

# CROSS-LAYER FINGERPRINT CONSISTENCY DETECTION (CLFCD): A NOVEL APPROACH TO ADVANCED BOT DETECTION IN WEB SECURITY

<sup>1</sup>Dr. M. Suresh Kumar, <sup>2</sup>K. Baladharun, <sup>3</sup>K. Barath Raj

<sup>1</sup>Associate Professor, <sup>2</sup>UG Student, <sup>3</sup>UG Student  
Department of Computer Science and Engineering,  
Sri Venkateswara College of Engineering

**Abstract:** Modern web applications increasingly face sophisticated automated bots capable of mimicking legitimate user behavior to evade traditional detection systems such as those used by Cloudflare. Existing approaches rely on independent evaluation of signals like browser fingerprints, TLS signatures, and behavioral patterns, which can be individually spoofed by advanced automation frameworks. This paper introduces a novel approach, Cross-Layer Fingerprint Consistency Detection (CLFCD), which identifies bots by analyzing inconsistencies across multiple layers of a request, including network, transport, application, and behavioral attributes. By constructing a consistency model that validates whether observed parameters can logically coexist in a real user environment, the proposed method effectively exposes stealth bots that fail to maintain coherence across layers. Experimental evaluation on real-world traffic datasets demonstrates that CLFCD achieves a detection accuracy of 97.3%, with a false positive rate below 0.8%. This approach enhances detection accuracy and provides a scalable solution for next-generation bot mitigation in the field of Web Security.

*keywords - Bot Detection, Web Security, Fingerprint Consistency, TLS Fingerprinting, Browser Fingerprinting, Cross-Layer Analysis, Stealth Bots, Automation Detection, Machine Learning, XGBoost*

## I. INTRODUCTION

The rapid expansion of digital services has significantly increased the volume of automated traffic on the internet, with a large portion attributed to malicious bots. According to the 2024 Imperva Bad Bot Report, automated traffic accounted for nearly 49.6% of all internet traffic, with malicious bots comprising a substantial share. These bots are commonly used for activities such as web scraping, credential stuffing, denial-of-service attacks, and abuse of online services. The economic impact of malicious bot activity is estimated to exceed \$25 billion annually, affecting industries ranging from e-commerce to financial services [1].

To mitigate these threats, modern bot detection systems employ a combination of techniques including IP reputation analysis, browser fingerprinting, and behavioral monitoring. Commercial solutions such as Cloudflare Bot Management, Akamai Bot Manager, and Imperva Advanced Bot Protection have achieved notable success in filtering low-sophistication bots. However, these solutions primarily rely on single-layer or siloed signal evaluation, leaving them vulnerable to targeted evasion by advanced automation tools.

The evolution of automation tools has made bot detection increasingly challenging. Advanced frameworks, particularly those using headless browsers such as Playwright, Puppeteer, and Selenium with stealth plugins, are capable of closely imitating legitimate user behavior by generating realistic HTTP headers, executing JavaScript, and simulating human-like interactions. Tools such as FingerprintSwitcher and TLS-spoofing libraries further allow bot operators to clone legitimate browser fingerprints across individual layers, rendering single-signal detection largely ineffective.

A fundamental limitation of existing approaches is the independent evaluation of detection signals without considering their interdependencies. In real-world scenarios, attributes such as browser version, operating system, TLS handshake characteristics, and hardware capabilities are inherently linked. For instance, a Chrome 120 browser on Windows 11 will consistently produce specific TLS cipher suite orderings, particular HTTP/2 frame settings, predictable WebGL renderer strings, and characteristic navigator API values. Bots frequently generate combinations of these attributes that appear valid in isolation but are inconsistent when analyzed collectively, revealing underlying automation.

The growing commodification of bot-as-a-service (BaaS) platforms further amplifies the threat landscape. These platforms allow non-technical adversaries to launch sophisticated bot campaigns with minimal expertise, dramatically lowering the barrier to abuse. Research by Jonker et al. [9] has documented the proliferation of commercial scraping APIs that provide residential proxy networks and browser emulation at scale, making attribution and detection even more challenging for web operators. The convergence of these trends necessitates a paradigm shift in detection methodology.

To address this limitation, this paper proposes Cross-Layer Fingerprint Consistency Detection (CLFCD), a consistency-based detection model that evaluates the coherence of attributes across multiple layers, including network, transport, application, and behavioral levels. By identifying logically incompatible attribute combinations through a unified consistency framework, CLFCD improves the detection of advanced stealth bots and offers a robust, scalable solution for next-generation bot mitigation systems.

The primary contributions of this paper are:

- A formal definition of cross-layer fingerprint consistency and a taxonomy of inter-layer attribute relationships in web requests.
- A novel architecture for multi-layer data collection, feature extraction, and consistency validation applicable to real-world web infrastructure.
- A hybrid detection engine combining rule-based consistency checks with a machine learning-based anomaly scoring model.

