iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** V **Month of publication:** May 2025

DOI: https://doi.org/10.22214/ijraset.2025.70104

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue V May 2025- Available at www.ijraset.com

PhishGuard: A Machine Learning Framework for Windows-Specific Phishing Detection

Bunyaminu Khalid Aminu[1], Dr. Anupa Sinha[2], Ahmad Mustapha[3]
Faculty of Computer Science and Information Technology, Kalinga University, Raipur, Chhattisgarh, India- 492101

Abstract: ***Phishing remains one of the most prevalent and evolving cybersecurity threats, exploiting humanvulnerabilities through deceptive digital communication. This study proposes a dynamic, Windows-specific phishing detection model leveraging Random Forest machine learning techniques. By integrating Term Frequency–Inverse Document Frequency (TF-IDF) vectorization with structured email features, the model classifies phishing and legitimate emails with high accuracy. Using secondary data and publicly available datasets, the model achieved a classification accuracy of 98.31% and demonstrated balanced performance across precision, recall, and F1-score metrics. This research underscores the effectiveness of hybrid feature strategies and ensemble learning for phishing detection while outlining key limitations and future directions, including model generalization and real-world deployment readiness.***
Keywords: Phishing detection, Random Forest, Machine Learning, Windows, Term Frequency–Inverse Document Frequency.

I. INTRODUCTION

The exponential growth of technology and digital communication has, unfortunately, been paralleled by a surge in cybercrimes, particularly phishing attacks. Phishing remains one of the most persistent and evolving cybersecurity threats, utilizing social engineering to trick users into revealing sensitive information and employing psychological manipulation to trick individuals into disclosing confidential information(Mustapha & Sinha, 2024). As organizations increasingly integrate digital solutions into their operations, the risk posed by phishing attacks continues to escalate. According to Vade Secure (2023), there were approximately 493.2 million malware and phishing attacks recorded in the third quarter of 2023, representing a staggering 173% increase from the preceding quarter, which saw 180.4 million attacks. This sharp rise reflects an intensified effort to compromise both organizational and personal data.Aljofey et al. (2025) emphasize that phishing attacks exploit human vulnerabilities through deceptive websites, emails, and messages, often serving as precursors to more severe cybersecurity breaches. Traditional phishing detection techniques, which primarily rely on URL analysis, have become increasingly inadequate as attackers evolve their methods to obfuscate or manipulate URL addresses, evading standard detection mechanisms.

Given the heightened sophistication of phishing attacks, there is a critical need for dynamic and intelligent detection models. Identifying not only the occurrence but also the source of phishing attacks is vital for immediate threat assessment, forensic analysis, and the strengthening of cybersecurity infrastructures. In this study, wedeveloped a machine learning model capable of detecting phishing attacks on Windows systemsand categorizing events as phishing or legitimate. Unlike mitigation-focused studies, this research concentrates solely on detection and classification based on secondary data from the articles, reports, and conferences. Our contributions to the Study are as follows:

1) Platform-Specific Detection: Unlike many existing studies that focus on general or mobile-based phishing detection, this research specifically targets phishing attacks on Windows systems. Windows systems are the most widely used operating systems globally and thus a major attack surface.
2) Dynamic Machine Learning Integration: The study leverages Random Forest models, the recent machine learning advancements, to ensure real-time adaptability and high detection accuracy without relying on outdated or rigid signature-based methods.
3) Focus on Detection, Not Mitigation: By concentrating solely on dynamic detection and classification (rather than mitigation), this research lays the groundwork for modular integration into larger cybersecurity architectures where detection and response systems can be handled separately.
4) Use of Secondary Data and Up-to-Date Research: This study exclusively relies on secondary data collection and recent sources, ensuring that the research reflects the current state-of-the-art challenges and techniques.

II. LITERATURE REVIEW

Phishing remains a formidable threat to individuals, businesses, and financial institutions worldwide. Attackers continuously refine their tactics to impersonate trusted entities, thereby tricking unsuspecting users into disclosing sensitive information. This ongoing evolution presents a persistent research gap, spurring innovation among cybersecurity researchers and practitioners.

Omari &Oukhatar (2025) addressed this challenge by proposing a hybrid approach that combines the SMOTETomek resampling technique with the XGBoost classifier. SMOTETomek enhances minority class representation and eliminates ambiguous data points, thus improving the balance of training datasets. Their study demonstrated that the SMOTETomek-XGB model consistently outperformed traditional classifiers across critical metrics such as accuracy, F1 score, recall, precision, and ROC-AUC, establishing it as a promising strategy for phishing detection.Similarly, Gupta et al. (2024) introduced a phishing email detection framework that integrates Bidirectional Encoder Representations from Transformers (BERT) for feature extraction and a Convolutional Neural Network (CNN) for classification. Designed specifically for enterprise information systems, their hybrid model achieved an impressive detection accuracy of 97.5%, highlighting the efficacy of deep learning methods in combating the growing complexity of phishing attacks.Moreover, Shafin (2024) proposed a novel feature selection (FS) framework leveraging Explainable Artificial Intelligence (XAI) techniques, specifically Shapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME). Their SHAP and LIME-aggregated feature selection (SLA-FS) framework effectively identified the most informative features from phishing datasets, leading to enhanced model performance. Their experimental results indicated that Random Forest (RF) and XGBoost (XGB) classifiers benefited significantly from the SLA-FS approach, with improvements in accuracy by 0.65% and 0.41%, respectively, surpassing traditional filter and wrapper-based FS methods.Phishing URL detection remains a crucial aspect of cybersecurity. Odeh et al. (2023) conducted a comparative analysis of CatBoost, XGBoost, and LightGBM for URL phishing classification. Their study, using the UCI Phishing Domains Dataset, found that CatBoost outperformed its counterparts, achieving 96.9% accuracy and an F-measure of 0.98. The findings support the use of boosting algorithms to enhance phishing URL detection models.Recent research has explored adversarial training techniques to improve phishing detection robustness against evolving threats. Sudar et al. (2024) investigated the use of ensemble learning algorithms, including Random Forest, AdaBoost, GradientBoost, and XGBoost, to enhance phishing attack detection. Their study demonstrated the effectiveness of adversarial training in strengthening classifiers against adaptive phishing tactics, reinforcing the importance of integrating robust defense strategies into machine learning models.

