EasyChair Preprint

№ 15168

Optimizing Phishing Detection with Advanced Feature Vectorization and Supervised Machine Learning Techniques

Jessica Comsie and Ayesha Noor

September 29, 2024

Optimizing Phishing Detection with Advanced Feature Vectorization and Supervised Machine Learning Techniques

Author: Jessica Comsie
Date:29th, Sep 2024

Abstract:

Phishing attacks are among the most common cybersecurity threats, taking advantage of users' trust to gain access to sensitive information. Detecting these attacks effectively is essential for protecting both individuals and organizations. This study focuses on improving phishing detection by developing an optimized framework for feature vectorization, combined with supervised machine learning techniques. By carefully selecting and designing features from email and website data, the goal is to enhance the accuracy of identifying phishing attempts. The analysis includes various text-based, URL-based, and metadata features, emphasizing their role in improving classification performance. Machine learning models such as Support Vector Machines (SVM), Random Forest, and Gradient Boosting are trained and tested on a dataset of legitimate and phishing samples. The study also examines the impact of feature scaling, selection, and dimensionality reduction methods like Principal Component Analysis (PCA) to determine which factors most effectively boost detection accuracy. Experimental findings show that an optimized feature set, combined with strong machine learning algorithms, greatly enhances phishing detection rates while reducing false positives. This approach highlights the potential for reliable, automated phishing detection systems, contributing to stronger cybersecurity defenses.

I. Introduction

A. Overview of Phishing Attacks

Phishing attacks have become a dominant form of cybercrime, where malicious entities attempt to deceive users into disclosing sensitive information such as passwords, credit card numbers, or personal identification details. These attacks often masquerade as legitimate emails, websites, or messages from trusted entities, exploiting social engineering tactics. Phishing schemes are evolving in sophistication, making them harder to detect and increasingly dangerous to individuals and organizations alike. The financial losses, reputation damage, and privacy violations associated with phishing attacks underscore the need for more effective detection mechanisms.

B. Importance of Detecting Phishing Attacks

As phishing techniques evolve, traditional detection methods, such as rule-based systems or blacklists, struggle to keep pace with new attack patterns. These methods often fail to detect novel or obfuscated phishing attempts, leading to high false negatives. Given the widespread impact of phishing on users across various sectors, from individuals to corporations, improving the accuracy and reliability of phishing detection systems is imperative. Efficient detection mechanisms can mitigate risks, prevent data breaches, and reduce the financial and reputational damage that phishing attacks cause.

C. Role of Machine Learning in Phishing Detection

Machine learning (ML) offers a promising solution to the growing challenge of phishing detection. By leveraging large datasets of phishing and legitimate instances, machine learning models can learn complex patterns and relationships, enabling them to differentiate between benign and malicious attempts. Supervised learning algorithms, such as Support Vector Machines (SVM), Random Forests, and Neural Networks, have been widely applied to this domain, offering enhanced detection accuracy. These models rely heavily on feature engineering, where relevant attributes are extracted from email or web content to enable precise classification. Machine learning techniques not only improve detection performance but also offer scalability and adaptability to new attack vectors.

D. Purpose of the Study

The purpose of this study is to enhance the detection of phishing attacks through the development of an optimized feature vectorization framework integrated with supervised machine learning algorithms. By focusing on the careful selection, transformation, and optimization of features—ranging from textual content to URL characteristics and email metadata—this study aims to improve the detection accuracy while reducing false positives. We will evaluate the effectiveness of various machine learning models in conjunction with feature engineering techniques such as feature selection and dimensionality reduction. This research seeks to provide a robust and scalable solution to phishing detection, thereby contributing to the ongoing efforts to strengthen cybersecurity.

II. Related Work

A. Traditional Phishing Detection Techniques

Traditional phishing detection techniques primarily rely on rule-based systems, blacklists, and heuristic analysis. Blacklists contain known phishing URLs or email addresses, preventing access to flagged content. However, blacklists are often reactive and fail to detect newly created phishing sites or those that frequently change their URL or IP addresses. Rule-based systems depend on predefined signatures or patterns, such as the presence of suspicious links or misleading domain names. While effective for certain cases, these methods are limited by their reliance on static rules, which makes them vulnerable to novel attack methods. Heuristic approaches extend rule-based methods by analyzing specific features such as spelling errors, missing HTTPS

security, and unusual email headers. Although somewhat more adaptable, traditional techniques often struggle with maintaining high detection accuracy and adapting to rapidly changing phishing strategies, leading to a higher rate of false positives and false negatives.

