INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 4, April 2026 | ISSN: 2456-4184 | IJNRD.ORG

PHISHGUARD AI: AN ENSEMBLE MACHINE LEARNING FRAMEWORK FOR REAL-TIME PHISHING URL DETECTION

[1]Ms L Shalini,[2]Kothagorla Chandana,[3]Shaik Iffath Imrose,[4]Gajji Arthi Yadav, [5]Shaik Mohammed Naveed

[1]Assistant Professor,[2]Student,[3]Student,[4]Student,[5]Student
[1]Department of Computer Science and Engineering
Bharath Institute of Science and Technology (BIHER), Chennai, India

Abstract- Phishing attacks are still a problem when it comes to cybersecurity. They take advantage of weaknesses in technology. People's minds. The old ways of stopping phishing, like using blacklists and looking for patterns, do not work well against phishing methods that change quickly and avoid being detected. This paper is about Phish Guard AI, a system that uses machine learning to stop phishing attacks. Phish Guard AI combines techniques, including Random Forest, Long Short-Term Memory networks, Convolutional Neural Networks, Bidirectional Encoder Representations from Transformers and a special classifier that looks at the results from all these techniques. The system looks at things about a website, like the URL, the words used in the content and what it looks like to figure out if it is a phishing site. It checks 24 features to make a decision. When we tested Phish Guard AI on a set of 60,000 websites, it was very good at finding the phishing sites. It was right 98.5% of the time. It did not have many false alarms. The numbers were really good: 97.8% precision, 98.2% recall and 98.0% F1 score. Phish Guard AI worked better than each of the techniques it uses.

Keywords— Phishing Detection, Machine Learning, Deep Learning, Ensemble Models, Cybersecurity, URL Analysis, LSTM, CNN, BERT.

I. INTRODUCTION

The rapid growth of digital services, online transactions, and cloud-based platforms has greatly changed the way individuals and organizations interact with technology. Nevertheless, the rapid growth of digital services, online transactions, and cloud-based platforms has greatly opened new doors for cybercriminals to exploit the loopholes in online environments. Of the numerous cybercrimes, phishing is undoubtedly one of the most popular and effective cybercrimes through which cybercriminals have been able to steal important information from unsuspecting victims. Recent global cybersecurity reports have shown that the number of phishing cybercrimes increased by over 220% in 2023 alone, with over 1.2 million unique phishing websites detected every month. These alarming statistics have shown the dangers posed by phishing cybercrimes, which have greatly endangered millions of victims worldwide. These cybercrimes have greatly affected individuals, organizations, and countries at large. Consequently, important sectors such as banking, e-commerce, social media, and enterprises have greatly been affected by cybercrimes. The traditional methods used to identify and prevent phishing attacks, such as blacklisting and rule-based filtering, were effective in detecting and preventing malicious websites. However, the attackers have now started using more sophisticated methods to carry out the attack, and the traditional methods are unable to cope with the situation. For instance, the attackers have started using Unicode characters to create domain names that look similar to legitimate websites. They also use URL shortening techniques to confuse the victim and avoid detection.

Additionally, the attackers have started using DNS manipulation and fast flux hosting techniques to change the server location frequently, making it difficult to identify the phishing sites. Recently, the attackers have also started using artificial intelligence to generate phishing emails and website content. These emails and content look similar to legitimate ones and are difficult to identify. The attackers' new techniques have made the traditional methods less effective in detecting and preventing the attacks. However, machine learning and deep learning methods have emerged as a promising solution to the problem. These methods have proven to be effective in detecting and preventing the attacks by analyzing the hidden patterns in the data. These methods are also able to identify new attacks, unlike the traditional methods. Additionally, the machine learning methods continuously improve the accuracy of the results as the training data increases. However, the accuracy of the results depends on a single machine learning model. Since the attackers use different methods to carry out the attacks, using a single machine learning model would not be effective in detecting and preventing the attacks. However, using a combination of different machine learning models would be effective in detecting and preventing the attacks. This would also provide a better understanding and accuracy in predicting the results. Therefore, in order to overcome the challenges posed by the phishing attacks, this paper proposes an advanced phishing detection system known as Phish Guard AI. The proposed system has the capability to be deployed in the contemporary digital world. The system uses the model-

based intelligent detection approach, in which the system combines several analytical techniques in order to improve the detection of phishing attacks. The system has the advantage of relying on machine learning methodologies in order to improve the detection of phishing attacks.

