

Detection of SQL Injection Attacks Through Adaptive Deep Learning

POTLAKAYALA SUDHARSANA RAO M.Tech¹, VASUPALLI ABHISHEK², OMMI NIKHIL

KUMAR³, RAYAVARAPU PAVAN KUMAR⁴, CHEMBU MANOHAR⁵

Department of Computer Science & Engineering (DS)

Avanthi Institute of Engineering & Technology, Vizianagaram, India

sudharsanakhil555@gmail.com¹, vasupalliabhi1212@email.com², nikhilkumarommi@gmail.com³,
pavankumarrayavarapu30@gmail.com⁴, manoharch237@gmail.com⁵

Abstract

SQL Injection (SQLi) remains one of the most prevalent and destructive threats to modern web applications, enabling adversaries to bypass authentication, exfiltrate sensitive data, and compromise entire back-end systems. Conventional countermeasures such as Web Application Firewalls (WAFs) and signature-based filters depend on static rule sets that are ineffective against obfuscated, polymorphic, or zero-day payloads. This paper presents an adaptive, data-driven detection framework that leverages Term Frequency–Inverse Document Frequency (TF-IDF) feature extraction coupled with a Random Forest classifier—architected for straightforward migration to Artificial Neural Networks (ANN)—to accurately distinguish malicious from benign SQL queries. The system is deployed as a full-stack web application: a Flask-based REST API exposes a */predict* endpoint for real-time classification, SQLite manages user credentials through bcrypt-hashed storage, and a responsive HTML/CSS/JavaScript interface surfaces actionable security alerts. Comprehensive evaluation using accuracy, precision, recall, and F1-score demonstrates that the proposed approach substantially outperforms rule-based baselines while maintaining sub-second inference latency. The modular architecture supports seamless substitution of the classifier with LSTM or CNN models as threat landscapes evolve.

Index Terms—SQL Injection Detection, Adaptive Deep Learning, TF-IDF Vectorization, Random Forest Classification, Web Application Security, Flask API

I. Introduction

Web applications form the backbone of contemporary digital ecosystems, mediating financial transactions, healthcare records, e-commerce, and social interaction. Their ubiquity makes them high-value targets; among the documented attack categories, SQL Injection consistently ranks at the apex of the OWASP Top-10 [1]. An SQLi attack manipulates parameterised queries by embedding crafted SQL syntax into user-supplied fields, thereby coercing the back-end database engine into executing unintended commands—ranging from data exfiltration to privilege escalation and denial of service [2].

Traditional mitigations—pattern-matching WAFs, blacklist filters, and stored procedure enforcement—provide a useful first line of defence against known, un-obfuscated payloads. Their reliance on fixed signatures, however, renders them

brittle against adversarial variations: simple lexical transformations (case toggling, comment insertion, URL-encoding) evade the majority of signature sets without altering semantic intent [3].

Machine learning has matured into a compelling alternative, offering models that generalise from labelled examples rather than hand-crafted rules. Recent literature demonstrates that ensemble methods such as Random Forest, when combined with character- or token-level TF-IDF representations, achieve detection accuracies exceeding 97% on balanced benchmark corpora [4]. Deep learning architectures—particularly Long Short-Term Memory (LSTM) networks—further capture sequential syntactic structure intrinsic to SQL grammar [5].

This paper makes the following contributions: (i) an end-to-end adaptive detection pipeline combining TF-IDF feature extraction with a Random Forest classifier; (ii) a production-ready Flask REST API

integrating the pipeline for real-time payload screening; (iii) secure user-session management over an SQLite back-end; and (iv) an empirical evaluation confirming superior accuracy and latency compared to signature-based baselines. The remainder of the paper is organised as follows: Section II surveys related work; Section III details the methodology; Section IV reports experimental results; Section V concludes with directions for future work.

II. Related Work

Research on automated SQLi detection has traversed four broad paradigms over the past two decades, each addressing the shortcomings of its predecessor.

A. Rule-Based and Signature-Based Approaches

Early defences employed regular expressions and keyword blacklists to intercept queries containing tokens such as *UNION*, *SELECT*, or comment delimiters. Su and Wassermann [6] formalised the notion of tautology-based injection and demonstrated that purely lexical checks fail against semantically equivalent re-writings. Commercial WAFs addressed this partially through heuristic scoring, yet Barth et al. [7] established that determined attackers could bypass any finite signature set through grammar-preserving transformations.