B. Machine Learning Approaches for Phishing Detection

In recent years, machine learning (ML) has emerged as a more dynamic and robust approach to phishing detection. Unlike traditional methods, machine learning models can learn from vast amounts of data, identifying complex patterns and distinguishing phishing attempts from legitimate communication. Supervised machine learning models such as Logistic Regression, Random Forest, Support Vector Machines (SVM), and Gradient Boosting have been extensively employed in phishing detection. These models leverage a wide range of features, including URL characteristics (e.g., domain age, length), textual content (e.g., language patterns, keyword frequency), and metadata (e.g., email headers, sender domain).

Deep learning approaches, such as Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), have also been applied, especially for detecting phishing in dynamic environments like social media and mobile platforms. Neural networks excel at processing unstructured data and discovering hidden patterns, allowing them to adapt to new phishing techniques. However, their complexity often requires larger datasets and computational power.

C. Challenges in Current Machine Learning Models

Despite the success of machine learning in phishing detection, several challenges remain. One of the primary issues is the selection of relevant and high-quality features. Feature engineering is critical for improving model performance, but identifying the most impactful features can be time-consuming and domain-specific. Moreover, phishing tactics evolve rapidly, which may cause models trained on outdated datasets to lose effectiveness against new attack vectors. Another challenge is the imbalance in phishing datasets—phishing instances are often much fewer than legitimate cases, which can lead to biased models that underperform in detecting rare phishing attacks.

Another challenge involves high false positive rates, where legitimate emails or websites are flagged as phishing, leading to a poor user experience. Additionally, the complexity of certain machine learning models, such as deep learning networks, can make them difficult to interpret, which hampers the understanding of how specific phishing features contribute to the final prediction. This lack of transparency also complicates model debugging and improvement efforts. Finally, real-time phishing detection, especially for rapidly evolving attacks, is computationally intensive, and optimizing models for quick decision-making without sacrificing accuracy is a challenge.

This study aims to address these challenges by optimizing feature vectorization and employing machine learning techniques that balance detection accuracy with interpretability and computational efficiency.

III. Feature Vectorization for Phishing Detection

A. Understanding Feature Vectorization

Feature vectorization is the process of converting raw data into a structured numerical format that machine learning models can process. In phishing detection, various attributes from emails, websites, or URLs are transformed into feature vectors, which represent the characteristics of phishing or legitimate instances. Effective feature vectorization is critical for model performance, as it directly impacts how well the model can discern between malicious and benign patterns. Features must capture the essential attributes of phishing attacks, such as the structure of URLs, email content, or sender information, and convert them into a format that enhances the ability of machine learning algorithms to identify phishing attempts.

The quality of the feature vectorization process, including the selection, transformation, and scaling of features, significantly influences the accuracy of the model. Poor feature representation can lead to misclassifications, increasing the number of false positives or false negatives. Therefore, a critical aspect of this study is optimizing feature vectorization to improve phishing detection performance.

B. Types of Features in Phishing Detection

Phishing detection typically involves extracting various types of features, each contributing to the model's ability to differentiate between legitimate and phishing content. These features fall into several broad categories:

URL-Based Features: Phishing URLs often exhibit unusual patterns such as long or obfuscated URLs, the presence of special characters, or the use of deceptive domain names. Common features include:

- URL length
- Domain age and registration information
- Presence of IP addresses in the URL
- Use of HTTPS or HTTP
- Number of subdomains
- Suspicious keywords in the URL (e.g., "login," "verify")

Content-Based Features: These features are derived from the text content of emails or websites and include:

- Frequency of suspicious words or phrases (e.g., "urgent," "account suspended")
- Presence of spelling or grammatical errors
- Embedded links or images that redirect to malicious websites
- Email body text length and structure
- Analysis of HTML tags or JavaScript used in webpages

