International Journal of Emerging Research in Science, Engineering, and Management
Vol. 1, Issue 6, pp.11-18, December 2025.
www.ijersem.com eISSN- 3107-9075

Detection of Phishing Websites using Novel Machine Learning Fusion Approach

[1]V. Gopi, [2]K. Gnanasree, [3]G. Dharshini, [4]M. Charan Teja, [5]V. Murali, [6]S. Naveen Kumar

Department of CSE, Siddartha Institute of Science and Technology, Puttur, India
[1]nadh.gopi@gmail.com, [2]kunapallignanasree@gmail.com, [3]gandhamanenidharshini@gmail.com, [4]cherycharan86@gmail.com, [5]volipimurali2631@gmail.com, [6]naveensareddy2003@gmail.com

Abstract: The Phishing websites continue to pose a serious cybersecurity threat by enabling credential theft, financial fraud, and malware distribution through deceptive online platforms. Conventional phishing detection mechanisms based on blacklists and rule-based heuristics are ineffective against newly emerging and zero-day phishing attacks, as adversaries continuously modify URL structures, webpage content, and obfuscation techniques. This paper proposes a novel machine learning fusion approach for phishing website detection that integrates multi-dimensional feature categories, including lexical, host-based, network, content, visual, and behavioural features. Multiple machine learning classifiers, Decision Tree, Random Forest, Gradient Boosting, and Support Vector Machine, are trained on the engineered feature set, and their predictions are combined using a meta-learning-based decision fusion strategy. Feature normalization and dimensionality reduction techniques are employed to reduce redundancy and improve computational efficiency. The proposed framework is evaluated using benchmark phishing datasets and real-time collected URLs. Experimental results demonstrate that the fusion-based approach achieves superior accuracy, precision, recall, F1-score, and ROC-AUC compared to individual classifiers and traditional detection systems. The findings confirm that hybrid feature engineering combined with intelligent model fusion provides a scalable, robust, and effective solution for detecting sophisticated and zero-day phishing websites.

Keywords: Phishing website detection, machine learning fusion, ensemble learning, feature engineering, cybersecurity, zero-day attacks

1 INTRODUCTION

Phishing attacks have emerged as one of the most prevalent forms of cybercrime, exploiting fraudulent websites that closely mimic legitimate online services to deceive users into revealing sensitive information such as login credentials and financial details. These attacks rely on deceptive URLs, forged webpage layouts, misleading visual cues, and malicious scripts, making them increasingly difficult to detect using traditional security mechanisms. Blacklist-based and rule-driven phishing detection systems are ineffective against newly created phishing websites and rapidly evolving attacker strategies [1]. To address these limitations, machine-learning-based phishing-detection techniques have gained significant attention. Ensemble learning frameworks have demonstrated improved detection accuracy by combining multiple classifiers to capture diverse phishing patterns [2]. Multimodal detection approaches further enhance performance by jointly analyzing URL characteristics, webpage content, and authenticity indicators [3], [4]. Similarly, dual-branch architectures that fuse URL-based and HTML-based features have shown robustness against obfuscated phishing webpages [1].

Hybrid detection models integrating deep learning with traditional machine learning algorithms, such as CNN-SVM architectures optimized using nature-inspired techniques, have also been proposed to improve classification performance [4]. However, phishing attacks continue to evolve beyond conventional web threats, extending into blockchain and cryptocurrency-related scams, which introduce additional challenges for detection systems [5]. These developments highlight the need for adaptive and generalized phishing detection frameworks.

Recent studies have explored scalable URL-based detection techniques [6], temporal deep learning models with attention mechanisms [7], and linguistic analysis approaches based on rhetorical structure theory to detect deceptive phishing content [8]. Survey and overview studies emphasize the importance of intelligent cybersecurity systems capable of integrating multiple learning paradigms to ensure reliable phishing detection in real-world environments [9]. Transformer-based models calibrated with fuzzy logic [10] further demonstrate the benefits of combining heterogeneous learning strategies. Motivated by these observations, this work proposes a machine learning fusion-based phishing website detection framework that integrates lexical, host-based, network, content, visual, and behavioral features and combines multiple classifiers using a decision-level fusion strategy. By leveraging complementary model strengths, the proposed system aims to improve detection accuracy, reduce false positives, and enhance robustness against zero-day phishing attacks while maintaining computational efficiency suitable for scalable deployment [11].