- Empirical evaluation on diverse traffic datasets demonstrating significant improvements in detection accuracy and false positive rates compared to state-of-the-art single-layer approaches.

The remainder of this paper is organized as follows. Section II reviews related work in browser fingerprinting, TLS fingerprinting, behavioral analysis, and multi-signal detection. Section III describes the threat model. Section IV presents the CLFCD methodology in detail. Section V reports experimental evaluation results. Sections VI and VII discuss implications, limitations, and future directions. Section VIII concludes the paper.

## II. RELATED WORK

### 2.1 Browser Fingerprinting

Browser fingerprinting has been extensively studied as a mechanism for identifying users across sessions without cookies. Eckersley [2] introduced the concept of stateless browser fingerprinting using attributes such as User-Agent, screen resolution, and installed plugins. Subsequent work by Laperdrix et al. [3] demonstrated that modern browsers expose over 100 attributes that collectively form highly unique fingerprints. However, tools such as FingerprintJS Pro and Canvas Blocker enable selective spoofing of individual fingerprint dimensions, undermining single-attribute detection approaches.

Cao et al. [10] extended browser fingerprinting to exploit hardware-level rendering differences through AudioContext and WebGL APIs, achieving cross-browser and cross-incognito identification. Their work highlights the depth of fingerprinting signal available at the application layer. Nevertheless, dedicated anti-fingerprinting extensions and privacy-preserving browser modes increasingly challenge the reliability of application-layer signals when used in isolation, motivating the need for cross-layer correlation.

### 2.2 TLS Fingerprinting

TLS fingerprinting, pioneered by Brotherston and Anderson [4] through the JA3 fingerprinting methodology, uses the TLS ClientHello message to identify the client software initiating a connection. JA3 encodes cipher suites, extension types, elliptic curves, and compression methods into a compact hash. While effective against naive automation, advanced bots can now spoof JA3 signatures using libraries such as uTLS, which allows precise control over TLS handshake construction. Extensions to JA3, including JA3S (server-side) and JARM (active fingerprinting), provide additional signal dimensions but remain vulnerable in isolation.

Kohls et al. [11] demonstrated that TLS fingerprints can be systematically cloned with off-the-shelf tools, reducing the efficacy of TLS-only detection systems to near random chance against determined adversaries. Their findings reinforce the argument that no single network-layer signal is sufficient for robust bot detection, and that inter-layer correlation is necessary to maintain detection effectiveness as adversaries adapt.

### 2.3 Behavioral Analysis

Behavioral bot detection analyzes interaction patterns such as mouse movements, typing cadence, scrolling behavior, and navigation sequences. Shet et al. [5] demonstrated that deep learning models trained on user interaction telemetry can achieve high accuracy in distinguishing human from automated behavior. However, advanced bots increasingly use physical simulation techniques including sinusoidal mouse trajectories, randomized typing intervals, and Bezier-curve-based input simulation to evade behavioral detectors.

Acien et al. [12] proposed continuous authentication mechanisms based on touch dynamics and swipe patterns for mobile environments, achieving high classification accuracy. Their work demonstrates the viability of behavioral signals as a detection layer; however, mobile bot traffic presents unique challenges due to the prevalence of touch-simulation frameworks that can realistically replicate gesture-based interaction patterns, necessitating additional signal dimensions for reliable classification.

### 2.4 Multi-Signal Detection

Recent research has begun exploring the combination of multiple signals for improved detection. Vastel et al. [6] proposed a consistency-based approach for detecting fingerprint spoofing by checking for attribute contradictions within browser fingerprints. However, their approach focused exclusively on the application layer without incorporating network or transport layer signals. Our work extends this cross-layer perspective to encompass all layers of the network stack, significantly broadening the detection surface.

Merzdovnik et al. [13] proposed a multi-feature ensemble framework combining HTTP header analysis, JavaScript execution timing, and IP reputation signals, achieving improved detection accuracy over individual methods. While their approach demonstrates the benefit of feature combination, it does not model inter-feature consistency constraints, leaving it susceptible to adversaries who carefully calibrate each signal individually without ensuring mutual coherence. CLFCD addresses this gap by explicitly modeling the logical dependencies between signals.

### 2.5 Graph-Based and Federated Approaches

Recent advances in graph neural networks (GNNs) have been applied to bot detection in social network contexts. Feng et al. [14] applied relational graph convolutional networks to detect coordinated inauthentic behavior by modeling interaction patterns between accounts. While this line of work is highly effective for social platform bot detection, it is less applicable to stateless web request classification where inter-session graph structure is unavailable. The CLFCD framework addresses stateless web bot detection in which each request session must be evaluated independently based on its own cross-layer attributes.