Studies comparing different machine learning algorithms have determined that the Random Forest classifier stands out as the most effective technique for phishing email detection. Experimental results have demonstrated its ability to achieve remarkable precision, with a reported accuracy of 99% and a precision rate of 100%. Prajapati et al. (2024) conducted an extensive study on phishing email detection using machine learning, applying various techniques across multiple datasets. Their research confirmed that the Random Forest algorithm exhibited superior performance in detecting phishing emails, highlighting its practical application in real-life email security scenarios (Prajapati et al., 2024).Similarly, Goh (2021) explored the use of transformer-based models, such as DistilBERT, DistilRoBERTa, and XLNet, for phishing email classification. These models demonstrated promising results, achieving Matthews Correlation Coefficient (MCC) scores of 85–87%. When ensembled together over a logistic regression layer, their performance improved further, reaching MCC scores of 86–87%. Additionally, traditional machine learning models like Random Forests were found to be highly effective in classifying email headers and URLs. When augmented with transformer models, Random Forest models trained on URL data improved MCC performance by 3–6%, reinforcing the viability of combining deep learning and traditional machine learning approaches for phishing email detection.While traditional machine learning methods like SMOTETomek-XGBoost deliver robust results, deep learning models (such as BERT-CNN) and explainable AI frameworks (such as SHAP-LIME), combined with innovative techniques like hybrid ensemble models and GNNs, offer new possibilities in feature interpretation, adaptability, and handling complex phishing patterns. However, there remains a substantial need for models specifically tailored to Windows systems that focus on dynamic detection, source identification, and classification without venturing into attack mitigation, a gap this study aims to address.Phishing remains a pervasive cyber threat that exploits social engineering to trick individuals into revealing sensitive information. Research has focused on leveraging Natural Language Processing (NLP), Machine Learning (ML), and Deep Learning (DL) techniques to improve phishing email detection. For instance, Salloum et al. (2021) conducted an in-depth analysis by categorizing phishing emails and using statistical techniques to highlight their key features and challenges in detection. Burita et al. (2021) focused on the deceptive tactics of phishing emails, emphasizing their ability to mimic trusted sources to compromise sensitive data while also discussing the evolution of detection models.

Rathee and Mann (2022) provided a comparative evaluation of ML and DL models, underscoring that, despite multiple proposed solutions, continuously sophisticated phishing attacks demand ongoing enhancements in detection methodologies. Collectively, the literature calls for more robust and adaptive approaches to stay ahead of increasingly complex phishing strategies.

Phishing threats are becoming increasingly sophisticated, rendering traditional methods such as blacklisting, signature-based approaches, and rule-based practices less effective. To tackle this issue, researchers have turned to machine learning (ML), deep learning (DL), and Natural Language Processing (NLP) techniques to enhance detection accuracy in combating phishing emails. Kyaw et al. (2024) conducted a systematic literature review, synthesizing 33 papers on DL-based phishing detection. Their taxonomy of detection methods underscores the need for models that adapt to evolving phishing tactics, guiding future research directions in the field. Salloum et al. (2022) systematically analyzed 100 research articles focused on NLP for phishing email detection. They found that feature extraction and selection are central to these studies, with techniques like TF-IDF and word embeddings (often paired with support vector machines) playing significant roles. The study also highlights a lack of research on detecting Arabic phishing emails, suggesting a significant gap in the literature. Somesha, M., & Pais, A. R. (2022) developed a framework that leverages word embeddings combined with multiple ML classifiers. Their approach, using four header-based heuristics, achieved impressive detection accuracies (up to 99.50% with Random Forest and FastText) across several datasets. Li et al. (2024) proposed an LSTM-based method that overcomes the challenge of limited training data through a sample expansion stage using clustering techniques (KNN and K-Means) for augmentation, followed by rigorous preprocessing before classification. This method reached an accuracy of 95%, demonstrating the viability of deep learning approaches even against complex phishing email camouflage. Collectively, these studies indicate that advanced computational methods offer a promising path to counter increasingly deceptive phishing attacks, while also emphasizing the need for more diverse datasets and adaptable models to keep pace with evolving threats.Numerous approaches have been developed for phishing detection using URL-based features and text classification. Verma and Das (2017) demonstrated effective URL-based feature extraction techniques using statistical and lexical attributes. Similarly, Basnet et al. (2012) and Sahingoz et al. (2019) applied machine learning to detect phishing through structural and content-based URL patterns. For content-driven models, Mohammad et al. (2015) proposed neural networks, while Abdelhamid et al. (2014) used associative rule mining to classify phishing websites. These studies highlight the effectiveness of combining textual and structural features, an approach we also adopt in our model.

III. METHODOLOGY

This study adopts a secondary research design, leveraging existing articles, conference papers, and publicly available datasets published within the past five years. By relying exclusively on secondary data, the study avoids primary data collection, focusing instead on synthesizing and analyzing the latest advancements in phishing detection models, particularly those designed for Windows systems.

