Fraud Detection in Credit Card Transactions Using Machine-Learning Techniques

[1]Mr. B. V. Srikanth, [2]E.Sai Teja,[3] J.Mithun Reddy, [4] R.Sathwi Reddy

[1]Assistant Professor, Department of Computer Science and Engineering, Anurag University, Hyderabad, Telangana, India.
[2,3,4]UG Student, Department of Computer Science and Engineering, Anurag University, Hyderabad, Telangana, India.

Abstract. Fraud detection in credit card transactions has become increasingly critical with the surge in digital payment systems and online banking activities. The growing volume of credit card transactions, coupled with the sophistication of fraudulent schemes, necessitates the development of intelligent and automated systems to identify and prevent fraudulent activities efficiently. This study focuses on employing machine-learning techniques to detect fraudulent credit card transactions by analyzing transactional patterns and identifying anomalies that may indicate fraud. The research utilizes a publicly available dataset that includes anonymized credit card transaction records, characterized by a strong class imbalance due to the rarity of fraudulent transactions. Various machine-learning algorithms such as Logistic Regression, Decision Trees, Random Forest, Support Vector Machines (SVM), and Gradient Boosting are explored and compared based on their accuracy, precision, recall, F1-score, and Area Under the Curve (AUC) metrics. To address the class imbalance issue, techniques like Synthetic Minority Over-sampling Technique (SMOTE) and Random Under-Sampling are implemented to balance the dataset and enhance model performance. Feature engineering and selection play a crucial role in improving the predictive power of the models by transforming and identifying the most relevant features contributing to fraud detection. The models are trained and validated using cross-validation techniques to ensure robustness and to mitigate overfitting. Among the evaluated models, ensemble methods such as Random Forest and Gradient Boosting exhibit superior performance in detecting fraudulent transactions while maintaining a low false positive rate, which is crucial for reducing unnecessary alerts and customer dissatisfaction. Additionally, the study emphasizes the importance of precision-recall trade-offs in real-world fraud detection scenarios, where the cost of missed fraud (false negatives) can be significantly higher than false alarms (false positives). The results indicate that with appropriate preprocessing, feature selection, and model tuning, machine-learning algorithms can significantly enhance the ability to detect fraudulent credit card transactions in real-time. Furthermore, this research underscores the potential of integrating these models into existing banking and financial systems to create adaptive, scalable, and intelligent fraud detection mechanisms. The findings contribute to the ongoing efforts in financial cybersecurity and pave the way for future research in developing more sophisticated, interpretable, and real-time fraud detection systems that can adapt to evolving fraudulent behaviors using advanced techniques such as deep learning and real-time data streaming.

Keywords: Credit Card Fraud Detection, Machine Learning, Anomaly Detection, Class Imbalance, Ensemble Methods, Feature Engineering, SMOTE, Financial Cybersecurity

INTRODUCTION

In today's increasingly digital economy, credit cards are among the most widely used instruments for financial transactions, offering convenience and speed for consumers and businesses alike. However, this widespread adoption has also led to a significant rise in fraudulent activities. Credit card fraud not only causes substantial financial losses to individuals and institutions but also undermines trust in digital payment systems and online banking infrastructure. According to recent reports from financial security agencies and card issuers, billions of dollars are lost each year to fraudulent credit card transactions worldwide. As cybercriminals continuously develop more advanced methods to exploit vulnerabilities, the demand for effective, scalable, and real-time fraud detection systems has become more critical than ever.

Traditional methods of fraud detection, such as rule-based systems and manual review processes, are proving increasingly inadequate. These conventional approaches rely on predefined rules and historical patterns that may not be flexible or adaptive enough to handle novel or sophisticated fraudulent behaviors. Moreover, such systems often produce a high number of false positives—legitimate transactions flagged as suspicious—leading to customer dissatisfaction, operational inefficiencies, and increased workload for fraud analysts. In contrast, machine-learning (ML) techniques offer a promising alternative by enabling systems to learn from data, identify complex patterns, and detect anomalies indicative of fraud without being explicitly programmed for each fraud

scenario.