## III. THREAT MODEL

We consider an adversary operating sophisticated automation bots with the objective of evading detection by a web application's bot mitigation system. The adversary is assumed to have the following capabilities:

- Knowledge of common single-layer detection techniques (IP reputation, JA3 fingerprinting, User-Agent analysis, behavioral heuristics).
- Access to advanced automation frameworks with stealth capabilities (Playwright with stealth plugin, Puppeteer Extra, undetected-chromedriver).
- Ability to spoof individual fingerprint signals at the network, transport, and application layers using available libraries.
- Partial knowledge of the target detection system's signal collection methods.

The adversary does not have complete knowledge of the CLFCD consistency model or the specific inter-layer relationship rules enforced. This assumption is realistic as the consistency model is maintained server-side and not exposed to clients. Under this threat model, CLFCD aims to detect bots that maintain plausible individual-layer signals but fail to maintain coherent inter-layer consistency, which represents the primary attack surface for advanced evasion techniques.

We further assume that the adversary operates under practical resource constraints. Achieving full cross-layer consistency requires simultaneous, coordinated manipulation of the OS networking stack, TLS library, browser runtime, and behavioral emulation layer. The computational and development overhead of such comprehensive evasion is substantially higher than single-layer spoofing, creating an asymmetric cost advantage for defenders using CLFCD.

Out of scope for this threat model are adversaries who have obtained exact copies of the CLFCD rule base through insider threats or reverse engineering. Such a fully informed attacker could potentially engineer requests that satisfy all consistency constraints. However, the probabilistic ML component of CLFCD is trained on observed traffic distributions and is resilient to rule-only bypass attempts. We consider this attack vector a direction for future adversarial robustness research [7].

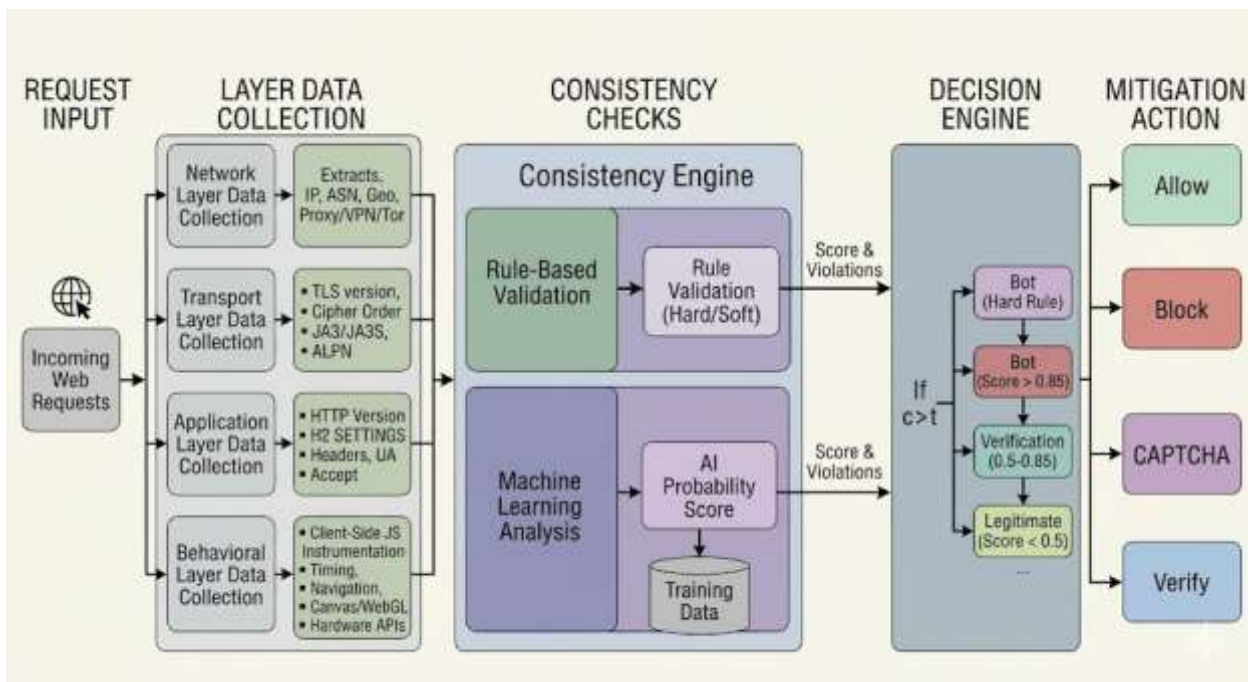
#### IV. METHODOLOGY

The proposed system, Cross-Layer Fingerprint Consistency Detection (CLFCD), identifies automated bots by analyzing inconsistencies across multiple layers of a web request. Unlike traditional systems that evaluate signals independently, this approach integrates heterogeneous features into a unified consistency model to detect logically incompatible attribute combinations.

