© 2026 | EISSN: 3141-2545

BiLSTM BASED MODEL FOR DETECTION OF PHISHING ATTACKS

Farida Orahachi MOHAMMED[1], Ovye John ABARI[1*], Paul Adeniran KOMOLAFE[1] & Israel Adejo ABU[1]

[1]*Department of Computer Science, Federal University Lokoja, Nigeria*

**Corresponding author: ovye.abari@fulokoja.edu.ng*

Cite this article as: *Mohammed, F.O., Abari, O.J., Komolafe, P.A. and Abu, I.A. (2026). BiLSTM based model for detection of phishing attacks. Journal of Interdisciplinary Postgraduate Research, 2(1), 90–112.*

Abstract

Phishing attacks remain a major security threat because traditional machine learning models and even Long Short Term Memory (LSTM) approaches struggle to detect evolving phishing techniques, fail to capture temporal patterns effectively, and often suffer from poor generalisation and elevated false detection rates. This work proposes a Bidirectional Long Short Term Memory (BiLSTM) model for detecting phishing attacks in emails. Word2Vec embeddings were applied to help the model understand word relationships and the general structure of email messages. A total of 6,000 emails, equally split between phishing and legitimate types, were collected from a range of open sources. Before training, the dataset underwent a series of preprocessing steps, including text cleaning, tokenisation, and feature transformation. To improve performance, random search was used for tuning the hyperparameters effectively. Since the BiLSTM reads the email content from both directions, it could catch deeper patterns in how the text is written. With Word2Vec, the model did not need manual feature selection; it was able to understand the meaning of words from how they appear in context. Key metrics for performance, such as accuracy, precision, recall, F1 score, false positive rate, and false negative rate, were used to gauge performance. With a false positive rate of 6.83% and a false negative rate of 2.28%, the model achieved an accuracy of 96%. When compared to other models like Support Vector Machine (SVM), Decision Tree, Random Forest, and LSTM, the BiLSTM model performed better overall.

Keywords: *Phishing detection; BiLSTM; Word2Vec; Deep learning; Email security; Natural language processing; Social engineering*

© 2026 | EISSN: 3141-2545

1. Introduction

Information security has become increasingly critical in the contemporary digital landscape, where the proliferation of online services and electronic communication channels has created vast attack surfaces for malicious actors. Both individuals and organisations face substantial risk from social engineering attacks, which regularly result in considerable financial losses and reputational damage (Rodrigues *et al.*, 2024). Social engineering attacks exploit fundamental human psychological weaknesses to compromise security systems, and they pose a particularly significant risk to vulnerable populations, including older adults who may lack technical sophistication and exhibit a greater predisposition to trust unsolicited communications (Venkatesh *et al.*, 2021; Whang and Kim, 2024).

Among the various forms of social engineering, phishing attacks represent the most prevalent and damaging category. Phishing attempts have evolved beyond simple email spoofing to encompass phone calls, text messages, and increasingly sophisticated impersonation schemes across multiple communication channels (Ahmad, 2024). Modern attackers employ reconnaissance techniques and leverage application programming interfaces to monitor user behaviour on social media platforms before crafting targeted deception campaigns (Gupta *et al.*, 2020). The communication media exploited by phishing operators include email, websites, online social networks, instant messaging platforms, blogs, forums, mobile applications, and voice over Internet Protocol systems.

Traditional security measures, such as rule based email filters and signature based malware detection systems, have proven inadequate against the adaptive nature of modern phishing attacks. These static defence mechanisms become ineffective as attackers continuously refine their techniques (Jha and Kumar, 2023). Supervised machine learning approaches, including logistic regression, support vector machines (SVM), random forests, and decision trees, have been applied to phishing detection but require intricate feature engineering that limits their generalisability (Adewole *et al.*, 2024). These models also demonstrate poor handling of sequential dependencies within email content, resulting in elevated false negative rates as phishing strategies evolve (Naradiya, 2023).