Metadata-Based Features: Metadata provides insights about the source and delivery of emails or websites. Common metadata features include:

- Sender's email domain (e.g., use of free email providers like Gmail, Yahoo)
- Mismatch between sender and reply-to addresses
- IP geolocation of the sender

- Email header anomalies (e.g., incorrect SPF records, DKIM signatures)
- Time-to-live (TTL) values for DNS records

Behavioral Features: These features analyze the behavior or interaction patterns of users with emails or websites:

- Click behavior (e.g., high click-through rates for malicious links)
- Mouse hover actions on phishing links
- Redirect patterns and frequency of redirection

Each of these feature types contributes unique information that can be useful in detecting phishing attacks. A key part of improving phishing detection involves combining these diverse features in a way that enhances the model's predictive ability.

C. Optimal Feature Selection and Engineering

Feature selection and engineering are critical processes for optimizing phishing detection models. Not all features contribute equally to model performance, and including irrelevant or redundant features can degrade accuracy and increase computational complexity. Feature selection techniques, such as recursive feature elimination (RFE), mutual information, and Chi-square tests, can help identify the most informative features for the task. These methods help reduce the feature space while retaining the most important attributes, improving both the efficiency and the accuracy of the model.

In addition to selection, feature engineering transforms raw features into more meaningful representations. For example:

- Logarithmic transformation of skewed features (e.g., URL length) to normalize the distribution.
- Feature scaling using methods like Min-Max scaling or standardization to ensure that numerical features are within similar ranges, preventing any one feature from dominating the learning process.
- Dimensionality reduction techniques such as Principal Component Analysis (PCA) can be used to reduce the feature set further by combining correlated features, thereby simplifying the model and speeding up computations without sacrificing accuracy.
- Combining well-engineered features with advanced selection methods ensures that the final feature set is both comprehensive and computationally efficient. By optimizing feature vectorization, phishing detection models can improve classification accuracy, reducing false positives and negatives, while remaining scalable and adaptable to new attack patterns.

This study focuses on identifying the most effective feature combinations and transformations to enhance phishing detection, ensuring a balanced trade-off between model complexity and performance.

IV. Supervised Machine Learning Models

A. Supervised Learning Overview

Supervised learning is a type of machine learning where models are trained using labeled data. In phishing detection, the dataset consists of instances that are clearly marked as either phishing (malicious) or legitimate (benign). The goal of the supervised learning process is to map input features (such as email content, URLs, or metadata) to the correct labels, enabling the model to predict whether future, unseen data is phishing or not.

Several supervised learning algorithms are commonly used in phishing detection, including:

- Support Vector Machines (SVM): A powerful classifier that separates data points using hyperplanes and is effective for high-dimensional datasets.
- Random Forest: An ensemble learning method that builds multiple decision trees and aggregates their outputs to improve accuracy and reduce overfitting.
- Logistic Regression: A simple but effective linear classifier that models the probability that an instance belongs to a particular class.
- Gradient Boosting Machines (GBMs): An ensemble method that combines weak learners (often decision trees) to create a strong classifier through iterative improvement.

These algorithms are trained on phishing datasets to learn patterns and relationships within the features, enabling them to classify new instances based on their learned knowledge. Each model has its strengths and trade-offs, which must be considered based on the nature of the data and the computational resources available.

B. Training Data and Labeling

The quality and quantity of training data are critical for the performance of supervised learning models. In phishing detection, training data typically includes a variety of features extracted from emails, websites, or URLs, alongside corresponding labels that indicate whether each instance is phishing or legitimate.

- Data Collection: Training data can be collected from multiple sources, such as public phishing databases, email providers, or web traffic logs. A diverse dataset that captures a wide range of phishing attacks is essential for creating robust models that generalize well to different types of phishing schemes.
- Data Preprocessing: Preprocessing steps may include data cleaning (e.g., removing duplicates, handling missing values), feature extraction, and vectorization. Since phishing data is often imbalanced (with fewer phishing instances compared to legitimate ones), techniques like oversampling, undersampling, or the use of synthetic data generation (e.g., SMOTE) can help balance the dataset.
- Labeling: Each instance in the dataset is labeled as either phishing (1) or legitimate (0). This binary classification enables supervised learning algorithms to learn from the labeled data. In some cases, manually labeled datasets are combined with automated methods (e.g., heuristic-based labeling) to enhance dataset size and diversity.