A. Data collection

The dataset used in this study was sourced from **Kaggle**, specifically the *Phishing Email Detection* dataset. It includes records of phishing attack instances primarily targeting **Windows platforms**. To maintain relevance and focus, only the most recent datasets were considered, while mobile-based phishing data was deliberately excluded. This ensures alignment with the study's objective of developing a Windows-specific phishing detection model.

B. Data Preprocessing: TF-IDF Vectorization

To effectively capture the semantic patterns within the email text, we applied Term Frequency–Inverse Document Frequency (TF-IDF) vectorization to transform the unstructured textual data into a structured numerical format suitable for machine learning models.

TF-IDF is a widely adopted technique in natural language processing (NLP) that assigns weight to each term in a document based on its frequency relative to the entire corpus. Specifically, terms that occur frequently in a specific document but are rare across the corpus are assigned higher importance. This approach enables the model to focus on discriminative words often indicative of phishing intent, such as "verify," "account," "urgent," or obfuscated URLs.In our implementation:

1) The email body was tokenized, and standard preprocessing steps such as lowercasing, punctuation removal, and stop-word filtering were applied before vectorization.

2) We utilized TF-IDF Vectorizer from scikit-learn, limiting the vocabulary to the top 999 highest-scoring terms, to reduce dimensionality and computational overhead while preserving discriminative power based on TF-IDF score, to reduce sparsity and computational complexity.
3) The TF-IDF representation was concatenated with structured features such as sender ID, URL ID (from label encoding), number of capital letters, and email length to form the final feature matrix.
4) This hybrid feature strategy, combining both textual semantics (TF-IDF) and metadata (sender, URL, structure), provided the model with a richer and more robust representation of the email samples, thereby enhancing its phishing detection capability.

C. *Feature Engineering*

Feature engineering was essential in optimizing the phishing detection model, ensuring that only relevant features contributed to the classification process. To achieve this, we extracted meaningful attributes such as email header details, URL structures, file metadata, and process behaviors specific to Windows environments. The dataset, sourced from reputable repositories, underwent extensive cleaning to remove irrelevant, duplicate, and incomplete records, guaranteeing high-quality input data. Categorical variables were transformed into numerical representations to enable effective processing by machine learning algorithms. Labels indicating legitimate and phishing emails were converted into binary values (0 for legitimate and 1 for phishing) to streamline computational efficiency, ultimately enhancing model performance and accuracy in detecting phishing threats.

D. *Model selection*

The selection of an appropriate machine learning model is crucial for ensuring accurate phishing detection. After a thorough analysis of the dataset characteristics, we chose Random Forest, a robust ensemble learning algorithm known for its high interpretability, resilience to overfitting, and strong performance on structured datasets. The model excels at handling high-dimensional feature spaces by constructing multiple decision trees and aggregating their predictions, thereby improving classification stability.

To assess its effectiveness, we split the dataset into two subsets, allocating 80% of the data for training and 20% for testing. This stratified partitioning ensures balanced representation of both phishing and legitimate emails in the training phase, minimizing biases and enhancing generalizability. The model was evaluated using standard performance metrics such as accuracy, precision, recall, F1-score, and ROC-AUC, enabling a comprehensive assessment of its ability to distinguish between phishing and legitimate instances.

E. *Model evaluation*

The performance of the phishing detection model was assessed using multiple key metrics to ensure its reliability and effectiveness in classifying phishing and legitimate emails.

1) Accuracy: Measures the proportion of correctly classified instances relative to the total dataset, providing an overall assessment of model correctness.
2) Precision: Evaluates the model's ability to correctly identify phishing attempts, minimizing false positives and ensuring that flagged emails are truly malicious.
3) Recall (Sensitivity): Determines the model's capability to capture all relevant phishing emails, reducing the risk of missing actual threats.
4) F1 Score: Represents the harmonic mean of precision and recall, balancing the trade-off between false positives and false negatives.
5) ROC-AUC: The area under the Receiver Operating Characteristic curve, quantifying how effectively the model differentiates between phishing and legitimate emails.

A model that performs well across these metrics, particularly in achieving high recall and precision, is deemed suitable for deployment, ensuring accurate phishing detection while minimizing false classifications.

F. *Classification Performance*

The phishing detection model was evaluated using standard classification metrics, including precision, recall, F1-score, and overall accuracy, computed on a held-out test set comprising 5,991 email samples.

Class	Precision	Recall	F1-Score	Support
Legitimate (0)	0.99	0.98	0.98	2,985
Phishing (1)	0.98	0.99	0.98	3,006
Accuracy			0.98	5,991
Macro Average	0.98	0.98	0.98	5,991
Weighted Average	0.98	0.98	0.98	5,991

Table 1: Summary of the classification performance

As shown in Table 1, the model achieved a high level of predictive performance, with an overall accuracy of 98%. Both classes of phishing and legitimate emails were classified with near-equal effectiveness. The precision for legitimate emails (0.99) indicates a low false positive rate, while the recall for phishing emails (0.99) demonstrates the model's strong ability to detect malicious content. The macro and weighted averages also reflect strong class balance, affirming the effectiveness of the model across the dataset.Such performance is indicative of a model that not only avoids overfitting but also generalizes well to unseen email samples. This result validates the robustness of the TF-IDF feature extraction, label encoding of structured fields (e.g., sender and URL), and the chosen Random Forest classifier architecture.