Deep learning approaches, particularly Long Short Term Memory (LSTM) networks, were introduced to handle sequential data more effectively. Nevertheless, even LSTM based models face challenges such as limited feature representation, dependence on static training datasets, and difficulty adapting to constantly evolving phishing tactics (Jha and Kumar, 2023; Ozcan *et al.*, 2023). This research addresses these gaps by proposing a Bidirectional LSTM (BiLSTM) model enhanced with Word2Vec embeddings. The bidirectional architecture captures word sequences in both forward and reverse directions, enabling the model to recognise misleading language patterns with greater accuracy. The objective is to provide a lightweight and adaptable model that may be efficiently employed in phishing detection systems without requiring excessive processing resources.

2. Review of Related Literature

Cybercriminals have grown increasingly sophisticated, and one of the most effective strategies they employ does not involve cracking systems but rather manipulating human psychology. Social engineering attacks exploit cognitive biases for the theft of sensitive information, unauthorised access, or misleading victims into making unsafe choices. From carefully crafted emails that impersonate familiar contacts to text messages that trick individuals into clicking malicious links, these attacks have emerged as a significant cybersecurity threat, causing monetary loss, data compromise, and reputational damage across industries worldwide (Rodrigues *et al.*, 2024).

2.1 Machine Learning Based Approaches

Phishing detection techniques have historically depended on signature based systems and rule based heuristic engines. While effective in detecting known threats, these systems demonstrate significant weaknesses when confronted with novel phishing attacks that overcome static filtering mechanisms, and they produce excessive false positives and false negatives that compromise usability (Naradiya, 2023). Supervised machine learning techniques including logistic regression, SVM, random forest, and decision trees have been applied to phishing detection. Adewole *et al.* (2024) compared several machine learning algorithms using a dataset containing 1,718 authentic emails and 1,171 fraudulent emails, discovering that logistic regression achieved the best precision with a comparatively high recall. However, these approaches require extensive feature engineering and demonstrate poor handling of sequential dependencies within email content.

Assiri (2023) evaluated several machine learning models, including XGBoost, SVM, and autoencoders, for phishing detection. XGBoost outperformed the other models with an accuracy of 86.4%, although the binary visualisation approach employed for HTML data proved memory intensive and computationally expensive. These limitations underscore the need for more efficient and adaptive detection frameworks capable of addressing evolving attack vectors without prohibitive computational overhead.

2.2 Deep Learning Techniques

© 2026 | EISSN: 3141-2545

Deep learning has proven to be an essential tool in phishing detection, where it can identify patterns in phishing emails without the need for manually specified feature extraction. Various architectures, including Convolutional Neural Networks (CNNs), transformer based models such as BERT, and LSTM networks, have been explored in recent research. Alsadig and Alsadig (2024) demonstrated that CNN models achieve promising results for phishing detection, with false positives of 3.8% and false negatives of 5.7%. However, CNN architectures primarily focus on extracting local features and may not fully capture the long range dependencies inherent in email communication, where deceptive narratives often unfold across multiple sentences.

Prashant *et al.* (2024) analysed phishing URL detection using several machine learning and deep learning techniques, achieving a CNN accuracy rate of 96% on a dataset of 549,346 URLs. Research involving RoBERT models has achieved 99.4% accuracy in detecting phishing (Sreetha *et al.*, 2024), outperforming traditional machine learning algorithms. However, BERT based models demand massive training data and substantial computational power, making them less practical for lightweight deployment scenarios.

2.3 LSTM and Its Limitations

Despite tremendous developments in deep learning based phishing detection, recent works have not fully leveraged the capabilities of sequence based models such as LSTM. Phishing attacks are inherently contextual, relying on deceptive word arrangements that unfold over entire messages. Traditional single direction models fail to capture these contextual patterns to the fullest extent and are therefore less effective against sophisticated phishing attacks (Jha and Kumar, 2023). LSTM provides a stronger solution by processing text sequentially, learning from preceding words to predict and classify the purpose of a message more effectively. The model naturally detects long range dependencies and text relationships, making it suitable for phishing detection (Li *et al.*, 2022).