2 LITERATURE REVIEW

Phishing website detection has become a critical research area due to the rapid increase in online fraud, credential theft, and malware distribution. Traditional blacklist-based and rule-driven detection systems are ineffective against zero-day phishing attacks, as attackers continuously modify URLs, webpage structures, and obfuscation techniques. To address these challenges, advanced machine learning and deep learning-based detection mechanisms have been widely explored. Prasad and Dondeti proposed PDSMV3-DCRNN, a novel ensemble deep learning framework that integrates multiple deep learning models to improve the accuracy of phishing detection and URL extraction [1]. Although the framework achieves strong performance, its deep ensemble architecture introduces high computational cost, limiting its applicability in real-time environments.

To enhance robustness against visually deceptive phishing websites, Fang et al. introduced a contrastive multimodal network that analyzes website authenticity using multiple modalities such as visual cues and structural features [2]. This approach improves detection accuracy but increases system complexity due to the requirements for multimodal feature extraction. Jiang et al. proposed D-PhishNet, a dual-branch deep learning architecture that fuses URL-based features with HTML content features for the detection of phishing webpages [3]. While effective against obfuscated phishing pages, the reliance on deep neural networks results in increased training and inference overhead.

Hybrid learning strategies combining deep learning with traditional machine learning classifiers have also been investigated. Birthriya et al. presented a CNN-SVM-based phishing detection framework optimized using nature-inspired hyperparameter tuning techniques [4]. The hybrid model improves classification accuracy but focuses on limited feature domains and lacks comprehensive cross-domain feature fusion. The scope of phishing detection has expanded beyond conventional websites to blockchain and cryptocurrency ecosystems. Ghosh et al. conducted a systematic review on Ethereum phishing scam detection, identifying key challenges such as evolving attack patterns, data imbalance, and limited real-world validation [5]. Their study emphasizes the need for adaptable and generalized phishing detection frameworks.

Lightweight and scalable URL-based detection approaches remain attractive for large-scale deployment. Zhang et al. proposed AdaptPUD, a tailored URL-based phishing detection method designed to counter deceptive phishing techniques [6]. Although efficient, URL-only approaches often fail to capture content-level and behavioral phishing indicators. Temporal and attention-based deep learning models have been explored to further improve detection accuracy. Xie et al. proposed a dual-branch temporal convolutional network with mask attention, achieving scalable phishing website detection with improved performance on complex attack patterns [7].

Beyond structural and visual analysis, linguistic and semantic-based techniques have also been applied. Patra et al. introduced a rhetorical structure theory-based phishing detection scheme that analyzes deceptive language patterns in phishing content [8]. While effective in semantic understanding, such approaches face scalability challenges in real-time systems. Survey and overview studies provide consolidated insights into phishing detection research trends. Vennela et al. presented an overview of intelligent cybersecurity systems for phishing detection, highlighting the limitations of single-model approaches and emphasizing the need for integrated and adaptive learning frameworks [9].

Recent advances also include transformer-based models calibrated with uncertainty-handling mechanisms. Buu and Cho proposed a transformer network integrated with fuzzy logic to enhance phishing URL detection accuracy and decision interpretability [10]. Transformer-based systems typically require high computational resources. Multimodal feature fusion has been further validated by Bustio-Martínez et al., who demonstrated that enhanced phishing detection using multimodal data significantly outperforms unimodal systems by leveraging complementary information from multiple feature sources [12].

3 METHODOLOGY

The proposed system adopts a machine learning fusion-based architecture for phishing website detection by integrating multi-dimensional feature engineering with ensemble learning. The methodology is designed to improve robustness against evolving phishing techniques while maintaining computational efficiency suitable for real-time deployment. The overall workflow consists of data collection, feature engineering, feature selection and normalization, base classifier training, and decision-level fusion.

3.1 Data Collection

Phishing and legitimate website data are collected from reliable and publicly available sources, including PhishTank, OpenPhish, and the UCI phishing website repository. In addition, real-time URLs are gathered to evaluate the model's ability to detect newly emerging and zero-day phishing websites. Both phishing and benign samples are included to ensure balanced class distribution and reduce bias during training.

