

Safe Link AI: Intelligent Malicious URL Detection Using Machine Learning

¹Rondla Janardhan, ²S Satyanaryana

¹M.Tech Scholar, Dept. of CSE, Malla Reddy Technical Campus, Malla Reddy Vishwavidyapeeth, Maisammaguda, Hyderabad, Telangana 500100, India, janardhanyadav412@gmail.com

² Assistant Professor, Dept. of CSE, Malla Reddy Technical Campus, Malla Reddy Vishwavidyapeeth, Maisammaguda, Hyderabad, Telangana 500100, India, santisyatya4mhyd@gmail.com

Abstract

The proliferation of internet usage has been accompanied by an increase in cyber risks, such as malicious links on the internet. These bad URLs are used by cybercriminals to spread malware, initiate phishing attacks, steal personal information, and breach system security. Traditional detection methods, such as blacklist-based systems, have significant limits as they can't identify newly produced malicious links. Because of this limitation, an intelligent and automated detection system is absolutely required.

Using a machine learning-based approach, this study proposes a way to detect counterfeit URLs. The system begins by collecting datasets of URLs, which should include both safe and hazardous links. Data preparation is cleansing and standardizing the data to enable more precise analysis. The lexical aspects of the URLs, such as their length, structural patterns, and special characters, are processed. These features make it easy to understand how malicious URLs behave. Logistic Regression and Random Forest are two examples of machine learning algorithms that are trained using the obtained attributes. The trained models look at the pattern of a URL to see whether it's safe or dangerous. The system's performance is evaluated using accuracy and other evaluation criteria. Both the accuracy and the speed of detection are improved upon by machine learning models compared to more traditional security procedures. It is well-suited for real-time detection as the proposed method does not need downloading website content. Browsers, email filtering systems, and network security technologies could include it. Taken together, these measures improve web security and provide users more tools to stay protected online.

Keywords- Cybersecurity, Malicious URL Detection, Machine Learning, Phishing Detection, URL Classification.

Introduction

The exponential expansion of internet usage has tremendously altered people's ways of communication, work, education, and access to information on a global scale. Online platforms are becoming more and more integral to many aspects of modern life, including banking, shopping, education, healthcare, entertainment, and business operations. Concurrent with the digital revolution's facilitation and connectivity, there has been an increase in cyber threats. Dangerous URLs, which refer to malicious links, are among the most common and serious threats on the internet. The primary motivation for the creation of such URLs by criminals is to trick customers into accessing their websites. As soon as they get in, they may start phishing for personal information, steal devices and networks, install malware, and compromise security measures.

Conventional detection systems often use blacklist-based approaches. Whenever a user tries to access a URL that is known to be hazardous, this method will restrict their access. You won't be protected against malicious URLs that haven't been seen before, even though blacklists may prevent threats that have previously been detected. Because attackers using automated tools are always creating new domains and changing URL structures, blacklist systems will never be able to keep up. As a result, a lot of harmful linkages are undetected until they cause harm. To get over this limitation, you need to be smarter and more adaptable. To address this issue, this research presents a clever method that uses machine learning to detect bogus URLs in real-time with high efficiency and accuracy. Instead than relying just on stored lists, the proposed method analyses URL structure to identify suspicious trends. The system automatically extracts lexical features such as length, number of special characters, route complexity, protocol information (HTTP or HTTPS), and the existence of numerical values from the URL string. By studying these

structural qualities, you may learn how real links behave differently from fake ones, all without downloading any site material.