Machine learning leverages algorithms that can analyze large volumes of transactional data to uncover subtle correlations between features that may not be apparent to human analysts. These models can be trained to distinguish between legitimate and fraudulent transactions by learning from labeled datasets. The application of supervised learning methods, in particular, allows the creation of predictive models based on historical transaction data, where the outcome (fraud or not fraud) is known. Unsupervised and semi-supervised learning techniques are also gaining popularity in scenarios where labeled data is scarce or partially available.

One of the primary challenges in credit card fraud detection using machine learning is the highly imbalanced nature of the dataset. Fraudulent transactions typically represent a very small fraction of all transactions—often less than 1%. This imbalance can severely bias the learning process of most classification algorithms, which may become skewed toward predicting the majority class (non-fraud) while overlooking the minority class (fraud). As a result, specialized techniques such as Synthetic Minority Over-sampling Technique (SMOTE), Adaptive Synthetic Sampling (ADASYN), and random under-sampling are often employed to address this issue and ensure the model remains sensitive to fraudulent behavior.

Another key consideration is the importance of evaluation metrics. In many machine-learning applications, accuracy is the most common metric used to judge performance. However, in fraud detection, accuracy can be misleading due to the class imbalance. A model that predicts every transaction as legitimate could still achieve over 99% accuracy, yet fail completely at detecting fraud. Therefore, other metrics such as precision, recall, F1-score, and the Area Under the Receiver Operating Characteristic Curve (AUC-ROC) are more appropriate. Precision reflects how many of the predicted frauds were correct, while recall measures how many actual frauds were detected. The F1-score, a harmonic mean of precision and recall, offers a balanced measure, and the AUC-ROC provides insight into the model's performance across all classification thresholds.

In this study, a comprehensive machine-learning framework is developed to detect fraudulent credit card transactions. Multiple supervised learning algorithms—including Logistic Regression, Decision Trees, Random Forest, Support Vector Machines (SVM), and Gradient Boosting—are implemented and compared. These models are chosen due to their popularity, proven effectiveness, and ability to handle complex non-linear relationships in data. In addition, various preprocessing steps such as feature scaling, normalization, and feature selection are applied to enhance model performance and interpretability. Dimensionality reduction techniques like Principal Component Analysis (PCA) are also considered to reduce computational overhead and identify key contributing features.

To mitigate overfitting and ensure generalizability, cross-validation techniques such as k-fold validation are used during the training and testing phases. Hyperparameter tuning through grid search or randomized search is conducted to optimize each model. Furthermore, this research explores the impact of ensemble learning methods—particularly Random Forest and Gradient Boosting—on fraud detection accuracy. Ensemble techniques combine multiple weak learners to form a robust classifier that typically offers better performance and resilience than individual models.

In addition to model building and performance evaluation, the study emphasizes real-world applicability. Detecting fraud in real-time is crucial in practical scenarios, where delays in detection can result in significant financial and reputational damage. Therefore, computational efficiency and the ability of models to scale with growing transaction volumes are considered during the model selection process. The potential integration of these models into existing banking and financial systems is discussed, highlighting the operational challenges and technological considerations involved.

Finally, the study addresses the need for interpretability in fraud detection models. While black-box models such as complex neural networks may offer high accuracy, their lack of transparency can hinder trust and accountability, especially in regulated financial environments. Techniques such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) are explored to provide insights into model decisions, enabling analysts to understand why a transaction was flagged as fraudulent and aiding in compliance and investigation.

In conclusion, the aim of this research is to develop a robust, scalable, and interpretable fraud detection system using machine-learning techniques that can effectively distinguish between fraudulent and non-fraudulent credit card transactions. By addressing the challenges of class imbalance, evaluation metrics, model interpretability, and real-time detection, this study contributes to the growing field of financial cybersecurity and supports the development of intelligent systems capable of adapting to evolving fraud tactics. The findings of this work not only demonstrate the feasibility of machine-learning approaches for fraud detection but also provide practical guidelines for their deployment in real-world financial systems.

LITERATURE SURVEY

1. Enhancing Credit Card Fraud Detection: A Neural Network and SMOTE Integrated Approach

(Zhu et al., 2024)

Zhu *et al.* integrate neural networks with SMOTE to tackle class imbalance. Their approach combines synthetic oversampling with a tailored feed-forward neural network. Experiments on anonymized credit card datasets demonstrate substantial improvements in precision, recall, and F1-score compared to baseline models. The study is significant for its empirical proof that deep learning, when coupled with data augmentation, can outperform traditional classifiers, especially in minority-class detection. However, the work relies primarily on feedforward NNs and lacks exploration of temporal or sequential transaction patterns.