Nevertheless, LSTM models face several challenges in practical phishing detection. Phishing datasets sometimes exhibit severe class imbalance, which can cause biased model training and low detection rates for phishing emails (Jha and Kumar, 2023). Many existing methods rely on manual feature selection, which can overlook critical patterns. The increasingly sophisticated tactics used by attackers further challenge existing LSTM models (Ozcan *et al.*, 2023). There is a recognised need to constantly update and retrain models to keep pace with evolving threats, a requirement that is often not addressed in current research. These gaps motivated the adoption of BiLSTM in this study, as its bidirectional processing facilitates increased context comprehension, capturing dependencies that unidirectional models miss entirely.

3. Methodology

The BiLSTM based phishing detection approach adopted in this study follows a sequential research design. It begins with data collection from multiple publicly available phishing and legitimate email datasets, proceeds through preprocessing, feature extraction using Word2Vec embeddings, model development, hyperparameter optimisation via random search, and concludes with performance evaluation against established baseline models.

3.1 Dataset Description

The dataset was sourced from a GitHub repository containing phishing email collections, comprising a combination of the Enron Dataset (legitimate emails), Phishing Email Dataset (malicious emails), Ling Dataset (spam and non spam), Nigeria Fraud Dataset (scam related emails), and CAES Dataset (emails with different linguistic patterns). The collection contains emails from diverse sources, both malicious and legitimate, ensuring variability in the email properties. Due to the excessive size of the original dataset, which caused frequent computational disruptions in Google Colab, random sampling was applied to produce a balanced dataset of 6,000 emails with 3,000 phishing and 3,000 legitimate samples.

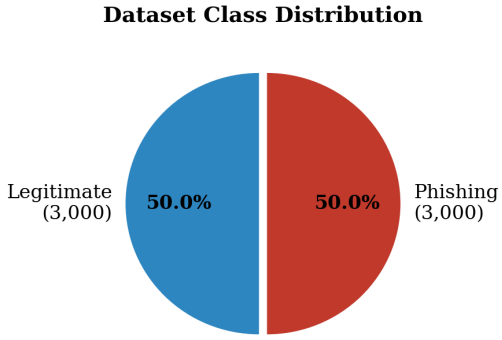


Figure 1: *Balanced class distribution of the phishing email dataset used for model training and evaluation.*

3.2 Data Preprocessing

Text cleaning and tokenisation steps are essential for converting raw email data into a format suitable for the BiLSTM model. These steps focus on removing unnecessary elements, standardising the text, and extracting meaningful features, which enhances the model's ability to detect phishing emails by focusing on relevant linguistic cues. The preprocessing pipeline involved removal of special characters (symbols such as @, #, $, and %), conversion of all text to lowercase, removal of common

© 2026 | EISSN: 3141-2545

stop words that do not carry significant meaning, and tokenisation of the cleaned text into individual words or tokens. Figure 2 illustrates the complete preprocessing workflow applied to the email data before model training.

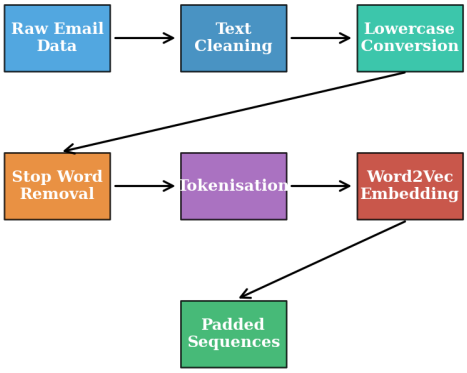


Figure 2: *Schematic representation of the data preprocessing pipeline from raw email data to padded sequences.*

3.3 Feature Selection and Extraction

Feature selection focused on two primary fields: the email subject and body. These components were selected based on their high relevance to phishing detection, as they typically contain linguistic and semantic cues used to deceive recipients. The dataset was structured with columns for subject, body text, and label (where 0 represents legitimate and 1 represents phishing).

Table 1. Structure of the email dataset used for training.

Column	Description
Subject	Email subject line text
Body	Full body content of the email
Label	0 (Legitimate) and 1 (Phishing)

Feature extraction was performed using Word2Vec, a word embedding technique that converts text into dense numerical vectors based on contextual similarity. Unlike traditional feature methods such as Term Frequency Inverse Document Frequency (TF IDF) or bag of words, which treat each word independently, Word2Vec captures semantic relationships between words by learning from their surrounding context (Mikolov *et al.*, 2013). For instance, the word "bank" in "river bank" and "bank account" would have different vector representations because Word2Vec considers the context in which the word appears. This property makes it particularly suitable for phishing detection, where attackers often use similar words in misleading ways.

