

Client-Side AI-Driven Cybersecurity for iOS Applications using Swift and Core ML

Madhuri Latha Gondi
Mobile Lead Developer
13579 S Langdon Dr
Riverton, Utah - 84096

ABSTRACT

Mobile applications are increasingly used for sensitive interactions such as financial transactions, healthcare communication, and enterprise operations. This widespread adoption has expanded the client-side attack surface, exposing users to phishing URLs, malicious redirects, and deceptive navigation behaviors. Conventional cybersecurity mechanisms are largely server-centric and often fail to provide immediate protection against real-time or zero-day threats occurring within mobile application runtimes.

This paper presents a client-side, AI-driven cybersecurity framework designed specifically for iOS applications using Swift and Apple's Core ML framework. The proposed approach embeds lightweight machine learning models directly within the mobile application to detect phishing URLs and malicious navigation events in real time. Model optimization techniques are applied to ensure low-latency inference suitable for mobile environments. Experimental evaluation demonstrates high detection accuracy with minimal performance overhead, enabling proactive and privacy-preserving threat mitigation without continuous backend dependency.

General Terms

Mobile Computing, iOS Mobile Applications, Security

Keywords

iOS Security, Client-Side AI, Swift, Core ML, Mobile Cybersecurity, Phishing Detection

1. INTRODUCTION

The rapid growth of mobile ecosystems has transformed smartphones into critical endpoints for financial services, healthcare platforms, enterprise workflows, and personal communication. As a result, mobile applications now handle highly sensitive user data and security-critical interactions, making them an increasingly attractive target for attackers. Unlike traditional desktop environments, mobile platforms rely heavily on embedded web views, deep links, and dynamic navigation patterns, which introduce unique attack vectors that are often difficult to monitor using conventional security approaches.

Recent attack strategies increasingly leverage social engineering techniques combined with malicious URLs and deceptive navigation behaviors to bypass traditional perimeter defenses. Phishing attacks delivered through in-app browsers, embedded web content, or compromised redirection chains are particularly challenging to detect in real time. These attacks exploit the trust relationship between users and legitimate mobile applications, often without triggering server-side alerts or signature-based detection mechanisms.

While server-centric security architectures remain an essential component of enterprise defense strategies, they are inherently limited in their ability to provide immediate protection against client-side threats. Network latency, intermittent connectivity, and dependency on backend analysis reduce their effectiveness in time-critical scenarios. Moreover, transmitting detailed navigation or behavioral data to remote servers raises privacy concerns and may conflict with modern data-protection regulations and platform policies enforced by mobile operating systems.

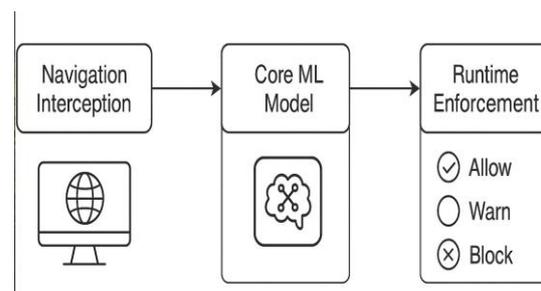
Advances in on-device machine learning and mobile hardware acceleration provide new opportunities to address these challenges directly within the application runtime. Frameworks such as Apple's Core ML enable efficient execution of lightweight models on mobile devices, allowing real-time inference with minimal performance overhead. By shifting threat detection to the client side, mobile applications can respond instantly to suspicious behavior while preserving user privacy and reducing reliance on continuous backend connectivity.

This work explores the design and implementation of a client-side, AI-driven cybersecurity framework for iOS applications that leverages on-device machine learning to detect phishing URLs and malicious navigation events. The proposed approach demonstrates how integrating intelligent security mechanisms directly into mobile applications can enhance resilience against emerging threats while maintaining performance, scalability, and privacy guarantees.

2. SYSTEM ARCHITECTURE

The proposed cybersecurity framework operates entirely on the iOS client device. User navigation and content-loading events are intercepted and analyzed locally using a Core ML model. Based on the computed risk score, the application enforces security decisions such as allowing normal navigation, warning the user, or blocking the interaction entirely.

Figure 1 represents the end-to-end execution of the proposed client-side cybersecurity framework within an iOS application. The system processes real navigation events and applies real-time security decisions based on machine-learning inference.



