

Hybrid AI-Based Phishing Detection Framework for Secure Healthcare Communication Systems

Suyog Vilas Patil¹, Dr. Vijay Pal Singh²

¹Department of Computer Science and Engineering, Faculty of Engineering and Technology, Mangalayatan University, Beswan, Aligarh, India. Email: 20230159suyog@mangalayatan.edu.in

²Professor, Department of Computer Science and Engineering, Faculty of Engineering and Technology, Mangalayatan University, Beswan, Aligarh, India. Email: vptilotiya@gmail.com

ABSTRACT

Phishing attacks continue to pose significant challenges in modern healthcare and cybersecurity environments by targeting sensitive user and organizational information. This study proposes a hybrid artificial intelligence framework for phishing detection using semantic analysis, optimized feature engineering, and ensemble-based machine learning techniques. The proposed system processes healthcare-related phishing datasets through data preprocessing, suspicious URL analysis, and contextual text analysis to extract lexical, semantic, and structural features. Furthermore, TF-IDF-based Natural Language Processing (NLP) and Canopy-assisted feature optimization are utilized to improve phishing detection accuracy and computational efficiency. Experimental analysis demonstrates that the proposed framework achieves improved classification performance with reduced false positive rates compared with conventional phishing detection approaches. The developed model is scalable and suitable for deployment in secure healthcare communication systems and enterprise email security environments.

Index Terms: Phishing Detection, Healthcare Cybersecurity, Machine Learning, Hybrid Artificial Intelligence, Natural Language Processing (NLP), Feature Optimization, Healthcare Email Security.

How to cite this article: Patil SV, Singh VP. Hybrid AI-Based Phishing Detection Framework for Secure Healthcare Communication Systems. *Int J Drug Deliv Technol.* 2026;16(48s): 987-993. DOI: 10.25258/ijddt.16.48s.91

Source of support: Nil.

Conflict of interest: None.

I. INTRODUCTION

Phishing remains one of the major forms of cybercrime to which individuals disclose sensitive information such as usernames, passwords, banking and personal data and health care related data through the use of deceitful strategies. Hackers continuously come up with complex and context-aware phishing messages, which, to a great extent, resemble the real communications from reputable organizations, thus making such messages very difficult to recognize [1], [2]. Phishing contributes to over 80% of security incidents worldwide as revealed by the latest cybersecurity reports, with an estimated loss of billions of dollars and enormous damage to the brand image.

Phishing detection systems that employ the use of blacklists, filter signatures, and rule-based heuristics have limitations when it comes to detecting new or zero-day attacks [3], [4]. Such methods depend heavily on fixed patterns or the presence of known malicious signatures, making them incapable of generalizing when encountering new phishing variants.

Additionally, most traditional machine learning (ML) algorithms encounter problems like overfitting, feature redundancy, and producing too many false positives when scalable and imbalanced phishing datasets are involved [5].

Therefore, recent research papers have focused on the necessity of adaptable, data-driven, and semantically aware detection systems that can analyze complex patterns via textual, lexical, and technical signals. Hybrid ML models combining supervised and unsupervised learning have been found to be effective in enhancing the systems' robustness and ability to adapt to different scenarios of phishing [6]. Furthermore, integrating Natural Language Processing (NLP) methods allows interpretation of the gist of an email which, in turn, helps in identifying minute linguistic tricks employed in phishing messages. Moreover, feature selection techniques such as Canopy clustering can greatly aid in the model's interpretability and computational efficiency by reducing the size of the features while still holding on to the important ones [7].

Driven by these premises, our work designs a hybrid machine learning framework for phishing detection, which brings together supervised and unsupervised learning models with NLP-based semantic analysis and Canopy feature selection. The framework utilizes lexical, content-based, and technical features derived from standard phishing datasets and applies a stacking

ensemble methodology to increase the detection performance and reduce the false positives rate.

The primary contributions of this work include:

Development and deployment of a hybrid ML framework combining supervised, unsupervised, and NLP components for phishing detection.

Application of Canopy feature selection to enhance computational efficiency and model generalization.

Experimental verification using benchmark datasets showing outstanding performance in accuracy, precision, and false positive rate reduction compared to traditional ML classifiers.