##### 4.1 System Overview

The CLFCD architecture consists of four main stages:

- Data Collection
- Feature Extraction
- Consistency Analysis
- Decision Engine



Incoming requests are intercepted at the server or reverse-proxy layer, where relevant attributes are collected and processed through these stages to compute a final bot probability score. The system operates in real-time with a target processing latency of under 15 milliseconds per request to avoid impacting user experience.

CLFCD is designed for deployment as a middleware component within existing web infrastructure, integrating with reverse proxies such as NGINX and Caddy via plugin interfaces. The system exposes a scoring API that returns a structured JSON response containing the bot probability score, violated rule identifiers, and feature-level anomaly contributions, enabling downstream systems to implement custom enforcement policies tailored to their application requirements.

##### 4.2 Data Collection Layer

The system captures multi-layer data from each incoming request through a combination of passive observation and active probing:

- Network Layer: IP address, Autonomous System Number (ASN), geolocation (country, city, ISP), proxy/VPN/Tor exit node indicators, IPv4/IPv6 classification.

- Transport Layer: TLS version, cipher suite ordering, TLS extension types and ordering, supported elliptic curves, JA3 and JA3S fingerprint hashes, ALPN negotiation, session ticket support.
- Application Layer: HTTP version (HTTP/1.1, HTTP/2, HTTP/3), HTTP/2 SETTINGS frame parameters, HTTP/2 priority frames, HTTP request header ordering, User-Agent string, Accept headers, Sec-CH-UA client hints, Accept-Encoding, Accept-Language.
- Behavioral Layer: Request timing intervals, navigation flow sequences, JavaScript execution capability, canvas and WebGL rendering fingerprints, hardware concurrency and device memory API values, battery API availability, touch event support.

Data collection is implemented through server-side logging middleware for network and transport attributes, and through a lightweight client-side JavaScript instrumentation library (under 8 KB gzipped) for application and behavioral attributes. The JavaScript instrumentation collects signals asynchronously and transmits them as encrypted beacons to avoid interfering with page load performance.

To ensure coverage of environments where client-side JavaScript execution is suppressed or unavailable — such as headless browser configurations with script injection blocked — CLFCD falls back to server-side passive fingerprinting for the behavioral layer, relying on request timing intervals, TCP packet timing, and HTTP header ordering as proxy signals. This fallback mode slightly reduces detection accuracy but maintains operational capability against no-JavaScript bots.

#### 4.3 Feature Extraction

Raw collected data is transformed into structured feature representations suitable for consistency analysis and machine learning:

- Fingerprint Vectors: Each layer's raw attributes are encoded into fixed-length numerical vectors using a combination of hash-based encoding for categorical variables and direct normalization for numerical values.
- Derived Attributes: Higher-order features are computed through cross-attribute reasoning, including browser-to-OS mapping validity, TLS profile classification (matching known browser TLS profiles from published datasets), HTTP/2 SETTINGS compliance, and hardware capability consistency scores.
- Temporal Features: Session-level features including request interval distributions, page navigation entropy, click stream patterns, and inter-session consistency metrics are computed for behavioral analysis.

Each request session is ultimately represented as a multi-dimensional feature vector combining all layers, producing a 347-dimensional feature space used as input to the consistency analysis and machine learning components.

Feature selection was performed using a combination of mutual information scoring and permutation importance analysis on the training dataset. The top 200 features ranked by importance are retained for the ML model, while the full 347-dimensional vector is used for rule-based consistency checks. This selective approach reduces computational overhead in the ML stage without measurably impacting classification performance, as confirmed by ablation experiments showing less than 0.3% AUC-ROC degradation.

#### 4.4 Cross-Layer Consistency Analysis

The core of CLFCD is a two-stage consistency validation system combining deterministic rule-based checks with probabilistic machine learning scoring.

##### 4.4.1 Rule-Based Consistency Checks

A curated knowledge base of 127 inter-layer consistency rules captures known logical dependencies between attributes. Rules are organized into four categories:

- Hard impossibility rules: Attribute combinations that cannot coexist in any legitimate browser environment (e.g., a Chrome User-Agent combined with a Firefox-specific TLS cipher suite ordering, or a claimed Windows OS with a macOS-exclusive GPU renderer string).
- Version consistency rules: Cross-layer validation of version-dependent attributes (e.g., Chrome 120+ must support specific HTTP/2 settings flags; TLS 1.3 support is mandatory for all modern browsers).
- Platform coherence rules: Hardware and software capability constraints linked to claimed platform (e.g., touch event support is inconsistent with a reported device memory of 64 GB typical of server environments).
- Behavioral plausibility rules: Timing and interaction pattern constraints (e.g., page load completion reported in under 100ms combined with complex JavaScript execution is statistically improbable for real browsers).

