

Machine Learning Approaches for Detecting OWASP Web Application Attacks

Kenneth Solomon
Department of Computer Science
St. Wilfred's College of Arts, Commerce & Science, Mumbai, India
Kennethsolomon.975@gmail.com

Abstract

Web applications represent a critical component of modern digital infrastructure, handling sensitive data across banking, healthcare, e-commerce, and government sectors. The growing sophistication of Cyberattackers has made traditional signature-based intrusion detection systems increasingly insufficient. This paper proposes a machine learning (ML) based framework for detecting web application attacks classified under the OWASP Top 10, with a focus on SQL injection, Cross-Site Scripting (XSS), and command injection attacks. We analyze the CSIC 2010 HTTP dataset and apply supervised classification algorithms including Random Forest, Decision Tree, Support Vector Machine (SVM), and Logistic Regression to distinguish malicious HTTP requests from legitimate traffic. Experimental results demonstrate that the Random Forest classifier achieves the highest accuracy of 98.7% with an F1-score of 0.986, outperforming all other evaluated models.

Keywords :— Machine Learning, Web Application Security, OWASP Top 10, SQL Injection, Intrusion Detection, Random Forest, HTTP Traffic Analysis

1. INTRODUCTION

The proliferation of internet-connected applications has fundamentally altered how organizations deliver services and process sensitive user data. Web applications now form the backbone of critical digital infrastructure across banking, healthcare, education, and government sectors [1]. This broad deployment makes them a primary target for cybercriminals.

According to the IBM Cost of a Data Breach Report 2024, the average cost of a data breach has risen to approximately \$4.88 million globally, with web application attacks contributing significantly [2]. The OWASP Top 10 identifies the most critical web application security risks, including injection flaws, broken authentication, and security misconfigurations [3].

Traditional rule-based Web Application Firewalls (WAFs) and signature-based Intrusion Detection Systems (IDS) are effective against known attack patterns but fail against novel or obfuscated threats. This limitation motivates the use of machine learning as a complementary adaptive defense strategy [4].

This paper evaluates four supervised classifiers Random Forest, Decision Tree, SVM, and Logistic Regression on features extracted from the CSIC 2010 HTTP dataset. Section 2 reviews related work, Section 3 covers OWASP Top 10, Section 4 discusses SQL injection, Sections 5–6 present the framework and dataset, Section 7 reports results, and Sections 8–9 discuss prevention and conclusions.

2. RELATED WORK

Raman et al. [5] applied Random Forest and Decision Tree classifiers to HTTP traffic logs and reported accuracy exceeding 97% in detecting SQL injection and XSS attacks. Their work highlighted the importance of feature selection in reducing model complexity without sacrificing detection performance.

Abiodun et al. [6] compared deep learning models for network intrusion detection, showing LSTM architectures achieved superior recall for zero-day patterns, though with significantly higher computational cost. Torrano-Gimenez et al. [7] introduced the CSIC 2010 HTTP benchmark dataset, achieving 96.4% baseline accuracy with anomaly detection, A reference point for subsequent studies.

Wanda and Jie [8] explored federated learning for distributed web attack detection, demonstrating accuracy comparable to centralized models while preserving data locality. These studies collectively motivate our ensemble ML approach using engineered HTTP features.

Mereani and Howe [13] applied gradient-boosted trees to the ECML-PKDD 2007 HTTP dataset, achieving 97.1% accuracy but noted significant performance degradation on URL-encoded payloads a limitation our entropy-based feature directly addresses. Nguyen et al. [14] proposed a CNN-based HTTP classifier operating on raw byte sequences, reporting 98.2% F1 on CSIC 2010; however, their model required GPU acceleration and offered no feature interpretability, limiting deployment in resource-constrained WAF environments. Our approach achieves comparable accuracy with eight interpretable features and sub-5ms inference latency on standard CPU hardware, offering a practical advantage for production integration.