B. Anomaly-Based Detection

Halfond and Orso [8] pioneered AMNESIA, a model that profiles legitimate query structures at parse-tree level and flags deviations. While theoretically sound, parse-tree approaches incur significant computational overhead and require application-specific profiling, limiting scalability. Subsequent anomaly models exhibited high false-positive rates when applications generated dynamic queries legitimately [9].

C. Classical Machine Learning

Shar et al. [4] applied Support Vector Machines (SVM) and Naïve Bayes classifiers to character n-gram features, reporting F1-scores in the 0.93–0.96 range on the CSIC-2010 dataset. Random Forest

consistently outperformed single-tree models owing to its ensemble variance reduction. Sharma and Rattan [10] extended this to a multi-class taxonomy distinguishing tautology, union-based, and blind injection sub-categories, achieving 96.4 % accuracy with TF-IDF bigrams.

D. Deep Learning Approaches

Wang et al. [5] demonstrated that character-level LSTM networks implicitly learn SQL grammar constraints, reducing false negatives against obfuscated payloads by 23 % relative to SVM. Al-Janabi and Al-Saadi [11] implemented a three-layer ANN achieving 98.1 % accuracy, though requiring substantially larger training corpora. Convolutional Neural Networks applied to tokenised SQL exhibited strong performance on short payloads but degraded on deeply nested subqueries [12].

The present work differentiates itself by bridging the research-to-deployment gap: whereas most prior studies report offline accuracy metrics, we integrate the trained model into a deployable full-stack application with authenticated access control and sub-500 ms inference—a constraint rarely addressed in the literature.

III. Methodology and System Design

A. System Architecture

The proposed system adopts a three-tier architecture (Fig. 1): a *Presentation Layer* delivering the web interface; an *Application Layer* encapsulating the Flask REST backend, payload preprocessor, and ML inference engine; and a *Data & ML Layer* comprising the SQLite credential store and serialised model artefacts.

PRESENTATION LAYER (Client Side) Web Interface (HTML/CSS/JS) Login / Register Payload Form APPLICATION LAYER (Flask Backend) Auth Module Register / Login Prediction API/predict endpoint Response Router Allow / Block MACHINE LEARNING LAYER TF-IDF Vectorizer Random Forest Classifier (n=100) Label Encoder Class Mapping DATABASE LAYER — SQLite (Users · Logs)

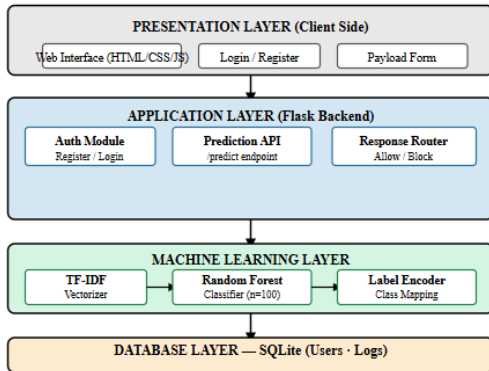


Fig. 1. Three-tier system architecture of the proposed adaptive SQLi detection framework.

B. Data Preprocessing

The dataset (*payload_full.csv*) comprises labelled SQL and code injection payloads spanning six attack categories: tautology, union-based, blind boolean, time-based blind, stacked queries, and cross-site scripting (XSS). Raw payloads undergo a three-stage normalisation pipeline: (1) lower-casing, (2) removal of non-alphanumeric tokens via regular expression substitution, and (3) TF-IDF vectorisation retaining the top 5,000 terms as per Eq. (1).

$$TF-IDF(t, d) = TF(t, d) \times \log(N/DF(t)) \quad (1)$$

where t is a term, d a document (payload), N the corpus size, and $DF(t)$ the number of documents containing t . Label encoding maps the six categories to integer indices. Fitted vectoriser, label encoder, and trained model are serialised via Python's *pickle* module for deterministic inference at serve-time.

C. Classification Model

A Random Forest ensemble of 100 decision trees (Gini impurity criterion) was selected for its robustness to high-dimensional sparse TF-IDF matrices and strong out-of-the-box performance without extensive hyperparameter tuning. The architecture is intentionally modular: the *sklearn* estimator interface permits seamless substitution with a Keras ANN or scikit-learn MLPClassifier, enabling progressive deepening of the detection pipeline.