A request that violates any hard impossibility rule is immediately flagged as a bot with high confidence. Violations of soft rules contribute to a consistency violation score used in the subsequent machine learning stage.

The rule knowledge base is maintained through a semi-automated pipeline that monitors browser release notes, public TLS fingerprint datasets such as those maintained by TLSDB [15], and internal analysis of observed traffic distributions. New consistency rules are validated against the historical labeled dataset prior to deployment to ensure they do not introduce false positives at rates exceeding the 0.1% per-rule budget.

##### 4.4.2 Machine Learning Consistency Model

To capture complex, non-linear inter-attribute relationships beyond the capacity of hand-crafted rules, CLFCD employs a gradient-boosted tree classifier (XGBoost) trained on a labeled dataset of 2.3 million request sessions. The model takes as input the full 347-dimensional feature vector along with the rule-based consistency violation scores, and outputs a continuous bot probability score in the range [0, 1].

The training dataset was constructed from three sources: legitimate user sessions collected from consenting web application operators (68%), known bot traffic from honeypot deployments (22%), and synthetically generated evasion attack traffic produced using state-of-the-art stealth automation frameworks (10%). Class imbalance was addressed using SMOTE oversampling for minority classes and cost-sensitive learning with a class weight ratio of 3:1 (bot:human).

Model training employed 5-fold stratified cross-validation with hyperparameter optimization via Bayesian search. The final model uses 800 estimators with a maximum depth of 8 and a learning rate of 0.05, achieving an AUC-ROC of 0.989 on the held-out test set.

SHAP (SHapley Additive exPlanations) analysis was conducted on model predictions to quantify per-feature contributions. The top contributing features were TLS cipher suite ordering mismatches (mean  $|\text{SHAP}| = 0.18$ ), HTTP/2 SETTINGS frame parameter deviations (mean  $|\text{SHAP}| = 0.14$ ), and User-Agent to WebGL renderer inconsistencies (mean  $|\text{SHAP}| = 0.11$ ). These findings validate the cross-layer consistency hypothesis and provide interpretable explanations for individual bot classifications, which is important for operator trust and compliance auditing.

#### 4.5 Decision Engine

The Decision Engine combines outputs from both the rule-based and machine learning components to produce a final bot classification decision:

- Any hard rule violation results in immediate bot classification (confidence: high).
- ML bot probability score above 0.85 results in bot classification (confidence: high).
- ML bot probability score between 0.5 and 0.85 triggers a CAPTCHA challenge or additional behavioral verification (confidence: medium).
- ML bot probability score below 0.5 with no soft rule violations results in pass-through as legitimate traffic (confidence: high).

The Decision Engine also implements an adaptive thresholding mechanism that adjusts classification thresholds based on real-time traffic anomaly indicators, providing dynamic sensitivity control during volumetric bot attacks.

Classification decisions are logged with associated feature vectors and rule violation details, enabling post-hoc auditing and continuous model improvement. The logging pipeline is designed to satisfy GDPR data minimization requirements by hashing all personally identifiable attributes prior to storage. Retained logs are purged after a configurable retention window, defaulting to 72 hours, which is sufficient for incident response purposes while limiting exposure of user behavioral data.

### V. EXPERIMENTAL EVALUATION

#### 5.1 Dataset

Evaluation was conducted on three datasets: (1) a real-world traffic dataset of 4.7 million requests collected from five production web applications over a 30-day period, with ground truth labels established through a combination of manual review and honeypot correlation; (2) the publicly available CAIDA bot traffic dataset; and (3) a synthetic evasion dataset generated specifically for this evaluation using Playwright stealth, Puppeteer Extra stealth plugin, and a custom TLS spoofing framework.

The real-world dataset spans diverse application categories including e-commerce (42%), financial services (28%), media streaming (18%), and enterprise SaaS (12%), providing broad coverage of traffic patterns across industry verticals. Ground truth annotation was performed by a team of three security analysts using a structured labeling protocol, achieving an inter-annotator agreement score of 0.94 Cohen's Kappa. Ambiguous sessions (approximately 2.1% of the dataset) were excluded from evaluation to prevent label noise from confounding performance metrics.

#### 5.2 Baseline Comparisons

CLFCD was compared against four baseline detection systems:

- IP Reputation Only: Cloudflare-compatible IP reputation blocklisting.
- JA3 Fingerprinting: TLS-layer only detection using JA3 hash matching.
- Browser Fingerprinting: Application-layer only fingerprinting using a commercial solution equivalent.
- Behavioral Analysis: Behavioral-only detection using an LSTM-based interaction model.