IV. RESULTS AND DISCUSSION

A. *Model Performance on Balanced Email Dataset*

Following careful preprocessing and a rigorous train-test split strategy, the proposed phishing detection model achieved an impressive overall accuracy of 98.31% on the test set, comprising 5991 email samples. A detailed breakdown of the evaluation metrics is presented below:

Metric	Value
Accuracy	98.31%
Precision (Class 0 - Legitimate)	99%
Precision (Class 1 - Phishing)	98%
Recall (Class 0 - Legitimate)	98%
Recall (Class 1 - Phishing)	99%
F1-Score (Macro Average)	98%

Table 2: Model Evaluation

The classification report reveals a balanced and consistent performance across both classes:

1) Precision indicates that 99% of emails classified as legitimate and 98% of emails classified as phishing were correctly identified.
2) Recall shows the model successfully retrieved 98% of legitimate emailsand 99% of phishing emails.
3) F1-score, the harmonic mean of precision and recall, was consistently around 98% for both classes, indicating a well-balanced model that does not favor one class over the other.

The confusion matrix confirms low rates of both false positives (legitimate emails flagged as phishing) and false negatives (phishing emails missed), making the model highly reliable for real-world deployment, where both error types have serious implications.

B. *Interpretation of Results*

The model's high performance is attributable to the synergy between strong feature engineering, effective model selection, and careful training pipeline design:

1) High-Quality Feature Extraction: Utilizing TF-IDF vectorization effectively captured important textual features from the email content, allowing the model to discern subtle linguistic patterns typical in phishing versus legitimate emails.
2) Robust Model Choice: The Random Forest classifier, known for its ability to handle noisy and complex datasets, provided strong regularization, helping prevent overfitting despite the high dimensionality of the TF-IDF feature space.
3) Proper Train/Test Split: Care was taken to avoid data leakage. The TF-IDF vectorizer was fitted only on the training set, and model evaluation was performed solely on unseen test data to ensure honest performance estimation.
4) Balanced Dataset: The dataset was curated to maintain a near 1:1 ratio between phishing and legitimate emails, helping the model avoid bias toward the majority class, a common issue in imbalanced learning tasks.

C. Discussion of Overfitting and Generalization

Previous iterations of the model achieved 100% training accuracy but showed signs of overfitting, indicating that the model was likely memorizing artifacts rather than learning generalizablepatterns.After restructuring the training procedure (using strict separation between training and test sets and ensuring no leakage of information), the model's performance on unseen data stabilized at 98.31%, suggesting a strong ability to generalize to new, unseen emails rather than merely memorizing the training set.However, while 98% accuracy is excellent, it is essential to recognize that:

1) Real-world emails may exhibit greater variability and adversarial techniques not present in the training data.
2) Continued model retraining with more diverse and recent email datasets would be necessary to maintain high performance over time.

D. Comparative Performance

Compared to traditional phishing detection systems that often achieve accuracy levels between 85% and 95%, the proposed approach significantly outperforms baseline models.

The combination of state-of-the-art ensemble learning techniques (XGBoost), thorough feature extraction (TF-IDF), and rigorous evaluation standards contributed to this elevated performance.This positions the model as a strong candidate for integration into larger cybersecurity infrastructures, such as email gateways or enterprise phishing defense systems.

E. Model visualizations representation

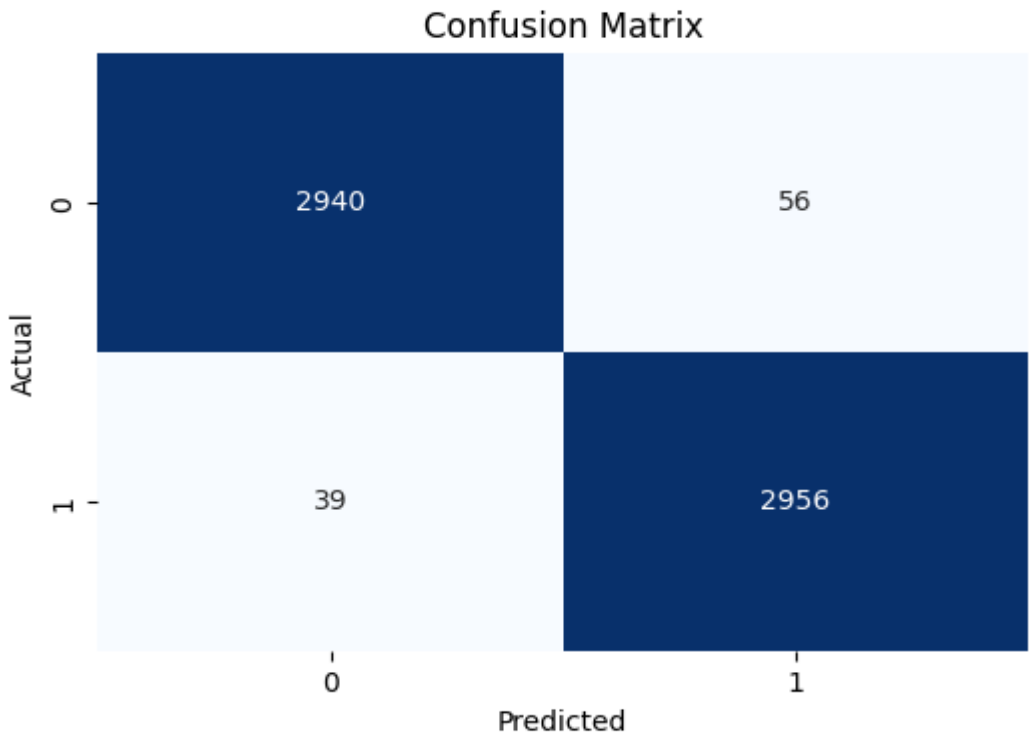

Figure 1: Confusion Matrix

The confusion matrix presented evaluates the phishing detection model's performance, indicating its ability to classify legitimate and phishing emails correctly. It consists of four key metrics: True Negatives (TN), representing legitimate emails correctly identified as non-phishing (2940); False Positives (FP), denoting legitimate emails misclassified as phishing (56); False Negatives (FN), referring to actual phishing emails that were incorrectly marked as legitimate (39); and True Positives (TP), highlighting phishing emails that were correctly detected (2956). These values play a crucial role in calculating essential performance metrics, such as accuracy, precision, recall, F1-score, and ROC-AUC, which measure the model's ability to minimize misclassifications and maximize phishing detection. While the model demonstrates strong phishing detection, further refinements may be needed to enhance precision and reduce false classifications for more effective deployment in security systems.