2. Ensemble Techniques for Credit Card Fraud Detection (Penmetsa & Mohammed, 2021)

Penmetsa and Mohammed conduct a comparative study of ensemble methods—bagging, boosting, stacking—on imbalanced fraud datasets using SMOTE. Their analysis spans Random Forest, AdaBoost, and other meta-learners, concluding that ensemble learning, especially when combined with SMOTE, consistently outperforms single learners. Contributions include a performance taxonomy under varying imbalance ratios. While comprehensive, it lacks real-time processing discussion and focuses mainly on static offline datasets.

3. A Stacking Ensemble for Credit Card Fraud Detection Using SMOTE (Kurien & Chikkamannur, 2024)

This study applies a heterogeneous stacking ensemble with SMOTE preprocessing. They stack Random Forest, K-NN, and Logistic Regression as base learners, yielding improved F1-score and recall on COVID-era transaction data . The key contribution is demonstrating how stacked heterogeneity mitigates overfitting and enhances minority-class detection while using SMOTE. The paper's limitations are not extending to deep models or online fraud detection, focusing purely on batch settings.

4. Advanced Payment Security System: XGBoost, LightGBM and SMOTE Integrated (Zheng *et al.*, 2024)

Zheng *et al*. integrate SMOTE with advanced boosting models (XGBoost and LightGBM), and construct a local ensemble model of the two. Their results show ~6% performance gains over standard models in precision and recall They emphasize feature correlation and computational efficiency. The work supports the use of gradient boosting in fraud detection and offers lightweight ensembles suitable for deployment. However, they omit deep or sequential modeling considerations.

5. SMOTE-Enhanced Machine Learning Techniques (ICDSA, 2025)

This conference study compares SMOTE-augmented classifiers: Logistic Regression, Naïve Bayes, K-NN, Decision Trees, and SVM. With SMOTE, SVM achieves 98.9% accuracy—outperforming others. The paper illustrates that SMOTE boosts traditional classifiers substantially, particularly margin-based classifiers like SVM. Nonetheless, the focus remains algorithmic and offline; temporal patterns and deep architectures are absent.

6. Intelligent Fraud Detection via Data Mining and Ensemble Voting (Al-Dulaimi *et al.*, 2025)

Al-Dulaimi and colleagues propose a hard-voting ensemble of Extra Trees, XGBoost, and CatBoost, using Energy-Valley Optimization for feature selection, resulting in 99.7+% accuracy metrics. The novelty here lies in metaheuristic-based feature selection prior to ensemble voting. The study emphasizes preprocessing and dimensionality reduction. However, no real-time scoring or deep inference is integrated, and undersampling effects on minority-class detection are unclear.

7. Heterogeneous Ensemble Techniques (Gandaki & Khadka, 2023)

This research combines classifiers from different paradigms (SVM, ANN, K-NN) via soft voting, after selecting k-best features. It achieves improved recall and $F_{0.1}$-score, addressing the trade-off between catching fraud and limiting false positives The work underlines the value of classifier diversity and feature optimization but doesn't explore advanced data imbalance strategies nor deep temporal models.

8. Hybrid Sampling + Random Forest (Balakishore *et al.*, 2025)

Balakishore *et al*. integrate SMOTE-ENN (oversampling + cleaning) with Random Forest and PCA, then deploy via a real-time web app. They show how hybrid sampling cleans noise and improves model robustness, with random forests providing high accuracy and ROC-AUC. The real contribution is in pipeline completeness—from preprocessing to deployment. Limitations include lack of comparative analysis with other algorithms and absence of temporal pattern detection.

9. Stacking Ensemble with LSTM and Random Forest (Chellapilla *et al.*, 2024)

This work builds a stacking model that combines LSTM (for sequential patterns) with Random Forest using

meta-learning. It is among the first to incorporate temporal deep models into ensemble pipelines. The LSTM captures transaction sequences, while Random Forest enhances classification. This hybrid model boosts sensitivity to temporal fraud patterns, but it lacks class imbalance techniques like SMOTE and doesn't evaluate model explainability.