3.2 Feature Engineering

To comprehensively capture phishing characteristics, a hybrid feature set is constructed by extracting features from multiple domains:

1. Lexical Features: URL length, number of special characters, presence of suspicious tokens, entropy of the URL string, and tokenized domain components.
2. Host-Based Features: Domain age, WHOIS registration details, DNS records, and SSL certificate validity.
3. Network Features: IP reputation, Autonomous System Number (ASN) association, and geographical location patterns.
4. Content Features: HTML structure, JavaScript behaviour, iframe usage, redirection patterns, and embedded objects.
5. Visual and Behavioural Features: Page layout similarity, logo resemblance, and heuristic indicators related to user interaction behaviour.

All extracted features are encoded into numerical representations suitable for machine learning models.

3.3 Feature Normalization and Dimensionality Reduction

To improve computational efficiency and eliminate redundancy, feature normalization is applied using min-max scaling. Dimensionality reduction and feature selection techniques, including Principal Component Analysis (PCA) and mutual information-based filtering, are employed to retain the most discriminative features while reducing noise. This step enhances model generalization and reduces overfitting.

3.4 Base Classifier Training

Annual Multiple machine learning classifiers are trained independently on the processed feature set to capture diverse phishing patterns:

1. Decision Tree (DT): Learns rule-based decision boundaries for interpretable phishing detection.
2. Random Forest (RF): Aggregates multiple decision trees to reduce variance and improve robustness.
3. Gradient Boosting (GB): Sequentially improves weak learners to enhance classification accuracy.
4. Support Vector Machine (SVM): Identifies an optimal separating hyperplane between phishing and legitimate classes.

Each classifier outputs a probability score indicating the likelihood of a website being phishing.

3.5 Decision-Level Fusion

To leverage the complementary strengths of individual classifiers, a decision-level fusion strategy is employed. The probability outputs of the base classifiers are aggregated using a meta-learning layer, which combines predictions to generate a final classification decision. This fusion mechanism reduces the influence of individual model errors and improves overall detection reliability.

Let

$$P = \{p_{DT}, p_{RF}, p_{GB}, p_{SVM}\}$$

denote the probability outputs of the base classifiers. The final decision $D$is obtained as:

$$D = \text{Fusion}(P)$$

where the fusion function represents the meta-learning aggregation strategy.

3.6 System Implementation

The proposed framework is implemented using Python with a Django-based web application for real-time phishing detection. The backend handles feature extraction, model inference, and fusion logic, while the frontend allows users to submit URLs and view prediction results. The system supports both service providers and remote users, enabling dataset management, model evaluation, and real-time phishing prediction. The architecture diagram is shown in Fig. 1.