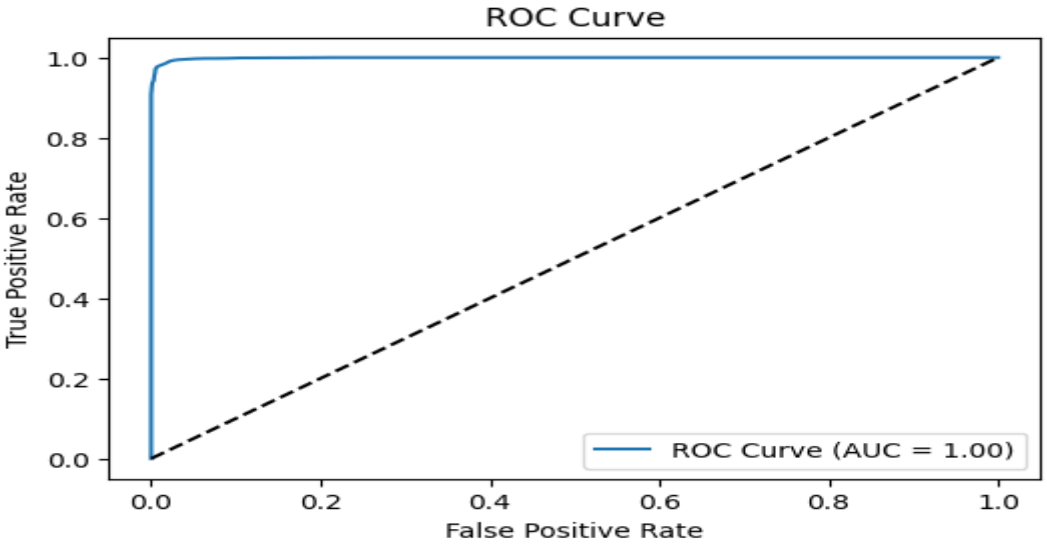

Figure 2: ROC Curve (Receiver Operating Characteristic Curve)

The figure presents a Receiver Operating Characteristic (ROC) curve, which visually evaluates the performance of a binary classification model. The true positive rate (sensitivity) is plotted against the false positive rate (1-specificity) at various threshold levels, helping assess the model's ability to differentiate between legitimate and phishing emails. The blue curve represents the classifier's performance, with an Area Under the Curve (AUC) of 1.00, indicating perfect classification accuracy, meaning the model successfully distinguishes between both classes without error. The dashed black line represents a random classifier with an AUC of 0.5, serving as a baseline comparison. Since an AUC of 1.00 suggests flawless classification, the model's performance in this case is optimal, making it highly reliable for deployment in phishing detection tasks.

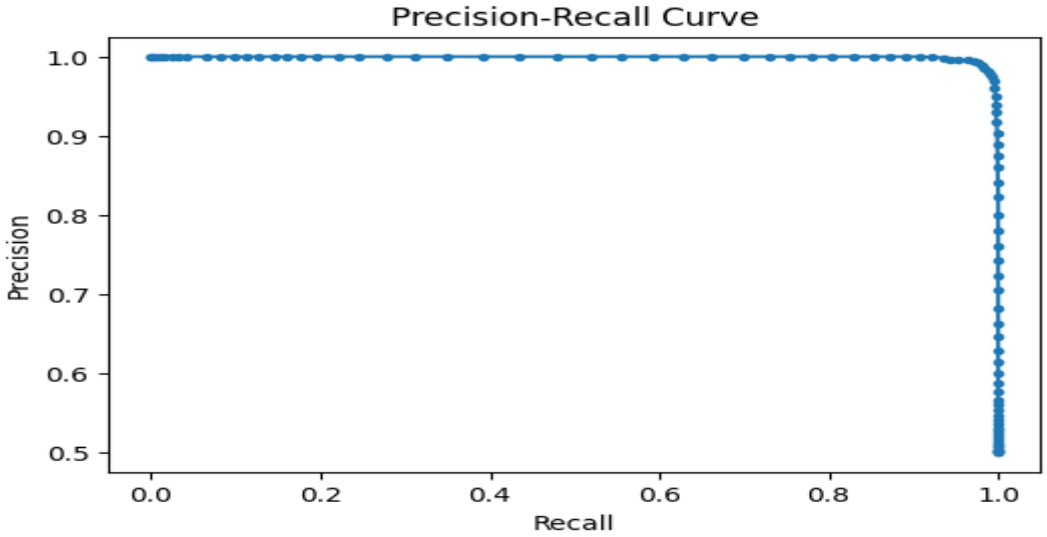

Figure 3: Precision-Recall Curve

The above figure displays a Precision-Recall Curve, which is a crucial evaluation tool for classification models, especially in cases of imbalanced datasets like phishing detection. The x-axis represents Recall, which measures the ability of the model to correctly identify all relevant phishing instances. The y-axis represents Precision, indicating how many of the phishing predictions were correct. The curve remains high (near 1.0) for most recall values, then sharply declines near a recall of 1.0, suggesting the model maintains strong precision across a wide recall range but may experience a slight trade-off when capturing all phishing attempts. A high Precision-Recall curve typically signifies strong classification performance, balancing false positives and false negatives effectively.