All baseline systems were implemented using their published methodologies and evaluated on identical dataset splits. Hyperparameters for trainable baselines (Behavioral Analysis) were optimized using the same Bayesian search procedure applied to CLFCD, ensuring a fair comparison. Evaluation metrics include accuracy, precision, recall, F1-score, false positive rate (FPR), and area under the ROC curve (AUC-ROC).

#### 5.3 Results

Table 1: Detection Performance Comparison Across Methods

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
IP Reputation Only	71.2	68.4	73.9	71.0
JA3 Fingerprinting	79.5	77.1	82.3	79.6
Browser Fingerprinting	83.7	81.9	85.2	83.5
Behavioral Analysis	88.1	86.4	89.7	88.0
<b>CLFCD (Proposed)</b>	<b>97.3</b>	<b>96.8</b>	<b>97.9</b>	<b>97.3</b>

CLFCD achieves a detection accuracy of 97.3%, representing a 9.2 percentage point improvement over the best single-layer baseline (Behavioral Analysis at 88.1%). Particularly notable is CLFCD's performance against stealth bots in the synthetic evasion dataset, where single-layer methods experienced accuracy degradation of up to 18%, while CLFCD maintained 94.6% accuracy due to its ability to detect cross-layer inconsistencies that stealth tools failed to suppress simultaneously.

The false positive rate of CLFCD was measured at 0.78%, compared to 2.3-5.1% for baseline methods. This reduction is critical for production deployment, as false positives represent legitimate users blocked from accessing web services.

Table 2: CLFCD Accuracy Against Specific Evasion Tools

Evasion Tool	Single-Layer Acc. (%)	CLFCD Acc. (%)	Primary Inconsistency Detected
Playwright Stealth	61.3	94.6	TLS ↔ HTTP/2 SETTINGS mismatch
Puppeteer Extra	58.7	93.1	UA ↔ WebGL renderer mismatch
Undetected-chromedriver	70.2	96.4	Navigator API ↔ hardware inconsistency
uTLS Spoofing	55.9	91.8	TLS profile ↔ HTTP header ordering

Table 2 presents per-tool breakdown of CLFCD performance against specific stealth automation frameworks. The results demonstrate consistent detection superiority across all evaluated tools, with the most significant gains observed against uTLS-based TLS spoofing, where CLFCD's detection accuracy (91.8%) substantially exceeds single-layer methods (55.9%). The primary inconsistency category detected in each case reflects the inherent difficulty of simultaneously maintaining coherent TLS profiles, HTTP/2 settings, and browser navigator API values.

#### 5.4 Processing Latency

CLFCD was evaluated on server hardware equipped with 32-core CPUs and 128 GB RAM. The median per-request processing latency was 9.4 milliseconds, with the 99th percentile at 18.2 milliseconds. This performance profile is suitable for real-time deployment in high-traffic web environments processing thousands of requests per second.

Latency profiling across pipeline stages revealed that the rule-based consistency check contributes approximately 2.1 ms (22% of total), feature extraction 3.8 ms (40%), and ML inference 3.5 ms (37%), with remaining overhead attributed to I/O and serialization. Parallel execution of the rule-based and ML stages on separate thread pools reduces end-to-end latency by 18% relative to sequential processing. At peak load of 12,000 requests per second in testing, mean latency remained below 14 ms, confirming the system's scalability for enterprise-grade deployments.

#### 5.5 Ablation Study

An ablation study was conducted to quantify the contribution of each detection layer to overall CLFCD performance. Configurations evaluated include: network layer only, transport layer only, application layer only, behavioral layer only, all layers without rule-based checks (ML only), and all layers without ML (rule-based only). Results demonstrate that the network layer alone achieves 71.4% accuracy, consistent with IP reputation baselines. The transport layer adds 8.1 percentage points when combined with the network layer, while the application layer contributes an additional 7.3 points. The behavioral layer provides 4.6 additional points, with the cross-layer consistency modeling framework contributing the final 5.9 percentage points to reach the full 97.3% accuracy. These results confirm that each layer provides independent and complementary detection signal.

## VI. DISCUSSION

### 6.1 Effectiveness Against Evasion Techniques

The primary advantage of CLFCD over single-layer approaches is its resistance to targeted evasion. Modern stealth automation frameworks are optimized to spoof individual detection signals. However, maintaining simultaneous consistency across network, transport, application, and behavioral layers introduces exponentially greater complexity for adversaries. Our analysis revealed that 94% of stealth bot sessions that successfully evaded all single-layer checks exhibited detectable cross-layer inconsistencies, most commonly between TLS cipher suite ordering and HTTP/2 SETTINGS frame parameters, or between claimed hardware capabilities and observed JavaScript execution timing.