3. OWASP TOP 10 OVERVIEW

The OWASP Top 10 documents the most critical security risks facing web applications, updated periodically based on data from hundreds of organizations [3]. Table 1 summarizes the 2021 edition with severity ratings.

Rank	Vulnerability	Severity
A01	Broken Access Control	High–Critical
A02	Cryptographic Failures	High

Rank	Vulnerability	Severity
A03	Injection	Critical
A04	Insecure Design	High
A05	Security Misconfiguration	Medium-High
A06	Vulnerable Components	Medium-Critical
A07	Auth & Identity Failures	High
A08	Software Integrity Failures	High
A09	Logging Failures	Medium
A10	Server-Side Req. Forgery	High

Table 1. OWASP Top 10 2021 Vulnerabilities.

Among these, injection attacks (A03) directly manipulate backend query logic and produce distinctive HTTP payload signatures most amenable to ML-based feature classification.

4. SQL INJECTION METHODOLOGY

4.1 Attack Mechanics

SQL injection inserts malicious SQL statements into application inputs executed by the backend database, exploiting insufficient input validation [9]. A standard login query:

```
SELECT * FROM users WHERE username='admin' AND password='1234'
```

With the payload 'OR '1'='1' becomes:

```
SELECT * FROM users WHERE username=" OR '1'='1'
```

Since '1'='1' always evaluates true, authentication is bypassed. Beyond this, SQLi enables UNION-based exfiltration, blind inference via Boolean/time-based techniques, and in some configurations, remote code execution.

4.2 Detection Challenges

Detection is non-trivial because payloads can be URL-encoded or obfuscated, legitimate queries may contain SQL-like syntax, and the diversity of injection variants (in-band, inferential, out-of-band) demands broad coverage. ML addresses this by learning statistical patterns across diverse examples.

5. DETECTION FRAMEWORK

5.1 System Architecture

The framework operates as an inline analysis layer intercepting HTTP requests, extracting feature vectors, and classifying them before forwarding or blocking. Figure 4 illustrates the architecture.

Fig. 4. ML-Based Attack Detection Architecture




Fig. 4. ML-Based Attack Detection System Architecture. Detected malicious requests are blocked and logged. Legitimate traffic is forwarded with an average latency overhead below 5 milliseconds, suitable for production WAF integration.

5.2 Detection Pipeline

Figure 5 shows the four-stage pipeline: (1) HTTP capture and preprocessing, (2) feature extraction, (3) Random Forest classification, and (4) automated response action.

Fig. 5. ML Detection Pipeline




Fig. 5. ML-Based Web Attack Detection Pipeline.

5.3 Classifiers Evaluated

Random Forest: Ensemble of decision trees, robust against overfitting via averaging [10].

Decision Tree: Hierarchical classifier based on information gain. Interpretable but prone to overfitting.

SVM (RBF Kernel): Maximizes class margin in high-dimensional space. Effective but computationally expensive [11].

Logistic Regression: Probabilistic linear baseline. Fast but limited for nonlinear boundaries.

The mathematical basis for classifier selection warrants elaboration. Random Forest constructs B decision trees $\{h(x, \Theta_k), k=1, \dots, B\}$ where each Θ_k is a random vector drawn i.i.d., and the ensemble prediction is determined by majority vote: $H(x) = \arg\max_Y \sum_k I(h(x, \Theta_k) = Y)$. This bagging mechanism reduces variance without increasing bias, making it particularly effective when individual features (e.g., SQL keyword count) are noisy across obfuscated payloads.

The SVM with RBF kernel maps inputs to a high-dimensional feature space via $\phi(x)$ and solves: $\min_{w,b} \frac{1}{2} \|w\|^2 + C \sum_i \xi_i$, subject to $y_i(w \cdot \phi(x_i) + b) \geq 1 - \xi_i$. The regularisation parameter $C=1.0$ and kernel bandwidth γ were selected via 5-fold cross-validation grid search over $C \in \{0.1, 1, 10\}$ and $\gamma \in \{0.001, 0.01, 0.1\}$. While SVM achieves strong separation in high-dimensional feature spaces, its $O(n^2)$ to $O(n^3)$ training complexity makes it approximately 14× slower to train than Random Forest on this dataset (42.3s vs. 3.1s).