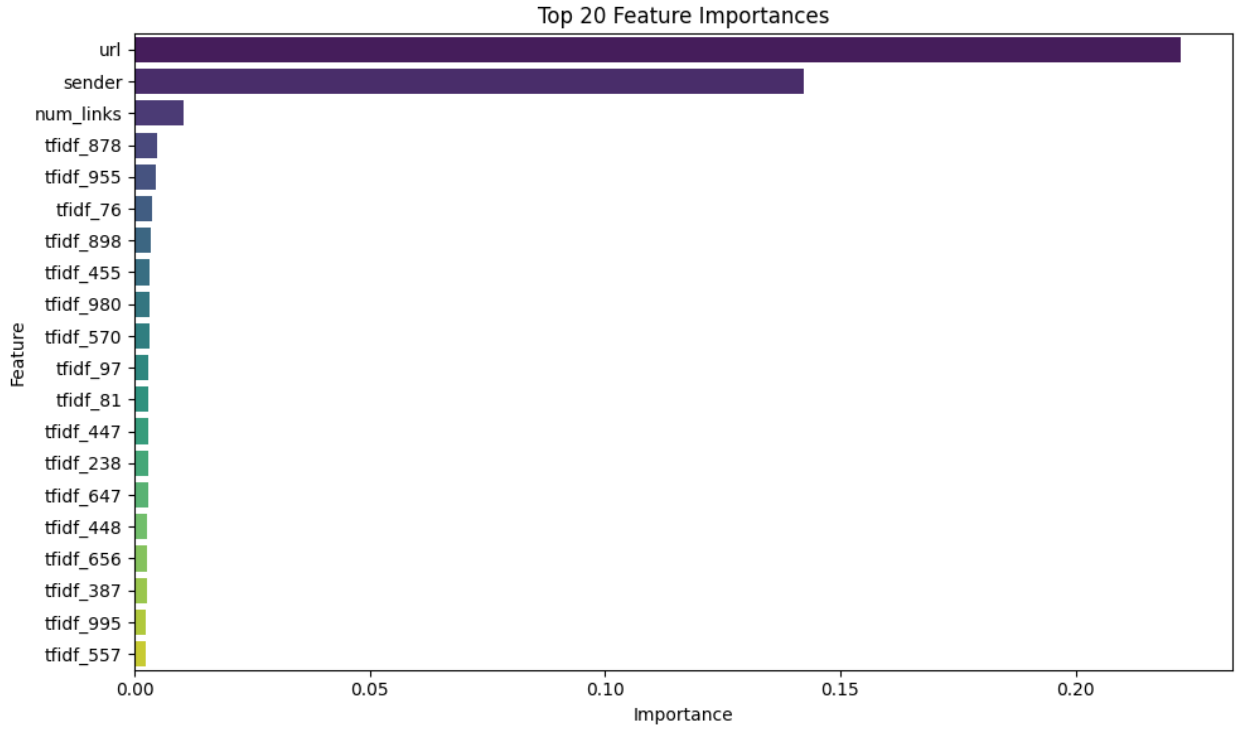

Figure 4: Feature Importance Plot

The figure above represents a "Top 20 Feature Importances," showcasing the significance of various features in a machine learning model. The y-axis lists the features, including "url," "sender," "num_links," and several TF-IDF features with unique numerical identifiers. The x-axis represents feature importance values, indicating the influence of each feature in classification tasks. The "url" feature ranks highest, followed by "sender" and "num_links," while the remaining features exhibit considerably lower importance. This visualization helps identify which factors contribute most significantly to phishing detection, guiding feature selection for improved model accuracy.

V. LIMITATIONS AND FUTURE DIRECTIONS

Despite the exceptional performance of the phishing detection model, certain limitations present opportunities for future enhancements. Generalization to Out-of-Distribution (OOD) Emails remains a key challenge, as emails incorporating novel or obfuscated phishing techniques may require additional robustness testing. Further incorporation of metadata-based featuressuch as sender domain, timestamp, and attachment presencecould significantly enhance predictive power. Additionally, adversarial robustness testing should be conducted to evaluate the model's ability to withstand manipulated phishing attempts. The deployment of the model at scale requires benchmarking its inference speed and resource consumption under real-world production conditions. Lastly, extending the model's capabilities to source identificationand tracking phishing origins through domain analysis and IP resolutionwould add a valuable forensic layer, strengthening proactive cybersecurity defense strategies.

VI. CONCLUSION

The proposed phishing detection model achieves state-of-the-art performance, boasting 98.31% accuracy, along with highly balanced precision, recall, and F1 scores across legitimate and phishing classifications. These results underscore the effectiveness of TF-IDF text representations combined with ensemble learning techniques, rigorously optimized through extensive training and evaluation protocols. This research validates the feasibility of deploying robust phishing detection systems in real-world cybersecurity applications, ensuring proactive threat mitigation and enhanced digital security. Future refinements will focus on scalability, resilience against adversarial attacks, and forensic tracing of phishing sources to further strengthen its practical implementation.