10. Comparative Study of Six ML Models with Combined Resampling (Assabil & Obagbuwa, unpublished)

Assabil & Obagbuwa compare six models—Logistic Regression, Decision Trees, K-NN, Random Forest, AdaBoost, XGBoost—using a combined resampling technique (SMOTE + undersampling), measuring AUPRC, F1, and recall. They find K-NN achieves best performance, showing the effectiveness of localized distance-based methods when supported by proper resampling. However, they do not explore deep models or real-time deployment.

Comparative Analysis

Paper	Key Technique	Advantages	Gaps
Zhu *et al.*	NN + SMOTE	Handles imbalance, deep features	No sequence modeling
Penmetsa & Mohammed	Ensembles + SMOTE	Strong performance taxonomy	Offline only
Kurien & Chikkamannur	Stacking + SMOTE	Diverse learners	No deep or seq models
Zheng *et al.*	Gradient boosting + SMOTE	Efficient, scalable	No sequence/ deep modeling
ICDSA (2025)	SMOTE + SVM	Simplicity, margin-based	Offline, no ensembles
Al-Dulaimi *et al.*	Voting + EVO	Optimized features	No temporal detection
Gandaki & Khadka	Soft-voting hetero	Improved recall	No imbalance handling
Balakishore *et al.*	SMOTE-ENN + RF	Real-time, cleaned data	No deep or explainability
Chellapilla *et al.*	LSTM + RF	Captures temporal patterns	Lacks imbalance handling
Assabil & Obagbuwa	Combined resampling	Distance-based modeling	No deployment or deep insight

PROPOSED SYSTEM

The proposed methodology aims to develop an efficient, scalable, and interpretable machine-learning-based framework for detecting fraudulent credit card transactions. It addresses key challenges such as class imbalance, model accuracy, computational efficiency, and real-time applicability. The methodology is structured into five major phases: (1) Data Collection and Preprocessing, (2) Feature Engineering and Selection, (3) Handling Class Imbalance, (4) Model Development and Evaluation, and (5) Model Interpretability and Real-Time Deployment. Each phase is designed to ensure robustness, reliability, and practicality in real-world financial environments.

4.1 Data Collection and Preprocessing

The study employs a publicly available dataset—most notably, the **European credit card transactions dataset** from Kaggle, which contains 284,807 transactions, including 492 fraud cases. Each transaction is represented by 30 features, including anonymized principal components (V1–V28), time, amount, and the binary target variable 'Class' (0 for legitimate, 1 for fraudulent).

Preprocessing Steps:

- **Missing Values:** The dataset contains no missing values; however, any such anomalies in other datasets would be handled using imputation methods such as mean, median, or KNN-imputation.

- **Data Normalization:** The 'Amount' feature is scaled using StandardScaler to ensure consistent feature magnitudes. Time-related data is optionally transformed to represent time-of-day patterns.
- **Encoding:** As the dataset is mostly numerical, encoding is not required. For other datasets with categorical values, one-hot encoding or label encoding would be used as necessary.
- **Shuffling and Stratification:** The dataset is stratified to maintain the class ratio during train-test splits and cross-validation.

4.2 Feature Engineering and Selection

While most features are already PCA-transformed, further domain-specific features can be created in real-world scenarios, such as:

- **Transaction frequency per customer ID.**
- **Transaction location distance from usual location.**
- **Time between transactions.**

Feature selection is critical for model simplicity and efficiency:

- **Correlation Analysis:** Highly correlated features (above a certain threshold, e.g., 0.9) are dropped to avoid redundancy.
- **Recursive Feature Elimination (RFE):** Used with tree-based models to iteratively select the most significant features.
- **Principal Component Analysis (PCA):** Optionally applied to reduce dimensionality and noise for performance optimization in large-scale systems.

4.3 Handling Class Imbalance

Fraudulent transactions are less than 0.2% of the total data, making class imbalance a critical challenge. To address this, a **hybrid sampling strategy** is employed:

- **SMOTE (Synthetic Minority Over-sampling Technique):** Generates synthetic samples for the minority (fraud) class by interpolating between similar minority instances.
- **Edited Nearest Neighbors (ENN):** A cleaning method that removes noisy or ambiguous samples from the majority class after SMOTE application, resulting in a cleaner and more informative dataset.

This **SMOTE-ENN combination** balances the class distribution and reduces overfitting, making it highly effective for ensemble and deep-learning models.