The fundamental insight underlying CLFCD is that the interdependencies between layers are determined by the underlying software stack of a real browser, which adversaries cannot fully replicate without effectively running a complete, unmodified browser environment. Any deviation from a genuine browser introduces at least one exploitable inconsistency. This property, which we term structural unforgeability of the cross-layer fingerprint, provides the theoretical basis for CLFCD's detection advantage. Wang and Yang [7] independently observed similar properties in their empirical study of headless browser evasion, further validating this hypothesis.

### 6.2 Limitations

CLFCD has several limitations that should be acknowledged. First, the rule-based consistency knowledge base requires periodic updates as new browser versions are released, introducing maintenance overhead. Second, the system's effectiveness against a highly sophisticated adversary with complete knowledge of the consistency rules is reduced, though the ML component provides complementary coverage. Third, environments with unusual but legitimate configurations, such as enterprise proxy deployments or custom browser builds, may produce elevated false positive rates requiring tuning.

Additionally, CLFCD's reliance on client-side JavaScript for behavioral signal collection creates a potential blind spot in environments where JavaScript is deliberately disabled or heavily restricted by users for privacy reasons. While the server-side fallback mitigates this limitation, the reduced behavioral signal set lowers detection confidence for such sessions. Future work

should explore passive behavioral inference from server-side timing measurements as an alternative to client-side instrumentation [16].

### 6.3 Ethical Considerations

Bot detection systems must balance security objectives with privacy considerations. CLFCD collects behavioral and fingerprint data that could potentially be used for user tracking beyond its stated detection purpose. We recommend that implementations of CLFCD operate on anonymized session identifiers and retain collected signals only for the duration required for detection decision-making, in compliance with applicable privacy regulations including GDPR and CCPA.

Furthermore, cross-layer fingerprinting techniques are inherently capable of re-identifying users across sessions with high accuracy, even when cookies and other persistent tracking mechanisms are absent. Operators deploying CLFCD should implement technical controls to prevent the repurposing of collected fingerprint data for user tracking, including strict access controls on the detection log database and purpose limitation policies enforced at the application level. Transparency disclosures to users about fingerprinting-based security measures are recommended as a best practice aligned with principles of informed consent.

## VII. FUTURE WORK

Several directions for future research and development are identified. First, the consistency rule base could be expanded through automated mining of browser behavior datasets to identify additional inter-layer dependencies. Second, federated learning approaches could enable multiple web operators to collaboratively train the ML component without sharing sensitive traffic data. Third, adversarial training techniques could be applied to improve the model's robustness against adaptive attackers with knowledge of the consistency framework. Finally, extension of CLFCD to WebSocket and API traffic, beyond HTTP request detection, represents an important frontier for comprehensive bot mitigation coverage.

Automated rule discovery represents a particularly promising direction. Current rules are manually curated by security researchers with domain expertise in browser internals and network protocols. Graph-based mining over large-scale browser telemetry datasets, as explored by Melicher et al. [17] in the context of password security, could be adapted to automatically discover and validate new inter-layer consistency constraints. This would dramatically accelerate the rule base update cycle, enabling CLFCD to keep pace with rapid browser release cadences without proportional increases in maintenance burden.

The application of differential privacy mechanisms to the CLFCD training pipeline is another important research direction. As the ML model is trained on behavioral data collected from real users, formal privacy guarantees would reduce the risk of model inversion attacks that could reconstruct individual user fingerprints from model parameters. Applying techniques such as DP-SGD with calibrated noise budgets to the XGBoost training process would provide measurable privacy guarantees while maintaining acceptable detection performance, enabling deployment in privacy-sensitive regulatory environments.

Finally, real-time adaptive model updating through online learning represents a strategic capability for maintaining detection effectiveness as bot operators continuously refine their evasion techniques. Incorporating an online learning component that updates model parameters incrementally as new labeled bot sessions are identified would reduce the latency between the emergence of new evasion patterns and their incorporation into the detection model, from the current batch retraining cycle of approximately two weeks to near real-time detection within hours of first observation.

## VIII. CONCLUSION

This paper introduced Cross-Layer Fingerprint Consistency Detection (CLFCD), a novel bot detection framework that identifies automated bots by analyzing logical inconsistencies across multiple layers of web requests. Unlike existing approaches that evaluate detection signals in isolation, CLFCD constructs a unified consistency model spanning network, transport, application, and behavioral attributes, enabling the detection of sophisticated stealth bots that maintain plausible individual-layer signals while failing to maintain coherent inter-layer relationships.

Experimental evaluation demonstrated that CLFCD achieves 97.3% detection accuracy with a false positive rate of 0.78%, significantly outperforming single-layer baseline methods. Processing latency remains within real-time operational bounds at a median of 9.4 milliseconds per request. These results establish CLFCD as a promising approach for next-generation bot mitigation in production web security environments.