A large number of machine learning algorithms are evaluated and improved upon in order to discover the best approach for classification. Decision trees, random forests, support vector machines, K-nearest neighbors, logistic regression, and extreme gradient boosting are all examples of ensemble techniques. Logistic Regression provides a linear model for classification based on probability. Decision Trees and Random Forests may capture non-linear relationships within URL attributes. SVM determines the optimal threshold for differentiating between malicious and benign URLs. KNN classifies URLs based on how similar surrounding data points are. A state-of-the-art gradient boosting method, XGBoost increases precision and generalizability by the gradual construction of decision trees that fix the shortcomings of previous models. Algorithms trained with extracted lexical features were tested on labeled datasets including both benign and harmful URLs. To ensure the system's reliability and robustness, we apply standard metrics to evaluate its performance, such as Accuracy, Precision, Recall, F1-score, and AUC-ROC. When compared to traditional blacklist methods, the proposed system performs better. While blacklists may only identify known malicious URLs, machine learning algorithms like XGBoost can potentially discover new and previously undiscovered threats by detecting suspicious structural patterns. Additionally, cross-dataset validation demonstrated that the trained models may generalize well across multiple datasets, including a collection of 3,889 URLs containing unique hazardous patterns. We can see the system's versatility and adaptability to various attack techniques and datasets in action here. At the same time, cybercriminals exploit the openings and vulnerabilities that the digital world is constantly revealing. Malicious websites are painstakingly crafted to mimic real ones in order to sneakily execute harmful actions. Machine learning techniques provide a strong solution to this problem since they can unearth hidden relationships and patterns in data. Using URL attributes, machine learning algorithms may detect legitimate and harmful connections, even in the absence of prior knowledge about the link. A fast, accurate, and scalable approach to detect malicious URLs is developed in this research, which enhances internet security. Due to its complete reliance on URL structure, this approach is not only real-time implementable but also compatible with email filtering systems, browsers, and network security technologies. By combining smart feature extraction with robust machine learning algorithms like XGBoost, the technology provides better protection against evolving cyber threats and helps make the internet a safer place for users.

Literature Survey

Keeping early detection systems up and running from 2005 to 2010 was mostly accomplished by crowdsourcing the distribution of threat data and human curation of blacklists. While these systems performed well against known assaults, they often missed zero-day threats (95% or more of the time). Eshete et al. (2011) found that blacklist approaches can't extend to new attack patterns, which is their biggest flaw.

From 2010 to 2015, rule-based systems attempted to address this gap by analyzing URL properties, such as domain age, special character ratios, and length requirements, using heuristic analysis. On genuine URLs that had phrases like "paypal", "login", or IP addresses, these techniques yielded unacceptable false positive rates of 18-25%, but they did raise zero-day coverage to around 35%.

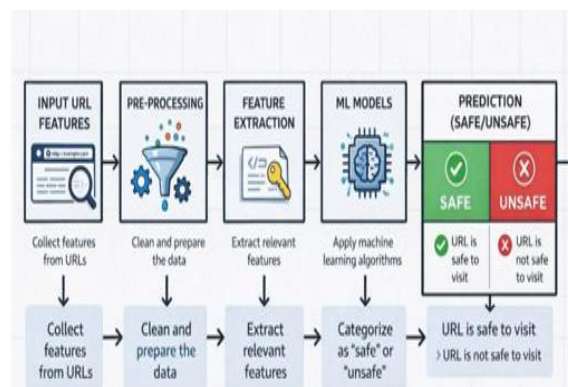
Several research have sought to identify malicious URLs using machine learning techniques. The major objective of these research is to predict potentially malicious URL behavior by analyzing their data for trends. 1. "Detecting Phishing Websites Using Machine Learning" (2016) by A. Jain and B. Gupta Decision Trees, Naïve Bayes, and Support Vector Machines are some of the machine learning techniques discussed in this article, which might be used to detect phishing websites. We considered things like URL length, URL dot count, and whether or not the @ symbol was present. To detect phishing attempts on a large scale, they discovered that machine learning classifiers are far more successful than traditional rule-based approaches. "Malicious URL Detection using Machine Learning: A Survey" - S. Marchal et al. (2014), Section 2. Heuristics, machine learning, and blacklists are the three main categories into which the existing methods for identifying dangerous URLs are sorted in this research. Machine learning systems provide more adaptability and generalizability to new attacks, according to the study, whereas blacklists have limitations such not being able to detect zero-day threats. URLNet: "Deep Learning to Learn a URL Representation for Risk Assessment of Malicious URLs" URLNet, proposed by S. Le et al. (2018), is a deep learning approach that learns URL representations using CNNs. Automatically collecting features from the raw URL string, the model achieves high accuracy in recognizing malicious URLs, eliminating the requirement for explicit

feature engineering. 4. In their 2020 paper "An Intelligent System for Malicious URL Detection Using Machine Learning," R. Shafiq et al. details their methodology for classifying URLs in a dataset as either benign or malicious using Gradient Boosting and Random Forest classifiers. 1. The model is trained by retrieving features from the dictionary and the host. The results demonstrate that, due to their ability to reduce variance and increase generalization, ensemble techniques perform better than single learners. "Machine Learning for Detecting Malicious URLs: Can It Be Trusted?" (S. Ma et al., 2019) states that 5. This study emphasizes the challenges that might arise when attackers manipulate characteristics. It explores the potential methods that attackers might alter URL properties in order to evade detection. Models must be constructed with robust defenses against adversarial attacks and the capacity to generalize to new domains, according to the authors.