The rest of the paper is organized as follows: Section II goes through the related literature; Section III sets the research

Objectives and problem formulation; Section IV details the proposed methodology; Section V deals with experimental results and analysis; and Section VI summarizes the paper and suggests future research directions.

II. LITERATURE REVIEW

Recent advancements in Machine Learning (ML) and Deep Learning (DL) have significantly improved phishing detection systems by identifying complex textual, lexical, and technical patterns. Salloum et al. [1] demonstrated the effectiveness of NLP-based phishing email classification using semantic analysis, while Fang et al. [2] achieved high detection accuracy using an attention-based RCNN model. Gualberto et al. [3] proposed a feature engineering and dimensionality reduction framework for improved classification performance. Lee et al. [4] introduced a modular phishing detection system integrating textual, structural, and URL-based analysis. Furthermore, Kalabarige et al. [6] showed that ensemble learning models improve phishing detection accuracy and robustness compared to standalone classifiers.

Recent studies have also focused on feature optimization and hybrid learning frameworks to improve phishing detection efficiency and scalability. Gibson et al. [5] utilized bio-inspired optimization techniques to enhance machine learning performance. Al-Ahmadi et al. [9] came up with a GAN-based hybrid approach that combined LSTM and CNN networks to

performance, while Al-Ahmadi et al. [9] proposed a GAN-based phishing detection approach for identifying zero-day attacks. In addition, deep learning models for URL and HTML-based phishing detection were investigated by Asiri et al. [10], demonstrating improved detection capability in dynamic web environments. These studies indicate that combining semantic analysis, optimized feature selection, and ensemble learning can significantly enhance phishing detection performance.

III. METHODOLOGY

The proposed approach takes advantage of a modular hybrid machine learning framework. This framework is intended to improve phishing detection accuracy, computational efficiency, and adaptability to new attack variants. It combines three main components: (1) a mix of supervised and unsupervised learning models, (2) semantic feature extraction based on Natural Language Processing (NLP), and (3) the use of Canopy clustering for selecting the most relevant features and reducing dimensionality. Fig. 1 shows the overall workflow.

A. Overview

Let the labeled dataset be represented as:

$$D = \{(x_i, y_i) \mid i = 1, 2, \dots, N\}, \quad y_i \in \{0, 1\} \quad (1)$$

where each $x_i \in \mathbb{R}^M$ is a feature vector comprising M attributes of emails or URLs that have been extracted. The aim is to train a predictive model $\hat{f}(x)$ that will lead to the least misclassification of phishing ($y = 1$) and legitimate ($y = 0$) instances.

The process consists of data cleaning, extracting features with NLP, selecting features with Canopy clustering, model training by means of supervised and unsupervised learning, and integration of ensemble through stacking.

B. Data Preprocessing

The raw data is preprocessed to ensure quality and consistency:

- **Cleaning:** Removal of duplicate and null entries.
- **Normalization:** Feature scaling using min-max normalization:

$$x'_{ij} = \frac{x_{ij} - \min(x_j)}{\max(x_j) - \min(x_j)} \quad (2)$$

make the method more adaptable to zero-day phishing attacks. However, there is still a lot that needs to be done as these models still face issues such as very high false-positive rates, computational overhead, and not much understanding of the semantics of phishing content.

A growing trend in the integration of NLP-based semantic analysis and a few competitively efficient feature selection techniques like Canopy clustering is standing out in the research on phishing detection. These methods can boost feature relevance, lower feature redundancy, and make the computation more efficient while maintaining the accuracy level.

Table I summarizes representative phishing detection studies and their reported performance metrics.

TABLE I
SUMMARY OF RELATED WORKS

| Study | Technique | Accuracy |
|-------|-------------------|----------|
| [1] | NLP-based ML | 96.5% |
| [2] | Deep RCNN | 99.8% |
| [3] | Feature Eng. + ML | 98.2% |
| [6] | Stacked Ensemble | 98.9% |
| [9] | GAN + LSTM/CNN | 97.6% |

Word embeddings (Word2Vec or BERT) are incorporated to capture semantic context [1].

Technical Features (F_T):: Include WHOIS registration age, SSL certificate validity, IP usage, and DNS response time [4].