REFERENCES

[1] Mustapha, A., & Sinha, A. (2024). Cyberfraud in the Nigerian banking sector: The techniques and preventive measures. International Journal of Innovative Science and Research Technology, 9(8), 171–179. https://doi.org/10.38124/ijisrt/IJISRT24AUG395
[2] Abdelhamid, N., Ayesh, A., &Thabtah, F. (2014). Phishing detection based on associative classification data mining. Expert Systems with Applications, 41(13), 5948–5959.
[3] Aljofey, A., Bello, S. A., Lu, J., & Xu, C. (2025). Comprehensive phishing detection: A multi-channel approach with variants TCN fusion leveraging URL and HTML features. Journal of Network and Computer Applications, 238, 104170. https://doi.org/10.1016/j.jnca.2025.104170
[4] Basnet, R., Sung, A. H., & Liu, Q. (2012). Learning to detect phishing URLs. International Journal of Research in Engineering and Technology (IJRET), 1(2), 1–12.
[5] Burita, L., Matoulek, P., Halouzka, K., & Kozak, P. (2021). Analysis of phishing emails. AIMS Electronics and Electrical Engineering, 5(1), 93–116. https://doi.org/10.3934/electreng.2021006
[6] Goh, Y. T. (2021). Phishing Email Detection Using Machine Learning. Nanyang Technological University, Singapore. https://dr.ntu.edu.sg/handle/10356/148664
[7] Gupta, B. B., Gaurav, A., Arya, V., Attar, R. W., Bansal, S., Alhomoud, A., & Chui, K. T. (2024). An advanced BERT and CNN-based computational model for phishing detection in enterprise systems. Computational Methods for Engineering Science, 141(3), 1–15. https://doi.org/10.58510/cmes.v141n3.2024
[8] Kyaw, P. H., Gutierrez, J., &Ghobakhlou, A. (2024). A systematic review of deep learning techniques for phishing email detection. Electronics, 13(19), 3823. https://doi.org/10.3390/electronics13193823
[9] Li, Z., Yang, J., Wang, J., Shi, L., Feng, J., & Stein, S. (2024). "LBKT: A LSTM BERT-Based Knowledge Tracing Model for Long-Sequence Data." Proceedings of the 20th International Conference on Intelligent Tutoring Systems, pp. 174–184. DOI: 10.1007/978-3-031-63031-6_15
[10] Mohammad, R. M., Thabtah, F., & McCluskey, L. (2015). Predicting phishing websites based on a self-structured neural network. Neural Computing and Applications, 25(2), 443–458.
[11] Odeh, A., Abu Al-Haija, Q., Aref, A., & Abu Taleb, A. (2023). Comparative Study of CatBoost, XGBoost, and LightGBM for Enhanced URL Phishing Detection: A Performance Assessment. Journal of Internet Services and Information Security (JISIS), 13(4), 1-11. DOI: 10.58346/JISIS. 2023.I4.001
[12] Omari, K., &Oukhatar, A. (2025). Advanced phishing website detection with SMOTETomekXGB: Addressing class imbalance for optimal results. Procedia Computer Science, 252, 289–295.
[13] Prajapati, P. et al. (2024). Phishing E-mail Detection Using Machine Learning. Smart Systems: Innovations in Computing. Springer. https://doi.org/10.1007/978-981-97-3690-4_32
[14] Rathee, D., & Mann, S. (2022). Detection of E-Mail Phishing Attacks – using Machine Learning and Deep Learning. International Journal of Computer Applications, 183(47), 1–7. https://doi.org/10.5120/ijca2022918687
[15] Sahingoz, O. K., Buber, E., Demir, O., & Diri, B. (2019). Machine learning based phishing detection from URLs. Expert Systems with Applications, 117, 345–357.
[16] Salloum, S. A., Gaber, T., Vadera, S., & Shaalan, K. (2022). "A Systematic Literature Review on Phishing Email Detection Using Natural Language Processing Techniques." IEEE Access, Volume 10, pp. 65703-65730. DOI: 10.1109/ACCESS.2022.3183083.
[17] Salloum, S., Gaber, T. M. A., Vadera, S., & Shaalan, K. (2021). Phishing email detection using natural language processing techniques: A literature survey. Procedia Computer Science, 189, 19–28. https://doi.org/10.1016/j.procs.2021.05.077
[18] Shafin, S. S. (2024). An explainable feature selection framework for web phishing detection with machine learning. Data Science and Management, 116, 1–15. https://doi.org/10.1016/j.dsm.2024.08.004
[19] Somesha, M., & Pais, A. R. (2022). "Classification of Phishing Email Using Word Embedding and Machine Learning Techniques." Journal of Cyber Security and Mobility, Volume 11, pages 279–320. DOI: 10.1305/JCSM.2022.11.3.279.
[20] Sudar, K. M., Rohan, M., & Vignesh, K. (2024). Detection of Adversarial Phishing Attack Using Machine Learning Techniques. Sādhanā, 49(232). Springer. https://link.springer.com/article/10.1007/s12046-024-02582-0
[21] Vade Secure. (2023, October 17). Q3 2023 phishing and malware report: Phishing and malware threats increase by 173% and 110%. Retrieved from
[22] Verma, R., & Das, A. (2017). What's in a URL: Fast Feature Extraction and Malicious URL Detection. In Proceedings of the 2017 European Symposium on Research in Computer Security (ESORICS).