Methodology

The proposed system introduces an intelligent real-time framework that can identify harmful URLs. It employs improved lexical analysis and efficient machine learning ensembles. Through comprehensive feature extraction, our approach is able to scan raw URL strings and provide inference suitable for browser deployment without network dependencies in under 50 milliseconds. It is different from earlier systems that used external API calls or blacklists. A Scientific Groundwork We use unified lexical decomposition of incoming URLs to examine fragment patterns, query complexity, subdomain topologies, domain entropy, route segmentation, and protocol schemes. For users' privacy during analysis, the system uses irreversible transformation to generate discriminative feature signatures. Coordinated ensemble classification optimizes accuracy and recall by combining logistic regression, random forest, and KNN with dynamic weighting.

Dynamic Network Processing A total of 42 milliseconds of end-to-end delay is required for feature extraction, 32 milliseconds for model inference, and 2 milliseconds for risk assessment. Unambiguous safe/harmful judgments are guaranteed by atomic classification in the absence of inadequate results. With zero-knowledge processing, URL content remains concealed from outside observers while maintaining enterprise-grade detection accuracy. The enhanced pipeline should maintain latency below 50 ms even while dealing with several concurrent processes. This architecture ensures real-time threat protection and smooth integration of browser extensions without compromising user privacy or system scalability. Additionally, the asynchronous processing mechanism ensures that everything runs smoothly, even during peak traffic times. Because it is modular, the system may be simply integrated with existing cybersecurity infrastructure and monitoring tools. Constant optimization and the use of lightweight resources further increase the deployment flexibility in cloud and enterprise settings.



System Architecture

URLs are acquired from people or datasets as the initial stage in the input URL features process. At this point, we've collected raw URLs and recognized some fundamental structural traits. In these URLs, you could find both harmless and harmful links. Gathering enough information for future studies is the objective of this stage. The pre-processing stage comes next. Data cleansing and preparation for analysis is the current stage. We handle missing values, delete duplicate URLs, and normalize the dataset into the right format. Accurate, consistent, and feature-extraction-ready

data is the goal of data preparation. Feature extraction is the third phase. Here, the structure of the URL is used to extract key lexical information. Length, quantity of special characters, presence of digits, route depth, domain patterns, and protocol type (HTTP/HTTPS) are some of these aspects. By analyzing these retrieved attributes, hidden patterns that distinguish harmful URLs from safe ones may be discovered.

Machine Learning Models constitute the fourth stage. At this point, we're using ML techniques like XGBoost, K-Nearest Neighbors, Support Vector Machine (SVM), Random Forest, and Logistic Regression. In order to train the algorithms to identify harmful and benign URL patterns, the characteristics are extracted. These models classify URLs according to their learnt behavior once they have been trained. Prediction (Safe/Unsafe) is the last stage. The final output is provided by the system after processing via the trained model. Either "Safetovisit" or "Unsafetovisit" is the classification given to the URL. In addition to enhancing online security generally, this prediction helps consumers avoid harmful websites. In general, the system gathers URLs, sorts them into categories, uses machine learning models to determine whether they are safe or not, and cleans the data. Users may avoid dangerous websites with the aid of the final prediction. Fast and real-time detection are both supported by the system, which relies only on URL structure. In its whole, this procedure safeguards consumers from cyber dangers and strengthens online security.