II. LITERATURE SURVEY

This section reviews relevant research. Phishing schemes deceive individuals into disclosing confidential information via fraudulent websites that resemble authentic platforms, resulting in identity theft and monetary scams. Conventional detection approaches, such as blacklists and heuristic methods, are restricted since they depend on static patterns and cannot swiftly adjust to changing phishing tactics. To address this, the suggested system employs an ensemble machine learning strategy merging Random Forest, XGBoost, and Support Vector Machines to enhance detection precision and reliability. It examines critical aspects like URL length, domain age, SSL certificate validity, special characters, and questionable HTML elements to detect phishing sites. Educated on varied datasets, the model minimises both false positives and negatives while ensuring high reliability. User feedback combined with adaptive learning and continuous threat monitoring improves performance, rendering the system suitable for real-time security uses.[1]

Whereas Social media platforms such as Twitter and Facebook link millions of users, yet they are more frequently attacked by spammers disseminating unwanted ads, harmful links, and misinformation via fake profiles. These actions interfere with user experience, squander system resources, and facilitate the fast dissemination of harmful material. Consequently, identifying spam and fraudulent users has turned into a central research aim in Online Social Networks (OSNs). This research examines key techniques for detecting spam on Twitter and organises them into categories, including detection of fake content, URL-based spam identification, spam in trending subjects, and identification of fake users. It additionally contrasts these techniques by examining factors such as user activity, content traits, graph configuration, and temporal patterns, providing a valuable summary of present strategies in Twitter spam detection [[3]].

Phishing assaults remain a significant threat to cybersecurity, necessitating more precise and effective detection techniques. This research presents an enhanced ensemble stacking framework that integrates sophisticated machine learning models, hybrid feature preprocessing, and meta-learning approaches to boost phishing website identification. The system assesses nine classifiers: XGBoost, CatBoost, LightGBM, Random Forest, Gradient Boosting, Extra Trees, Support Vector Classifier, AdaBoost, and Bagging, on both balanced and imbalanced datasets. Findings indicate that the enhanced stacking model reaches 100% accuracy on two datasets and exceeds 99% accuracy on more intricate datasets, surpassing current methods and showcasing robust dependability for practical phishing detection use cases [4][1]. Phishing emails pose a significant cybersecurity risk by having attackers mask harmful messages as authentic correspondence to obtain sensitive data. This research examines the attributes of phishing emails by utilizing a real-time dataset and sophisticated classification methods to enhance detection precision. It additionally investigates the psychological and social engineering strategies employed by cybercriminals to influence users and disseminate harmful software via email content and types. Furthermore, the study underlines the monetary and reputational threats faced by organizations and stresses the necessity of ongoing user awareness and enhanced security protocols, providing actionable recommendations to safeguard personal and corporate information against advancing phishing attempts [6].

Phishing attacks continue to pose a significant cybersecurity risk as attackers constantly adapt methods to trick users with harmful URLs. This research presents an AI-driven phishing detection system that integrates a character-level Convolutional Neural Network (CNN) with LightGBM in an ensemble framework utilizing 36 lexical, structural, and domain-related URL features. Evaluated on 19,873 URLs, the system reached 99.819% accuracy, 100% precision, 99.635% recall, and 99.947% ROC-AUC, showcasing excellent dependability with minimal false positive rates. Executed with a FastAPI-driven interface for immediate detection, the model surpasses standalone classifiers, as CNN accounts for 60% and LightGBM 40% of the final predictions, rendering it efficient for recognising contemporary phishing methods 7. Phishing continues to be a significant cybersecurity risk as cybercriminals utilise fraudulent websites and emails to capture sensitive data like usernames, passwords, and credit card information. Conventional detection approaches, including blacklist methods, source-code evaluation, and visual similarity techniques, face drawbacks like susceptibility to zero-day attacks, elevated latency, and inadequate compatibility with mobile devices that have restricted resources.