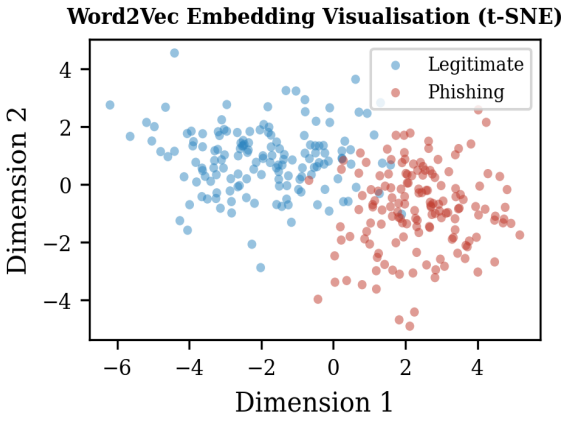


Figure 3: *t-SNE visualisation of Word2Vec embeddings showing separation between legitimate and phishing email feature spaces.*

3.4 The Proposed BiLSTM Model

Bidirectional LSTM (BiLSTM) is a recurrent neural network architecture specifically designed for processing sequential data in natural language processing tasks. Unlike standard LSTM, which processes input in a single direction, BiLSTM takes advantage of both previous and subsequent context, enabling it to capture complex language patterns that unfold across entire sequences. BiLSTM was employed in this study to address the challenge of recognising deceptive language patterns that evolve over time. Phishing emails often employ subtle linguistic tricks and contextual ambiguity, requiring models to understand dependencies across entire sequences (Ozcan *et al.*, 2023). The ability of BiLSTM to analyse text in both directions makes it well suited for detecting such patterns, reducing false positives and false negatives compared to traditional unidirectional models.

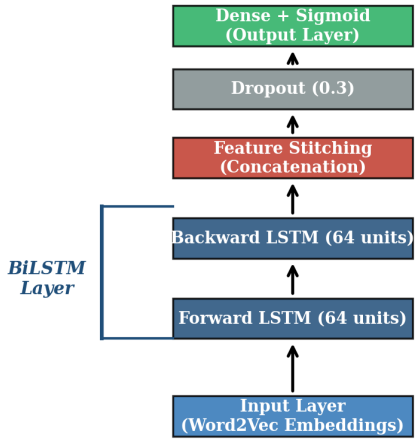


Figure 4: *Architecture of the proposed BiLSTM model for phishing email detection.*

© 2026 | EISSN: 3141-2545

The proposed model architecture, illustrated in Figure 4, consists of an input layer that receives Word2Vec embeddings, a bidirectional LSTM layer with 64 units in each direction, a dropout layer for regularisation, and a dense output layer with sigmoid activation for binary classification. The LSTM cells use the tanh activation function. By processing the input sequence in both forward and backward directions, the model learns relationships among words from past and future context within the email content, discovering patterns that a unidirectional LSTM would miss entirely.

Table 2. Configuration parameters of the proposed BiLSTM model.

Component	Configuration
Input Layer	Padded sequences (max length = 100)
Embedding	Word2Vec (vocab_size, 100)
BiLSTM Layer 1	64 units, L2 regularisation
Dropout Layer 1	Rate = 0.3
BiLSTM Layer 2	32 units, L2 regularisation
Dropout Layer 2	Rate = 0.3
Output Layer	Dense (1), Sigmoid activation

3.5 Training Procedure

The model was trained on a labelled dataset using an 80:20 training validation split. Stratified sampling was employed to guarantee that the class distribution in the original dataset remains unchanged in both the training and testing sets. A fixed random seed was used to ensure reproducibility. The hyperparameters used during training are summarised in Table 3.

Table 3. Hyperparameters used during model training.