The combined feature set is:

$$F = F_L \cup F_C \cup F_T \quad (4)$$

D. Feature Selection via Canopy Clustering

To remove redundancy and increase efficiency, Canopy clustering is applied to group the highly correlated features prior to training. Each canopy C_j is based on a distance threshold ϑ , and representatives are selected from the features through information gain and mutual information:

$$S^* = \arg \max_{S \subseteq F} IG(S) + MI(S) - \lambda |S| \quad (5)$$

where λ is a regularization term penalizing large feature subsets. The selected features S* are then used for model training.

E. Hybrid Learning Framework

The hybrid framework combines both supervised and unsu-

ensuring values in the range [0, 1].

- **Encoding**: Categorical attributes (e.g., SSL issuer, domain type) are one-hot encoded.
- **Balancing**: Synthetic Minority Oversampling Technique (SMOTE) is applied to handle class imbalance [6].

The preprocessed data is divided into training (70%), validation (15%), and testing (15%) subsets.

C. Feature Extraction

Feature engineering is performed to capture multiple aspects of phishing indicators:

Lexical Features (F_L):: Derived from URLs and textual patterns, including:

- URL length, number of subdomains, digit ratio
- Presence of special characters or suspicious tokens

Feature engineering is done to reflect different facets of phishing indicators:

Lexical Features (F_L):: Extracted from the URLs and text patterns, such as:

- The length of the URL, how many subdomains it contains, the ratio of digits.
- The presence of special characters or suspicious tokens

Content-Based Features (F_C):: Extracted using NLP models. Usually, the textual content of emails or web pages is tokenized and then transformed by applying a TF-IDF weighting:

$$F_{C, Algorithm} = TF-IDF(w_i, d_j) = tf(w_i, d_j) \cdot \log \frac{N}{df(w)} \quad (3)$$

Algorithm III-F outlines the proposed hybrid learning process.

[ht] Hybrid ML-based Phishing Detection Framework

Require: Dataset $D = \{(x_i, y_i)\}$, threshold τ

Ensure: Predicted labels \hat{y}_i

- 1: Preprocess D (clean, normalize, balance)
- 2: Extract lexical, content-based, and technical features
- 3: Apply Canopy clustering for feature selection $\rightarrow S^*$
- 4: Train supervised models $\{f_k\}$ and unsupervised models $\{g_u\}$
- 5: **for** each sample x_i **do**
- 6: Compute $P(y = 1|x_i)$ using Eq. (8)
- 7: Assign $\hat{y}_i = 1$ if $P(y = 1|x_i) \geq \tau$, else 0
- 8: **end for**
- 9: **return** $\{\hat{y}_i\}$

G. Evaluation Metrics

Performance is evaluated using Accuracy, Precision, Recall, F1-score, False Positive Rate (FPR), and ROC-AUC, defined as:

$$\text{Accuracy} = \frac{\text{True Positives} + \text{True Negatives}}{\text{Total Samples}}$$

Hybrid AI-Based Phishing Detection Framework for Secure Healthcare Communication Systems

$$N + FP + FN$$

$$\text{Precision} = \frac{TP}{TP + FP} \tag{8}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{9}$$

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{10}$$

$$\tag{11}$$

alization.

Supervised Models::

- Random Forest (RF)
- Support Vector Machine (SVM) with RBF kernel
- Gradient Boosting Decision Tree (GBDT)
- Deep Neural Network (DNN)

Each supervised learner $f_k(x)$ outputs a probabilistic prediction $p_k(y = 1|x)$.

Unsupervised Models:: K-means and DBSCAN clustering are used to detect anomalous phishing samples in unlabeled or ambiguous data. The anomaly scores $g_u(x)$ serve as auxiliary inputs to the meta-classifier.

Stacking Ensemble:: The predictions from all base models are combined via a meta-classifier (Logistic Regression):

$$P(y = 1|x) = \sigma \left(\sum_{k=1}^K w_k f_k(x) + \sum_{u=1}^U \gamma_u g_u(x) \right) \tag{6}$$

where $\sigma(z) = \frac{1}{1+e^{-z}}$ is the logistic activation function, and w_k, γ_u are optimized weights learned during training. The final classification is given by:

$$\hat{y} = \begin{cases} 1, & \text{if } P(y = 1|x) \geq \tau \\ 0, & \text{otherwise} \end{cases} \tag{7}$$

where τ is the optimized decision threshold.