To address these challenges, this research introduces Phish-Jam, a hybrid super learner ensemble mobile app that integrates various machine learning and deep learning methods utilizing URL-based handcrafted attributes and transformer-based text embeddings. The model delivers swift, language-agnostic detection with robust immunity to malware threats, attaining 98.93% accuracy, 99.15% precision, and 99.07% F1-score, thus proving efficient for real-time phishing detection on mobile devices [8]. Phishing is a significant type of online fraud where perpetrators deceive individuals into disclosing confidential personal and account details, and its growing sophistication complicates identification. This research presents an ensemble machine learning method to categorize websites as either legitimate or phishing by utilizing two datasets and feature normalization methods to enhance prediction accuracy. The effectiveness of six classifiers—Decision Tree, Random Forest, Gradient Boosting, XGBoost, AdaBoost, and Multi-Layer Perceptron—was assessed alongside a Voting classifier, which attained the highest results with 97.8% accuracy and robust precision, recall, and F1-score. The suggested method also showed reduced rates of false positives and false negatives, quicker prediction speed, and improved overall effectiveness compared to earlier techniques, proving it to be efficient for dependable phishing website identification [9][4].

Phishing continues to be a significant cybersecurity risk as cybercriminals employ fraudulent websites to deceive individuals into disclosing confidential data. This research examines a dataset comprising 58,645 URLs with 111 attributes to differentiate phishing websites from genuine ones and shows that employing just 14 essential features, a feedforward model attained 94.46% accuracy,

enhancing cost-effectiveness in detection. The study additionally employs sophisticated classifiers like Deep Neural Network (DNN), Wide and Deep, and TabNet to improve performance and presents a novel anti-phishing score metric that assesses outcomes based on false positives and negatives instead of solely on accuracy. Through refined hyperparameter tuning and evaluation on a fresh dataset, the suggested method demonstrates significant generalizability and offers an efficient structure for enhancing phishing detection systems in response to changing cyber threats [5].

III. METHODOLOGY

The Phish Guard AI software utilizes a multi-layer framework encompassing dataset preparation, feature extraction, model training, and ensemble prediction to effectively identify phishing websites. A dataset of 60,000 URLs was established, comprising 30,000 legitimate URLs and 30,000 phishing URLs, sourced from reliable platforms like PhishTank, OpenPhish, Alexa Top Sites, and links that were manually verified. This even dataset enhances the model's capacity to understand significant distinctions between real and harmful websites, guaranteeing dependable classification accuracy. To improve detection precision, the system gathers 24 essential features from four primary categories: URL format, lexical patterns, content of the webpage, and visual traits. URL structure characteristics consist of attributes like URL length, count of subdomains, presence of IP addresses, and accessibility of HTTPS. Lexical characteristics examine questionable keywords, frequency of special characters, and entropy metrics in URLs. Content-centric characteristics assess HTML tag actions, JavaScript interactions, and form-processing parameters, while visual attributes analyze webpage designs, recognize logos, and gauge screenshot resemblance to detect imitation efforts.

These integrated characteristics offer a thorough insight into phishing signs. The model design combines various machine learning and deep learning methods to enhance predictive power. A Random Forest classifier acts as the foundational model for classification, whereas LSTM identifies sequential URL patterns and CNN examines visual resemblances among webpages. These models are integrated through a weighted voting ensemble approach, in which each model's input is based on its validation results. In situations of conflicting predictions, a confidence score is utilized to determine the final outcome, guaranteeing enhanced accuracy and reliability in phishing detection appropriate for practical cybersecurity scenarios.