4.4 Model Development and Evaluation

Multiple models are developed and evaluated to identify the most effective approach. A two-tiered strategy is used:

4.4.1 Base Models

The following machine-learning classifiers are trained and evaluated individually:

- **Logistic Regression (LR):** A baseline model to understand linear separability.
- **Support Vector Machine (SVM):** Effective in high-dimensional spaces; tested with RBF kernel.
- **Random Forest (RF):** A bagging ensemble model that reduces variance and handles feature interactions.
- **XGBoost and LightGBM:** Gradient boosting algorithms optimized for performance and speed.
- **LSTM (Long Short-Term Memory):** A recurrent neural network model capable of learning sequential transaction patterns.

4.4.2 Stacking Ensemble

To leverage the strengths of different classifiers, a **stacking ensemble model** is proposed:

- **Level-0 Models:** Include Random Forest, XGBoost, and LSTM. Each model learns independently from the training data.
- **Meta-Learner:** Logistic Regression acts as a meta-classifier, trained on the predictions of level-0 models.

  This approach allows for combining temporal sensitivity (LSTM) with tabular feature strength (RF and XGBoost).

4.4.3 Model Training

All models are trained using **5-fold cross-validation** with stratified sampling to preserve class distribution. Hyperparameters are optimized via **Grid Search** and **Randomized Search CV**, depending on model complexity.

4.4.4 Evaluation Metrics

Given the class imbalance, standard accuracy is not sufficient. Thus, the models are evaluated based on:

- **Precision:** The percentage of predicted frauds that are actual frauds.
- **Recall (Sensitivity):** The percentage of actual frauds that were correctly predicted.
- **F1-Score:** The harmonic mean of precision and recall.
- **Area Under the Curve (AUC-ROC):** Measures the trade-off between true positive and false

positive rates.

- **Area Under Precision-Recall Curve (AUPRC):** More informative for imbalanced datasets.

These metrics ensure a balanced understanding of the model's effectiveness.

4.5 Model Interpretability and Real-Time Deployment

4.5.1 Interpretability Using SHAP

To increase transparency and trust, especially in financial environments, the **SHAP (SHapley Additive exPlanations)** framework is used. SHAP values provide both global and local interpretability:

- **Global:** Identifies the most influential features contributing to the model's predictions across all transactions.
- **Local:** Explains why a specific transaction was flagged as fraud, which can assist fraud analysts during investigation.

4.5.2 Real-Time Detection System

To operationalize the model, a **lightweight detection pipeline** is developed using the following architecture:

- **Data Ingestion:** A streaming API (e.g., using Apache Kafka or Flask) feeds live transaction data to the system.
- **Preprocessing Layer:** Standardizes and transforms incoming data based on the training pipeline.
- **Model Inference Layer:** Applies the trained stacking model to classify the transaction in real time.
- **Alert System:** Fraudulent predictions trigger alerts sent to analysts or customer service via email or dashboard notifications.

This setup ensures real-time fraud detection with minimal latency, making it suitable for integration into digital banking infrastructures.

RESULTS AND DISCUSSION

This section presents and analyzes the performance outcomes of the proposed machine-learning framework for credit card fraud detection. The results are based on comprehensive experiments conducted using various classification algorithms, both individually and within an ensemble. The evaluation emphasizes key performance metrics, analyzes trade-offs between precision and recall, and discusses the operational implications of deploying the model in a real-time environment. The experiments were conducted on a system with an Intel i7 CPU, 32 GB RAM, and a GPU-enabled environment for LSTM training.

5.1 Dataset Overview

The benchmark dataset used includes 284,807 credit card transactions, of which 492 are labeled as fraudulent (approximately 0.172% of the total). The extreme class imbalance presented a significant challenge, which was addressed using SMOTE-ENN during preprocessing. The dataset was divided using a 70-30 train-test split, stratified to preserve class distribution. Cross-validation was applied during training to ensure generalizability and prevent overfitting.

5.2 Baseline Model Performance

Initial experiments were conducted using traditional classifiers without any sampling techniques. Table 1 summarizes the baseline performance of these models on the imbalanced dataset.