A well-labeled and representative dataset is crucial to ensure that the machine learning model can accurately detect phishing attempts and avoid overfitting to specific attack patterns.

C. Model Evaluation Metrics

Once the supervised learning models are trained, their performance must be evaluated using appropriate metrics. In phishing detection, accuracy alone is not sufficient, as it may not account for the imbalance between phishing and legitimate instances. Instead, a range of metrics is used to assess the model's effectiveness:

- **Accuracy:** The percentage of correctly classified instances (both phishing and legitimate) out of the total instances. While accuracy provides a general sense of model performance, it can be misleading in imbalanced datasets where most instances are legitimate.

- **Precision:** The proportion of true phishing detections (true positives) out of all instances predicted as phishing. High precision indicates that the model has a low false positive rate.

- **Recall (Sensitivity or True Positive Rate):** The proportion of actual phishing instances that the model correctly identifies. High recall means the model catches most phishing attacks, even at the cost of flagging some legitimate instances.

- **F1 Score:** The harmonic mean of precision and recall, providing a single measure of model performance that balances false positives and false negatives. The F1 score is particularly useful when dealing with imbalanced datasets.

- **AUC-ROC (Area Under the Receiver Operating Characteristic Curve):** This metric evaluates the trade-off between the true positive rate (recall) and the false positive rate at different thresholds. A higher AUC score indicates better model discrimination between phishing and legitimate instances.

- **False Positive Rate (FPR):** The proportion of legitimate instances that are incorrectly classified as phishing. A low false positive rate is crucial to minimize disruptions for legitimate users.

Choosing the right evaluation metrics depends on the application context. For example, in phishing detection, minimizing false negatives (i.e., undetected phishing attempts) is critical to security, while balancing false positives to avoid unnecessary disruption.

In this study, the effectiveness of different supervised learning models will be evaluated using these metrics, ensuring that the chosen model offers the best trade-off between detection accuracy, false positives, and computational efficiency.

V. Experimental Results and Analysis

A. Datasets and Experimental Setup

To evaluate the effectiveness of phishing detection models, this study uses publicly available phishing datasets combined with real-world legitimate instances from various sources. The datasets consist of features extracted from phishing and legitimate emails, websites, and URLs, including URL-based, content-based, and metadata features.

Datasets: The experiment utilizes a combination of widely recognized datasets such as the PhishTank database for phishing URLs and the UCI Machine Learning Repository's phishing websites dataset. Legitimate instances are gathered from trusted sources like Alexa's top websites and email service providers. The final dataset includes balanced samples of phishing and legitimate data for improved comparison.

Data Preprocessing: Before training, the data undergoes several preprocessing steps:

- Feature extraction: URL characteristics (e.g., length, domain age), content-based features (e.g., keyword frequencies, HTML tags), and metadata (e.g., sender information, email headers) are extracted.
- Normalization: Numerical features are normalized using Min-Max scaling to ensure uniformity across different feature types.
- Train-Test Split: The dataset is split into training (70%) and testing sets (30%) using stratified sampling to maintain the balance between phishing and legitimate instances.

Models and Training: The study trains and evaluates various supervised machine learning models, including:

- Support Vector Machines (SVM)
- Random Forest
- Logistic Regression
- Gradient Boosting Machines (GBM)
- Neural Networks

These models are tuned using cross-validation and grid search to optimize hyperparameters such as learning rate, regularization strength, and tree depth (for ensemble methods). Each model is trained on the preprocessed training set and evaluated on the testing set using performance metrics.