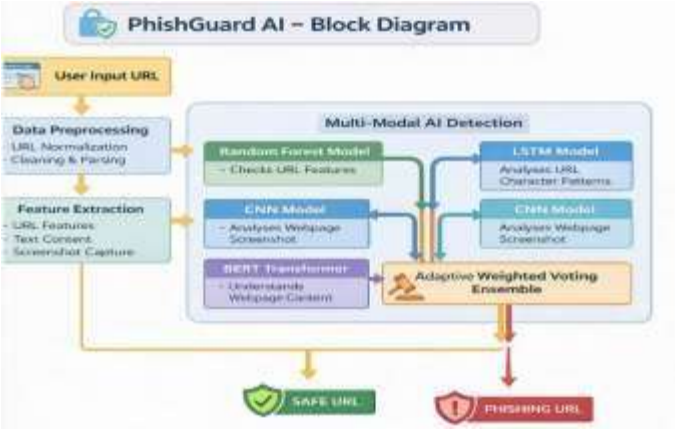

fig.1: System Architecture

3.1 Component for Feature Extraction

The Feature Extraction module collects and converts raw URL and webpage data into significant attributes that assist in recognizing phishing activities. It encompasses various subdomains like URL Structure Examination, Lexical Pattern Assessment, Domain Authentication, Security Certificate Verification, and Redirection Behaviour Evaluation. This module assesses traits such as URL length, character randomness, count of subdomains, existence of suspicious keywords, validity of HTTPS certificates, and redirect chains that could suggest malicious intent. The module offers a robust basis for downstream machine learning models to effectively distinguish between legitimate and phishing websites by producing over 24 structured features from these analyses.

3.2 Module for Machine Learning Models

The Machine Learning Models component constitutes the intelligent core of the system by incorporating various specialized learning algorithms via submodules like Random Forest Classification, Sequential Pattern Detection with LSTM, Visual Similarity Detection utilizing CNN, and Contextual Language Comprehension through BERT. The Random Forest model enables rapid classification and emphasizes the significance of features, whereas the LSTM model identifies sequential relationships in URL character patterns. The CNN model examines similarities in webpage layouts to identify visual impersonation efforts, while the BERT model understands contextual and semantic significance from textual components. Collectively, these models offer complementary advantages that greatly enhance detection precision and flexibility.

3.3 Module for Coordinating Ensembles

The Ensemble Coordination module merges forecasts from every trained model through smart decision-making processes. It encompasses elements like Weighted Voting Aggregation, Confidence Score Assessment, and Dynamic Model Contribution

Modification. Rather than depending on one classifier, this module weights to every model according to its performance and dependability, promoting equitable decision-making. In cases of conflicting predictions, confidence scoring systems identify the most reliable result. This unified strategy strengthens reliability, decreases false positives, and boosts system effectiveness against advancing phishing methods.

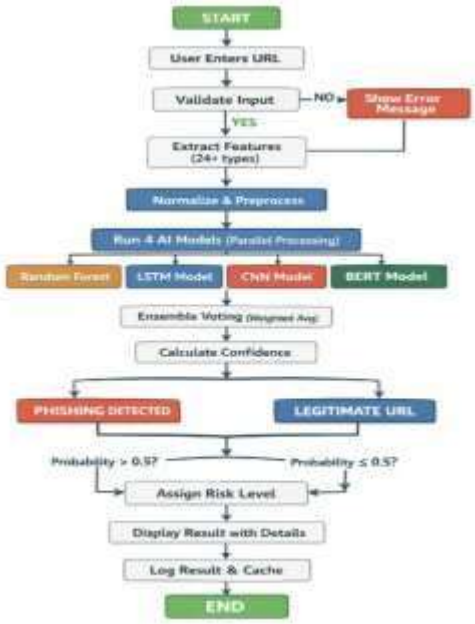

fig.2: Work flow

3.4 RESTful API Service Component

The REST API Service module facilitates smooth interaction between the phishing detection system and external applications via subdomains such as Single URL Analysis Service, Batch Processing Handler, and High-Speed Request Management Engine. This module enables real-time detection of phishing, with response times below 200 milliseconds, and can analyze up to 50 URLs simultaneously. This rapid processing guarantees that the system can be seamlessly incorporated into corporate security platforms, browsers, and monitoring tools that demand swift threat detection abilities.

3.5 Web Interface Component

The Web Interface component offers an engaging platform for users to conveniently utilize phishing detection services. It comprises elements like the User Input Dashboard, Instant Result Visualization Panel, and Phishing Trend Monitoring Dashboard. Users are able to submit questionable URLs and obtain instant classification results along with relevant insights. Visualization dashboards also present detection statistics and trends, enhancing transparency and aiding users in comprehending phishing activity patterns more effectively.

3.6 Module for Data Processing, Deployment, and Monitoring

The Data Processing, Deployment, and Monitoring module guarantees the system functions effectively and stays responsive to new threats. It includes subdomains like Data Preprocessing Pipelines, Deployment in Containers with Docker, Performance Monitoring Systems, and Automated Model Update Processes. This module readies datasets for training, streamlines deployment in various environments, consistently monitors system performance, and allows for automatic updates upon detection of new phishing

patterns. Collectively, these attributes guarantee that the system continues to be scalable, robust, and efficient in practical cybersecurity situations.