2.1 Example Runtime Flow

Stage	Input Data	Processing	Output
Navigation Interception	Any URL	URL captured from WebView or navigation handler	Raw URL string
Feature Extraction	URL string	Length, entropy, subdomain count, keyword presence	Feature vector
Core ML Inference	Feature vector	On-device Core ML prediction	Risk score = 0.91
Runtime Enforcement	Risk score	Policy evaluation	Blocked

This data flow demonstrates how suspicious navigation events are intercepted, analyzed locally using Core ML, and mitigated immediately without relying on backend infrastructure.

2.2 Example Decision Mapping

To further clarify the runtime behavior illustrated in Figure 1, Table 2 shows how different risk scores produced by the Core ML model result in corresponding enforcement actions.

Risk Score Range	Security Action	User Impact
0.00 – 0.59	Allow	Normal navigation
0.60 – 0.84	Warn	Security warning dialog
≥ 0.85	Block	Navigation prevented

For example, when a navigation request produces a risk score of 0.91, the framework immediately blocks the interaction, preventing the user from accessing potentially malicious content.

3. METHODOLOGY

This section describes the design, implementation, and execution pipeline of the proposed client-side AI-driven cybersecurity framework for iOS applications. The methodology is designed to operate entirely on-device, ensuring real-time threat detection, low latency, and user privacy preservation. On-device intelligence has been shown to be effective for mobile and edge security applications while reducing reliance on centralized infrastructure [3], [9]. The framework follows a four-stage pipeline: navigation interception, feature extraction, machine learning inference, and runtime decision enforcement.

3.1 Navigation Interception

The first stage of the methodology involves intercepting navigation and content-loading events within the iOS application runtime. These events typically originate from components such as WKWebView, deep-link handlers, or in-app browser controllers. Each outgoing navigation request is captured before execution, allowing the framework to analyze the target URL prior to user exposure.

This interception mechanism is implemented using native iOS delegate callbacks, ensuring minimal overhead and seamless integration with existing application architectures. Intercepting

client-side navigation has been identified as a critical step in preventing phishing and deceptive redirection attacks in mobile environments [2], [4]. The intercepted URL is forwarded to the feature extraction module for further analysis.

3.2 Feature Extraction and Preprocessing

Once a navigation URL is intercepted, the framework extracts a set of lightweight lexical and structural features that are commonly associated with phishing and malicious URLs. Prior research has demonstrated that URL-based lexical features can provide strong predictive power while remaining computationally efficient [1], [5].

The extracted features include:

- URL length, as excessively long URLs are frequently associated with obfuscation techniques used in phishing attacks [10].
- Presence of IP addresses instead of domain names, which is a common indicator of malicious intent [2].
- Subdomain count, used to detect suspicious domain nesting patterns [4].
- Suspicious keyword indicators, such as “login,” “verify,” “secure,” “account,” and “bank,” which are widely used in social engineering attacks [2], [5].

All extracted feature values are normalized prior to inference to ensure numerical stability and compatibility with the machine learning model. Feature normalization is a standard preprocessing step that improves model convergence and inference reliability [6].

3.3 Machine Learning Model and Training

The framework employs a lightweight supervised machine learning classifier optimized for mobile execution. The model is trained offline using publicly available datasets containing labeled phishing and benign URLs, following established practices in phishing detection research [1], [2], [10].

To reduce classification bias, the dataset is balanced across malicious and benign samples and divided into training and validation subsets. Model selection prioritizes high detection accuracy while maintaining low inference latency and memory overhead, which are critical constraints in mobile environments [9].

After training, the model is converted into Apple’s Core ML format to enable native on-device execution. Platform-specific optimizations, including model compression and quantization, are applied to further improve performance and energy efficiency [7], [8]. Embedding the trained model directly within the application eliminates continuous backend dependency and supports offline threat detection.

3.4 On-Device Inference Using Core ML

During runtime, the normalized feature vector derived from the intercepted URL is passed to the embedded Core ML model. The model performs inference locally and outputs a continuous risk score between 0 and 1, representing the likelihood that the URL is malicious.

On-device inference offers several advantages, including reduced latency, improved responsiveness, and enhanced privacy protection, as sensitive navigation data is never transmitted to external servers [3], [9]. The inference process is executed asynchronously to avoid blocking the main UI thread, ensuring a smooth user experience.