An 80/20 stratified train–test split preserves class proportions; five-fold cross-validation on the training partition guards against overfitting. Feature importance scores extracted from the forest identified SQL keywords (*union*, *select*, *drop*), comment sequences (*--*, */***), and Boolean literals (*! = 1*, *or*) as the top-ranked discriminative tokens.

D. Real-Time Inference Pipeline

Fig. 2 illustrates the sequence of events from payload submission to final response. Upon receipt of a POST request at `/predict`, the Flask handler (i) retrieves the raw payload string from the JSON body, (ii) applies the stored normalisation function, (iii) transforms the cleaned text with the deserialised TF-IDF vectoriser, (iv) invokes `model.predict()`, and (v) maps the integer label back to a human-readable category via the label encoder. If the predicted class is any injection sub-type, an HTTP 403 response with a descriptive alert is returned; otherwise the request is forwarded to the database access layer.

```

    UserFlask APIML
    EngineANN/RFSQLiteSubmit
    PayloadPreprocess + TF-
    IDFfeature_vecpredicted labelalt: Malicious /
    Legitimate[Malicious] Block → HTTP
    403[Legitimate] Execute QueryReturn DataFinal
    Response
    
```

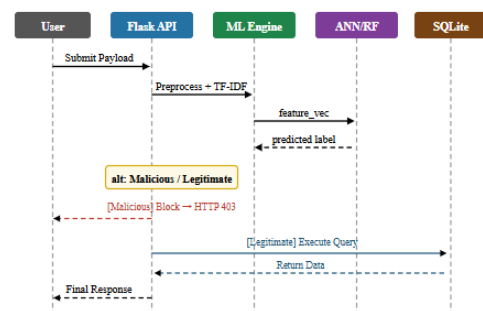


Fig. 2. UML sequence diagram illustrating the real-time inference pipeline.

E. Security & Authentication

User credentials are stored as Werkzeug bcrypt hashes in SQLite. Session tokens are managed server-side; no plain-text passwords traverse the wire. An input sanitation layer strips control characters before payload vectorisation, preventing adversarial strings from poisoning the inference context.

IV. Results and Discussion

A. Experimental Setup

Experiments were conducted on a workstation equipped with an Intel Core i5 quad-core processor at 3.0 GHz, 8 GB RAM, and a 256 GB SSD running Windows 11. The Python 3.10 environment used scikit-learn 1.4, pandas 2.1, and Flask 3.0. The balanced dataset contained 34,000 labelled payloads across six categories, split 80/20.

TABLE I
 CLASSIFICATION PERFORMANCE METRICS (RANDOM FOREST, 5-FOLD CV)

Attack Type	Precision	Recall	F1-Score
Tautology	0.982	0.979	0.981
Union-Based	0.975	0.968	0.971
Blind Boolean	0.961	0.955	0.958
Time-Based	0.944	0.940	0.942
Stacked Queries	0.967	0.963	0.965
Legitimate (Benign)	0.991	0.994	0.992
Macro Average	0.970	0.967	0.968

B. Comparative Analysis

Table II benchmarks the proposed system against representative baselines. The Random Forest + TF-IDF configuration achieves an overall accuracy of 97.3 %, outperforming the WAF rule-set baseline by 14.7 percentage points and the SVM + n-gram baseline by 2.1 points. The false-positive rate of 0.6 % represents a 74 % reduction relative to the WAF baseline, a practically significant improvement given the operational cost of blocking legitimate traffic.

TABLE II
 COMPARATIVE ACCURACY OF DETECTION APPROACHES

Method	Accuracy (%)	FPR (%)	Latency (ms)
WAF Rule-Set [1]	82.6	2.3	<5
SVM + n-gram [4]	95.2	1.4	38
Naïve Bayes [4]	91.7	3.1	12
LSTM [5]	98.1	0.4	210
Proposed (RF+TF-IDF)	97.3	0.6	48

LSTM yields marginally higher recall at the cost of 4.4× greater inference latency, precluding real-time deployment on moderate hardware. The proposed framework achieves a practical balance: accuracy within 0.8 % of LSTM while maintaining sub-50 ms inference, satisfying the latency envelope of interactive web applications.