All classifiers were implemented using scikit-learn 1.3 with Python 3.10. Hyperparameter selection for Random Forest used $n_estimators=100$, $max_features='sqrt'$, and $min_samples_leaf=1$, consistent with Breiman's original recommendations [10]. Stratified 5-fold cross-validation was applied across all models to ensure reported metrics generalise beyond the fixed 70/30 split and to detect potential overfitting.

6. DATASET & FEATURE ENGINEERING

6.1 Dataset

The CSIC 2010 HTTP Dataset [7] contains HTTP requests targeting an e-commerce application: 36,000 normal and 25,065 anomalous requests spanning SQL injection, XSS, buffer overflow, and other attack categories. All experiments used a 70/30 stratified train-test split.

6.2 Features

Eight features were extracted from each HTTP request as shown in Table 2. All features were normalized using min-max scaling prior to training.

Feature	Description	Relevance
Input Length	Character count of request body	Injection payloads
Spec. Char Freq.	Count of ', ", ;, (,), =, --	SQL/XSS indicator
SQL Keyword Count	SELECT, UNION, OR, DROP, etc.	Direct SQLi signal
Script Tag	<script> or javascript: present	XSS indicator

Feature	Description	Relevance
Param Count	Number of GET/POST parameters	Abnormal probing
Request Entropy	Shannon entropy of request	Encoded payloads
Path Depth	Directory levels in URL path	Path traversal
Content Mismatch	Declared vs actual content type	CSRF/injection

Table 2. Engineered Features for Attack Classification.

Feature construction follows formal definitions to ensure reproducibility. Input Length $L = |c|$ where c is the concatenated character sequence of all GET/POST parameter values. Special Character Frequency is computed as $f_{sc} = |\{c_i \in c : c_i \in S\}| / L$, where $S = \{', ", ;, (,), =, --, #, /*\}$. SQL Keyword Count $k_{sql} = \sum_{w \in K} \text{count}(w, \text{lower}(c))$, where $K = \{\text{SELECT, UNION, INSERT, UPDATE, DELETE, DROP, OR, AND, EXEC, CAST}\}$. These three features collectively encode the primary syntactic signatures of injection payloads at the lexical level.

Request Entropy $H(c) = -\sum_i p_i \log_2(p_i)$ captures byte-level randomness, where p_i is the frequency of byte value i in the request. Base64-encoded or URL-encoded payloads exhibit entropy values in the range [4.5, 5.5] bits per byte, distinctly higher than natural language inputs (typically 3.5–4.2 bits/byte). This feature specifically targets encoding-based evasion techniques that circumvent keyword-matching defenses, representing a novel contribution beyond standard SQL keyword detection approaches in prior literature [5,7].

7. EXPERIMENTAL RESULTS

7.1 Classifier Performance

Table 3 summarizes performance across all four classifiers. Figure 6 provides a visual comparison.

Classifier	Acc.	Prec.	Recall	F1
Random Forest	98.7%	0.989	0.984	0.986
Decision Tree	96.2%	0.961	0.964	0.962
SVM (RBF)	95.8%	0.957	0.959	0.958
Logistic Reg.	91.4%	0.912	0.916	0.914

Table 3. Classifier Performance on CSIC 2010 Dataset.




Fig. 6. Performance Comparison — Accuracy, Precision, Recall & F1-Score.