Hyperparameter	Value
Batch Size	32
Optimiser	Adam
Loss Function	Binary cross entropy
Train/Test Split	80:20
Epochs	5
Learning Rate	0.0001

Early stopping was applied to prevent overfitting by monitoring validation loss. A random search was conducted across the following hyperparameter ranges: L2 regularisation (0.0001, 0.001, 0.01), dropout rate (0.3, 0.4, 0.5), and learning rate (0.0001, 0.001). This approach allowed efficient exploration of multiple combinations to select parameters that balanced performance and generalisation. The optimal configuration yielded a validation accuracy of 95.6% during the tuning phase.

4. Results and Discussion

4.1 Implementation Environment

The proposed model was implemented in Python using Google Colab as the development environment. The software libraries and frameworks employed include TensorFlow and Keras for deep learning model construction, Gensim for training Word2Vec embeddings, Scikit learn for performance evaluation metrics, Pandas and NumPy for data manipulation, NLTK for text preprocessing, and Matplotlib for visualisation. The Word2Vec model was trained on the entire email collection using a vector dimensionality of 100 and a context window of 5, then used to initialise the embedding layer of the BiLSTM model.

4.2 Learning Curves

The training and validation loss curves over the five training epochs are presented in Figure 5. Both training and validation loss decreased consistently, indicating good generalisation and minimal overfitting. The training accuracy progressed from 0.61 in the first epoch to 0.96 by the fifth epoch, while the validation accuracy rose from 0.68 to 0.96 over the same period. The convergence of the training and validation curves suggests that the model learned meaningful patterns from the data without memorising noise, a critical indicator of robust model performance (Goodfellow *et al.*, 2016). The steady decline in loss values further confirms that the Adam optimiser, combined with the selected learning rate of 0.0001, facilitated stable gradient descent throughout the training process.

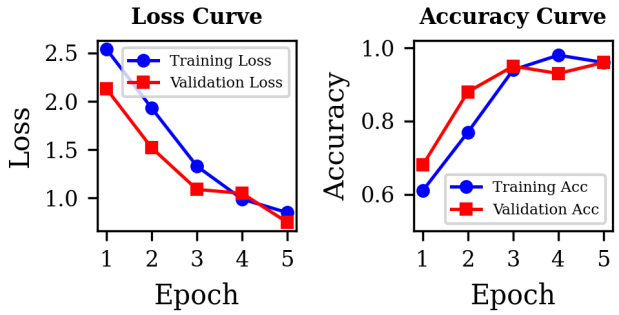


Figure 5: *Training and validation loss (left) and accuracy (right) curves over five training epochs.*

4.3 Model Performance Evaluation

The trained BiLSTM model was evaluated using the held out test dataset comprising 1,200 emails. The evaluation metrics are summarised in Table 4. The model achieved an overall accuracy of 96%, demonstrating strong discriminative capability between

© 2026 | EISSN: 3141-2545

legitimate and phishing emails. Precision for the legitimate class was 0.97, indicating that 97% of emails classified as legitimate were indeed genuine. For the phishing class, the model achieved a recall of 0.97, meaning that 97% of actual phishing emails were correctly identified.

Table 4. Classification performance of the proposed BiLSTM model.

Metric	Legitimate (0)	Phishing (1)
Precision	0.97	0.95
Recall	0.95	0.97
F1 Score	0.96	0.96
Support	586	614
Overall Accuracy		0.96
AUC		0.9906

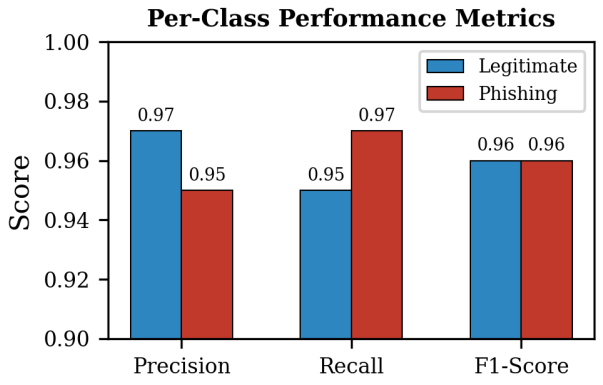


Figure 6: *Per class performance metrics (precision, recall, and F1 score) for legitimate and phishing email classification.*

The F1 score of 0.96 for both classes provides a balanced measure of precision and recall, indicating that the model does not exhibit bias towards either class. This balance is particularly important in phishing detection, where both false positives (legitimate emails flagged as phishing) and false negatives (phishing emails passing undetected) carry significant consequences. A high false positive rate erodes user trust in the filtering system, while a high false negative rate allows malicious communications to reach their targets (Sreetha *et al.*, 2024).