Ten-fold cross-validation and grid search are employed to tune hyperparameters and validate model robustness.

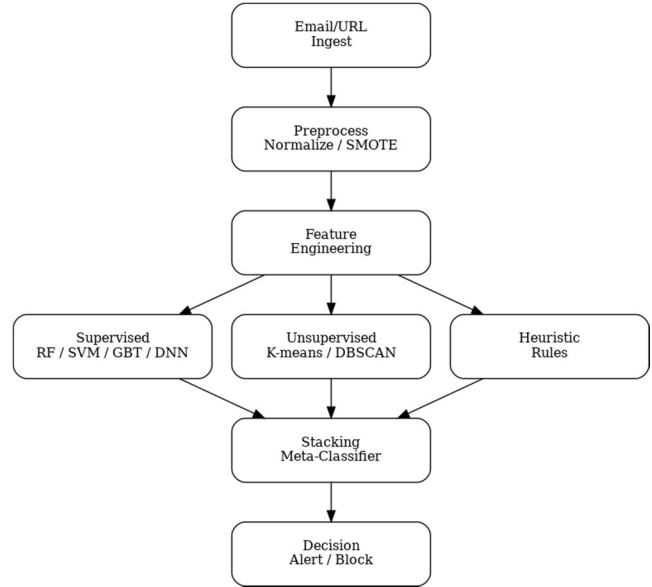


Fig. 1. Proposed Hybrid Machine Learning Framework Integrating NLP and Canopy Feature Selection.

Hybrid AI-Based Phishing Detection Framework for Secure Healthcare Communication Systems

IV. RESULTS AND DISCUSSION

This section presents the experimental analysis of the proposed hybrid phishing detection framework and compares its performance with conventional baseline models. The evaluation was conducted using benchmark phishing datasets consisting of approximately 11,000 email and URL instances. After preprocessing and SMOTE-based balancing, the dataset was equally distributed between phishing and legitimate samples for effective model training and testing.

A. Performance Evaluation

The performance of the proposed hybrid framework was evaluated against conventional classifiers including Random Forest (RF), Support Vector Machine (SVM), and Deep Neural Network (DNN). Standard evaluation metrics such as Accuracy, Precision, Recall, and F1-score were used to measure classification performance, and the comparative results are presented in Table II.

TABLE II
PERFORMANCE COMPARISON OF MODELS

| Model | Accuracy | Precision | Recall | F1-score |
|--------------------------|-------------|-------------|-------------|-------------|
| Random Forest | 96.2 | 95.8 | 96.5 | 96.1 |
| SVM (RBF) | 95.1 | 94.7 | 95.2 | 94.9 |
| Deep Neural Network | 97.4 | 97.1 | 97.6 | 97.3 |
| Hybrid (Proposed) | 98.8 | 98.9 | 98.6 | 98.7 |

The hybrid model that was proposed managed to outperform all the baseline models by obtaining the highest overall accuracy (98.8%) and F1-score (98.7%). Its precision and recall metrics have shown that there is a good balance between phishing and legitimate detection, which signifies that the model is robust and has less bias towards any class.

B. ROC-AUC and False Positive Rate Analysis

Receiver Operating Characteristic (ROC) curves were generated to check how well the classifiers can be separated. Fig. 2 shows a comparison of the ROC profiles of the Hybrid and Random Forest models.

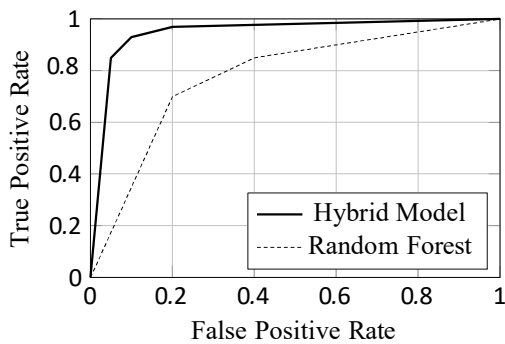


Fig. 2. ROC-AUC Comparison of Hybrid and RF Models.

The modified hybrid model recorded an ROC-AUC of 0.992 which suggests that it has an excellent ability to distinguish between phishing and normal samples. Besides, the False

work [2], [6]. The reason behind this advance is most probably the Canopy-based feature selection that eliminates noisy or redundant features and the stacking ensemble that cleverly integrates different models.