B. Model Performance Comparison

The performance of each model is compared across key evaluation metrics, including accuracy, precision, recall, F1 score, and AUC-ROC. The following table summarizes the results of the experiments:

Model	Accuracy (%)	Precision		Recall	F1 Score	AUC-ROC
Support Vector Machines	95.2	0.93	0.89	0.91	0.97	
Random Forest	97.1	0.96	0.93	0.94	0.98	
Logistic Regression	93.5	0.91	0.86	0.88	0.95	
Gradient Boosting Machines	96.5	0.95	0.92	0.93	0.97	
Neural Networks	96.8	0.94	0.94	0.94	0.98	

Random Forest outperformed other models in terms of overall accuracy (97.1%) and AUC-ROC (0.98), making it the best model for phishing detection in this experiment.
Neural Networks showed comparable results, particularly excelling in recall (0.94), which indicates its strong ability to catch phishing attacks.
Support Vector Machines (SVM) had a strong AUC-ROC (0.97) but slightly lower recall compared to Random Forest and Neural Networks.
Logistic Regression, although simpler, achieved a respectable performance but trailed behind more complex models in both recall (0.86) and F1 score (0.88).
The experiments reveal that ensemble methods like Random Forest and Gradient Boosting Machines perform exceptionally well in phishing detection, leveraging the diverse feature sets effectively.

C. Statistical Significance and Analysis

To ensure that the differences in model performance are statistically significant, the study conducts a paired t-test between the top-performing models (Random Forest, Gradient Boosting Machines, and Neural Networks) across multiple cross-validation folds. The null hypothesis states that there is no significant difference in the performance of these models, while the alternative hypothesis asserts that there is a statistically significant difference.

- Null Hypothesis (H_0): There is no significant difference in model performance.
- Alternative Hypothesis (H_1): There is a significant difference in model performance.
- After performing the t-test, the p-values obtained for the Random Forest versus Gradient Boosting Machines and Random Forest versus Neural Networks are both less than 0.05, indicating that the differences in their performance metrics are statistically significant at a 95% confidence level.

- Further analysis of feature importance reveals that URL-based features, particularly domain age, URL length, and the presence of suspicious keywords, were the most critical factors influencing phishing detection across all models. This insight highlights the importance of focusing on these high-impact features for future model refinement.

In conclusion, Random Forest demonstrated the best overall performance with statistically significant results, balancing high precision, recall, and interpretability. The findings confirm that optimizing feature vectorization and selecting the right supervised learning model can significantly enhance phishing detection systems, improving accuracy and reducing false positives.

VI. Challenges and Future Directions

A. Challenges in Phishing Detection using Machine Learning

Phishing detection using machine learning (ML) presents several technical and practical challenges, which must be addressed to improve the efficacy and adaptability of these systems.

- Evolving Phishing Tactics: Phishing attacks continuously evolve, adopting new tactics to bypass detection systems. These include the use of sophisticated social engineering techniques, highly targeted spear-phishing attacks, and dynamically generated phishing sites. ML models trained on historical data may struggle to detect these novel forms of phishing, leading to a high number of false negatives. The need for constant updates and retraining of models to accommodate these evolving tactics remains a significant challenge.

- Data Imbalance: In phishing detection, legitimate emails or websites far outnumber phishing instances, leading to highly imbalanced datasets. This imbalance can bias machine learning models, causing them to underperform in detecting rare phishing cases while maintaining high accuracy on legitimate instances. Techniques like oversampling, undersampling, or synthetic data generation (e.g., SMOTE) can mitigate this issue, but they are not always sufficient to fully address the imbalance problem.

- High False Positive Rate: Many phishing detection models, particularly those tuned for high recall, tend to flag legitimate emails or websites as phishing, leading to a high false positive rate. This reduces user trust and increases the workload for security teams tasked with manually reviewing flagged instances. Reducing false positives while maintaining a high detection rate remains a significant challenge.

- Feature Engineering Complexity: While feature selection and engineering are crucial for effective phishing detection, determining the most relevant and impactful features can be time-consuming and domain-specific. The need for frequent updates to feature sets in response to emerging attack vectors requires ongoing investment in feature engineering, which adds complexity to system maintenance.

- Real-Time Detection and Scalability: Phishing detection systems must operate in real-time to effectively block or flag phishing attempts before users fall victim. However, many machine learning models, especially deep learning models, are computationally intensive and require significant processing power. Balancing the need for real-time detection with model complexity and scalability is a key challenge, particularly for systems deployed at scale in large organizations.