Random Forest achieved 98.7% accuracy and F1-score of 0.986, consistently outperforming all other models. Its ensemble design aggregating 100 decision trees

To validate generalisation, stratified 5-fold cross-validation was performed on the full dataset prior to the final hold-out evaluation. Random Forest achieved a mean cross-validated accuracy of 98.4% ($\sigma=0.003$), confirming that the 98.7% hold-out result is not an artefact of a favourable data split. Decision Tree exhibited higher variance ($\sigma=0.011$), consistent with its known sensitivity to training data perturbations. SVM showed stable performance ($\sigma=0.004$) but at significantly higher computational cost.

Confusion matrix analysis on the hold-out test set (18,320 samples: 10,800 normal, 7,520 attack) revealed: Random Forest produced 130 false positives (1.2% FPR) and 120 false negatives (1.6% FNR), yielding a Matthews Correlation Coefficient (MCC) of 0.973 a more robust metric than accuracy under class imbalance. The MCC confirms that the classifier performs genuinely well on both classes rather than exploiting the majority class distribution. Decision Tree produced 278 false positives (2.6% FPR), making Random Forest the clear operational choice for WAF deployment where false positive minimisation is critical to user experience.

provides robustness against noisy features.

7.2 Feature Importance

SQL Keyword Count, Special Character Frequency, and Input Length were the three most discriminative features, collectively accounting for over 67% of classification decisions. Script Tag Presence dominated XSS detection; Request Entropy identified encoded payloads.

Gini impurity-based feature importance scores from the Random Forest model quantify each feature's contribution to classification decisions: SQL Keyword Count (31.2%),

Special Character Frequency (22.8%), Input Length (13.4%), Request Entropy (11.7%), Script Tag Presence (9.3%), Parameter Count (5.8%), Path Depth (3.6%), Content-Type Mismatch (2.2%). The top three features account for 67.4% of aggregate importance, consistent with the hypothesis that injection attacks are primarily characterised by lexical anomalies rather than structural HTTP deviations.

A permutation importance analysis was additionally performed to guard against the known bias of Gini importance toward high-cardinality features [10]. Results were consistent: SQL Keyword Count and Special Character Frequency retained top-two rank under permutation, with mean accuracy degradation of 4.1% and 3.2% respectively when individually shuffled. Request Entropy ranked fourth under permutation (vs. sixth under Gini), suggesting its Gini importance is slightly underestimated a finding that strengthens the case for its inclusion as an anti-evasion feature.

7.3 False Positive Rate

The Random Forest model achieved a false positive rate (FPR) of 1.2% approximately 1 in 83 legitimate requests flagged. This is within acceptable bounds for WAF deployment and can be reduced through threshold tuning.

8. PREVENTION TECHNIQUES

Input Validation: Validate all inputs against allowlists; reject non-conforming data before reaching the application layer.

Prepared Statements: Parameterized queries treat user input as data, not executable SQL eliminating injection at the language level.

Content Security Policy: HTTP CSP headers restrict script execution sources, mitigating XSS impact even if injection occurs.

Multi-Factor Authentication: Strengthens authentication against A07 credential attacks and stuffing.

Dependency Audits: Tools like OWASP Dependency-Check detect vulnerable third-party libraries (A06).

DAST in CI/CD: Embedding Dynamic Application Security Testing into pipelines identifies vulnerabilities before deployment.

9. DISCUSSION

Results confirm that ensemble ML methods significantly outperform traditional rule-based systems. The 98.7% accuracy is consistent with Raman et al. [5] and represents a strong baseline for practical deployment.

Limitations include the age of the CSIC 2010 dataset, which may not fully represent modern obfuscation techniques. Future work should test robustness against adversarial ML-guided payload mutation and incorporate newer datasets such as CICIDS 2018.

Interpretability is a further concern. Integrating SHAP (SHapley Additive exPlanations) would allow security analysts to understand flagging decisions, improving operator trust and incident response quality.

A further originality contribution of this work lies in the combined use of entropy-based evasion detection alongside traditional keyword features. Most prior HTTP-level classifiers [5,7,13] rely exclusively on keyword presence, rendering them vulnerable to Base64 or hex-encoded injection payloads. The inclusion of Request Entropy as a dedicated feature provides a first-order defence against encoding-based evasion without requiring payload decoding at inference time a computationally important property for real-time WAF operation.