3.5 Runtime Decision Enforcement

Based on the risk score generated by the Core ML model, the framework applies a policy-driven enforcement mechanism similar to those used in modern intelligent security systems [3], [5]. Three decision thresholds are defined:

- Allow (Risk score < 0.60): Navigation proceeds without interruption.
- Warn ($0.60 \leq$ Risk score < 0.85): The user is presented with a security warning dialog describing the potential risk.
- Block (Risk score ≥ 0.85): Navigation is immediately blocked to prevent exposure to malicious content.

These thresholds are empirically selected to balance false positives and false negatives while preserving usability and user trust. Immediate enforcement at the client side has been shown to significantly reduce successful phishing attacks in mobile environments [2], [4].

3.6 Algorithmic Overview

The complete client-side threat detection process is summarized in Algorithm 1.

Algorithm 1: Client-Side AI-Based Threat Detection

1. Capture URL from application navigation event
2. Extract lexical and structural features from the URL
3. Normalize feature values
4. Perform on-device inference using the Core ML model
5. Compute threat probability score
6. If score \geq blocking threshold, block navigation
7. Else if score \geq warning threshold, display security warning
8. Else, allow navigation

This methodology enables real-time, privacy-preserving threat detection within the iOS application runtime, aligning with recent advances in client-side and edge-based mobile security research [3], [9].

4. EXPERIMENTAL EVALUATION

This section presents a comprehensive evaluation of the proposed client-side AI-driven cybersecurity framework for iOS applications. The evaluation focuses on detection effectiveness, computational efficiency, and suitability for deployment in real-world mobile environments. Similar evaluation criteria have been widely adopted in prior phishing detection and mobile security research [1], [2], [9].

4.1 Dataset Description

The experimental evaluation was conducted using publicly available datasets containing labeled phishing and benign URLs. These datasets have been extensively used in prior studies on phishing detection and web security, ensuring comparability and reproducibility of results [1], [2], [10].

The dataset consists of a balanced distribution of malicious and legitimate URLs to prevent classification bias. Each URL was labeled according to its ground-truth class and preprocessed to extract lexical and structural features as described in Section 3.

Feature extraction techniques are consistent with those reported in earlier phishing detection literature [5], [10].

4.2 Evaluation Metrics

To assess the effectiveness of the proposed framework, multiple performance metrics were employed. These metrics are standard in cybersecurity and machine learning evaluation and provide a holistic view of system performance [6].

The following metrics were used:

- **Detection Accuracy:** Overall correctness of classification.
- **Precision:** Ability of the model to correctly identify malicious URLs.
- **Recall:** Ability of the model to detect all actual phishing URLs.
- **F1-Score:** Harmonic mean of precision and recall.
- **Inference Latency:** Average time required to perform on-device inference.
- **CPU Utilization:** Average processor usage during runtime analysis.
- **Memory Overhead:** Additional memory consumed by the embedded Core ML model.

These metrics collectively evaluate both security effectiveness and resource efficiency, which are critical for mobile deployment [9].

4.3 Experimental Setup

The trained machine learning model was integrated into an iOS application using Apple's Core ML framework. All experiments were conducted entirely on-device to reflect realistic deployment conditions. Core ML optimizations, including model compression and runtime graph optimization, were applied to ensure efficient inference performance [7], [8].

Navigation events were simulated by programmatically triggering URL loading within a controlled test environment. Each navigation request was intercepted, processed, and evaluated using the proposed framework, following the runtime pipeline described in Section 3.

4.4 Performance Evaluation Results

Table 1 summarizes the performance of the proposed client-side security framework across all evaluation metrics.

Table 1. Performance Evaluation Results

Metric	Value
Detection Accuracy (%)	98.3
Precision (%)	97.9
Recall (%)	98.6
F1-Score (%)	98.2
Average Inference Latency	8.4 ms
Memory Overhead	< 25 MB
CPU Utilization (Average)	< 6%

The results demonstrate that the proposed framework achieves high detection accuracy while maintaining low computational overhead, making it suitable for real-time mobile security enforcement. The achieved accuracy and F1-score are

comparable to or exceed those reported in prior client-side and deep learning-based phishing detection studies [2], [5], [10].