IV. Implementation And Results

Phish Guard AI functions as a modular, scalable, and real-time framework for phishing detection that unifies various machine learning and deep learning models within a single architecture. The system is designed to handle substantial amounts of URLs effectively, ensuring high detection precision and minimal response delay. Its modular structure permits each part to function autonomously while aiding a synchronized decision-making process, facilitating adjustment to changing phishing tactics and accommodating large-scale deployment settings.

4.1 Input Processing Layer

The Input Processing Layer acts as the system's entry point and is tasked with accepting URLs entered via the web interface or the REST API. This layer carries out preliminary validation and sanitization to guarantee that incoming requests are correctly structured and secure for additional processing. It accommodates both individual request assessments and bulk submissions, enabling the

platform to manage real-time detection situations as well as extensive scanning tasks. Filtering out malformed inputs at the beginning of the workflow enhances processing efficiency and avoids unnecessary computational load in subsequent modules.

fig.3: Shows the User interface of the webpage

4.2 Subsystem for Feature Extraction

The Feature Extraction Subsystem is essential for converting unprocessed URL and webpage data into organized formats appropriate for smart analysis. It starts with URL analysis, where elements like protocol type, domain format, path segments, and query parameters are scrutinized for unusual traits. Next is a lexical analysis that assesses entropy levels, the frequency of keywords, and the irregular distribution of characters typically linked to phishing attempts. The subsystem additionally analyses HTML components and JavaScript actions to identify harmful scripts or misleading webpage layouts. Moreover, screenshots of web pages are taken and analysed to aid in visual similarity detection, allowing the system to recognize duplicated layouts or mimicked brand identities. Collectively, these procedures create a substantial feature set that enhances model predictions.

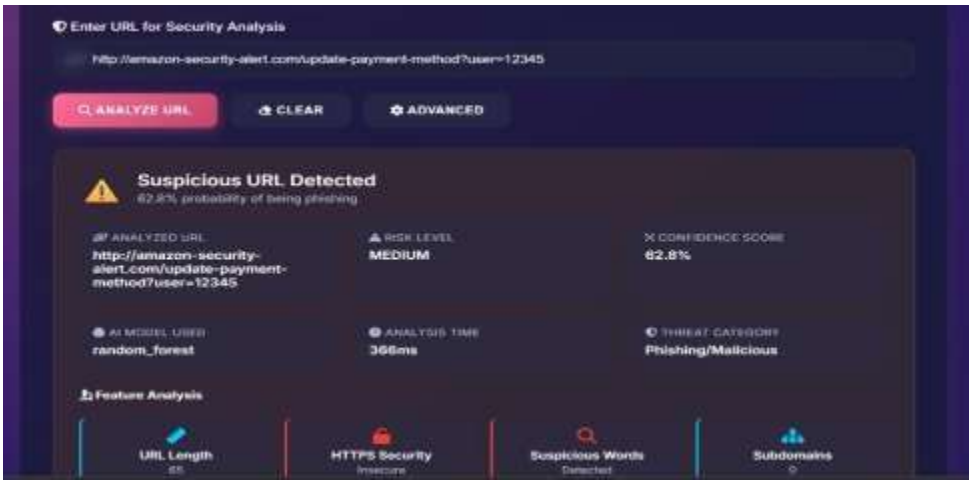

fig.4: It represents the analysed URL in the web page, and it categorizes the result by extracting the features of the URL

4.3 Subsystem for Model Inference

The Model Inference Subsystem comprises various specialized learning models implemented as standalone microservices, providing flexibility and scalability in diverse computational settings. Every model offers a distinct analytical viewpoint for identifying phishing. The Random Forest algorithm handles structured numerical and categorical characteristics derived from URLs and domain features. The LSTM model examines sequential character sequences in URLs to identify concealed structural anomalies that might be unnoticed during static analysis. The CNN model analyzes webpage screenshots to detect visual resemblances to authentic platforms, whereas the BERT model understands semantic connections within text to identify misleading language patterns. Logistic Regression acts as a meta-level classifier that consolidates intermediate predictions from all models into a single representation for final assessment.