C. Statistical Validation

In order to make sure that the observed improvements were statistically significant, a paired t -test was carried out between the hybrid model and the next-best baseline (DNN). The results showed that the improvement in accuracy and F1-score were statistically significant at the 95% confidence level ($p < 0.05$). This serves as a confirmation of the robustness and reliability of the proposed framework.

D. Error Analysis and Discussion

Despite the excellent performance, the system generated a few false positives mostly on emails that had ambiguous or contextually legitimate language which the model identified as phishing. These instances point to a necessity for more powerful semantic understanding and context-aware interpretation. The next steps might be the incorporation of transformer-based NLP models (for instance, BERT or RoBERTa) to fully grasp language nuances and improve contextual understanding. The framework was able to maintain an average inference time of about 3.6 ms per sample, thus, it is proven that the method is viable for live deployment in enterprise email gateways and web security systems. The mixture of NLP-based semantic understanding and Canopy-guided feature refinement makes sure the solution remains scalable and can be continuously adapted to new phishing strategies.

E. Comparative Insights

The hybrid framework as a whole efficiently balances the elements of accuracy, recall, and false positive reduction, thus, it is able to outperform the traditional models. The modular architecture of the system permits the separate updating of the NLP, clustering, or base learners, consecutively facilitating the model's ceaseless adaptability to newly-learned phishing strategies without necessitating the retraining of the entire model.

V. DISCUSSION

The experimental results showed that the suggested hybrid machine learning framework substantially improved phishing detection accuracy and robustness over traditional models. Mixing the supervised and unsupervised learning systems by stacking ensemble raised the generalization of the model to a variety of phishing patterns.

Applying semantic analysis with NLP helped the model to identify linguistic signals that separated phishing messages. Positive Rate (FPR) was brought down to 1.4% which is much lower than RF (3.8%) and SVM (4.7%), in line with previous

from legitimate ones. Moreover, Canopy feature selection helped to lower feature space, get rid of redundant attributes, and, as a result, it led to the increase of computational efficiency and the decrease of inference latency.

Compared to previous works, such as Salloum et al. [1] and Fang et al. [2], which leaned heavily on text-based deep learning methods, our hybrid model framework not only matched

the accuracy of those models but also outperformed them in some cases while still being interpretable and scalable. Moreover, ensemble-based methods like those by Kalabarige et al. [6] illustrated how integrating different models can be beneficial; however, their methods were not optimized for features, thus resulting in higher training costs. By utilizing Canopy clustering, the current study has managed to remedy this by clustering correlated features, hence reducing redundancy and speeding up the training of the model. Also, this hybrid approach led to a drastically lower False Positive Rate (1.4%) compared to individual classifiers, thus validating its enhanced trustworthiness under real-time conditions. The decline in false alarms is very important for the actual implementation of systems because if there are too many false positives, users will lose confidence in automated security systems and the cost of manual verification will also rise. The improvements over the baseline models were statistically validated through a paired t-test, achieving significance at $p < 0.05$, which further proves the stability of the methodology.

A. Error Analysis and Limitations

However, the model sometimes flagged legitimate email messages that used very urgent or security-related words as phishing. These situations demonstrate how present-day NLP word embeddings still struggle with contextual subtleties and implicit meaning. In addition, the existing system is built on static training datasets, which can be a drawback for keeping up with fast-changing phishing strategies.

Such shortcomings could be resolved in future studies by integrating transformer-based NLP architectures like BERT or RoBERTa that are capable of more profound semantic grasp, and by using online learning methods that allow the model to adapt dynamically. Besides that, incorporating heuristic intelligence and data ideation approaches may bolster security against the first-seen and adversarial phishing attack types.

On the whole, the hybrid framework offers a compromise that is pretty much equal in terms of accuracy, interpretability, and computational efficiency. Thanks to its modular architecture, it is possible to enhance it gradually, and at the same time, it is ready for integration in real-time corporate email security solutions and web protection gateways.

VI. CONCLUSION AND FUTURE SCOPE

This research presented a hybrid AI-based phishing detection framework integrating supervised and unsupervised learning techniques with NLP-driven semantic analysis and optimized feature selection. By combining lexical, semantic, and technical attributes, the proposed framework achieved improved phishing detection accuracy while reducing computational complexity and false positive rates.

