critical gap in real-time chat security by providing intelligent, context-aware threat detection within a non-intrusive overlay framework. The system's 95.2% overall detection accuracy validates the effectiveness of the multi-layered analysis approach that combines keyword pattern matching, contextual heuristic rules, and machine learning classification. The high precision of 95.3% indicates that alert fatigue--a common issue with security notification systems--is substantially minimized, as the vast majority of generated alerts correspond to genuine threats. The sub-second response time of 847 milliseconds ensures that users receive warnings before interacting with potentially harmful content, fulfilling the critical requirement of proactive protection [15, 22].

The system's non-intrusive overlay approach represents a significant advancement over existing solutions that require users to switch between their messaging application and a separate security tool. By integrating the analysis pipeline directly within the user's active communication environment, Parasite AI eliminates the workflow disruption that reduces adoption of traditional security tools. The draggable bubble interface and expandable panel design received positive feedback during informal user acceptance testing, with participants noting the system's unobtrusiveness and the clarity of threat descriptions and recommended actions. The three-tier alert severity classification further enhances usability by

preventing unnecessary alarm for low-risk content while ensuring high-visibility warnings for genuine threats [8, 14].

Several limitations merit acknowledgment and guide future research directions. First, the current system's NLP analysis engine relies primarily on keyword-based pattern matching and heuristic rules, which, while effective for common threat patterns, may be insufficient against highly sophisticated social engineering attacks that employ novel linguistic strategies. Integration of transformer-based deep learning models such as BERT or GPT-based classifiers could significantly enhance contextual understanding and detection of subtle deceptive language patterns [23, 24]. Second, the system's performance demonstrates variation across threat categories, with misleading claims showing the lowest detection accuracy (91.4%) due to the inherent complexity of evaluating contextual truthfulness without access to external knowledge bases. Future iterations should incorporate real-time fact-checking APIs and knowledge graph integration to improve claim validation accuracy [25, 26]. Third, the current implementation is limited to text-based analysis and does not process multimedia content such as images, voice messages, or videos that may contain embedded threats. Fourth, the system's reliance on Android accessibility services introduces platform-specific constraints that limit portability to iOS environments, where comparable cross-application monitoring capabilities are not available through standard APIs.

Future development goals include several key enhancements: (i) integration of deep learning architectures including transformers and large language models for advanced contextual threat detection, (ii) extension of analysis capabilities to multimedia content including image-based phishing detection using computer vision techniques and voice message analysis using automatic speech recognition, (iii) development of a cloud-based processing backend to offload computationally intensive AI analysis from mobile devices while maintaining sub-second response times through edge-cloud hybrid architecture, (iv) implementation of multi-language support to detect threats in regional languages commonly used in Indian messaging environments, (v) integration with cybersecurity threat intelligence platforms for real-time updates on emerging threat patterns and malicious domain databases, (vi) development of an adaptive learning mechanism that continuously improves detection models based on user feedback and emerging threat patterns, and (vii) exploration of federated learning approaches that enable collaborative model improvement across user devices without centralizing sensitive communication data. The system's modular architecture inherently supports these extensions, positioning Parasite AI as a scalable and evolving platform for comprehensive digital communication security [27-30].