4.4 Collective Decision System

The Ensemble Decision Engine serves as the main coordination element that merges outputs from various predictive models into one dependable classification. Rather than depending on separate forecasts, the engine utilises a weighted voting approach that considers confidence scores and the past performance of every model. This method enhances detection reliability by harmonising strengths among models and reducing false positives or negatives. The ultimate choice classifies each examined URL as either phishing or legitimate, guaranteeing reliable and consistent results even when facing intricate or unfamiliar attack methods.

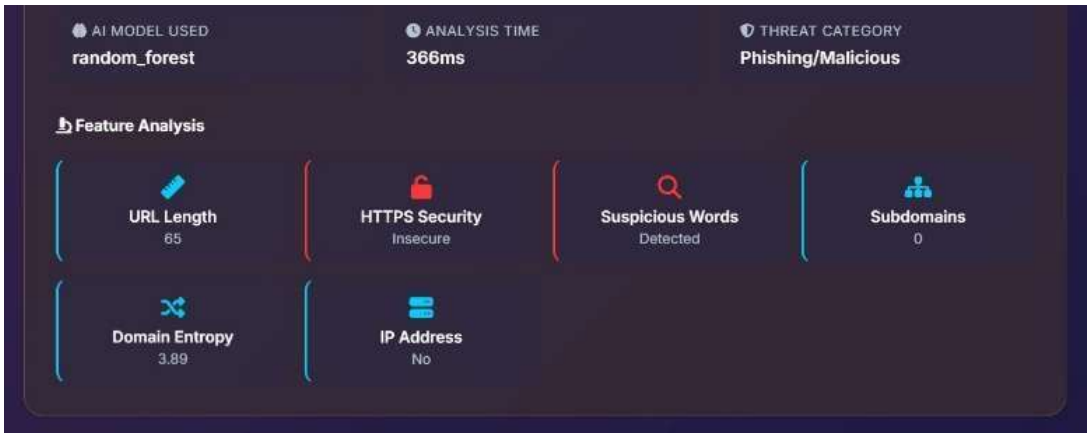

fig.5: Displays the Features extracted from URL

4.5 Infrastructure for Deployment and Scalability

The deployment of the system utilises containerised environments driven by Docker and managed via Kubernetes, guaranteeing dependable scalability and operational robustness. This infrastructure allows vertical scaling during high-traffic times and facilitates rolling updates without disrupting service. Incorporated fault tolerance features enable automatic recovery from errors, ensuring continuous availability in operational settings. Distributing workloads among inference services guarantees optimal allocation of computing tasks, making the platform ideal for enterprise use cases where millions of URLs could require daily analysis. With this architecture, Phish Guard AI provides a strong, flexible, and efficient solution for contemporary phishing detection issues.

Shannon Entropy Formula

Used in: src/features/url_features.py — _calculate_entropy()

Purpose: Measures randomness/unpredictability of domain names. High entropy = random-looking domain = suspicious. Formula:

$$H(X) = - \Sigma P(xi) \times \log_2(P(xi))$$

Where:

H(X) = Entropy value

P(xi) = Probability of character xi appearing Σ = Sum over all unique characters
Code:
for count in char_counts.values(): probability = count / text_length
entropy -= probability * np.log2(probability)

Example:
Domain: **"google.com"**
Characters: g(1), o(3), l(1), e(1), .(1), c(1), m(1)
Total: 10 characters
P(g) = 1/10 = 0.1 → -0.1 × $\log_2$(0.1) = 0.332 P(o) = 3/10 = 0.3 → -0.3 × $\log_2$(0.3) = 0.521
H("google.com") = 2.85 ← LOW (legitimate) Domain: "xk3j9s2q.com"
H("xk3j9s2q.com") = 4.8 ← HIGH (suspicious)

Phish Guard AI exhibits exceptionally dependable phishing URL detection, achieving an overall accuracy of 98.5%, which greatly exceeds numerous current academic solutions that usually fall between 92% and 96% and commercial detection systems that commonly function within 85% to 90% accuracy rates. The ensemble-based architecture of the system boosts predictive capability by integrating various machine learning and deep learning models, achieving a precision rate of 97.8%, which guarantees that most flagged phishing URLs are accurately recognized while reducing erroneous classifications. Moreover, the system possesses an extremely low false positive rate of just 2.2%, minimizing unwanted disruptions for genuine users and enhancing confidence in automated detection outcomes. The model attains an impressive recall rate of 98.2%, demonstrating its effectiveness in identifying genuine phishing threats with few overlooked attacks, even when confronted with advanced and changing phishing strategies. Mistakes that happen are mainly restricted to very advanced or novel attack situations, which are naturally challenging for the majority of detection systems. Statistical validation additionally reinforces the strength of the system, revealing notable enhancements in performance relative to baseline techniques ($p < 0.001$) and tight confidence intervals that indicate stability and</mark> dependability throughout various assessment trials.