- Interpretability and Transparency: More complex models like deep learning networks often act as "black boxes," making it difficult to interpret their decision-making processes. This lack of transparency complicates efforts to understand why a particular instance was flagged as phishing and hampers model debugging, improvement, and user trust. Developing interpretable models without sacrificing detection performance is an ongoing challenge in phishing detection.

B. Future Trends in Feature Vectorization and Detection

- Automated Feature Engineering: As phishing attacks become more sophisticated, manual feature engineering may struggle to keep up with emerging attack vectors. Automated feature engineering, driven by advanced algorithms like neural architecture search (NAS) and AutoML, is expected to become more prominent. These techniques can dynamically identify and optimize features without

requiring domain-specific knowledge, improving the adaptability of phishing detection systems.

- Use of Deep Learning and NLP: The increasing complexity of phishing attacks calls for more advanced detection methods. Deep learning models, especially those that leverage natural language processing (NLP), are expected to play a greater role in phishing detection. These models can process unstructured data, such as email content or website text, more effectively than traditional methods. NLP techniques like transformers (e.g., BERT, GPT) can analyze the context and semantics of phishing messages to improve detection accuracy.

- Graph-Based Detection Methods: Future trends may include the adoption of graph-based models that map relationships between various entities (e.g., URLs, email domains, IP addresses) in phishing attacks. Graph neural networks (GNNs) can analyze these relationships and detect hidden patterns in phishing campaigns that rely on interconnected systems. This approach would improve the detection of sophisticated, network-based attacks, such as spear-phishing and business email compromise (BEC).

- Adversarial Machine Learning: As phishing attackers increasingly target ML-based detection systems, the need for adversarial machine learning defenses will grow. Attackers may craft subtle, adversarial examples that deceive detection models by exploiting their weaknesses. Developing models that are resistant to such adversarial attacks is a key future direction. This could involve integrating adversarial training techniques or building more robust defenses against model manipulation.

- Transfer Learning and Incremental Learning: To address the challenge of evolving phishing techniques, transfer learning and incremental learning offer promising solutions. These methods allow models to adapt to new data without requiring full retraining, reducing the computational burden and enabling more timely updates. Transfer learning can apply knowledge from previously seen phishing patterns to detect new variations, enhancing the model's adaptability.

- Hybrid Models: Future phishing detection systems are likely to employ hybrid models that combine the strengths of multiple approaches, such as rule-based systems, machine learning, and deep learning. Hybrid systems can balance the interpretability of traditional models with the adaptability and accuracy of ML models, offering a more robust detection framework. This integration can improve detection across different phishing attack types and contexts.

- Behavioral-Based Detection: In addition to content and metadata features, future systems may increasingly rely on behavioral-based detection. Monitoring user interaction patterns, such as click behavior or time spent on suspicious websites, can provide additional signals for phishing detection. Behavioral analysis could be particularly useful in detecting phishing-as-a-service (PhaaS) campaigns that continuously evolve their attack methods.

- Real-Time Detection with Edge Computing: To improve the scalability and speed of phishing detection, edge computing could play a key role in future systems. By

processing data at the network edge (closer to the user), real-time phishing detection becomes more feasible without overloading centralized servers. Edge-based models can analyze phishing threats locally, enabling faster detection and response times, especially in large-scale deployments.

In summary, the future of phishing detection lies in leveraging advanced machine learning techniques, automated feature engineering, and more sophisticated models that adapt to the evolving landscape of phishing attacks. By addressing current challenges, future detection systems will offer improved accuracy, real-time detection, and robustness, contributing to stronger cybersecurity defenses.

VII. Conclusion

A. Summary of Key Findings

This study explored the use of machine learning (ML) techniques for enhancing phishing attack detection through optimal feature vectorization and supervised learning models. Key findings include:

- Optimal Feature Vectorization: Proper selection and engineering of features, such as URL-based, content-based, and metadata attributes, significantly impact model performance. URL length, domain age, and suspicious keyword presence were identified as critical features.
- Supervised Learning Models: Among the models evaluated, Random Forest emerged as the top-performing algorithm, achieving high accuracy, recall, and AUC-ROC scores. Ensemble methods, including Gradient Boosting Machines (GBM), also performed well, while simpler models like Logistic Regression provided good but lower performance.
- Challenges: The study highlighted challenges such as data imbalance, evolving phishing tactics, high false positive rates, and real-time detection constraints, which must be addressed to improve the robustness of phishing detection systems.