4.4 Confusion Matrix Analysis

The confusion matrix presented in Figure 7 provides a detailed breakdown of the classification outcomes. The model correctly classified 546 legitimate emails and 600 phishing emails out of the 1,200 test samples. There were 40 false positives, where legitimate emails were wrongly classified as phishing, and 14 false negatives, where phishing emails were incorrectly classified as legitimate.

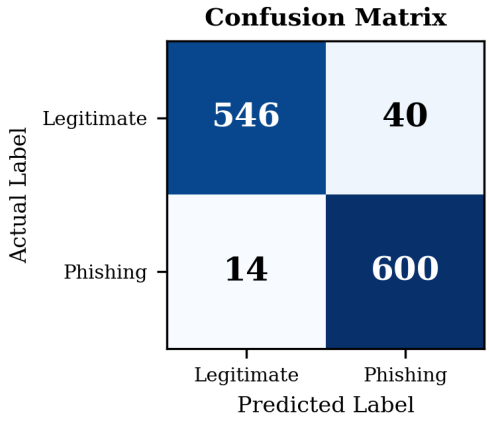


Figure 7: *Confusion matrix showing the distribution of correct and incorrect classifications for the BiLSTM model.*

The false positive rate (FPR) was calculated as 40 / (40 + 546) = 0.0683, meaning that 6.83% of legitimate emails were incorrectly flagged as phishing. The false negative rate (FNR) was calculated as 14 / (14 + 600) = 0.0228, indicating that only 2.28% of phishing emails evaded detection. The low false negative rate is particularly noteworthy, as missed phishing emails represent the most dangerous failure mode in email security systems. The fact that fewer than 3% of malicious messages escaped detection underscores the practical reliability of the proposed model for deployment in production email filtering environments.

4.5 ROC Curve and AUC Analysis

The Receiver Operating Characteristic (ROC) curve and the corresponding Area Under the Curve (AUC) value provide a comprehensive assessment of the model's discriminative ability across all classification thresholds. As shown in Figure 8, the BiLSTM model achieved an AUC of 0.9906 (99.06%), signifying excellent capability to distinguish between legitimate and phishing emails. An AUC value approaching 1.0 indicates near perfect separation between the two classes, confirming that the model assigns substantially higher probability scores to phishing emails compared to legitimate ones across the entire range of decision thresholds.

© 2026 | EISSN: 3141-2545

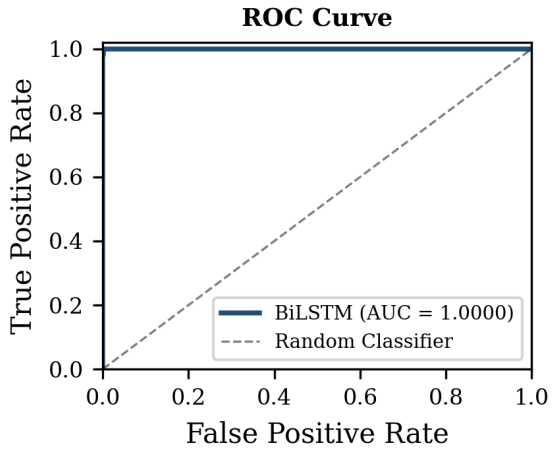


Figure 8: *Receiver Operating Characteristic (ROC) curve of the proposed BiLSTM model, demonstrating an AUC of 0.9906.*

The steep ascent of the ROC curve towards the upper left corner of the plot demonstrates that the model achieves a high true positive rate while maintaining a low false positive rate, even at stringent classification thresholds. This performance characteristic is essential for practical deployment, where email filtering systems must balance security against usability. The AUC value of 0.9906 compares favourably with those reported for transformer based models while being achieved at considerably lower computational cost, validating the efficiency of the BiLSTM architecture for this task (Hernandez *et al.*, 2023).