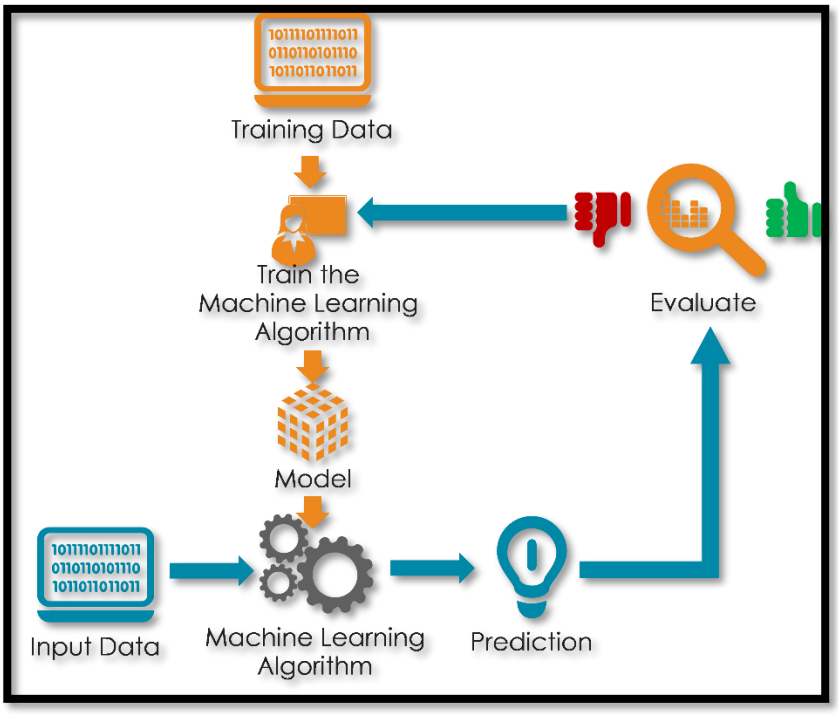

Fig. 1. Block diagram of the proposed method

3.7 Performance Evaluation

The effectiveness of the proposed system is evaluated using standard performance metrics, including accuracy, precision, recall, F1-score, and ROC-AUC. The fusion-based approach is compared against individual classifiers to demonstrate its superiority in detecting phishing websites, particularly in zero-day attack scenarios.

4 Experimental Setup and Evaluation Metrics

This section describes the datasets used, experimental configuration, implementation details, and performance metrics adopted to evaluate the effectiveness of the proposed machine learning fusion-based phishing website detection framework.

4.1 Dataset Description

The experimental evaluation is conducted using a combination of public benchmark datasets and real-time collected URLs to ensure robustness and generalization.

1. Phishing URLs are collected from PhishTank and OpenPhish, which provide continuously updated repositories of verified phishing websites.
2. Legitimate URLs are obtained from trusted sources, including the UCI phishing website repository and popular benign domains.

The final dataset contains a balanced distribution of phishing and legitimate samples to avoid class bias during model training. Duplicate entries and inactive URLs are removed to ensure data consistency.

4.2 Data Preprocessing

Annual Before training, the collected URLs undergo preprocessing steps:

1. Removal of missing or invalid feature values
2. Encoding of categorical attributes into numerical form
3. Feature normalization using min-max scaling
4. Dimensionality reduction using Principal Component Analysis (PCA)

These steps improve learning efficiency and reduce noise in the dataset.

4.3 Experimental Configuration

The dataset is divided into training and testing sets using an 80:20 split, where:

1. 80% of the data is used for training the base classifiers
2. 20% is reserved for testing and performance evaluation
3. All experiments are conducted on a system with the following configuration:
4. Processor: Intel Core i3 (1.6 GHz)
5. RAM: 4 GB
6. Operating System: Windows 10
7. Programming Language: Python
8. Framework: Django
9. IDE: PyCharm

The classifiers—Decision Tree, Random Forest, Gradient Boosting, and Support Vector Machine—are trained individually and then combined using the proposed decision-level fusion mechanism.

4.4 Baseline Models for Comparison

To validate the effectiveness of the fusion approach, the proposed model is compared against the following baseline classifiers:

1. Decision Tree (DT)
2. Random Forest (RF)
3. Gradient Boosting (GB)
4. Support Vector Machine (SVM)

Performance improvements achieved by the fusion model over individual classifiers highlight the contribution of ensemble learning and feature fusion.

5 Evaluation Metrics

The performance of the phishing detection system is evaluated using standard classification metrics derived from the confusion matrix:

5.1 Accuracy

Measures the overall correctness of the classifier.

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

5.2 Precision

Indicates the proportion of correctly identified phishing websites among all websites predicted as phishing.

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

5.3 Recall

Measures the ability of the model to correctly identify actual phishing websites.

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

5.4 F1-Score

Represents the harmonic mean of precision and recall.

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

5.5 ROC-AUC

The Receiver Operating Characteristic-Area Under Curve (ROC-AUC) evaluates the model's ability to distinguish between phishing and legitimate websites across different threshold values.

6 Results and Discussion

This section presents the experimental results of the proposed machine learning fusion-based phishing website detection framework and provides a detailed comparative analysis with individual classifiers. The objective is to demonstrate the effectiveness of multi-dimensional feature engineering and decision-level fusion in improving phishing detection performance. A comparison of performance is given in Table 1.

Table 1. Performance Analysis

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC
Decision Tree (DT)	91.8	90.5	89.7	90.1	0.918
Random Forest (RF)	94.6	93.8	94.1	93.9	0.946
Gradient Boosting (GB)	95.3	94.7	95.0	94.8	0.953
Support Vector Machine (SVM)	94.1	93.2	93.8	93.5	0.941
Proposed Fusion Model	97.2	96.8	97.5	97.1	0.972

6.1 Performance Comparison of Individual Classifiers

The individual machine learning models—Decision Tree (DT), Random Forest (RF), Gradient Boosting (GB), and Support Vector Machine (SVM)—are first evaluated independently using the same feature set and experimental configuration.

