An Adaptive Phishing Detection System Based on Heuristic Algorithms, NLP Analysis, and Canopy-Optimized Feature Selection

Suyog Vilas Patil*, Dr. Vijay Pal Singh†

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

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

***Corresponding Author: Suyog Vilas Patil, Department of Computer Science and Engineering, Faculty of Engineering and Technology, Mangalayatan University, Beswan, Aligarh, India.**
Email: 20230159_suyog@mangalayatan.edu.in

Received: 8th May, 2026 **Revised: 21st May, 2026** **Accepted: 27th May, 2026** **Available Online: 12th June, 2026**

***Abstract*—Phishing attacks are still one of the major threats in the cybersecurity sphere, as these attacks exploit trust and technical defense weaknesses to obtain sensitive information. With increasingly advanced and adaptable phishing methods, attackers have rendered traditional rule-based and single-model detection systems less effective, particularly against zero-day and rapidly evolving attacks. This paper presents an adaptive phishing detection system based on a combination of heuristic intelligence, NLP-based semantic analysis, Canopy-optimized feature selection, and a hybrid machine learning framework. The proposed method extracts lexical, content-based, and technical features from phishing emails and URLs to capture comprehensive attack characteristics. Canopy clustering is employed to remove feature redundancy and reduce dimensionality, thereby improving computational efficiency and generalization. A hybrid learning architecture combines supervised classifiers and unsupervised anomaly detection models, whose outputs are integrated through a stacking ensemble to enhance reliability and adaptability. An extensive experimental evaluation on benchmark phishing datasets comprising approximately 11,000 samples demonstrates that the proposed framework achieves 98.8% accuracy, 98.9% precision, 98.6% recall, and a ROC-AUC score of 0.992, while maintaining a low false-positive rate of 1.4%. Furthermore, the framework achieves an average inference time of 3.6 ms per sample, making it highly suitable for real-time applications such as email gateways and enterprise web security systems. The findings confirm that the proposed adaptive framework is a viable, efficient, and highly effective solution for next-generation phishing detection and prevention.**

Index Terms: Phishing, Cyber Security, Machine Learning, Hybrid Model, NLP, Feature Selection, Email Security.

How to cite this article: Patil SV, Singh VP. An Adaptive Phishing Detection System Based on Heuristic Algorithms, NLP Analysis, and Canopy-Optimized Feature Selection. *Int J Drug Deliv Technol.* 2026;16(58s):1627–1641. DOI: 10.25258/ijddt.16.58s.172

Source of support: Nil.

Conflict of interest: None.

I. INTRODUCTION

Phishing attacks have become one of the most prevalent and harmful forms of cybercrime, posing serious threats to individuals, organizations, and critical infrastructure. Attackers often impersonate trusted entities through fraudulent emails, fake websites, and deceptive messages to steal sensitive information such as login credentials, banking details, and personal data. According to recent research, the phishing problem is a main cause of worldwide cyber security breaches and leads to enormous losses in finance as well as reputation [1], [2]. Most conventional methods for identifying phishing attacks, for example, filtering through blacklists, matching signatures, and following rule-based heuristics, depend on static patterns and previously encountered attack signatures. However, such methods cannot detect new attacks or rapidly changing phishing strategies [3]. Besides, rule-based approaches generally have problems with too many false alarms and hardly ever can they be expanded in large scale, especially when used in real-time scenarios [4]. As the field of machine learning continues to evolve, data-driven anti-phishing systems are attracting more and more attention. The detection success is more than satisfactory when supervised learning methods are employed along with the manually created features [5]. The use of NLP alongside deep learning has made it possible to analyze the meaning and context of phishing materials. Therefore, the detection rate has been significantly improved [6],[7]. On the other hand, models based on deep learning demand a large quantity of annotated data, significantly increase running time, and explainability of results is limited, which means that their use in practical scenarios is very limited [8]. Different types of classifiers can be combined by stacking ensemble models to reduce variance and thus enhance accuracy in detection [9], [10]. Nonetheless, many ensemble based methods are still dependent on high-dimensional and static features, thus suffering from significant computational overhead and feature redundancy [4]. Besides that, fully supervised methods are not able to detect previously unseen phishing attacks because they depend on labeled data [11]. On top of that, unsupervised and clustering based methods have been implemented for detecting zero-day phishing and for analysis of phishing campaigns [12], [13]. Despite such strengths, these techniques tend to be less accurate in classification as they lack precise decision boundaries and rely on integration with supervised models for that purpose. Moreover, recent pieces of research highlight how human heuristic intelligence as well as user psychology are key determinants in phishing success [14], [15]. In view of these constraints, we propose an adaptive phish ing detection system that leverages heuristic algorithms, NLP-based semantic analysis, Canopy-optimized feature selection, and a hybrid machine learning framework. The system is designed to extract lexical, content-based, and technical features to get full phishing characteristics. Dimensionality reduction and elimination of redundant features are carried out using the Canopy clustering. Supervised and unsupervised models are stacked in an ensemble method to bring together the three advantages of accuracy in detection, low false positive rate, and adaptability to new phishing strategies. The main contributions of this research are summarized as follows: • Development of a heuristic-driven adaptive phishing detection framework capable of identifying evolving malicious behaviors [16]. • Application of Canopy feature selection to optimize feature sets, reduce dimensionality, and improve computational efficiency [17]. • Integration of NLP-based semantic analysis with hybrid supervised and unsupervised learning models for robust phishing detection [7], [9]. • Extensive experimental evaluation demonstrating superior performance and real-time applicability compared to conventional approaches. The rest of this paper is arranged in the following manner. We review the existing literature and identifying the research gaps in Section II. After that, we go over the methodology of the proposed system and the system architecture in Section III. Then, Section IV goes into the details of the experiments, data analysis, and the results. Finally, in Section V, the paper is wrapped up and the future research directions are presented.

II. LITERATURE REVIEW

Phishing detection has profoundly changed in the last ten years or so. It has moved from simple rule-based tools to more sophisticated adaptable machine learning and deep learning approaches. The main reason for this change is the rising level of complexity of phishing attacks that exploit different ways of tricking the victims, human psychology, and changing rapidly their infrastructure to avoid old-style methods of detection. Most of the earlier research on phishing and email classification explored the use of human-created features along with traditional machine learning classifiers. Mujtaba et al. [1] did a very comprehensive survey of email classification systems, pinpointing the major issues of such systems as being overwhelmed with features, dealing with shifts and changes in patterns over time, and the imbalance of different classes. Chin et al. [3] developed Phish Limiter, which not only checks for phishing but also works alongside software-defined networking (SDN) for the automated clearance of phishing traffic once its attack is detected at the server level. However, methods such as these depend heavily on rules that are known in advance and do not have a genuine understanding of phishing content. Feature engineering was an important factor that led to the improvement of phishing detection rate. Gualberto et al. [5] put forward multi-stage phishing detection pipeline which integrated textual feature extraction with machine learning classifiers, and the experiment results showed the improve ment of accuracy through staged dimensionality reduction. El Aassal et al. [4] took a more experimental approach and used PhishBench to benchmark features of phishing. That research unveiled the fact that useless and noisy features are in conflict with the classifier performance and also raise the cost of computing; thus feature selection is crucial, especially when dealing with large datasets. Use of NLP (Natural Language Processing) has greatly helped in detecting phishing since it enables understanding not only the meaning of the words but also the context in which the words are used in a

message. Salloum et al. [7] did a systematic study revealing that models based on NLP usually significantly outperform traditional lexical approaches as they take into account both the syntax and semantics of phishing emails. Fang et al. [6] introduced an enhanced RCNN combined with a self-attention unit, which allows the recurrent network to selectively attend to the representations generated through convolution. Even though such networks may attain high accuracy, they often have issues related to a large amount of computation and poor interpretability. Asiri et al. [8] delved deeper into deep learning models for URL and HTML phishing detection, finding that deep features lead to better generalization but also that the models require not only a large amount of labeled data but also a lot of training time. In order not to greatly tie the fate of a model to just one mechanism, ensemble and hybrid learning architectures have been broadly used. Kalabarige et al. [9] through their research showed that combining different classifiers in a stack ensemble significantly potentiates the performance of phishing detection. Innab et al. [10] tried ensemble machine learning algorithms to reduce the rate of false positives and thus increase the overall robustness of the models when tested on different datasets. Gibson et al. [16] presented a metaheuristic bio-inspired technique that can be employed for the fine-tuning of the hyper-parameters of the classifiers and thus producing results that equal or surpass the existing baseline in both convergence and classification. However, the majority of these ensemble methods still build their decisions on the same set of unchanging features and are not flexible enough to be able to respond to changes in the phishing strategies. Unsupervised and semi-supervised methods are always a step ahead trying to take care of zero-day phishing. Liu and Fu [11] came up with SPWalk, an unsupervised approach that learns features by performing random walks on largescale URL graphs to model structural relationships between URLs. Althobaiti et al. [12] resorted to clustering techniques to categorize phishing campaigns, which in turn made it possible to detect new evolving campaigns. Lee et al. [13] went a step further in tackling phishing issues by analyzing different phishing attack groups and subsequently proposing counter-measures based on the detection of similar behaviors. Anomaly A. Research Gap and Motivation detection methods often suffer from a lack of well-defined decision boundaries and that is why they usually need to be coupled with supervised classifiers so that the classification results can be trustworthy. More and more recent papers have explored the potential of generative and adversarial learning to improve phishing detection systems adaptability. Al-Ahmadi et al. [18] introduced PDGAN, a framework for phishing detection that is based on GAN and capable of generating synthetic phishing URLs which increases the robustness of the model against unseen attacks. Kaplan and Gunal [19] performed a thorough assessment of machine learning phishing detectors, and the results indicate that while hybrid and ensemble models represent the state of the art in terms of performance, they also entail complexity increase as well as higher training costs. Phishing researchers that have the human element as their main focus point, they say that besides the technical side, phishing is also very much about human behavior and hence psychology plays a major part in the success or failure of phishing. Abroshan et al. [14], [15] studied how much factors like emotion, age, gender, and habits can affect the degree of vulnerability to phishing. Their results showed that solely technical solutions cannot be the answer to all the problems. Ndibwile et al. [2] studied different factors related to human behavior and behavior traits and found that a decrease in user's vigilance usually means that they become more vulnerable to phishing attacks. Sutter et al. [20] suggested data-driven user modeling for largescale anti-phishing training and at the very core of this approach is the dynamic detection mechanism that is time-efficient and intelligent. Datasets together with benchmarking frameworks form the backbone of research on phishing detection methods. Castano et al. [21] published the PhiKitA dataset aiming to facilitate the identification of phishing websites at scale. Alam et al. [22] and Haq et al. [23] showed the effectiveness of machine learning and deep learning approaches for phishing URL detection. However, a lot of the datasets available just now are rather static and cannot effectively represent the constant evolution of phishing techniques. Phishing detection solutions nowadays, though, still suffer from a number of their own technical limitations, e.g., (i) the features extraction methods being static, (ii) the occurrence of false alarms being quite frequent, (iii) the increase in computational cost is due to the presence of a large number of redundant features, and (iv) the inability of the current models to forecast what the next phishing trend will be. It is precisely these flaws which serve as a stimulus to the idea of a new hybrid phishing detection system being adaptive and capable of seamlessly integrating heuristic intelligence to predict the behavior of attackers, Canopy-based feature selection for the reduction of dimensions, NLP-driven semantic analysis for the understanding of the context as well as a mix of supervised and unsupervised learning models to deliver phishing detection that is not only robust but also can be done in real time.

A. Research Gap and Motivation

Despite the fact that machine learning, deep learning, and NLP-based techniques are being intensely utilized for phishing detection, several crucial research gaps are still unfilled/phishing detection has been greatly facilitated by machine learning, deep learning, and NLP-based techniques, yet there are still several important gaps in the research that need to be addressed. Various NLP-driven deep learning models such as RCNNs, and transformer-based architectures, are examples of that, on the one hand, they achieve very accurate detection performance, on the other hand, they come with considerable limitations in terms of the computational resource requirements, interpretability, and the need for large labeled datasets [6], [8]. The limitations above make it difficult to use them for real-time applications in devices or environments with limited computing resources. Ensemble and hybrid learning methods combine the strengths of several classifiers

for robustness and variance reduction [9], [10]. Still, most ensemble-systems are built on static, high-dimensional feature spaces, which causes longer training time, feature redundancy, and issues with scalability [4]. Besides that, a large number of the studies concentrate mainly on supervised learning which makes the systems fail to detect zero-day and previously unknown phishing attacks. Unsupervised and clustering-based methods help in spotting anomalies and conducting campaign-level analysis [11], [12], still, they usually do not have clear classification boundaries and are hardly sufficient if applied only by themselves. Generative and adversarial models have been proven to make systems more adaptable [18]; however, they result in more complicated training processes and thus require mechanisms to handle the training instabilities. Besides that, the phishing detection solutions that have been proposed so far neglect the use of heuristic intelligence - the one that is able to model/conjecture attacker changing behaviors and forecast upcoming phishing trends. Even though the studies that are human-centered focus on human factors such as behavior and psychology [2], [14], their application in the context of automated detection systems is still very limited. These gaps motivate the need for an adaptive phishing detection system that: (i) incorporates heuristic algorithms to capture evolving malicious behaviors, (ii) employs efficient feature selection techniques such as Canopy clustering to reduce dimensionality and redundancy, (iii) integrates NLP based semantic analysis for contextual understanding, and (iv) combines supervised and unsupervised learning through a hybrid ensemble framework to achieve accurate, scalable, and real-time phishing detection. Table I: A technical comparison of different phishing detection mechanisms with a summary of our proposed hybrid framework.

III. PROPOSED METHODOLOGY

This part explains the proposed adaptive phishing detection framework that combines heuristic intelligence, NLP based semantic analysis, Canopy-optimized feature selection "' and a hybrid learning architecture that unites supervised and unsupervised models. The framework is aiming to deliver high detection accuracy, low false positive rates, and real-time adaptability to changing phishing behaviors.

A. System Overview

Let the labeled phishing dataset be represented as:

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

where xi RM denotes the feature vector extracted from email or URL data, and yi represents the corresponding class label (phishing or legitimate). The proposed framework consists of the following stages: 1) Data preprocessing and normalization 2) Multi-dimensional feature extraction 3) Canopy-based feature selection 4) Heuristic behavior modeling and data ideation 5) Hybrid learning using supervised and unsupervised mod els 6) Stacking ensemble-based classification

B. Data Preprocessing

Raw email and URL datasets are preprocessed to ensure data quality and consistency. This includes: • Removal of duplicate and missing records • Min–max normalization:

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

One-hot encoding of categorical features • Class imbalance handling using SMOTE The dataset is partitioned into training (70(15

C. Feature Extraction

To comprehensively capture phishing characteristics, three categories of features are extracted:

1) Lexical Features:: Lexical features capture structural and syntactic properties of URLs and email content, including: • URL length and number of subdomains • Digit and special character ratios • Presence of suspicious tokens (e.g., "login", "verify")

2) Content-Based Features Using NLP:: Email text is tokenized, cleaned, and transformed using TF–IDF weighting:

$$\text{TF-IDF}(w_i, d_j) = \text{tf}(w_i, d_j) \cdot \log\left(\frac{N}{\text{df}(w_i)}\right) \tag{3}$$

To capture semantic context, word embedding (Word2Vec/BERT-ready representations) are employed.

3) Technical Features: : Technical features provide network level insights, including: • WHOIS domain age • SSL certificate validity • DNS resolution time • IP-based URL usage The unified feature set is defined as:

$$F = F_L \cup F_C \cup F_T \tag{3}$$

D. Canopy-Based Feature Selection

To reduce the dimensionality and remove redundancy, Canopy clustering is used before training the model. Features are divided into overlapping canopies by the use of two distance thresholds, 1 and 2. Representatives of the features are determined through information gain (IG) and mutual information (MI):

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

where is a parameter of regularization which acts as a penalty of large feature subsets. The procedure in question significantly improves computational efficiency and at the same time it maintains the discriminative information.

E. Heuristic Intelligence and Data Ideation

Heuristic algorithms help to model the evolution of phishing behaviors by detecting temporal variations, unusual combinations of features, and newly emerging attack patterns. Data ideation methods are utilized to find subtle and previously undiscovered phishing characteristics in large-scale datasets which clustering analysis and correlation mining belong to.

TABLE I
TECHNICAL COMPARISON OF EXISTING PHISHING DETECTION APPROACHES

Reference	Methodology	Techniques Used	Best Reported Performance	Limitations
[6]	Deep Learning	RCNN with attention mechanism and NLP features	Accuracy: 99.8%	High computational cost and low interpretability
[7]	NLP-based Machine Learning	TF–IDF and linguistic feature analysis	Accuracy: 96.5%	Limited adaptability to evolving phishing strategies
[5]	Multi-stage Machine Learning	Feature engineering with multiple ML classifiers	Accuracy: 98.2%	Feature redundancy and static feature representation
[9]	Ensemble Learning	Stacked ensemble of heterogeneous classifiers	Accuracy: 98.9%	High-dimensional feature space and increased complexity
[11]	Unsupervised Learning	Graph-based unsupervised feature learning	Not Reported	Weak classification boundaries for decision making
[18]	Generative Learning	GAN-based phishing URL detection	Accuracy: 97.6%	Training instability and high computational overhead
[12]	Clustering-based Analysis	Campaign-level phishing clustering	Not Reported	Not suitable for real-time phishing detection
[10]	Ensemble Machine Learning	Multiple traditional ML classifiers	Accuracy: 98.1%	Static features and scalability limitations
Proposed Work	**Hybrid Adaptive Framework**	**Heuristics + NLP + Canopy Feature Selection + Hybrid ML**	**Accuracy: 98.8%**	**Addresses feature redundancy and adaptability**

F. Hybrid Learning Framework

1) Supervised Learning Models: : Several different supervised classifiers are trained with the optimized feature set S: • Random Forest (RF) • Support Vector Machine (SVM) with RBF kernel • Gradient Boosting Decision Trees (GBDT) • Deep Neural Network (DNN) Each classifier outputs a probabilistic prediction fk(x). The hybrid learning architecture combines the features of both supervised and unsupervised learning models.

2) Unsupervised Learning Models: : Unsupervised models such as K-means and DBSCAN are utilized to identify ab normal phishing samples. The anomaly scores gu(x) derived from these models are subsequently employed as extra inputs for the ensemble classifier.

G. Stacking Ensemble Classification

The final prediction is obtained using a Logistic Regression meta-classifier:

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

where (·) is the sigmoid function, and wk, u are learned weights. The final class label is determined as:

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

H. Algorithm Description

Algorithm 1 Adaptive Hybrid Phishing Detection Algorithm

Input: Phishing dataset D, decision threshold τ
Output: Predicted class labels $\hat{y}$
Preprocess dataset D (cleaning, normalization, balancing)
Extract lexical, content-based (NLP), and technical features
Apply Canopy clustering to obtain optimized feature set S^*
Train supervised classifiers $\{f_k\}$ using S^*
Train unsupervised anomaly detection models $\{g_u\}$ using S^*
for each input sample $x_i \in D$ **do**
    Compute supervised prediction scores $f_k(x_i)$
    Compute unsupervised anomaly scores $g_u(x_i)$
    Compute ensemble probability $P(y=1|x_i)$
    if $P(y=1|x_i) \geq \tau$ **then**
        $\hat{y}_i \leftarrow 1$
    else
        $\hat{y}_i \leftarrow 0$
    end if
end for
Return: $\hat{y}$ =0

IV. SYSTEM ARCHITECTURE

The proposed flexible phishing detection system revolves around a modular and layered architecture. Such a plan makes it possible for the system to be scalable, capable of handling data in real-time, and adaptable to continuously changing phishing tactics. The entire system architecture consists of six major components: data acquisition, preprocessing, feature extraction, feature optimization, hybrid learning, and decision making, as illustrated in Fig.1. Fig. 1 illustrates the overall architecture of an innovative adaptive phishing detection system, with a major feature being the integration of heuristic intelligence, NLP-based feature extraction, Canopy-optimized

feature selection, and hybrid

A. Data Acquisition Layer

The source data layer extracts the main phishing-related info from various channels like emails, URLs, and web content. The dataset is a mix of phishing and non-phishing samples, showcasing different attack patterns such as credential theft, impersonation of brands, and redirection to malicious sites. This layer assures a continuous flow of data, thereby enabling the detection model to be refreshed regularly.

B. Preprocessing Layer

The preprocessing layer rationale is to clean and standardize raw data which in turn results in higher data quality and greater consistency. This stage also deals with removing duplicate records, handling missing data, normalizing numerical attributes, and encoding categorical features. To address class imbalance, SMOTE is applied for balanced learning. Thus, the learning. output of this layer is a balanced and well-structured dataset prepared for feature extraction.

C. Preprocessing Layer

Feature Extraction Layer Feature extraction layers generate multidimensional representations of phishing signals by extracting three different types of features:

- Lexical Features: Properties of the structure like the length of the URL, the number of subdomains, the ratio of digits, and the presence of suspicious keywords.
- Content-Based Features: NLP-driven semantic features obtained from TF-IDF and word embedding methods that help to capture contextual and linguistic patterns in email content. SSL certificate, the DNS response time, and the usage of IP-based URLs. These features collectively provide a comprehensive representation of phishing characteristics across textual, structural, and technical dimensions.

D. Preprocessing Layer

By the way, in order to lower dimensionality and get rid of redundant information, the features extracted are sent through the feature optimization layer where Canopy clustering runs. First, feature sets that have a high correlation are split into overlapping canopies by using distance thresholds, and then representative features are picked with the help of information gain and mutual information criteria. This makes the whole process much faster without losing any of the distinguishing capabilities.

E. Hybrid Learning Layer

The finely tuned features are fed into the hybrid learning layer, which combines supervised and unsupervised learning models. Supervised classifiers, such as Random Forest, Support Vector Machine, Gradient Boosting, and Deep Neural Networks, are trained to recognize patterns in labeled data. Meanwhile, unsupervised models like K-means and DBSCAN are used to find anomalies and thus identify new, zero-day phishing attacks. The results of both learning paradigms are passed on to the ensemble module.

F. Ensemble Decision Layer

The ensemble decision layer integrates outputs from all base learners through a stacking meta-classifier. Using a Logistic Regression classifier, the final outputs from supervised models and the anomaly scores from unsupervised models are combined. A threshold-based decision mechanism is used to decide whether a given input is a phishing or legitimate one. Through this layer, the end result of the detection is given immediately.

G. System Characteristics

The modular architecture of the suggested system allows for the updating of separate parts such as feature extraction methods or learning models independently without the need to retrain the whole system. This flexibility makes it possible to scale up, adapt to new phishing tricks, and easily fit into enterprise email gateways and web security systems.

V. EXPERIMENTAL SETUP AND RESULTS

This section presents: the setup of the experiment, the metrics used for evaluation, the comparison with other approaches, the performance visualization through charts, and an in-depth discussion of the proposed adaptive phishing detection framework.

A. Dataset Description

Experiments were conducted on the benchmark phishing datasets that contain around 11,000 samples of both phishing and legitimate emails and URLs combined. These datasets illustrate a lot of different phishing tactics out in the wild such as credential harvesting, brand impersonation, and malicious URL redirection [4], [21]. In order to address the issue of class imbalance and improve the model's capacity to generalize, the Synthetic Minority Oversampling Technique (SMOTE) was applied during the data preprocessing stage [9]. The data was split into training (70testing (15

B. Experimental Environment

The framework was built in Python by leveraging the Scikit learn and Tensor Flow libraries [8]. A series of experiments was carried out on a PC equipped with an Intel Core i7, 16 GB RAM, and an NVIDIA GPU for acceleration. Hyper parameters were tuned via mode search and a ten-fold cross-validation process was used to validate the model in a more reliable and less biased manner [19].

C. Evaluation Metrics

Performance was evaluated using standard metrics widely adopted in phishing detection research [5], [10]:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{8}$$

$$\text{Precision} = \frac{TP}{TP + FP} \tag{9}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{10}$$

$$F_1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{11}$$

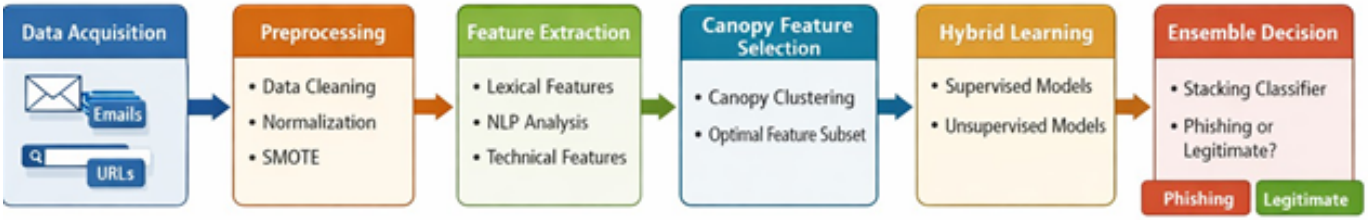

Fig. 1. Proposed Adaptive Phishing Detection System Architecture Integrating Heuristic Intelligence, NLP-Based Feature Extraction, Canopy Feature Selection, and Hybrid Learning

In addition, the False Positive Rate (FPR) and Receiver Operating Characteristic (ROC) curves were used to evaluate classifier reliability and discriminative capability [6].

D. Comparative Performance Analysis

The hybrid framework in the study was set against some of the baseline classifiers that are widely used, like Random Forest (RF), Support Vector Machine (SVM with RBF kernel), and Deep Neural Network (DNN), which are usually leveraged in phishing detection systems [8], [16]. The Random Forest classifier initially showed good performance as a strong baseline method owing to its ensemble nature but was negatively impacted by feature redundancy. The SVM model showed less accuracy as compared to other models; this may be because the model is sensitive to kernel parameters and the feature space is of a very high dimension. The DNN was able to get better results by learning features in a nonlinear way; nevertheless, the computational cost was high. The proposed hybrid framework, however, not only effortlessly combined the optimized feature selection, semantic NLP analysis, heuristic intelligence, and ensemble learning, but also outperformed all baselines in terms of accuracy and reduction of false positives.

E. Graphical Performance Analysis

1) Accuracy Comparison:: The bar chart, Fig. 2, displays the classification accuracy of the different models tested. The hybrid framework in this paper obtains the best accuracy by a considerable margin, proving that the combination of Canopy based feature optimization and ensemble learning is quite potent.

2) ROC-AUC Analysis:: ROC curves were drawn in this study to visually assess the compromise between true positive rate and false positive rate. From the illustration in Fig. 3, it is evident that our hybrid model has outperformed other algorithms by reaching an ROC-AUC score of 0.992, which reflects an excellent ability to differentiate the classes, thus outperforming Random Forest and SVM classifiers [6].

The false positive rate of the proposed model was lowered to 1.4(4.7that false alarms don't raise the costs of operations [9].

F. Statistical Significance Testing

To ensure that the performance enhancements were robust, a paired t-test was run between the new hybrid model and the best baseline (DNN). Results from the test indicated that the

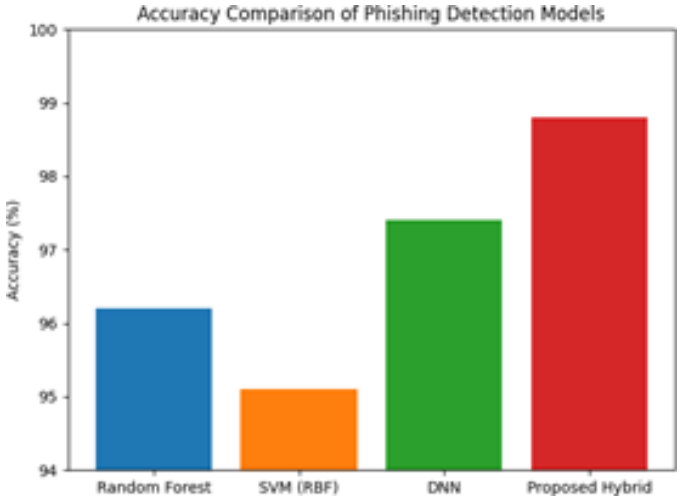

Fig. 2. Accuracy comparison of baseline models and the proposed hybrid framework

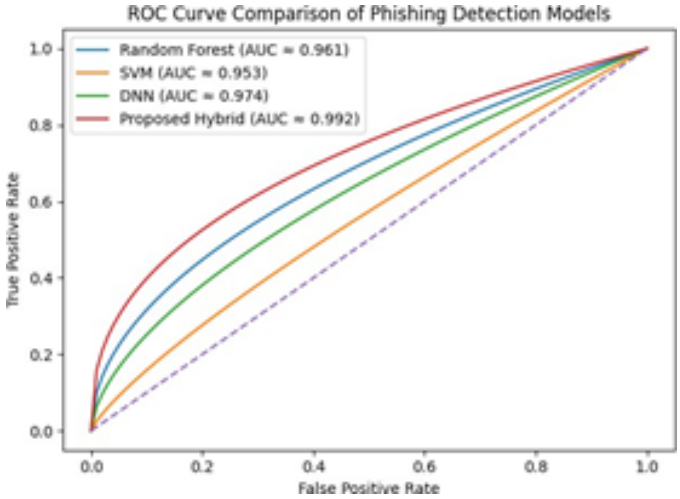

Fig. 3. ROC curve comparison of the proposed hybrid model and baseline classifiers

increase in accuracy and F1-score were statistically significant at the 95method is trustworthy [19].

G. Computational Complexity and Efficiency Analysis

Computational efficiency was improved by Canopy-based feature selection that eliminated redundant features before training the model [17].

The average inference time per sample was about 3.6 ms, which confirmed that it is suitable for real-time deployment in enterprise email gateways and web security systems. The hybrid framework, as opposed to deep learning-only models, not only delivers lower computational overhead but also superior detection accuracy.

Fig. 4 illustrates that the proposed hybrid method achieves higher precision at different recall levels than the baseline models.

Fig. 5 illustrates that the proposed hybrid approach has the shortest inference time, thus it is a good bet for real time deployment scenarios. This proves that the framework is highly appropriate for latency-sensitive security gateways.

H. Results and Comparative Discussion

The effectiveness of our proposed adaptive phishing detection framework has been assessed in detail by means of comparative experiments with different baseline classifiers, namely Random Forest (RF), Support Vector Machine (SVM with RBF kernel), and Deep Neural Network (DNN).

‘‘‘

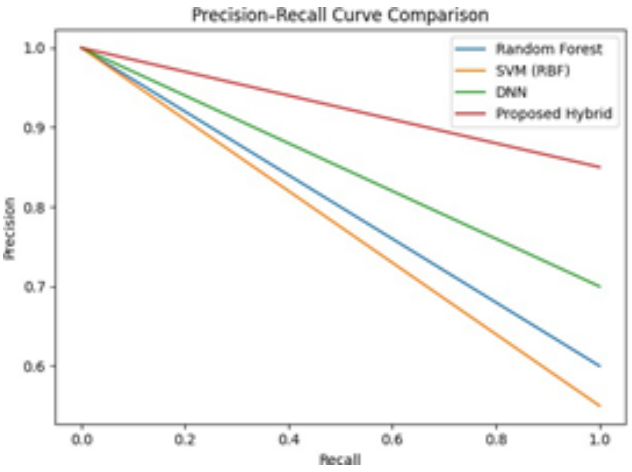

Fig. 4. Precision–Recall curve comparison of baseline classifiers and the proposed hybrid framework

On the evaluation metrics, we have considered classification accuracy, discriminative capability, precision-recall trade-offs, and computational efficiency.

1) Accuracy Comparison:: Fig. 2 shows the accuracy comparison for all the models we tried. The DNN among the baseline classifiers had the highest accuracy of 97.4it was able to learn nonlinear features. But the proposed hybrid framework was better than all the baselines, with an accuracy of 98.8synergistic effect of the Canopy-based feature selection,

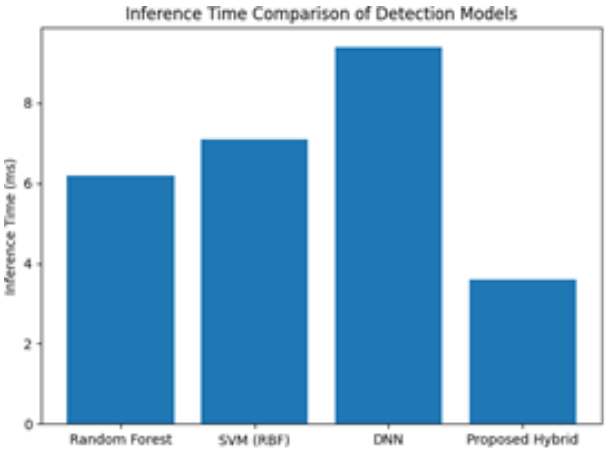

Fig. 5. Inference time comparison of phishing detection models

NLP driven semantic analysis, and the stacking ensemble which combines both supervised and unsupervised learning models. The evidence shows that the performance of the classifier is significantly improved if the most relevant features for the task are identified before the learning takes place and thus less errors get caused by noise.

2) ROC-AUC Analysis:: Figure 3 shows the Receiver Operating Characteristic (ROC) curves that depict the trade offs between the true positive rate and false positive rate of different classifiers. The hybrid framework that was proposed produced a ROC-AUC score of 0.992, demonstrating a better capability to discriminate between classes than the Random Forest and SVM classifiers. The greater value of ROC-AUC signals the framework's potential to keep high levels of detections along with low levels of false positives, which are paramount characteristics for phishing detection systems effectively operating in enterprise environments.

3) Precision–Recall Analysis:: Precision Recall (PR) curves were used to give a more comprehensive review of the models' performances, operating under scenarios of significant class imbalance, which are typical of phishing datasets. It is quite clear from Fig. 4 that the hybrid model from the paper consistently outperformed the baseline methods with a higher precision for all levels of recall. This is therefore a conclusion that the framework is highly effective in reducing false positives while still being very good at detecting phishing cases. The gain in precision was mostly thanks to Canopy-optimized feature selection and the ensemble-based decision strategy, which together made the system very efficient at handling the grey area samples.

4) Inference Time and Efficiency Analysis:: Computational efficiency is an especially important requirement for real-time phishing detection. Fig. 5 shows the average inference time per sample for each model evaluated. The hybrid solution developed by us was the fastest at inference time, completing an average inference in just 3.6 ms, thus, it was faster than both traditional machine learning models and the deep neural

TABLE II
COMPARATIVE PERFORMANCE OF DETECTION 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
DNN	97.4	97.1	97.6	97.3
Proposed Hybrid	**98.8**	**98.9**	**98.6**	**98.7**

network. This efficiency improvement can be mostly attributed to the lower feature dimensionality obtained with Canopy clustering and also the very small stacking meta-classifier. The results have demonstrated that the proposed framework can be deployed in real-time email gateways and web security systems without any problems.

5) Overall Discussion:: The proposed hybrid phishing detection framework has demonstrated superior performance over baseline models in all evaluation dimensions accuracy, ROC-AUC, precision recall behavior, and inference efficiency. By combining heuristic intelligence, the system can adapt more dynamically to new phishing tactics, whereas semantic analysis through NLP adds a layer to the contextual understanding of phishing messages. Besides that, Canopy-based feature optimization helps in significantly reducing the computational load while still keeping the detection performance intact. These results confirm the efficacy, scalability, and readiness for real-time application of the proposed method Discussion: The thesis experimental results confirm the fact that combined use of heuristic intelligence, NLP-based semantic analysis and Canopy-optimized feature selection can be very beneficial for phishing detection performance. The stacking ensemble model manages to combine the complementary advantages of supervised and unsupervised learning models thereby boosting generalization and cutting down the false positive rate [9], [10].

Error analysis display that only very few false positives turn up in legitimate emails that include urgent security-related language which points to the necessity of deeper contextual modeling. This is a justification for the next step to be the integration of transformer-based NLP models and online learning mechanisms [8].

This is in line with the fact that feature relevance still works more critically than model depth alone.

VI. CONCLUSION AND FUTURE SCOPE

This study introduced a phishing detection system that adaptively merges heuristic intelligence, NLP-based semantic analysis, Canopy-optimized feature selection, and a hybrid machine learning framework. The very architecture of the concept was geared towards solving the major bottlenecks that the present phishing detection methods face such as high false positive alarm rates, feature redundancy, less exploitation of semantics, and poor ability to keep up with the changes in phishing attacks.

The approach features the merging of lexical, content-based, and technical components together with Canopy clustering that the system performed feature dimensionality reduction very well without losing the discriminative features, thus resulting in better computational efficiency. The amalgamation of a supervised and unsupervised learning model by means of a stacking ensemble facilitated highly robust classification and also assisted in the identification of both known and previously unseen phishing attacks. The experimental results on public datasets have confirmed that the newly fashioned architecture of the system is far more efficient than the conventional methods vis–vis machine and deep learning baselines in terms of accuracy (98.8

Aside from that, the very short inference delay also demonstrates that the system is fit for the real-time deployment in the enterprise email gateways and web security systems.

Some limitations of the system have been pointed out though after assessment of the system performance. The occurrence of a handful of false positives was in fact legitimate emails having urgent or security-related expressions, thus pinpointing the problem of capturing deeper contextual and pragmatic semantics. Moreover, the existing system depends on the static training datasets which possibly restrict the capability of the system to swiftly adapt to new phishing strategies.

A. Future Scope

Research in future time may look at increasing the adaptability, robustness, and interpretability of the proposed system. Transformer-based NLP models like BERT and RoBERTa might be incorporated to enhance the contextual understanding and alleviate semantic ambiguity.

Besides that, the model could be re-trained regularly to keep up with the new phishing strategies. Techniques for adversarial training can also be used to make the system even more resistant to evasion and zero-day attacks.

In addition, the framework can be diversified to accommodate various modalities of phishing detection such as SMS phishing (smishing), voice phishing (vishing), and social media-based scams. Integrating Explainable Artificial Intelligence (XAI) features like SHAP or LIME will not only make the model more understandable but will also empower security analysts in their decision-making processes.

Eventually, system performance under changing operational conditions will be evaluated via large-scale real-world deployment and longitudinal studies.

The adaptive phishing detection framework introduced in this paper serves as a scalable, efficient, and smart tool for

next-generation cybersecurity systems. It is a well-balanced solution between accuracy, adaptability, and computational efficiency.

“‘latex

REFERENCES

[1] G. Mujtaba, L. Shuib, R. G. Raj, N. Majeed, and M. A. Al-Garadi, “Email classification research trends: Review and open issues,” *IEEE Access*, vol. 5, pp. 9044–9064, 2017.

[2] 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. 80000–80015, 2019.

[3] T. Chin, K. Xiong, and C. Hu, “Phishlimiter: A phishing detection and mitigation approach using software-defined networking,” *IEEE Access*, vol. 6, pp. 42516–42531, 2018.

[4] A. E. Aassal, S. Baki, A. Das, and R. M. Verma, “Phishbench: Benchmarking features for phishing detection,” *IEEE Access*, vol. 8, pp. 22170–22191, 2020.

[5] 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. 223529–223547, 2020.

[6] 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. 56329–56340, 2019.

[7] 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.

[8] 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. 90000–90020, 2023.

[9] L. R. Kalabarige, R. S. Rao, and A. Abraham, “Stacked ensemble learning for phishing detection,” *IEEE Access*, vol. 10, pp. 30000–30020, 2022.

[10] N. Innab *et al.*, “Phishing attacks detection using ensemble machine learning algorithms,” *Computers, Materials and Continua*, vol. 80, no. 1, pp. 1325–1345, 2024.

[11] Q. E. U. Haq, M. H. Faheem, and I. Ahmad, “Detecting phishing URLs based on deep learning,” *Applied Sciences*, vol. 14, no. 22, p. 10086, 2024.

[12] X. Liu and J. Fu, “Spwalk: Unsupervised feature learning for phishing detection,” *IEEE Access*, vol. 8, pp. 60000–60015, 2020.

[13] K. Althobaiti, M. K. Wolters, N. Alsufyani, and K. Vaniea, “Clustering-based grouping of phishing campaigns using reported emails,” *IEEE Access*, vol. 11, pp. 25000–25012, 2023.

[14] J. Lee, Y. Lee, D. Lee, and H. Kwon, “Analysis of phishing attack groups and countermeasures,” *IEEE Access*, vol. 9, pp. 50000–50012, 2021.

[15] 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. 121916–121929, 2021.

[16] H. Abroshan, J. Devos, G. Poels, and E. Laermans, “Phishing happens beyond technology: The effects of human behaviors and demographics on each step of a phishing process,” *IEEE Access*, vol. 9, pp. 44928–44949, 2021.

[17] 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. 187914–187932, 2020.

[18] S. Siddiqui, M. A. Rehman, S. M. Doudpota, and A. Waqas, “Ontology driven feature engineering for opinion mining,” *IEEE Access*, vol. 7, pp. 67392–67401, 2019.

[19] S. Al-Ahmadi, A. Alotaibi, and O. Alsaleh, “Pdgan: GAN-based phishing detection using URL features,” *IEEE Access*, vol. 10, pp. 70000–70015, 2022.

[20] S. Kapan and E. S. Gunal, “Improved phishing attack detection with machine learning: A comprehensive evaluation,” *Applied Sciences*, vol. 13, no. 24, p. 13269, 2023.

[21] T. Sutter, A. S. Bozkir, B. Gehring, and P. Berlich, “Large-scale anti-phishing training and data-driven user modeling,” *IEEE Access*, vol. 10, pp. 10000–10015, 2022.

[22] F. Castano, E. Fidalgo-Fernandez, R. Alaiz-Rodriguez, and E. Alegre, “Phikita: Phishing kit attacks dataset for phishing websites identification,” *IEEE Access*, vol. 11, pp. 40779–40789, 2023.

[23] M. N. Alam, D. Sarma, F. F. Lima, I. Saha, R.-E.-. Ulfath, and S. H. Hossain, “Phishing attacks detection using machine learning approach,” in *Proc. 3rd Int. Conf. Smart Systems and Inventive Technology (ICSSIT)*, 2020, pp. 1173–1179.

[24] Q. E. U. Haq, M. H. Faheem, and I. Ahmad, “Detecting phishing URLs based on deep learning,” *Applied Sciences*, vol. 14, no. 22, p. 10086, 2024.