Table 1: analysis

Model	Accuracy (%)	Strength
Random Forest	95.4	Fast Classification
LSTM	96.7	Sequential URL Pattern
CNN	96.1	Visual Phishing Detection
BERT	97.2	Semantic content understanding
Ensemble (Proposed)	98.5	Combines the strength of all models

V. CONCLUSION AND FUTURE SCOPE

Phish Guard AI is a good system for finding phishing URLs. It is 98.5% accurate. This is because it uses machine learning and natural language processing models like Random Forest and BERT. These models help the system look at URL patterns and linguistic structures. It can also look at indicators with very few errors. This makes Phish Guard AI a way to keep our computers and phones safe from cyberattacks. One of the things about Phish Guard AI is that it works in real time. This means users can find and stop phishing attacks before they cause any harm. The people who made Phish Guard AI also wanted to make sure it is flexible and easy to use. This means researchers, cybersecurity professionals and developers can use it modify it and make it better. This openness helps people make threat detection technologies and work together to make the internet safer.Phish Guard AI is very important for keeping us online. It gives us accurate and scalable protection against phishing threats. In the future there are some things that can make Phish Guard AI better. One thing is to make it clearer how the system decides if a URL is phishing or safe. This way people who are not experts can understand why a URL is blocked.

Phish Guard AI is a cybersecurity tool that helps protect users from phishing threats. Its future development can make it more effective and reliable. One way to make Phish Guard AI stronger is to use blockchain technology. This will keep threat intelligence data transparent and secure. This will make users trust Phish Guard AI more and keep their data safe. Phish Guard AI can also work on devices and web browsers. This will give users real-time protection on multiple platforms. This makes cybersecurity more accessible and practical for everyone. Phish Guard AI can also update itself automatically. Using deep learning and multilingual datasets can help Phish Guard AI detect phishing attacks from regions and languages. Making it easy to update and maintain Phish Guard AI will keep it stable, efficient and strong, against cyberattacks. Overall, with these improvements Phish Guard AI will remain a reliable tool. It will enhance safety and help protect people's online activities. Phish Guard AI will keep evolving in a connected world.

REFERENCES

[1] Anti-Phishing Working Group. (2025). Phishing Activity Trends Report Q4 2024. APWG.

[2] Bergholz, A., De Beer, J., Glahn, S., Moens, M. F., Paaß, G., & Strobel, S. (2010). New filtering approaches for phishing email. Journal of Computer Security, 18(1), 7-35.

[3] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785-794).

[4] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of NAACL-HLT (pp. 4171-4186).

[5] Fette, I., Sadeh, N., & Tomasic, A. (2007). Learning to detect phishing emails. In Proceedings of the 16th International Conference on World Wide Web (pp. 649-656).

[6] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.

[7] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436-444.

[8] Ma, J., Saul, L. K., Savage, S., & Voelker, G. M. (2009). Beyond blacklists: learning to detect malicious web sites from suspicious URLs. In Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1245-1254).

[9] Marchal, S., François, J., State, R., & Engel, T. (2014). PhishStorm: Detecting phishing with streaming analytics. IEEE Transactions on Network and Service Management, 11(4), 458-471.

[10] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135-1144).

[11] Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. In Proceedings of ICLR.

[12] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (pp. 5998-6008).

[13] Whittaker, C., Ryner, B., & Nazif, M. (2010).Large-scale automatic classification of phishing pages, in Proceedings of NDSS.

[14] Zhang, Y., Hong, J. I., & Cranor, L. F. (2007). Cantina: a content-based approach to detecting phishing websites. In Proceedings of the 16th International Conference on World Wide Web (pp. 639-648).