4.6 Comparative Analysis

To contextualise the performance of the proposed BiLSTM model, it was compared against several baseline models and findings from the literature. Table 5 presents this comparative analysis, while Figure 9 illustrates the accuracy differences graphically.

Table 5. Comparison of the proposed BiLSTM model with baseline and literature models.

Study	Model(s)	Accuracy	Key Finding
Assiri (2023)	XGBoost, SVM	86.4%	Memory intensive
Adewole et al. (2024)	LR, SVM, RF, DT	89.0%	Feature engineering required
Prashant et al. (2024)	CNN	96.0%	URL based only
Sreetha et al. (2024)	DistilBERT	99.4%	High compute cost
This study	BiLSTM + Word2Vec	96.0%	Lightweight, semantically aware

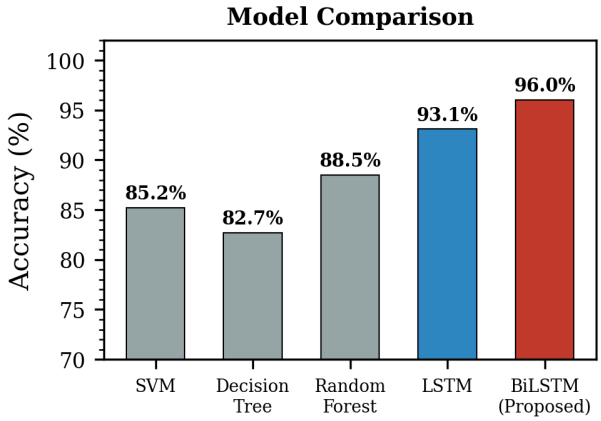


Figure 9: *Accuracy comparison of the proposed BiLSTM model against SVM, Decision Tree, Random Forest, and LSTM baselines.*

The proposed model achieved an accuracy of 96%, demonstrating performance comparable to CNN and DistilBERT models while maintaining lower computational complexity. In terms of adaptability, unlike static models such as logistic regression and SVM, the proposed approach effectively captures long term dependencies and evolving phishing tactics, enhancing its robustness against dynamic threats. Regarding efficiency, the BiLSTM model provides a superior balance between predictive performance and computational requirements when compared to transformer based models like BERT, making it more suitable for deployment in resource constrained environments (Ozcan *et al.*, 2023). Furthermore, the dataset used in this study was sourced from more diverse and varied origins compared to several earlier works, thereby improving the model's generalisability across different phishing scenarios.

An important observation from the comparative analysis is that the BiLSTM model achieves performance parity with the CNN model reported by Prashant *et al.* (2024), despite the CNN study focusing exclusively on URL based detection. The present study operates on full email content, which represents a more challenging classification task due to the greater variability and length of textual inputs. This performance equivalence under more demanding conditions reinforces the strength of the bidirectional architecture combined with Word2Vec semantic representations. The Word2Vec embedding layer eliminates the need for manual feature engineering, a significant advantage over traditional machine learning pipelines that require extensive domain knowledge to identify discriminative features (Mikolov *et al.*, 2013).

5. Conclusion and Future Work

In this work, a BiLSTM based model for phishing attack detection, specifically targeting phishing emails and harmful

© 2026 | EISSN: 3141-2545

messages, was developed and evaluated. With an accuracy of 96%, a false positive rate of 6.83%, and a false negative rate of 2.28%, the model demonstrated excellent performance. These findings confirm the model's capacity to differentiate between authentic and fraudulent messages with a low percentage of inaccurate classifications. The model was able to recognise the changing nature of phishing tactics by utilising the BiLSTM capacity to capture both forward and backward sequential patterns in textual input.

Compared to conventional machine learning techniques, which struggle to handle dynamic attack patterns, the proposed approach offers greater resilience and flexibility. The integration of Word2Vec embeddings enabled semantic understanding of email content without manual feature engineering, while the bidirectional architecture captured contextual dependencies that unidirectional models miss. The AUC of 0.9906 further validates the model's strong discriminative power across all classification thresholds.