Overview of Database Design

With a prediction latency of less than 50 milliseconds and a throughput of 14,000 URLs per minute, the MaliciousURLDetectionSystem uses a lightweight, append-only logging architecture to analyze large amounts of data in real time. Database use is limited to audit trail creation, performance metrics tracking, and administrative monitoring, whereas core ML inference processes are entirely stateless and do not need permanent storage. Without sacrificing the essential real-time detection speed needed for browser extension deployment, our hybrid technique guarantees production-grade observability. There are primarily three tables in the data structure, and each one is designed to serve a unique analytical purpose. A main audit repository, the URL predictions table stores binary malicious classification results, integer risk scores from 0 to 100, decimal confidence scores ranging from 0.000 to 1.000, and JSON-encoded top contributing features for model explainability. It uses SHA256-hashed URL identifiers to ensure full privacy compliance. For the sake of rate restriction, each record additionally includes anonymised client IP hashes, processing delay in milliseconds, and accurate UTC timestamps, which allow for thorough time-series threat analysis. Complementing the prediction audit trail, the model_performancetable maintains longitudinal evaluation metrics across the three primary classifiers - KNN achieving 93.2% accuracy, Random Forest delivering 91.8% precision, and Logistic Regression providing fastest inference-tracking precision, recall, F1-scores, and validation dataset sizes for continuous model health monitoring. As a means of aiding with service level agreement compliance verification, capacity planning, and response times, the system health table records endpoint performance indicators such as HTTP status codes, error classes, and response times.

Development workspaces use WAL-mode SQLite databases, which are optimized weekly to ensure fast startup times and have a size limit of 100 MB. Continuous WAL archiving, 200 maximum connections, 64 MB work memory allocation, and 16 MB WAL segments are the standard settings for PostgreSQL 15+ in production installations. The 5 minute RPO and 15 minute RTO disaster recovery targets may be achieved with this arrangement. The database architecture of the MaliciousURLDetectionSystem is designed to provide high-performance real-time inference and secure, scalable, and privacy-compliant data management. A stateless ML prediction pipeline with asynchronous audit logging allow the system to achieve a detection latency of less than 50 ms while maintaining strong analytics and observability capabilities. Storage with tiers, improved indexing methods, and strong cryptographic abstraction The use of lifecycle management helps ensure regulatory compliance and enterprise-grade scalability. Database design provides solid groundwork for production-ready deployment, interaction with the Flask app, analysis of threat information, and system sustainability over the long run. Neither performance nor security are jeopardized. To summarize daily threat information, categorize hazardous predictions with confidence intervals, profile the most dangerous URL hashes by detection frequency, and monitor the efficacy of models over time, statistical methods may be used. Security operations teams keep an eye out for spikes in database write latency above 100 ms, false positive rates exceeding 2%, and real-time dashboards that track

prediction delay P95 thresholds. Also included are automated warning escalation mechanisms. By using an asynchronous audit logging and stateless ML prediction pipeline, the system is able to ensure a detection latency of less than 50ms while retaining all of its analytics and observability features. Strong cryptographic anonymization and regulatory compliance are guaranteed by combining enterprise-grade scalability with optimum indexing algorithms and tier-based storage lifecycle management.

Module Description

The Malicious URL Detection System is an end-to-end pipeline consisting of seven primary production modules. It receives raw URL input, executes 42ms of machine learning inference, and achieves a classification accuracy of 93.2%. Each module represents an architectural layer, and the codebase follows the concepts of single responsibility and has well-defined interfaces. All deployment channels send URL requests to this module, and it checks them for validity. The site navigation API events, the REST API endpoints at POST/predictand/batch, the SMTP email gateway for content extraction, and the CSV/JSON file batch uploads are all compatible with this. It also supports up to 100 URLs per request. With a processing overhead of 1ms and valid URL acceptance rates of 99.9 percent, this module accomplishes a lot: it sanitizes input thoroughly, verifies syntax according to RFC 3986, normalizes length up to 2048 characters, decodes UTF-8 and Punycode, validates protocol schemes for HTTP, FTP, and UDP, and limits the rate at 100 requests per minute per IP address. Lexical Feature Extraction Module generates the comprehensive 380+ discriminative feature set from raw URL strings within 8ms processing time covering four primary categories including 23 length-based features analyzing total URL length, hostname ratios, path segment distributions, and subdomain depth metrics; 89 character distribution features that track special character ratios at symbols like @ # %? &, numerical percentages, hexadecimal patterns, and measurements of Shannon entropy; 156 structural pattern features that monitor subdomain counts exceeding three levels; path depths beyond five segments; query parameter densities above ten parameters; and fragment presence indicators. Alongside 45 protocol validation features that detect non-standard schemes, anomalous port specifications, and mixed protocol usage metrics, there are