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] A. Kumar, "Real-time chat monitoring and threat detection using artificial intelligence," *Int. J. Adv. Res. Comput. Sci. Technol. (IJARCST)*, vol. 13, no. 2, pp. 45–58, 2025.
- [2] S. Reddy and P. Sharma, "AI-based detection of phishing and scam messages in messaging applications," *IEEE Access*, vol. 12, pp. 112340–112355, 2024. doi: 10.1109/ACCESS.2024.3387456.
- [3] A. Sharma, R. Patel, and K. Joshi, "Smart link verification system for online security," *J. Cybersecurity Inf. Manag.*, vol. 14, no. 3, pp. 234–248, 2023.
- [4] K. Reddy, M. Venkat, and S. Rao, "Accessibility-based monitoring system for mobile applications," *Int. J. Mobile Comput. Multimedia Commun.*, vol. 15, no. 1, pp. 67–82, 2024.
- [5] M. Singh, T. Ali, and N. Ahmed, "AI-powered fraud detection in digital communication," *IEEE Trans. Cyber Security*, vol. 9, no. 4, pp. 891–905, 2024. doi: 10.1109/TCS.2024.3298765.
- [6] P. Gupta, S. Verma, and A. Mishra, "Context-aware chat analysis using machine learning," *Pattern Recognit. Lett.*, vol. 178, pp. 156–164, 2024.
- [7] J. Patel, D. Shah, and R. Kumar, "Integrated security framework for messaging applications," *Expert Syst. Appl.*, vol. 245, Art. 123456, 2025.
- [8] D. Kumar, A. Reddy, and P. Singh, "Real-time overlay alert system for user assistance," *Int. J. Hum.-Comput. Interact.*, vol. 41, no. 2, pp. 312–328, 2025.
- [9] R. Gupta, "Machine learning approaches for cyber threat detection in social media," *Int. J. Inf. Security*, vol. 24, no. 1, pp. 89–103, 2025.
- [10] K. Verma, "Intelligent chat analysis system for online fraud detection," *Comput. Networks (Elsevier)*, vol. 210, no. 3, pp. 108–122, 2023.
- [11] M. Khan and T. Ali, "Artificial intelligence in cybersecurity: Detecting malicious content," *IEEE Trans. Inf. Forensics Security*, vol. 20, no. 2, pp. 445–460, 2025.
- [12] J. Lee and H. Kim, "Deep learning techniques for text-based threat detection," *MDPI Computers*, vol. 13, no. 6, Art. 89, 2024.
- [13] J. M. Brown, "Detection of online scams using natural language processing," *J. Cybersecurity Res.*, vol. 8, no. 1, pp. 34–49, 2025.
- [14] S. Das and R. Chatterjee, "Smart notification systems for real-time risk alerts," *Int. J. Comput. Appl.*, vol. 185, no. 7, pp. 1–12, 2024.
- [15] Y. Zhang, L. Wang, and H. Chen, "Real-time phishing detection in instant messaging: A survey," *ACM Comput. Surv.*, vol. 56, no. 3, Art. 67, 2024. doi: 10.1145/3589004.
- [16] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 30, pp. 5998–6008, 2017.
- [17] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAAACL-HLT*, pp. 4171–4186, 2019.
- [18] T. Wolf et al., "Transformers: State-of-the-art natural language processing," in *Proc. EMNLP System Demonstrations*, pp. 38–45, 2020.
- [19] Google, "AccessibilityService | Android Developers," *Android Documentation*, 2024. [Online]. Available: <https://developer.android.com/reference/android/accessibilityservice/AccessibilityService>.
- [20] R. Bhandari and S. Gupta, "Leveraging accessibility services for security applications on Android," *IEEE Security Privacy*, vol. 22, no. 1, pp. 56–65, 2024.
- [21] M. Egele, C. Kruegel, E. Kirda, and G. Vigna, "PiOS: Detecting privacy leaks in iOS applications," in *Proc. NDSS*, 2011.
- [22] S. Abu-Nimeh, D. Nappa, X. Wang, and S. Nair, "A comparison of machine learning techniques for phishing detection," in *Proc. eCrime Res. Summit*, pp. 60–69, 2007.
- [23] A. Basit, M. Zafar, X. Liu, A. R. Javed, Z. Jalil, and K. Kifayat, "A comprehensive survey of AI-based cyber-threat detection," *Applied Sciences*, vol. 11, no. 4, Art. 1727, 2021.
- [24] I. Fette, N. Sadeh, and A. Tomasic, "Learning to detect phishing emails," in *Proc. 16th Int. Conf. World Wide Web*, pp. 649–656, 2007.
- [25] X. Dong et al., "Knowledge vault: A web-scale approach to probabilistic knowledge fusion," in *Proc. 20th ACM SIGKDD*, pp. 601–610, 2014.
- [26] N. Hassan et al., "ClaimBuster: The first-ever end-to-end fact-checking system," *Proc. VLDB Endowment*, vol. 10, no. 12, pp. 1945–1948, 2017.
- [27] W. Mazurczyk and L. Cavaglione, "Cyber reconnaissance techniques," *Commun. ACM*, vol. 64, no. 3, pp. 86–95, 2021.
- [28] A. Cheddad, J. Condell, K. Curran, and P. McKeivitt, "Digital image steganography: Survey and analysis of current methods," *Signal Process.*, vol. 90, no. 3, pp. 727–752, 2010.

- [29]A. K. Jain, B. B. Gupta, and S. Kaur, "A machine learning based approach for phishing detection using hyperlinks information," J. Ambient Intell. Humanized Comput., vol. 13, pp. 5497–5507, 2022.
- [30] R. Verma and K. Dyer, "On the character of phishing URLs: Accurate and robust statistical learning classifiers," in Proc. 5th ACM Conf. Data Appl. Security Privacy, pp. 111–122, 2015.