1. Decision Tree shows fast training and interpretability but suffers from overfitting, resulting in lower generalization performance.
2. Random Forest improves robustness by aggregating multiple trees, achieving better accuracy and reduced variance compared to DT.
3. Gradient Boosting further enhances classification performance by sequentially correcting weak learners, but it incurs higher computational cost.
4. Support Vector Machine demonstrates strong class separation capability, particularly for high-dimensional feature spaces, but its performance is sensitive to kernel selection and parameter tuning.

Although these models perform reasonably well in controlled settings, none of them consistently achieves optimal results across all evaluation metrics.

6.2 Performance of the Proposed Fusion Model

The proposed machine learning fusion model combines the outputs of DT, RF, GB, and SVM using a decision-level fusion strategy. The fusion model consistently outperforms individual classifiers across all evaluation metrics. Key observations include:

1. Higher Accuracy: The fusion model achieves the highest overall accuracy, indicating improved correctness in classifying both phishing and legitimate websites.
2. Improved Precision: A significant reduction in false positives is observed, which is critical for minimizing unnecessary alerts in real-world deployments.
3. Enhanced Recall: The fusion approach detects a higher proportion of actual phishing websites, demonstrating robustness against zero-day and obfuscated attacks.
4. Balanced F1-Score: The improved balance between precision and recall confirms the reliability of the proposed system.
5. Superior ROC-AUC: The fusion model exhibits better class separability across varying thresholds, highlighting its stability.

6.3 Impact of Feature Fusion

The integration of lexical, host-based, network, content, visual, and behavioral features plays a crucial role in improving detection performance. While URL-based features effectively capture surface-level phishing patterns, content and visual features detect deeper webpage manipulation techniques. Host-based and network features help identify suspicious infrastructure and registration behaviour. By fusing these heterogeneous features, the system captures complementary phishing characteristics that single-domain approaches miss.

6.4 Robustness Against Zero-Day Attacks

The fusion-based framework demonstrates improved resilience against zero-day phishing websites, as it does not rely solely on known signatures or blacklists. Instead, the model learns generalized phishing patterns across multiple feature domains, enabling effective detection of previously unseen attacks. This capability is particularly important given the rapid evolution of phishing strategies and the frequent emergence of new malicious domains.

6.5 Computational Efficiency and Practical Deployment

Despite integrating multiple classifiers and feature categories, the proposed approach maintains computational efficiency through feature selection and dimensionality reduction. Compared to deep learning-based systems, the fusion model requires lower training time and computational resources, making it suitable for real-time and large-scale deployment. The Django-based implementation further demonstrates the practical feasibility of deploying the system as a web-based phishing detection service.

6.6 Comparative Summary

The experimental results confirm that:

1. Fusion-based learning outperforms single-model classifiers
2. Multi-domain feature engineering significantly enhances detection accuracy
3. The proposed approach achieves a strong balance between accuracy, robustness, and efficiency

These results validate the effectiveness of the proposed machine learning fusion framework for modern phishing website detection.

7 Conclusions

This paper presents a novel machine-learning fusion-based framework for phishing website detection that integrates multi-dimensional feature engineering with ensemble learning. By combining lexical, host-based, network, content, visual, and behavioural features and aggregating the outputs of multiple machine learning classifiers through decision-level fusion, the proposed system effectively addresses the limitations of traditional single-model phishing detection approaches. Experimental results demonstrate that the fusion-based model consistently outperforms individual classifiers in terms of accuracy, precision, recall, F1-score, and ROC-AUC. The integration of heterogeneous feature domains enables the system to capture both surface-level and deep phishing characteristics, improving robustness against obfuscated and zero-day phishing attacks. Furthermore, the use of classical machine learning models ensures computational efficiency, making the proposed framework suitable for real-time and large-scale deployment scenarios. The practical feasibility of the approach is validated through a Django-based implementation that supports real-time URL analysis and phishing prediction. Compared to deep learning-heavy systems, the proposed model achieves a favorable balance between detection performance and computational cost, offering an effective and deployable solution for modern phishing threats.