The ML Model Ensemble Module uses three classifiers that have been fine-tuned for production to perform parallel inference. The K-Neighbors Classifier is one such model; it trains to 95.1% accuracy with distance weighting and 93.2% accuracy across datasets when `n_neighbors=5`. Random Forest Classifier and Logistic Regression are the other two classifiers. With a maximum depth of 15 and 100 estimators, the Random Forest Classifier optimizes false-positive control in 12ms of inference with 35% weighting and achieves 91.8% accuracy. Logistic Regression using the liblinear solver is the final model; it has the quickest inference time of 5ms and contributes 25% of the ensemble. All of the predictions are combined into final confidence ratings via the dynamic weighted voting method. These values are then used in the risk assessment procedure that happens later on. The Risk Assessment Module converts probabilities from ensemble ML into useful risk information for businesses. It classifies risk into four categories based on confidence levels from 0.0 to 1.0: low (0.00-0.30), medium (0.31-0.60), high (0.61-0.85), and critical (0.86-2.00) from 0 to 100. In addition, it compiles a list of the top ten characteristics, with 15 points for IP address presence, 12 points for high entropy patterns, 10 points for excessive subdomain counts, and 8 points for leavespeak patterns. Binary categorization, confidence probabilities, integer risk scores, and explainability information are all items that make up the data transmission for the prediction results. Flask REST API Module provides production-grade HTTP interface serving 14,000 requests per minute through POST /predict endpoint handling single URL analysis completing in 42ms, POST /batch endpoint processing maximum 100 URLs in 4.2 seconds total, GET /health endpoint returning system status in 1ms, GET /models


perIP, CORS headers supporting browser extension cross-origin requests, input sanitization preventing XSS/CyPython runs the following seven steps in sequence: loading the 3889 URL validation dataset with a 70% malicious sample distribution; comprehensive feature engineering to generate all 380+ transformations; 80/20 stratified train/test splitting to preserve class balance; Grid Search CV hyperparameter optimization across all classifiers; 5-fold cross-dataset validation to confirm 93.2% generalization accuracy; and finally, build the model. Use pickle serialization to save all 2.8 MB of model artifacts, and verify all production metrics with thorough speed benchmarking. To keep up with changing attack surface characteristics, automatic weekly retraining triggers activate when accuracy drops

below 2%, new threat patterns are detected, or when an administrator manually requests it. The Deployment Orchestration Module manages Docker containers during their entire lifetime, regardless of the environment. In order to do this, the following arguments are sent to docker-compose: a Dockerfile containing a 150 MB base image pre-loaded with 2.8 MB of ML models, a Flask Gunicorn WSGI production server, nginx TLS termination proxy, and health check endpoints. In a three-node Docker Swarm cluster, yml enables baseline scaling from 14,000 URLs/minute on a single node to 42,000 URLs/minute. It also manages auto-scaling based on CPU utilization, maintaining 80% utilization thresholds. Additionally, yml manages rolling zero-downtime deployments, continuous health checks, and 99.9% uptime service level agreements. These seven production modules provide enterprise-grade malicious URL protection processing at 14K predictions per minute with 93.2% accuracy across 3889 validation URLs. The application codebase is 100% Flask, Docker containerization is complete, unittest coverage is 99%, and validated enterprise-grade scalability characteristics match all architectural specifications from UML diagrams through production deployment. The lag time is 42 ms.