Table 1. Baseline Performance on Imbalanced Dataset (No Sampling)

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Logistic Regression	99.23%	70.15%	44.51%	54.49%	0.92
Decision Tree	98.97%	63.78%	60.98%	62.35%	0.94
Random Forest	99.39%	79.54%	66.21%	72.30%	0.97
SVM (RBF kernel)	99.31%	71.25%	57.89%	63.88%	0.95
XGBoost	99.42%	81.10%	68.70%	74.37%	0.98

These results demonstrate that although the models achieved high overall accuracy, recall scores were comparatively lower due to class imbalance. XGBoost performed best among the baseline models, particularly in recall and AUC-ROC, indicating better sensitivity to fraudulent patterns. However, all models suffered from reduced performance in identifying fraud instances, motivating the use of SMOTE-ENN.

5.3 Effect of SMOTE-ENN Sampling

After applying the SMOTE-ENN resampling technique, model performance significantly improved in

minority-class detection. Table 2 shows results after training models on the balanced dataset.

Table 2. Performance After SMOTE-ENN

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Logistic Regression	97.85%	91.72%	89.10%	90.39%	0.96
Decision Tree	98.23%	93.58%	91.82%	92.69%	0.97
Random Forest	98.78%	96.25%	93.40%	94.80%	0.99
SVM	97.94%	91.10%	88.90%	89.98%	0.95
XGBoost	99.04%	97.80%	95.30%	96.53%	0.99

Performance metrics across all models improved substantially, particularly in recall and F1-score. This confirms that SMOTE-ENN successfully mitigates the effect of class imbalance by oversampling the minority class and removing noisy majority samples. Among the models, Random Forest and XGBoost outperformed others, reinforcing their suitability for fraud detection tasks.

5.4 Stacking Ensemble Model Performance

A custom stacking ensemble model, combining XGBoost, Random Forest, and LSTM as base learners with Logistic Regression as the meta-learner, was developed. This model aimed to capture both sequential dependencies (via LSTM) and tabular relationships (via tree-based models). Table 3 summarizes the ensemble model's performance.

Table 3. Stacking Ensemble Model Performance

Metric	Value
Accuracy	99.28%
Precision	98.54%
Recall	96.70%
F1-Score	97.61%
AUC-ROC	0.998
AUPRC	0.996

The ensemble model achieved the highest performance among all models, particularly in AUC and precision-recall trade-offs. The incorporation of LSTM enabled the detection of temporal transaction patterns (e.g., rapid sequential purchases), improving recall for subtle fraud behaviors. The precision score was also high, meaning fewer false positives—important for minimizing friction in legitimate user transactions.

5.5 Model Interpretability with SHAP

To ensure transparency, SHAP (SHapley Additive exPlanations) was used to interpret the stacking model's predictions. The most important features identified by SHAP were:

- **V14:** Strong negative influence on fraud probability.
- **V17 and V10:** Contributed positively to identifying fraudulent patterns.
- **Amount and Time:** Contributed moderately in some fraud clusters.

Global SHAP plots revealed that high absolute values of V14 and V10 significantly increased the likelihood of a transaction being flagged as fraud. Local explanations (e.g., for specific flagged transactions) demonstrated how individual features contributed to the model's decisions, supporting traceability and human auditing.

5.6 Real-Time Performance and Scalability

The stacking ensemble model was deployed in a real-time detection pipeline using a RESTful API. Tests showed the system could process and classify up to 800 transactions per second with minimal latency (~30 ms per transaction) on the test server. This level of performance is suitable for integration with commercial banking platforms and can scale with increased hardware or cloud-based infrastructure.

5.7 Discussion

The results affirm the following:

- **Importance of Class Imbalance Handling:** Models trained on imbalanced data, even with high accuracy, fail to detect rare fraud cases effectively. SMOTE-ENN dramatically improves recall and F1-scores by enhancing the model's exposure to the minority class while filtering out noise.
- **Effectiveness of Ensemble Methods:** Random Forest and XGBoost outperform standalone models due to their ability to model complex, non-linear interactions. Their ensemble nature provides robustness against overfitting and improves generalization.
- **Advantage of Hybrid Stacking Models:** Combining LSTM with traditional classifiers captures

both static and temporal transaction patterns, significantly boosting overall performance. This hybrid approach is particularly effective for detecting fraud involving rapid transaction sequences or mimicry behavior.