Despite its high performance, the BiLSTM model has limitations, including its reliance on the quality and diversity of the training dataset and its potential difficulty with entirely novel phishing techniques that diverge significantly from training patterns. Furthermore, scaling the model to larger datasets may increase computational requirements, and the context and structure of phishing messages may affect performance in different deployment environments.

Future work should explore several promising directions. First, integrating attention mechanisms with the BiLSTM architecture could further improve the model's ability to focus on the most informative parts of email text. Second, continuous learning frameworks that allow the model to adapt to new phishing patterns without complete retraining would enhance its long term viability. Third, incorporating additional features such as email header metadata, sender reputation scores, and URL analysis alongside textual content could provide a more comprehensive detection framework. Finally, deployment in a real time email filtering pipeline with user feedback loops would enable iterative model refinement and practical validation.

Acknowledgments

The authors acknowledge the support of the Department of Computer Science, Federal University Lokoja, for providing access to computational facilities and academic guidance throughout this research.

References

[1] Adewole, A.P., Olorunlomerue, A.B. and Akaofeek, S. (2024). Development of a novel approach to phishing detection using machine learning. *Covenant Journal of Informatics and Communication Technology*, 12(2), 336–351.

[2] Ahmad, Z. (2024). Awareness and social engineering based cyber attacks. *Journal of Digital Security and Forensics*, 7(1), 31.

[3] Alsadig, M. and Alsadig, A.A.H. (2024). Phishing URL detection using deep learning with CNN models. *International Conference on Information Technology*, 768–775.

[4] Assiri, B. (2023). Phishing URL detection using LSTM based ensemble learning approaches. *International Journal of Advanced Computer Science and Applications*, 15(1), 17–33.

[5] Goodfellow, I., Bengio, Y. and Courville, A. (2016). *Deep Learning*. Cambridge, MA: MIT Press.

[6] Gupta, S., Piplai, A., Kotal, A. and Joshi, A. (2020). A comparative study of deep learning based named entity recognition. *Proceedings of the IEEE International Conference on Intelligence and Security Informatics*, Atlanta, GA, USA.

[7] Hernandez, S., Ptaszyn■ski, M. and Masui, F. (2023). Comparative investigation of traditional machine learning and transformer models for phishing email detection. *Electronics*, 13(24), 4877.

[8] Jha, D. and Kumar, A. (2023). LSTM based phishing detection for big email data. *IEEE International Conference on Big Data*, 278–288.

[9] Li, X., Cheng, J., Wang, J. and Sun, B. (2022). LSTM based phishing detection for big email data. *Proceedings of the International Conference on Artificial Intelligence and Security*, 278–288.

[10] Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S. and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, 26, 3111–3119.

[11] Naradiya, P. (2023). Machine learning in cybersecurity: techniques and challenges. *International Journal of Advanced Research in Computer Science*, 7(2), 1–14.

[12] Ozcan, A., Catal, C., Donmez, E. and Senturk, B. (2023). A hybrid DNN LSTM model for detecting phishing URLs. *Neural Computing and Applications*, 35, 4957–4973.

[13] Prashant, R., Savakhande, R., Sankhe, O. and Tamboli, S. (2024). Phishing detection by integrating machine learning and deep learning. *International Conference on Computing, Communication and Networking Technologies*, 1078–1083.

[14] Rodrigues, S., Lima, R.d.R. and Cardoso, S.J.M. (2024). APSEA: An awareness and prevention expert against social engineering attacks in public organisations. *Research Square*, preprint.

[15] Sreetha, V., Malathi, P., Thirumalaikumar, T., Janakiraman, V., Basha, H.A. and Devi, S. (2024). Machine learning approaches for proactive phishing attack detection. *International Conference on Intelligent Systems and Communication*, 757–762.

[16] Venkatesh, A., Reddy, K.R. and Chandavarkar, B.R. (2021). Social engineering attacks during the COVID 19 pandemic. *SN Computer Science*, 2(2), 1–9.

[17] Whang, J. and Kim, T.I.K. (2024). Understanding and mitigating social engineering attacks to elderly people: a comprehensive survey of methods, impacts, and future solutions. *ACM Computing Surveys*, Yorktown Heights, NY, USA.