B. Implications for Cybersecurity

The findings of this research have important implications for the broader field of cybersecurity:

- Improved Phishing Detection: By enhancing phishing detection systems through optimal feature engineering and leveraging advanced machine learning models, organizations can better defend against phishing attacks, which remain a top cybersecurity threat.
- Reduced False Positives: Reducing false positives while maintaining high detection accuracy is essential for maintaining user trust and minimizing the workload for security teams. Advanced techniques, such as ensemble models and feature optimization, can help achieve this balance.
- Adaptability to Emerging Threats: The continuous evolution of phishing tactics calls for dynamic and adaptable detection systems. The application of techniques like automated feature engineering and transfer learning can enhance the adaptability of phishing detection models to new attack vectors.

C. Recommendations for Further Research

The study identifies several areas for future research to further enhance phishing detection:

- Adversarial Machine Learning: Future work should focus on developing models resistant to adversarial attacks, where attackers attempt to manipulate machine learning systems to evade detection.
- Automated Feature Engineering: Exploring automated feature engineering methods, such as AutoML and neural architecture search (NAS), can help continuously optimize phishing detection systems without manual intervention.
- Real-Time Detection: Further research on deploying real-time detection systems using edge computing and lightweight machine learning models could improve the scalability and responsiveness of phishing defenses.
- Hybrid Detection Models: Future research could also focus on integrating hybrid models, combining rule-based, machine learning, and deep learning approaches to enhance robustness across various phishing attack types.

In conclusion, this study demonstrates that optimizing feature vectorization and selecting appropriate machine learning models can significantly enhance phishing detection capabilities, contributing to stronger defenses against cyber threats.

References:

1. Ghedabna, L., Ghedabna, R., Imtiaz, Q., Faheem, M. A., Alkhayyat, A., & Hosen, M. S. (2024). Artificial Intelligence in Human Resource Management: Revolutionizing Recruitment, Performance, and Employee Development. *Nanotechnology Perceptions*, 52-68.Tamal, M. A., Islam, M. K., Bhuiyan, T., Sattar, A., & Prince, N. U. (2024). Unveiling suspicious phishing attacks: enhancing detection with an optimal feature vectorization algorithm and supervised machine learning. *Frontiers in Computer Science*, *6*, 1428013.
2. Chowdhury, R. H., Prince, N. U., Abdullah, S. M., & Mim, L. A. (2024). The role of predictive analytics in cybersecurity: Detecting and preventing threats. *World Journal of Advanced Research and Reviews*, *23*(2), 1615-1623.
3. Chowdhury, N. R. H., Prince, N. N. U., Abdullah, N. S. M., & Mim, N. L. A. (2024d). The role of predictive analytics in cybersecurity: Detecting and preventing threats. *World Journal of Advanced Research and Reviews*, *23*(2), 1615–1623. https://doi.org/10.30574/wjarr.2024.23.2.2494
4. Prince, N. U., Faheem, M. A., Khan, O. U., Hossain, K., Alkhayyat, A., Hamdache, A., & Elmouki, I. (2024). AI-Powered Data-Driven Cybersecurity Techniques: Boosting Threat Identification and Reaction. *Nanotechnology Perceptions*, 332-353.
5. Faheem, M. A., Zafar, N., Kumar, P., Melon, M. M. H., Prince, N. U., & Al Mamun, M. A. (2024). AI AND ROBOTIC: ABOUT THE TRANSFORMATION OF CONSTRUCTION INDUSTRY AUTOMATION AS WELL AS LABOR PRODUCTIVITY. *Remittances Review*, *9*(S3 (July 2024)), 871-888.
6. Priyadharshini, S. L., Al Mamun, M. A., Khandakar, S., Prince, N. N. U., Shnain, A. H., Abdelghafour, Z. A., & Brahim, S. M. (2024). Unlocking Cybersecurity Value through Advance Technology and Analytics from Data to Insight. *Nanotechnology Perceptions*, 202-210.