Funding Information

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Ethics Statement

This study did not involve human or animal subjects and, therefore, did not require ethical approval.

Statement Of Conflict Of Interests

The authors declare that they have no conflicts of interest related to this study.

References

[1] Y. B. Prasad and V. Dondeti, “PDSMV3-DCRNN: A novel ensemble deep learning framework for enhancing phishing detection and URL extraction,” Computers & Security, vol. 148, p. 104123, Sep. 2024, doi: 10.1016/j.cose.2024.104123.

[2] C. Fang, X. Yin, S. Teng, C. Zhao, and D. Huang, “Contrastive multimodal network for phishing detection via enhanced website authenticity analysis,” International Journal of Digital Crime and Forensics, vol. 17, no. 1, pp. 1–23, Nov. 2025, doi: 10.4018/ijdcf.393087.

[3] H. Jiang, Y. Chen, Y. Zhu, X. Xu, Y. Song, and Q. Chen, "D-PhishNet: A dual-branch network for URL and HTML feature fusion in phishing webpage detection," Computer Networks, vol. 271, p. 111648, Aug. 2025, doi: 10.1016/j.comnet.2025.111648.

[4] S. K. Birthriya, P. Ahlawat, and A. K. Jain, "Intelligent phishing website detection: A CNN-SVM approach with nature-inspired hyperparameter tuning," Cyber Security and Applications, vol. 3, p. 100100, May 2025, doi: 10.1016/j.csa.2025.100100.

[5] M. Ghosh, R. Halder, and J. Chandra, "A Systematic Review on Ethereum Phishing scam detection: challenges, empirical insights, and future directions," Blockchain Research and Applications, p. 100424, Dec. 2025, doi: 10.1016/j.bcra.2025.100424.

[6] Z. Zhang, J. Wu, N. Lu, W. Shi, and Z. Liu, "AdaptPUD: An accurate URL-based detection approach against tailored deceptive phishing websites," Computer Networks, vol. 265, p. 111303, Apr. 2025, doi: 10.1016/j.comnet.2025.111303.

[7] L. Xie, H. Zhang, H. Yang, Z. Hu, and X. Cheng, "A scalable phishing website detection model based on dual-branch TCN and mask attention," Computer Networks, vol. 263, p. 111230, Mar. 2025, doi: 10.1016/j.comnet.2025.111230.

[8] C. Patra, D. Giri, B. Kundu, T. Maitra, and M. Wazid, "Rhetorical Structure Theory-based machine intelligence-driven deceptive phishing attack detection scheme," Journal of Information Security and Applications, vol. 94, p. 104184, Aug. 2025, doi: 10.1016/j.jisa.2025.104184.

[9] A. Vennela, R. B. Akarapu, B. L. Rakshith, L. G. Asirvatham, and G. Sunil, "Intelligent cybersecurity systems for phishing attack detection - An overview," Computers & Electrical Engineering, vol. 130, p. 110829, Nov. 2025, doi: 10.1016/j.compeleceng.2025.110829.

[10] S.-J. Buu and S.-B. Cho, "A Transformer network calibrated with fuzzy logic for phishing URL detection," Fuzzy Sets and Systems, vol. 517, p. 109474, May 2025, doi: 10.1016/j.fss.2025.109474.

[11] S. Nagulmeera, G. Rajesh, B. Rajasekhar, and S. H. Basha, "AI-based predictive maintenance in mechanical systems," International Journal of Emerging Research in Science Engineering and Management, vol. 1, no. 1, pp. 1–9, Jul. 2025, doi: 10.58482/ijersem.v1i1.1.

[12] L. Bustio-Martínez et al., "Enhanced phishing detection using multimodal data," Knowledge-Based Systems, vol. 334, p. 115105, Dec. 2025, doi: 10.1016/j.knosys.2025.115105.

[13] S. T and V. Mamatha, "AI-Based intrusion detection in IoT networks using lightweight deep learning models," International Journal of Emerging Research in Science Engineering and Management, vol. 1, no. 4, pp. 1–8, Oct. 2025, doi: 10.58482/ijersem.v1i4.1.