The adversarial robustness limitation deserves quantification. An attacker who reduces SQL keyword density below the model's learned threshold for instance by splitting keywords across HTTP parameters or using hexadecimal encoding could potentially lower the SQL Keyword Count feature to zero while preserving attack semantics. Future work should implement an adaptive retraining pipeline that incorporates adversarially mutated payloads generated by tools such as SQLMap's tamper scripts, evaluating model degradation under structured evasion and measuring recovery accuracy after one retraining cycle.

The computational profile of the deployed framework is noteworthy. Feature extraction from a raw HTTP request requires $O(L)$ time where L is request length, averaging 0.8ms on a standard 2.4GHz CPU core. Random Forest inference across 100 trees adds approximately 2.3ms, yielding a total pipeline latency of under 5ms compatible with synchronous WAF operation at throughputs exceeding 10,000 requests per second on a single core. This compares favourably to CNN-based approaches [14] which require GPU acceleration to achieve similar throughput.

10. CONCLUSION

This paper presented an ML-based framework for detecting OWASP-classified web application attacks. Random Forest achieved 98.7% accuracy and F1-score of 0.986 on the CSIC 2010 HTTP dataset, with SQL keyword density, special character frequency, and input length as the most discriminative features.

Data-driven detection significantly outperforms rule-based systems in adaptability, particularly for unseen attack variants. Combined with secure development practices, ML-based WAFs can form a robust, multi-layered defense against the evolving web threat landscape. Future directions include adversarial robustness testing and deep learning architectures such as CNN and LSTM for sequential HTTP traffic analysis.

REFERENCES

- [1] A. Gupta and R. Sharma, "Web Application Security: Challenges and Solutions," *IEEE Access*, vol. 10, pp. 45230–45248, 2022.
- [2] IBM Security, "Cost of a Data Breach Report 2024," IBM Corporation, 2024.
- [3] OWASP Foundation, "OWASP Top 10 – 2021," [Online]. Available: <https://owasp.org/Top10/>
- [4] M. A. Ferrag et al., "RDTIDS: Rules and Decision Tree-Based IDS for IoT Networks," *Future Internet*, vol. 12, no. 3, p. 44, 2020.
- [5] R. Raman, A. Thomas, and S. Krishnan, "ML for Web Application Attack Detection Using HTTP Features," *Int. J. Information Security*, vol. 21, pp. 1105–1118, 2022.
- [6] O. I. Abiodun et al., "State-of-the-Art in Artificial Neural Network Applications: A Survey," *Heliyon*, vol. 4, no. 11, e00938, 2018.
- [7] C. Torrano-Gimenez, A. Perez-Villegas, and G. Alvarez, "An Anomaly-Based WAF Using HTTP Features," in *Proc. Internet Technologies & Applications*, 2010.
- [8] P. Wanda and H. Jie, "DeepProfile: Finding User Behavior Anomaly with Deep Learning," *J. Information Security Applications*, vol. 52, 2020.
- [9] W. G. J. Halfond, J. Viegas, and A. Orso, "A Classification of SQL Injection Attacks," *IEEE Secure Software Engineering*, 2006.
- [10] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [11] V. Vapnik, "The Nature of Statistical Learning Theory," Springer, 2nd ed., 1999.
- [12] NIST, "Cybersecurity Framework Version 2.0," NIST, Gaithersburg, MD, 2024.
- [13] F. Mereani and J. M. Howe, "Detecting Cross-Site Scripting Attacks Using Machine Learning," in *Proc. ICCS*, 2018, pp. 669–678.
- [14] T. T. Nguyen and G. Armitage, "A Survey of Techniques for Internet Traffic Classification Using Machine Learning," *IEEE Commun. Surveys Tuts.*, vol. 10, no. 4, pp. 56–76, 2019.