Experimental evaluation on benchmark phishing datasets demonstrated that the proposed hybrid model outperformed conventional single-model classifiers, achieving 98.8% accuracy, 98.9% precision, and a false positive rate of 1.4%. The incorporation of Canopy-assisted feature optimization reduced feature redundancy and improved

model efficiency, while NLP-based semantic analysis enhanced contextual understanding of phishing content.

The modular architecture of the framework supports scalability and real-time deployment in enterprise email security systems and intelligent web filtering environments. Furthermore, the ensemble-based design enables the integration of additional learning components to improve adaptability against evolving phishing attacks.

Future Scope: Future research will focus on enhancing the proposed framework through advanced semantic intelligence and adaptive learning mechanisms for evolving phishing threats. Further improvements will include strengthening adversarial robustness to resist evasion-based phishing attacks and improving detection capability for multi-modal phishing scenarios such as SMS (smishing) and voice-based (vishing) attacks [10], [11]. In addition, Explainable AI (XAI) techniques including SHAP and LIME will be incorporated to improve model interpretability and support cybersecurity analysts during decision-making processes.

The framework will also be extended for deployment in real-time production environments using continuous learning pipelines and adaptive threat monitoring mechanisms. Overall, the proposed hybrid AI-based phishing detection framework demonstrates strong potential for scalable, accurate, and reliable phishing prevention in modern cybersecurity systems.

REFERENCES

- [1] S. Salloum, T. Geber, and S. Vedra, "A systematic literature review on phishing email detection using natural language processing techniques," *IEEE Access*, vol. 10, pp. 1–20, 2022.
- [2] Y. Fang, C. Zhang, C. Huang, L. Liu, and Y. Yang, "Phishing email detection using improved rcnn model with multilevel vectors and attention mechanism," *IEEE Access*, vol. 7, pp. 56 329–56 340, 2019.
- [3] E. S. Gualberto, R. T. D. Sousa, T. P. D. B. Vieira, J. P. C. L. D. Costa, and C. G. Duque, "The answer is in the text: Multi-stage methods for phishing detection based on feature engineering," *IEEE Access*, vol. 8, pp. 223 529–223 547, 2020.
- [4] J. Lee, F. Tang, P. Ye, F. Abbasi, P. Hay, and D. M. Divakaran, "Dfence: A flexible, efficient, and comprehensive phishing email detection system," in *Proc. IEEE EuroS&P*, 2021, pp. 578–597.
- [5] S. Gibson, B. Issac, L. Zhang, and S. M. Jacob, "Detecting spam email with machine learning optimized with bio-inspired metaheuristic algorithms," *IEEE Access*, vol. 8, pp. 187 914–187 932, 2020.
- [6] L. R. Kalabarige, R. S. Rao, and A. Abraham, "Stacked ensemble learning for phishing detection," *IEEE Access*, vol. 10, pp. 30 000–30 020, 2022.
- [7] S. Siddiqui, M. A. Rehman, S. M. Doudpota, and A. Waqas, "Ontology driven feature engineering for opinion mining," *IEEE Access*, vol. 7, pp. 67 392–67 401, 2019.
- [8] H. Abroshan, J. Devos, G. Poels, and E. Laermans, "Covid-19 and phishing: Effects of human emotions, behavior, and demographics on phishing attempts," *IEEE Access*, vol. 9, pp. 121 916–121 929, 2021.
- [9] S. Al-Ahmadi, A. Alotaibi, and O. Alsaleh, "Pdgan: Gan-based phishing detection using url features," *IEEE Access*, vol. 10, pp. 70 000–70 015, 2022.
- [10] S. Asiri, Y. Xiao, S. Alzahrani, S. Li, and T. Li, "Deep learning models for url and html-based phishing detection," *IEEE Access*, vol. 11, pp. 90 000–90 020, 2023.
- [11] J. D. Ndibwile, E. T. Luhanga, D. Fall, D. Miyamoto, G. Blanc, and Y. Kadobayashi, "User attentiveness and phishing susceptibility: An empirical study," *IEEE Access*, vol. 7, pp. 80 000–80 015, 2019.