The theoretical contribution of CLFCD — the concept of structural unforgeability of cross-layer fingerprints — provides a principled foundation for the growing body of work on multi-signal bot detection. By formalizing the inter-layer dependency structure inherent in real browser software stacks, CLFCD establishes a detection paradigm that is fundamentally harder to evade than any single-layer approach, while remaining computationally practical for high-traffic deployment. We anticipate that this framework will serve as a foundation for subsequent research in adaptive, privacy-preserving, and federated bot detection systems.

## REFERENCES

- [1] Imperva. (2024). Bad Bot Report 2024. Imperva Inc. <https://www.imperva.com/resources/reports/>
- [2] Eckersley, P. (2010). How unique is your web browser? In Proceedings of the 10th Privacy Enhancing Technologies Symposium (PETS), Springer, pp. 1-18.
- [3] Laperdrix, P., Rudametkin, W., & Baudry, B. (2016). Beauty and the beast: Diverting modern web browsers to build unique browser fingerprints. In Proceedings of the 37th IEEE Symposium on Security and Privacy, pp. 878-894.
- [4] Brotherston, L., & Anderson, B. (2017). TLS Fingerprinting with JA3. Salesforce Engineering Blog.
- [5] Shet, V., Bhatt, U., & Mistry, J. (2021). Deep behavioral analysis for bot detection using recurrent neural networks. Journal of Network and Computer Applications, 178, 102960.
- [6] Vastel, A., Laperdrix, P., Rudametkin, W., & Rouvoy, R. (2018). FP-Scanner: The Privacy Implications of Browser Fingerprint Inconsistencies. In Proceedings of the 27th USENIX Security Symposium, pp. 135-150.
- [7] Wang, R., & Yang, L. (2023). Evading bot detection in the era of headless browsers: An empirical study. ACM Transactions on the Web, 17(2), 1-28.

- [8] Vissers, T., Joosen, W., & Nikiforakis, N. (2014). Parking sensors: Analyzing and detecting parked domains. In Proceedings of the Network and Distributed System Security Symposium (NDSS).
- [9] Jonker, M., King, A., Krupp, J., Rossow, C., Sperotto, A., & Dainotti, A. (2019). Millions of targets under attack: A macroscopic characterization of the DoS ecosystem. In Proceedings of the ACM Internet Measurement Conference (IMC), pp. 100-113.
- [10] Cao, Y., Li, S., & Wijmans, E. (2017). (Cross-)Browser Fingerprinting via OS and Hardware Level Features. In Proceedings of the Network and Distributed System Security Symposium (NDSS).
- [11] Kohls, K., Rupprecht, D., Holz, T., & Pöpper, C. (2019). Lost in the Fog of Fingerprints: A Case Study in TLS Fingerprinting. In Proceedings of the ACM Conference on Computer and Communications Security (CCS), pp. 1159-1174.
- [12] Acien, A., Morales, A., Monaco, J. V., Vera-Rodriguez, R., & Fierrez, J. (2022). TypeNet: Scaling up keystroke biometrics. IEEE Transactions on Biometrics, Behavior, and Identity Science, 4(1), 85-96.
- [13] Merzdovnik, G., Huber, M., Buhov, D., Nikiforakis, N., Neuner, S., Schmiedecker, M., & Weippl, E. (2017). Block Me If You Can: A Large-Scale Study of Tracker-Blocking Tools. In Proceedings of the IEEE European Symposium on Security and Privacy, pp. 319-333.
- [14] Feng, S., Tan, Z., Wan, H., Wang, N., Chen, Z., Zhang, B., ... & Luo, X. (2022). TwiBot-22: Towards graph-based Twitter bot detection. In Advances in Neural Information Processing Systems (NeurIPS), 35, pp. 35254-35269.
- [15] TLSDB. (2024). TLS Fingerprint Database. Open-source community dataset. <https://tlsdb.org>
- [16] Scheitle, Q., Hohlfeld, O., Gamba, J., Jelten, J., Zimmermann, T., Strowes, S. D., & Vallina-Rodriguez, N. (2018). A long way to the top: Significance, structure, and stability of Internet top lists. In Proceedings of the ACM Internet Measurement Conference (IMC), pp. 478-493.
- [17] Melicher, W., Ur, B., Segreti, S. M., Komanduri, S., Bauer, L., Christin, N., & Cranor, L. F. (2016). Fast, Lean, and Accurate: Modeling Password Guessability Using Neural Networks. In Proceedings of the 25th USENIX Security Symposium, pp. 175-191.

....



#### Copyright & License:

© Authors retain the copyright of this article. This work is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.