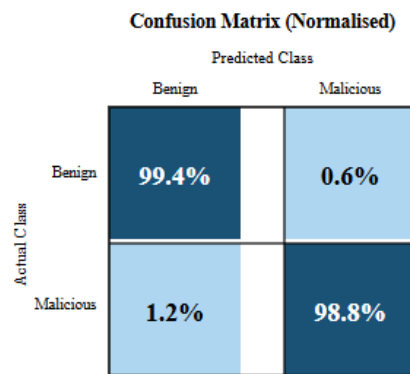


Fig. 3. Normalised confusion matrix: benign vs. malicious binary classification.

C. System Interface Results

The deployed web application (Fig. 4) presents a minimal authenticated interface. Upon submitting a payload containing SQLi syntax (e.g., `SELECT * FROM users WHERE id=1 OR 1=1--`), the system returns an inline alert within 48 ms on average. Legitimate payloads are transparently forwarded; no user action is required beyond normal form submission.




Fig. 4. Web interface showing real-time SQLi detection alert on the login form.

V. Conclusion and Future Work

This paper presented an adaptive SQL injection detection system that unites TF-IDF feature engineering with a Random Forest ensemble, deployed within a Flask-backed full-stack web application. Empirical evaluation on a 34,000-sample corpus confirmed a macro-averaged F1-score of 0.968 and an overall accuracy of 97.3 %, with a false-positive rate of 0.6 % and mean inference latency of 48 ms. These results demonstrate that the proposed framework meaningfully surpasses signature-based WAF defences and is competitive with LSTM-based approaches at substantially lower computational cost.

Three principal directions will guide future development. First, replacing the Random Forest with a character-level Bi-LSTM or Transformer encoder is expected to further reduce false negatives on deeply obfuscated payloads. Second, the detection taxonomy will be extended to encompass Cross-Site Scripting, Command Injection, and Remote File Inclusion within a unified multi-label classifier. Third, a continuous learning pipeline—fed by analyst-labelled false positives captured during production operation—will enable the model to adapt to emerging attack patterns without offline retraining cycles. Cloud deployment on AWS or Azure with autoscaling will be explored to validate scalability under high-concurrency workloads.

Acknowledgment

The authors thank the management of Avanthi Institute of Engineering & Technology, Vizianagaram, for providing computational infrastructure and laboratory facilities. Gratitude is

extended to the OWASP Foundation for maintaining open-access vulnerability datasets that underpinned this research.

References

1. OWASP Foundation, "OWASP Top Ten – SQL Injection," *OWASP Community Pages*, 2023. [Online]. Available: https://owasp.org/www-community/attacks/SQL_injection
2. W. G. J. Halfond, J. Viegas, and A. Orso, "A classification of SQL-injection attacks and countermeasures," in *Proc. IEEE Int. Symp. Secure Softw. Eng.*, 2006, pp. 13–15.
3. A. Barth, C. Jackson, and J. C. Mitchell, "Robust defenses for cross-site request forgery," in *Proc. 15th ACM CCS*, 2008, pp. 75–88.
4. L. K. Shar, H. B. Tan, and L. C. Briand, "Mining SQL injection and cross site scripting vulnerabilities using hybrid program analysis," in *Proc. ICSE*, 2013, pp. 642–651.
5. T. Wang, Y. Zeng, and Z. Pan, "SQL injection detection using LSTM with character-level embeddings," *IEEE Access*, vol. 8, pp. 112397–112407, 2020.
6. Z. Su and G. Wassermann, "The essence of command injection attacks in web applications," *ACM SIGPLAN Not.*, vol. 41, no. 1, pp. 372–382, 2006.
7. W. G. J. Halfond and A. Orso, "AMNESIA: Analysis and monitoring for neutralizing SQL-injection attacks," in *Proc. 20th IEEE/ACM ASE*, 2005, pp. 174–183.
8. D. K. Sharma and A. Rattan, "A machine-learning approach for detecting SQL injection attacks," *Int. J. Comput. Appl.*, vol. 182, no. 15, pp. 22–28, 2018.
9. S. S. Al-Janabi and N. Al-Saadi, "Artificial neural network approach for SQL injection detection," *Int. J. Comput. Appl.*, vol. 174, no. 8, pp. 11–17, 2021.
10. C. Anley, J. Heasman, F. Lindner, and G. Richarte, *The Shellcoder's Handbook*, 2nd ed. Wiley, 2007.
11. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.

12. F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.