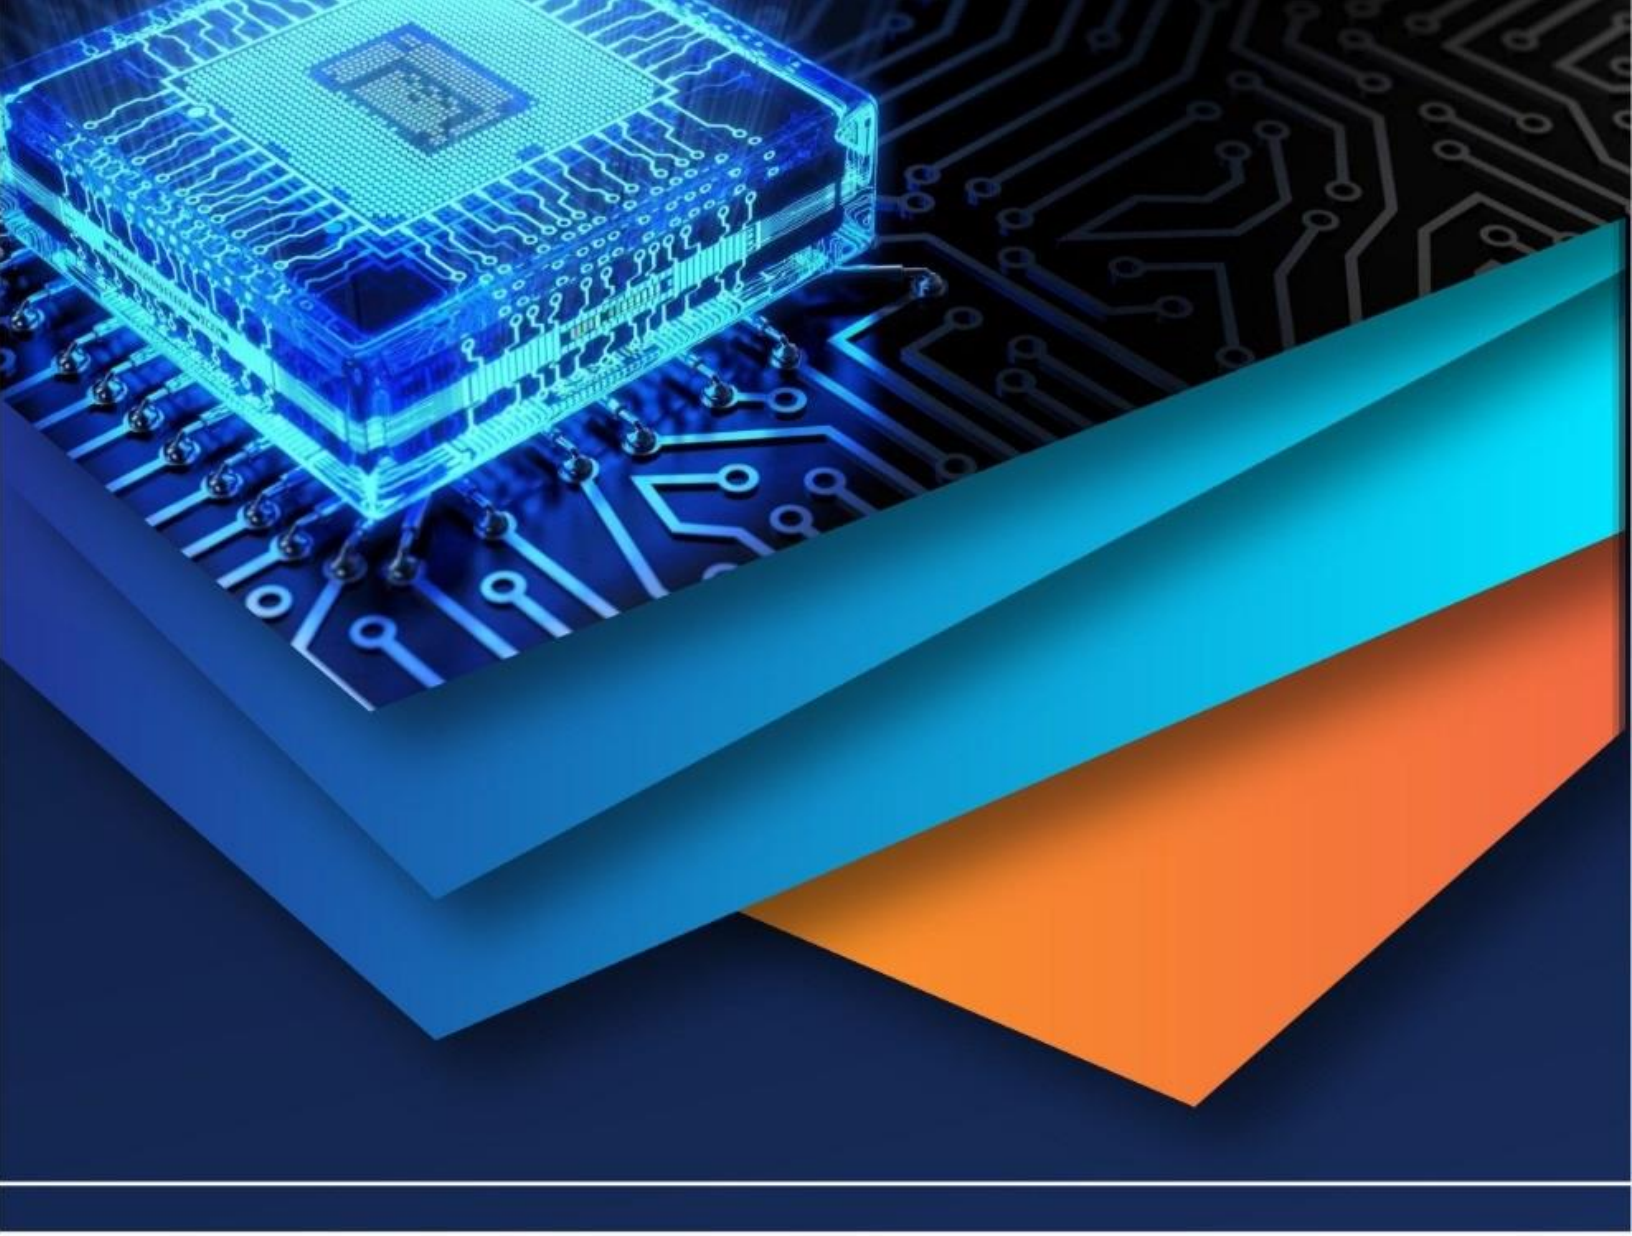

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY