

# Phishy?

## Detecting Phishing Emails Using ML and NLP

Md Fazle Rabbi      Arifa I. Champa      Minhaz F. Zibran  
Department of Computer Science, Idaho State University, USA  
{mdfazlerabbi, arifaislamchampa, minhazzibran}@isu.edu

**Abstract**—Phishing emails, a type of cyberattack using fake emails, are difficult to recognize due to sophisticated techniques employed by attackers. In this paper, we use a natural language processing (NLP) and machine learning (ML) based approach for detecting phishing emails. We compare the efficacy of six different ML algorithms for the purpose. An empirical evaluation on two public datasets demonstrates that our approach detects phishing emails with high accuracy, precision, and recall. The findings from this work are useful in devising more efficient techniques for recognizing and preventing phishing attacks.

**Index Terms**—phishing, spam, email, classification, detection, machine learning, natural language processing

### I. INTRODUCTION

Phishing is a common approach that cybercriminals use to steal sensitive information from individuals and organizations. Messaging apps, phone calls, social media, and emails are just a few of the many platforms that phishing attacks target. The outbreak of the coronavirus epidemic resulted in an increase in phishing attacks that was 220% higher than the yearly average [1]. According to the Anti-Phishing Working Group, there were a record-breaking total of 1,270,883 unique phishing incidents reported in the third quarter of 2022, demonstrating the prevalence and persistence of this threat [2].

Email has long been arguably the most used platform for electronic communication in both formal and informal settings. Therefore, email platforms have been the most frequent targets for phishing attacks where cybercriminals create scam emails that appear genuine but contain traps for the end-user for becoming a victim of information leaks or other malicious attacks.

According to Forrester Research, 91% of all hacking attacks begin with a phishing email [3]. In 2022, 96% of organizations were targeted by email-related phishing attempts [4]. Phishing attacks are becoming increasingly sophisticated and can deceive even experienced users, as illustrated by the case of the John Podesta phishing attack [5]. This incident, involving a fake email that tricked the former chairman of the Hillary Clinton presidential campaign, demonstrated that anyone is susceptible to an email phishing attack.

Phishing often takes advantage of the naivety of end-users, and combating it requires human intervention in the form of awareness and training campaigns [6]–[8]. Despite the ongoing efforts in employee training and social awareness of email phishing attacks, these attacks continue to increase, resulting in an increasing number of email phishing victims. This indicates

that we need more sophisticated approaches for the automatic detection of phishing emails before those emails are presented to the end-users, who may or may not be technically well-educated.

To develop an effective approach for detecting phishing emails, it is crucial to understand the characteristics of phishing emails that distinguish them from legitimate ones. To accomplish this, we examine the characteristics and traits of phishing emails to identify patterns that can be used to develop an effective detection model. In particular, we address the following research questions:

**RQ1:** *What is the most effective machine learning (ML) algorithm for detecting phishing emails?*

— Since different ML algorithms have their strengths and weaknesses in different application contexts, it is critical to identify the most effective algorithm for detecting phishing emails. We measure effectiveness in terms of accuracy, precision, recall, and F-score.

**RQ2:** *What impact does ML algorithm training time have on the performance of phishing email detection?*

— Lengthy training times may not be feasible in some operational settings, and if longer training times do not produce significant improvements in the performance of the model, the additional time and resources invested would just be wasted. Hence, a deep understanding of the relationship between training time and performance can assist in the development of more efficient phishing detection systems for adoption in practical settings.

**RQ3:** *To what extent does the ‘subject’ feature of an email contribute to the accurate detection of phishing emails?*

— One of an email’s most noticeable and prominent features is the ‘subject’, and the content of that line has a significant impact on how the recipient perceives the email’s authenticity. Investigating the ‘subject’ feature’s contribution to accurately detecting phishing emails can provide insights into the creation and identification of phishing emails, which will aid in the development of more effective phishing detection systems.

To address the aforementioned research questions, we use two phishing email datasets, which we process using natural language processing (NLP) and then we compare six machine learning algorithms by separately operating them on the two datasets.

The rest of the paper is structured as follows. Section II describes the data sets used in this work. Section III describes

the methodology of this study. The findings from this study are presented and discussed in Section IV. In Section V, we describe the limitations of our work. Previous work in the literature that are related to ours are discussed in Section VI. Finally, Section VII concludes the paper.

## II. DATASETS

For our work, we use two publicly available datasets. Table I presents the number of phishing and legitimate emails in each of the datasets, which are further described below.

TABLE I  
DATASETS USED IN OUR WORK

Dataset	Ling [9]	TREC [10]
<b>Email</b>		
<b>Legitimate</b>	2,412	25,220
<b>Phishing</b>	481	50,199
<b>Total</b>	2,893	75,419

### A. Ling Dataset

The Ling Spam dataset [9] is a collection of emails specifically focused on topics of interest to linguists. The dataset is organized around two main attributes: ‘subject’ and ‘message,’ as well as a ‘label’ indicating whether each email is spam or legitimate. The subject and body of an email are captured in ‘subject’ and ‘message’ respectively.

### B. TREC Dataset

The TREC Public Corpus [10] is a collection of email messages collected between April 8 and July 6, 2007. For our study, a preprocessed TREC Public Corpus Dataset is collected from Kaggle [11]. This dataset includes four attributes: ‘subject,’ ‘email\_to,’ ‘email\_from,’ and ‘message,’ as well as a ‘label’ indicating whether the email is phishing or legitimate. The ‘email\_to’ and ‘email\_from’ attributes respectively capture the email addresses of the recipients and sender of the email.

## III. METHODOLOGY

We use a three-phase procedure for recognizing phishing emails. First, we sanitize the dataset through a preprocessing phase. Second, subsets of the preprocessed data are then separately fed to six ML algorithms [12] to create ML models. Finally, the ML models are operated on separate subsets of the preprocessed datasets. The models’ performances are measured using a set of metrics described in Section III-C.

To avoid the overhead of installing and configuring different tools and libraries on a local machine for operating the ML algorithms, we leverage a cloud-based platform, Google Colab [13], for the purpose. However, we use the processing power of our local machine rather than Colab’s GPU or TPU. Our local machine, which runs a Windows 10 operating system, is equipped with an Intel Core i7-8650U processor with a base clock speed of 1.90 GHz and turbo boost up to 4.2 GHz, and 16GB of DDR4 memory.

### A. Data Preprocessing

In the preprocessing phase, we handle missing attribute values, eliminate duplicates/redundancy, clean data, and convert them to a format for feeding to the ML algorithms.

1) *Handling Missing Attribute Values:* In the email datasets, we notice that some emails do not include a ‘subject’ but have a message body while some emails include a ‘subject’ with an empty message body. In cases where the ‘subject’ attribute is missing for an email, we exclude the attribute from consideration and only use the ‘message’ body for that particular email. If the ‘message’ attribute is missing for a particular email, we exclude that entire email from our study. Table II presents the number of instances we encountered the missing attribute values in the two datasets used in our work.

TABLE II  
MISSING ATTRIBUTE VALUES

Dataset	Ling [9]	TREC [10]
<b>Attribute</b>		
subject	62	793
message	0	1487

2) *Removing Duplicates:* In the TREC dataset, the ‘message’ attribute contains 14,885 instances of duplicate data, and when we combine the ‘subject’ and ‘message’ attributes, we identify 11,107 instances of duplicate emails. In the Ling dataset, the ‘message’ feature contains 34 instances of duplicate data, and when we merge the ‘subject’ and ‘message’ features, we discover 17 instances of duplicate emails. We eliminate the duplicates keeping only one instance from each duplicate set.

3) *Data Cleansing:* To prepare data for ML algorithms, we perform further clean then as follows:

- We convert all text to lowercase, to ensure consistency.
- We remove any leading and trailing white spaces from each email body.
- We replace actual email addresses, URLs, currency symbols, and contact numbers with the placeholders ‘MAILID’, ‘LINKS’, ‘MONEY’, and ‘contact number’, respectively.
- Any non-alphanumeric character found in the email’s subject or body/message is replaced with a white space.
- We remove stop words (e.g., ‘the’, ‘is’) from both the subject and body of each email to reduce noise.

4) *Vectorization:* To convert email information into numerical vectors, we use the term frequency-inverse document frequency (TF-IDF). This well-known NLP technique multiplies a word’s term frequency (TF) by its inverse document frequency (IDF). The TF refers to how many times a word appears in a document, whereas the IDF refers to how frequently a word appears across the entire corpus of documents. Table III presents the total number of phishing and legitimate email instances we obtain in each of the datasets after the completion of the preprocessing step.

TABLE III  
DATASETS AFTER PREPROCESSING

	Ling	TREC
Legitimate	2408	24673
Phishing	468	38152
Total	2876	62825

### B. Classification Algorithms

For the detection of phishing emails, we separately attempt with six ML algorithms [12], as listed below.

- 1) Logistic Regression (LR)
- 2) K-Nearest Neighbors (KNN)
- 3) AdaBoost (AB)
- 4) Multinomial Naive Bayes (MNB)
- 5) Gradient Boosting (GB)
- 6) Random Forest (RF)

Details of all these well-known ML algorithms can be found elsewhere [12]. The attributes in the datasets, as described in Section II, are used as features while operating the ML algorithms on the datasets.

TABLE IV  
CONFUSION MATRIX

Actual \ Predict	Legitimate	Phishing
Legitimate	TN	FP
Phishing	FN	TP

### C. Evaluation Metrics

For each ML algorithm operated on each dataset, we record the number of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). We then measure the performance of each algorithm in terms of accuracy, recall, precision, and F-score, which are characterized according to the confusion matrix shown in Table IV. We then measure the performance of each algorithm in terms of accuracy, recall, precision, and F-score as defined as follows.

- **Accuracy** is measured by the percentage of correctly categorized instances in the dataset. Mathematically

$$\text{accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

- **Precision** measures the number of actual positive cases among all those that are predicted to be positive. Mathematically,

$$\text{precision} = \frac{TP}{TP + FP} \quad (2)$$

- **Recall** calculates the proportion of true positive cases among all positive examples. Mathematically defined as,

$$\text{recall} = \frac{TP}{TP + FN} \quad (3)$$

- **F-score** is the harmonic mean of precision and recall. Mathematically defined as follows:

$$\text{F-score} = 2 \times \frac{p \times r}{p + r} \quad (4)$$

True positive rate (TPR), also known as sensitivity, indicates how accurately phishing email detection systems identify true positive instances. Another essential metric is the false positive rate (FPR), also known as specificity, which represents the percentage of true negative cases that a detection system incorrectly interprets as positive.

For a typical detection system, when TPR increases, FPR also goes high. But high TPR and low FPR are desirable. Thus, by plotting TPR versus FPR at various classification levels, the receiver operating characteristic (ROC) curve is used to assess the overall effectiveness of a phishing email detection system. These evaluation indicators assist pinpoint problem areas and offer insights into how well a classification model is performing.

### D. Applying the ML Algorithms

We operate ML algorithms on the two datasets, Ling and TREC, using only their common ‘subject’ and ‘message’ features, and excluding TREC’s additional ‘email\_to’ and ‘email\_from’ features.

From the Ling dataset, we first randomly pick 80% of phishing emails and 80% of legitimate emails and combine them to create our training subset of the Ling dataset. The rest 20% phishing and legitimate emails are combined to create the testing subset of the Ling dataset. Using the same procedure, we also create separate training and testing subsets of the TREC dataset. The training subsets are used to train the ML algorithms and the testing subsets are used to evaluate the algorithms’ performances.

## IV. ANALYSIS AND FINDINGS

We now describe our analyses for addressing each of the research questions and derive their answers.




Fig. 1. TN, FP, FN, and TP on the Ling dataset

### A. Performance of ML Algorithms (RQ1)

Figure 1 and Figure 2 present the number of TP, TN, FP, and FN in the ML algorithms’ detection of phishing emails in the Ling and TREC datasets respectively. In the current context, it is important to minimize FN predictions, which happen when the algorithm mistakenly classifies a phishing email as a legitimate one. With respect to the number of FNs, the RF algorithm appears to have outperformed others with zero FN in the Ling dataset and only 45 in the TREC dataset.




Fig. 2. TN, FP, FN, and TP on the TREC dataset

TABLE V  
PERFORMANCE OF ML ALGORITHMS ON THE Ling DATASET

ML Algorithms	Accuracy (%)	Precision (%)	Recall (%)	ROC	F-score (%)
LR	96.70	97.26	96.70	0.9755	96.84
KNN	97.22	97.21	97.22	0.9343	97.17
AB	98.61	98.63	98.61	0.9785	98.62
MNB	87.67	96.98	87.67	0.9358	90.99
GB	97.92	98.06	97.92	0.9782	97.95
RF	97.92	98.18	97.92	0.9879	97.97

The performance ML algorithms on the Ling dataset are presented in Table V. As seen in the table, the RF algorithm achieves the highest ROC score (0.9879) and second highest F-score (97.97%), which is slightly (0.65%) lower than that of Adaboost. Adaboost is found to have achieved slightly higher accuracy, precision, recall, and F-score compared to those of RF but the differences always remain less than 0.7% only.

TABLE VI  
PERFORMANCE OF ML ALGORITHMS ON THE TREC DATASET

ML Algorithms	Accuracy (%)	Precision (%)	Recall (%)	ROC	F-score (%)
LR	98.32	98.32	98.32	0.9813	98.32
KNN	96.86	96.93	96.86	0.9637	96.84
AB	96.80	96.80	96.80	0.9654	96.80
MNB	97.86	97.86	97.86	0.9769	97.86
GB	96.98	96.98	96.98	0.9687	96.98
RF	98.42	98.43	98.42	0.9817	98.41

The classification performance of the ML algorithms on the TREC dataset is presented in Table VI. As seen in the table, the RF algorithm clearly outperforms the other five algorithms achieving the highest scores in all the evaluation metrics. This is possibly because RF is effective in identifying phishing emails in large datasets such as TREC as it can improve the model's generalization performance and reduce overfitting by incorporating the predictions of multiple decision trees.

Figure 3 and Figure 4 present the ROC curves for the ML algorithms operated on the Ling and TREC datasets respectively. The closer an ROC curve reaches to the upper-left corner, the better the performance of the corresponding classifier. On both datasets, RF performs the best, with its ROC curve being the closest to the upper left corner.

Based on the observations discussed above, we now derive the answer to the research question RQ1 as follows.

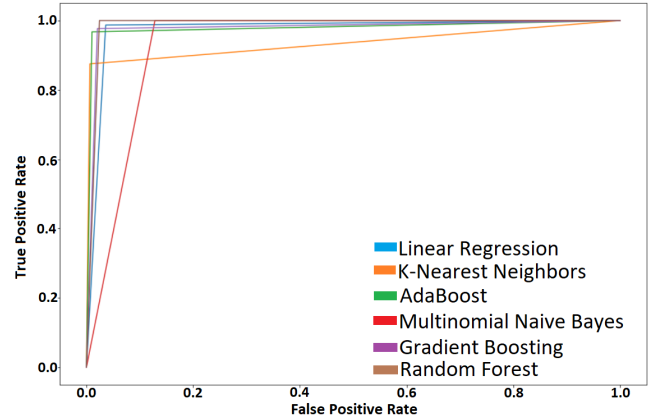


Fig. 3. ROC Curve for the Ling Dataset




Fig. 4. ROC Curve for the TREC Dataset

**Ans. to RQ1:** In detecting phishing emails, the classification performance of the RF algorithm appears superior to the other ML algorithms in our study. This superiority is more evident when the algorithms are operated on the larger dataset. Especially the low number of FNs makes RF particularly suitable for phishing email detection.

### B. ML Algorithms' Training time (RQ2)

We calculate the training time of six ML algorithms while separately training the datasets. Figure 5 and Figure 6 respectively present the training time taken by the ML algorithms when trained on the Ling and TREC datasets.

As seen in the figures, on both datasets, KNN and MNB algorithms have significantly shorter training times compared to the other ML algorithms. This is because these two algorithms are computationally simpler than the others.

The MNB algorithm uses probability theory to determine the likelihood that a new data point is a member of a certain class based on the feature values, whereas the KNN algorithm compares the new data point to all existing data points and classifies it based on the nearest neighbors. Both of these algorithms require less computation compared to others such as GB or RF. Both GB and RF utilize complex decision trees, and ensemble methods, and thus take more time to




Fig. 5. Training time (in seconds) on the *Ling* dataset

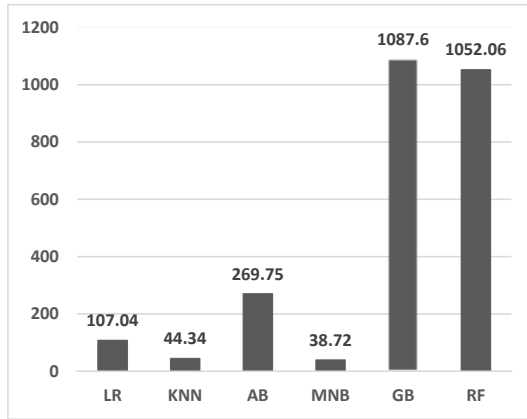


Fig. 6. Training time (in seconds) on the *TREC* dataset

execute, which is also reflected in the figures. Based on the observations, the answer to the research question RQ2 is derived as follows.

**Ans. to RQ2:** *KNN and MNB algorithms have much shorter training times compared to the other ML algorithms studied.*

### C. Impact of the ‘subject’ Feature (RQ3)

To investigate the importance of the ‘subject’ feature in accurately detecting phishing emails, we operate the ML algorithms on the dataset according to the procedure described in Section III-D, but at this phase, we make the ML algorithms disregard the ‘subject’ feature in the datasets.

In Figure 7, we present the number of TN, FP, FN, and TP found in the ML algorithms’ detection of phishing emails in the *Ling* dataset excluding the ‘subject’ feature. If we compare the results in this figure with those of Figure 1, we notice that, when the ‘subject’ feature is excluded, there is a slight decrease in TP for all the algorithms. Figure 8 presents the number of TN, FP, FN, and TP found in the ML algorithms’ detection of phishing emails in the *TREC* dataset excluding the ‘subject’ feature. Now, if we compare the results in this figure with those in Figure 2, we again see that all the ML algorithms’ TP also decreases for the *TREC* dataset when the ‘subject’ feature is disregarded.

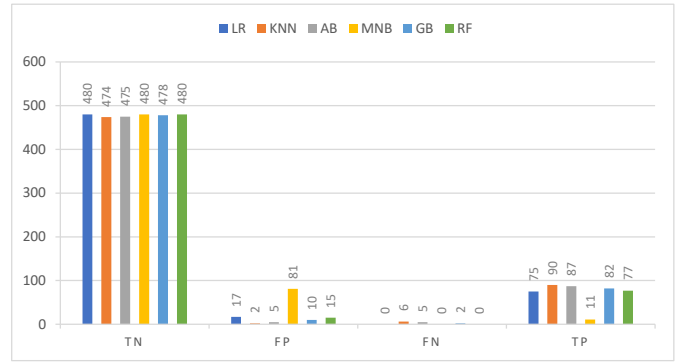


Fig. 7. TN, FP, FN, and TP on the *Ling* dataset excluding the ‘subject’ feature




Fig. 8. TN, FP, FN, and TP on *TREC* dataset excluding the ‘subject’ feature

Table VII presents the accuracy, precision, recall, and F-score of the ML algorithms’ detection of phishing emails in the *Ling* dataset while disregarding the ‘subject’ feature. A comparison of the results in this table with those in Table V reveals the following. In this relatively smaller dataset, when the ‘subject’ feature is excluded, LR and KNN perform better, while AB, MNB, GB, and RF perform worse. Because the *Ling* dataset is relatively small, exclusion of the ‘subject’ feature might not have a substantial impact on the algorithms’ performances. Alternatively, the ‘subject’ may not really have a distinguishable impact in the determination of whether an email is phishing or legitimate in general, irrespective of the algorithms in use. We further examine these possibilities by looking at the performances of the ML algorithms on the larger *TREC* dataset.

In Table VIII, we present the performance of the ML algorithms’ phishing email detection in the *TREC* dataset without taking into account the ‘subject’ feature. If we compare

TABLE VII  
PERFORMANCE ON *Ling* DATASET, EXCLUDING THE ‘SUBJECT’ FEATURE

ML Algorithms	Accuracy (%)	Precision (%)	Recall (%)	ROC	F-score (%)
LR	97.03	97.58	97.03	0.9829	97.15
KNN	98.60	98.59	98.60	0.9666	98.59
AB	98.25	98.25	98.25	0.9676	98.25
MNB	85.84	98.31	85.84	0.9278	90.86
GB	97.90	98.05	97.90	0.9779	97.94
RF	97.38	97.81	97.38	0.9848	97.47

TABLE VIII

PERFORMANCE ON TREC DATASET, EXCLUDING THE ‘SUBJECT’ FEATURE

ML Algorithms	Accuracy (%)	Precision (%)	Recall (%)	ROC	F-score (%)
LR	98.31	98.32	98.31	0.9817	98.31
KNN	83.16	88.55	83.16	0.8730	84.04
AB	96.29	96.29	96.29	0.9608	96.29
MNB	97.71	97.72	97.71	0.9753	97.71
GB	96.78	96.78	96.78	0.9670	96.78
RF	98.38	98.40	98.38	0.9820	98.38

the results of this table with those in Table VI, we notice that accuracy, precision, recall, and F-score decreases for all the algorithms when the ‘subject’ feature is excluded. This larger TREC dataset includes a varied collection of emails, making it more challenging for the ML algorithms to correctly distinguish phishing emails based only on message bodies without taking the ‘subject’ feature into account. Based on the observations and discussion above, we, therefore, derive the answer to the research question RQ3 as follows.

**Ans. to RQ3:** *Inclusion of the ‘subject’ feature increases accuracy of the ML algorithms in detection of phishing emails. However, this impact of the ‘subject’ feature is negligible when ML algorithms are applied on small datasets, but significant when the algorithms operate on larger datasets having a diverse collection of phishing and legitimate emails.*

## V. THREATS TO VALIDITY

Our phishing email detection model is developed using only two datasets (i.e., TREC and Ling datasets). Thus the generalizability of our findings can be argued limited in scope.

We solely focus on the ‘subject’ and ‘message’ (i.e., email body) features of the dataset in this work. However, the sender of an email, along with other header details, email attachments, and URLs in the email body can offer useful information and boost a classification model’s performance. We plan to include all these in our future work.

We examine the effectiveness of six ML algorithms in detecting phishing emails. There are a number of other ML algorithms, which are popular but not included in our work. We plan to include them in future extension to this work.

## VI. RELATED WORK

ML and NLP have been used in many software engineering studies [14]–[18] while many other studies explored different security aspects of software systems [19]–[25]. Sattar et al. [26] applied predictive model analysis for detecting web-spams in web-graphs. Several recent attempts to detect phishing emails have tried with machine learning [27]–[29] and deep learning [27], [30]–[33] techniques. These studies have used various strategies to improve the performance of detection algorithms. Some studies used a single algorithm [29]–[31], while others have compared several algorithms to identify the most effective approaches [27], [28]. Researchers have also tried using single [28], [29] and multiple datasets [27], [30], [31], as well as combination of multiple datasets [27].

Islam et al. [27] compared four ML algorithms on TREC and Ling datasets, using seven features of TREC and two features of Ling. Similar to our work, they also used the Ling and TREC datasets, and RF was reported to have the highest accuracy compared to the four ML algorithms included in their work. However, in their work, the accuracy of RF did not exceed 95% on the TREC dataset (and 96% on the Ling dataset) and while in our study RF achieves more than 98% accuracy on the TREC dataset (and over 97% on the Ling dataset).

Agarwal and Kumar [29] combined Naive Bayes (NB) with Particle Swarm Optimization and evaluated its performance on the Ling dataset. They found that while using NB alone resulted in an accuracy *below* 90%, the integrated approach achieved an accuracy above 90%. We have also used the Ling dataset and achieved almost 98% accuracy on this dataset while the Multinomial Naive Bayes (MNB) algorithm in our work has been able to achieve more than 97% accuracy on the TREC dataset.

Pan et al. [31] used a Semantic Graph Neural Network on the TREC dataset but did not compare their method with any other algorithms. Dhavale [30] employed a convolutional neural network-based approach on the TREC dataset without comparing their approach with any other algorithms. Similar to ours, Dhavale’s method also achieved approximately 98% accuracy. However, unlike the work of Dhavale [30] or Pan et al. [31], we compared the performances of six ML algorithms.

Unlike any of the work discussed above, in our study, we examined the impact of a particular email feature on the six ML algorithms’ detection of phishing emails. This also makes our work unique from the earlier work in the literature.

## VII. CONCLUSION AND FUTURE WORK

In this paper, we have presented a performance comparison of six different machine learning (ML) algorithms in distinguishing phishing emails from legitimate ones. The algorithms are trained and evaluated on two publicly available datasets (i.e., Ling and TREC datasets). Classification performances are measured in terms of accuracy, precision, recall, F-score, and ROC. The training time required for the ML algorithms is also compared for measuring their applicability in a practical context. We have also examined the importance of a particular feature in the identification of phishing emails.

Our study reveals that the Random Forest (RF) algorithm performs better than other ML algorithms, particularly for larger datasets. Due to its low false negatives, RF can be particularly suitable for pragmatic application in phishing email detection, if training time is not a limiting factor. Classification performance of the ML algorithms generally decreases when the ‘subject’ feature is disregarded. The impact of this particular feature is found more substantial when the ML algorithms are operated on larger datasets. KNN (K-nearest Neighbours) and MNB (Multinomial Naive Bayes) algorithms took much less time in training compared to other ML algorithms in our study. In our work, Gradient Boosting (GB) has taken the longest time in training, but this algorithm has not produced



the best results in terms of classification performances. This suggests that simply choosing the algorithm with the longest training time may not necessarily lead to better results. It is important to choose an ML algorithm based on the size and nature of the datasets as well as classification performance requirements.

Our future research will address the limitations of this work and will investigate the use of additional features to improve the performance of ML algorithms in phishing email detection. There are other benchmark datasets for phishing emails, such as the Enron [34], Nazario [35], and SpamAssassin [36] corpora. These datasets are available in different formats, such as mdir and mbox. They can be combined into a single dataset in a single format, such as CSV, to create a bigger, more diversified dataset. This remains within our plans for future work. The planned extension of this work will also include additional ML algorithms. We hypothesize that applying deep learning approaches, such as convolutional neural networks (CNN), could improve the results by automatically recognizing complex patterns and features that might be difficult to detect with conventional ML approaches. We also plan to extend this work in this direction in the future.

## REFERENCES

- [1] D. Warburton, *2020 Phishing and Fraud Report*. <https://www.f5.com/labs/articles/threat-intelligence/2020-phishing-and-fraud-report>, Verified: March 2023.
- [2] APWG, *Phishing Activity Trends Reports*. <https://www.f5.com/labs/articles/threat-intelligence/2020-phishing-and-fraud-report>, Verified: March 2023.
- [3] MimeCast, *How to Stop Phishing Attacks (Whitepaper)*. <https://www.mimecast.com/resources/white-papers/how-to-stop-phishing-attacks/>, Verified: March 2023.
- [4] MimeCast, *The State of Email Security 2023 (E-book)*. <https://www.mimecast.com/state-of-email-security/>, Verified: March 2023.
- [5] J. Uchill, "Typo led to podesta email hack: Report," *The Hill*, vol. 13, 2016.
- [6] Z. Dou, I. Khalil, A. Khreishah, A. Al-Fuqaha, and M. Guizani, "Systematization of knowledge (sok): A systematic review of software-based web phishing detection," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2797–2819, 2017.
- [7] A. Das, S. Baki, A. El Aassal, R. Verma, and A. Dunbar, "Sok: a comprehensive reexamination of phishing research from the security perspective," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 1, pp. 671–708, 2019.
- [8] A. Mukherjee, N. Agarwal, and S. Gupta, "A survey on automatic phishing email detection using natural language processing techniques," *Int. Res. J. Eng. Technol*, vol. 6, no. 11, pp. 1881–1886, 2019.
- [9] G. Sakkis, I. Androutsopoulos, G. Paliouras, V. Karkaletsis, C. D. Spyropoulos, and P. Stamatopoulos, "A memory-based approach to anti-spam filtering for mailing lists," *Information retrieval*, vol. 6, pp. 49–73, 2003.
- [10] G. V. Cormack, "Trec 2007 spam track overview," in *Proc. of the 16th Text REtrieval Conference (TREC)*, vol. 500, 2007, p. 274.
- [11] A. Chatterjee, *Preprocessed TREC 2007 Public Corpus Dataset*. <https://www.kaggle.com/datasets/imdeepmind/preprocessed-trec-2007-public-corpus-dataset>, Verified: March 2023.
- [12] A. Géron, *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow*. O'Reilly Media, Inc., 2022.
- [13] E. Bisong, *Building machine learning and deep learning models on Google cloud platform*. Springer, 2019.
- [14] M. Zibran, "On the effectiveness of labeled latent dirichlet allocation in automatic bug-report categorization," in *38th ACM/IEEE International Conference on Software Engineering (ICSE)*, 2016, pp. 713–715.
- [15] M. Islam and M. Zibran, "Leveraging automated sentiment analysis in software engineering," in *MSR*, 2017, pp. 203–214.
- [16] M. Islam and M. Zibran, "Sentistrength-se: Exploiting domain specificity for improved sentiment analysis in software engineering text," *Journal of System and Software*, vol. 145, pp. 125–146, 2018.
- [17] M. Islam and M. Zibran, "Deva: Sensing emotions in the valence arousal space in software engineering text," in *SAC*, 2018, pp. 1536–1543.
- [18] M. Islam, M. Ahmmad, and M. Zibran, "Marvalous: Machine learning based detection of emotions in the valence-arousal space in software engineering text," in *34th ACM/SIGAPP Symposium on Applied Computing (SAC)*, 2019, pp. 1786–1793.
- [19] M. Islam and M. Zibran, "A comparative study on vulnerabilities in categories of clones and non-cloned code," in *10th IEEE Intl. Workshop on Software Clones*, 2016, pp. 8–14.
- [20] M. Islam, M. Zibran, and A. Nagpal, "Security vulnerabilities in categories of clones and non-cloned code: An empirical study," in *11th ACM/IEEE Intl. Symposium on Empirical Software Engineering and Measurement*, 2017, pp. 20–29.
- [21] R. Joseph, M. Zibran, and F. Eishita., "Choosing the weapon: A comparative study of security analyzers for android applications," in *Intl. Conference on Software Engineering, Management and Applications*, 2021, pp. 51–57.
- [22] D. Murphy, M. Zibran, and F. Eishita, "Plugins to detect vulnerable plugins: An empirical assessment of the security scanner plugins for wordpress," in *Intl. Conference on Software Engineering, Management and Applications*, 2021, pp. 39–44.
- [23] A. Rajbhandari, M. Zibran, and F. Eishita, "Security versus performance bugs: How bugs are handled in the chromium project," in *Intl. Conference on Software Engineering, Management and Applications*, 2022, pp. 70–76.
- [24] J. Rodriguez, M. Zibran, and F. Eishita, "Finding the middle ground: Measuring passwords for security and memorability," in *20th IEEE International Conference on Software Engineering, Management and Applications*, 2022, pp. 77–82.
- [25] A. Champa, M. Rabbi, F. Eishita, and M. Zibran, "Are we aware? an empirical study on the privacy and security awareness of smartphone sensors," in *21st IEEE International Conference on Software Engineering, Management and Applications (SERA)*, 2023, pp. 1–8 (to appear).
- [26] N. Sattar, S. Arifuzzaman, M. Zibran, and M. Sakib, "Detecting web spam in webgraphs with predictive model analysis," in *3rd International Workshop on Big Data Analytic for Cyber Crime Investigation and Prevention*, 2019, pp. 4299–4308.
- [27] M. Islam, M. Al Amin, M. Islam, M. Mahbub, M. Showrov, and C. Kaushal, "Spam-detection with comparative analysis and spamming words extractions," in *9th International Conference on Reliability, Information Technologies and Optimization*, 2021, pp. 1–9.
- [28] N. A. Unnithan, N. Harikrishnan, R. Vinayakumar, K. Soman, and S. Sundarakrishna, "Detecting phishing e-mail using machine learning techniques," in *Proc. 1st anti-phishing shared task pilot 4th acm iwspa co-located 8th acm conf. data appl. secur. privacy (codaspy)*, 2018, pp. 51–54.
- [29] K. Agarwal and T. Kumar, "Email spam detection using integrated approach of naïve bayes and particle swarm optimization," in *2018 Second International Conference on Intelligent Computing and Control Systems (ICICCS)*. IEEE, 2018, pp. 685–690.
- [30] S. Dhavale, "C-asft: convolutional neural networks-based anti-spam filtering technique," in *Proceeding of International Conference on Computational Science and Applications: ICCSA 2019*, 2020, pp. 49–55.
- [31] W. Pan, J. Li, L. Gao, L. Yue, Y. Yang, L. Deng, and C. Deng, "Semantic graph neural network: a conversion from spam email classification to graph classification," *Scientific Programming*, vol. 2022, pp. 1–8, 2022.
- [32] L. Halgaš, I. Agrafiotis, and J. R. Nurse, "Catching the phish: Detecting phishing attacks using recurrent neural networks (rnns)," in *20th International Conference on Information Security Applications*, 2020, pp. 219–233.
- [33] S. Magdy, Y. Abouelseoud, and M. Mikhail, "Efficient spam and phishing emails filtering based on deep learning," *Computer Networks*, vol. 206, p. 108826, 2022.
- [34] B. Klimt and Y. Yang, "The enron corpus: A new dataset for email classification research," in *Machine Learning: ECML 2004: 15th European Conference on Machine Learning, Pisa, Italy, September 20-24, 2004. Proceedings 15*. Springer, 2004, pp. 217–226.
- [35] J. Nazario, *The online phishing corpus*. <http://monkey.org/jose/wiki/doku.php>, Verified: March 2023.
- [36] A. Schwartz, *Apache SpamAssassin*. <https://spamassassin.apache.org/>, Verified: March 2023.