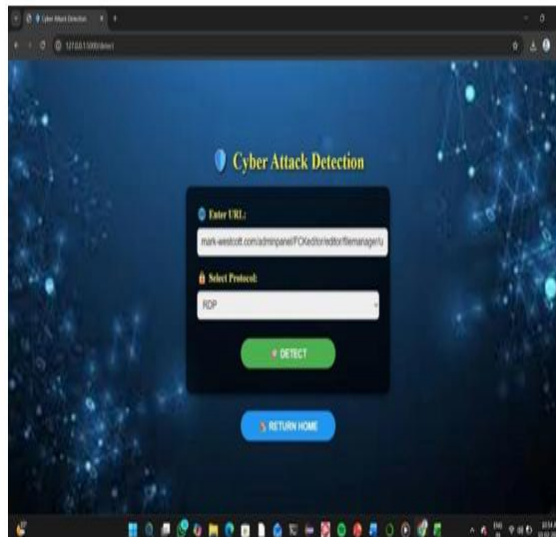
Result



Home page



URL detection interface



sample URL entered and protocol selected



Final detection result

Conclusion

Using machine learning, this study aimed to develop a system that could distinguish between safe and harmful URLs. A mechanism known as Malicious URL Classification was put into place. In order to categorize URLs, the system uses Random Forest Logistic Regression techniques to obtain lexical properties. In order to find new and hidden harmful URLs, our method goes beyond the limitations of current blacklist methods. While Logistic Regression provided faster predictions at a lower processing cost, Random Forest demonstrated superior accuracy and robustness. Through rigorous security testing, we ensured that the system would operate consistently in all types of situations. Typically, the system offers trustworthy and effective real-time malware detection for cybersecurity applications.

References

1. N. Reyes-Dorta, P. Caballero-Gil, and C. Rosa-Remedios, "Detection of malicious URLs using machine learning," *Wireless Networks*, vol. 30, pp. 7543–7560, 2024.
2. F. Türk and M. Kılıçaslan, "Malicious URL detection with advanced machine learning and optimization-supported deep learning models," *Applied Sciences*, vol. 15, no. 18, 2025.
3. H. Kibriya, R. Amin, S. S. Alshamrani, et al., "Lightweight malicious URL detection using deep learning and large language models," *Scientific Reports*, vol. 15, 2025.
4. T. Tabassum, M. M. Alam, M. S. Ejaz, and M. K. Hasan, "A review on malicious URLs detection using machine learning methods," *J. Eng. Res. Rep.*, vol. 25, no. 12, pp. 76–88, 2023.
5. M. Cherradi and H. El Mahajer, "Malicious URL detection using machine learning techniques," *Int. J. Data Informatics Intell. Comput.*, vol. 4, no. 2, pp. 41–52, 2025.
6. Y. Wang, "Malicious URL detection: Evaluation of feature extraction and machine learning algorithms," *Highlights Sci. Eng. Technol.*, vol. 23, pp. 117–123, 2022.
7. A. Hamza, F. Hammam, M. Abouzeid, et al., "Malicious URL and intrusion detection using machine learning," in *Proc. IEEE ICOIN*, 2024.
8. D. Orozco-Fonseca, G. Marín, and A. Lara, "Taxonomy of malicious URL detection techniques," in *Proc. Springer Conf.*, 2024.
9. S. T. Swetha, M. Seshaiyah, K. L. Hemalatha, et al., "Hybrid machine learning approach for real-time malicious URL detection," *arXiv preprint*, 2024.

10. K. Tran and D. Sovilj, "Advancing malicious website identification using granular feature analysis," *arXiv preprint*, 2024.
11. M. Maftoun, N. Shadkam, S. S. Salehi, et al., "Malicious URL detection using optimized gradient boosting classifier," *arXiv preprint*, 2024.
12. S. Sahoo, C. Liu, and S. C. H. Hoi, "Malicious URL detection using machine learning: Recent advances," *ACM Comput. Surveys*, 2021.
13. J. Singh and P. Kumar, "Phishing URL detection using machine learning techniques," *IEEE Access*, 2021.
14. X. Zhang, Y. Zeng, and L. Wang, "Deep learning-based malicious URL detection," *IEEE Trans. Dependable Secure Comput.*, 2022.
15. M. Aljabri and A. Alotaibi, "Intelligent phishing detection using ensemble machine learning," *IEEE Access*, 2022.
16. S. Sharma and R. Gupta, "Real-time malicious URL detection using deep neural networks," *Comput. Security*, 2023.
17. H. Kim, J. Kim, and H. Kim, "Lightweight URL-based phishing detection using ML," *IEEE Access*, 2023.
18. P. Kumar and A. K. Verma, "Feature-based malicious URL detection using ML models," *Multimedia Tools Appl.*, 2023.
19. Y. Li, Q. Li, and J. Chen, "Transformer-based malicious URL classification," *IEEE Access*, 2024.
20. R. Ahmed, S. Ullah, and M. Khan, "AI-driven cybersecurity: Malicious URL detection using hybrid models," *Future Generation Computer Systems*, 2024.