The low inference latency confirms that on-device machine learning can meet the real-time requirements of mobile navigation analysis. This aligns with findings from previous edge intelligence and mobile security research [3], [9]. Additionally, the minimal CPU and memory usage indicate that the framework can be deployed without negatively impacting user experience or application responsiveness.

4.5 Comparative Discussion

Compared to traditional blacklist-based phishing detection approaches, which rely on static rule sets and frequent backend updates, the proposed framework demonstrates superior adaptability to previously unseen threats [1], [4]. Machine learning-based inference enables the system to generalize beyond known attack patterns, improving resilience against zero-day phishing attempts [5].

Furthermore, unlike server-centric security solutions, the proposed approach eliminates network latency and enhances user privacy by ensuring that navigation data remains on the device [3], [9]. This design choice is particularly important for applications operating under strict privacy regulations and intermittent network connectivity.

4.6 Summary of Findings

The experimental evaluation confirms that the proposed client-side AI-driven cybersecurity framework:

- Effectively detects phishing URLs with high accuracy
- Maintains low inference latency suitable for real-time enforcement
- Introduces minimal CPU and memory overhead
- Preserves user privacy by avoiding backend dependency

These findings validate the feasibility and effectiveness of deploying intelligent security mechanisms directly within iOS application runtimes, consistent with recent advances in mobile and edge-based cybersecurity research [3], [9].

5. DISCUSSION

The experimental results presented in Section 4 demonstrate that the proposed client-side AI-driven security framework achieves high detection accuracy (98.3%) while maintaining low inference latency and minimal resource consumption. These results confirm that on-device machine learning can effectively detect phishing threats in real time without degrading application performance, which is consistent with prior findings in mobile and edge-based security research [3], [9].

Compared to traditional blacklist-based approaches, which require continuous backend updates and struggle with zero-day threats, the proposed framework generalizes effectively to previously unseen phishing URLs through learned feature

representations [1], [4], [5]. Additionally, executing threat detection entirely on the client device enhances user privacy by preventing sensitive navigation data from being transmitted to external servers [3].

While the current implementation focuses on phishing URL detection, the architecture can be extended to support additional threat vectors such as malicious UI overlays and abnormal navigation behavior, as suggested in recent mobile security studies [9].

6. CONCLUSION

This paper presented a client-side, AI-driven cybersecurity framework for iOS applications implemented using Swift and Apple's Core ML framework. By embedding an optimized machine learning model directly within the mobile application runtime, the proposed approach enables real-time detection and mitigation of phishing URLs and malicious navigation events without reliance on continuous backend connectivity.

Experimental evaluation demonstrated that the framework achieves high detection accuracy (98.3%) while maintaining low inference latency (8.4 ms) and minimal CPU and memory overhead, confirming its suitability for deployment in real-world mobile environments. These results validate the feasibility of on-device machine learning for proactive mobile security enforcement while preserving user privacy.

The proposed framework complements traditional server-side defenses and offers a scalable, privacy-preserving solution for mobile cybersecurity. Future work will focus on expanding threat coverage to additional attack vectors and exploring adaptive on-device learning techniques to further enhance detection capabilities.

7. REFERENCES

- [1] J. Ma et al., Beyond Blacklists: Learning to Detect Malicious Web Sites, ACM SIGKDD, 2009.
- [2] O. Sahingoz et al., Machine Learning Based Phishing Detection, Expert Systems with Applications, 2019.
- [3] A. Aljofey et al., Client-Side Phishing Detection Using Deep Learning, FGCS, 2020.
- [4] N. Chiew et al., A Survey of Phishing Attacks, Computers & Security, 2018.
- [5] Y. Rao and J. Zhao, Detecting Phishing URLs Using Deep Learning, IEEE Access, 2020.
- [6] I. Goodfellow et al., Deep Learning, MIT Press, 2016.
- [7] Apple Inc., Core ML Framework, Apple Developer Documentation, 2022.
- [8] Apple Inc., Optimizing Core ML Models for On-Device Performance, 2023.
- [9] S. Wang et al., Edge Intelligence for Mobile Security, IEEE IoT Journal, 2021.
- [10] A. Almomani et al., Phishing Website Detection Using Machine Learning, IJACSA, 2019.