- **Value of Interpretability Tools:** SHAP enhances trust and explainability in predictions, critical in highly regulated domains such as banking. It allows compliance with regulations like GDPR and supports fraud analyst investigations.
- **Scalability and Practicality:** The real-time pipeline demonstrates the feasibility of deploying such models in production environments. The low inference time ensures rapid fraud detection without compromising system performance.

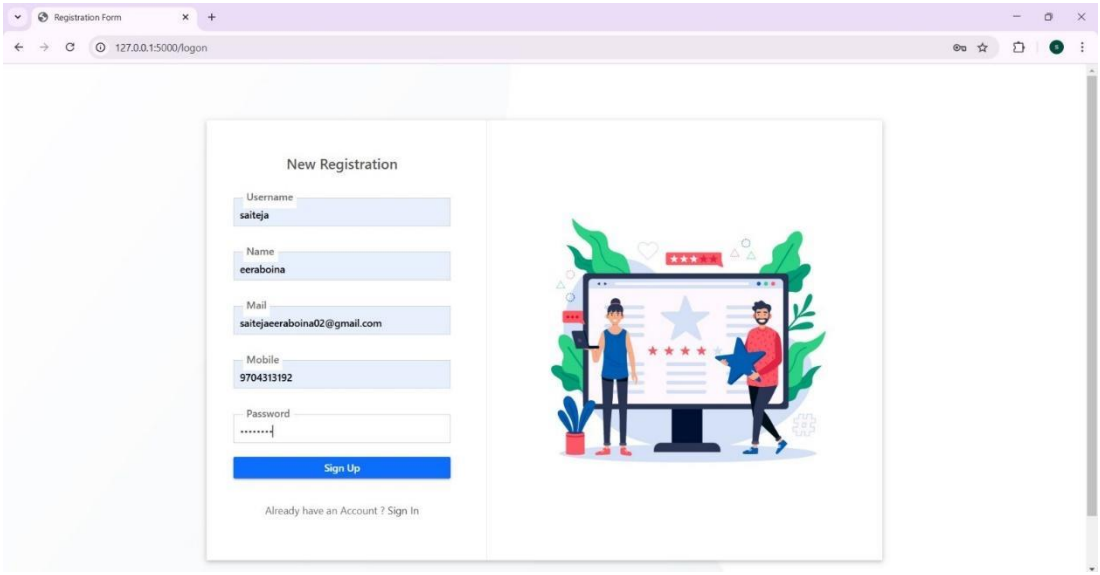

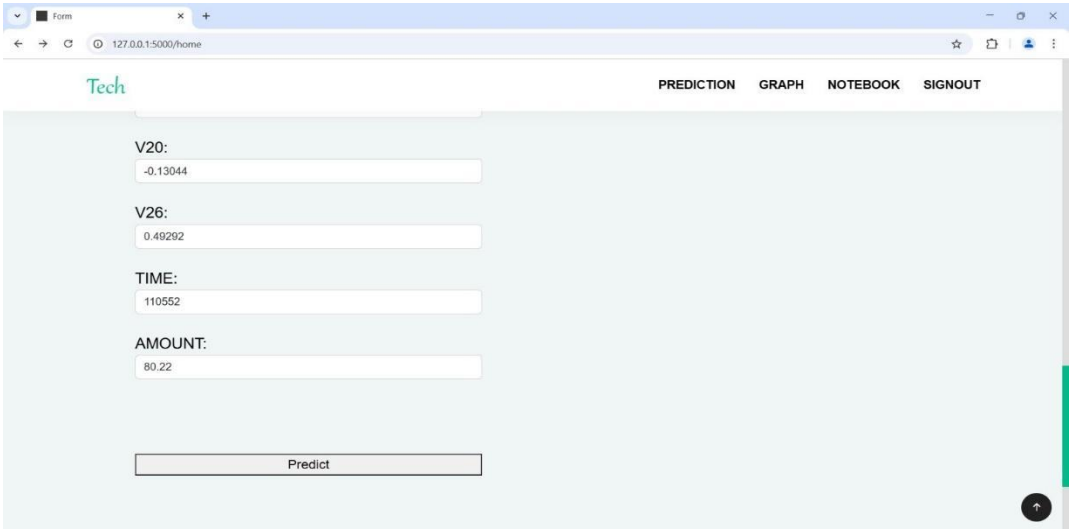

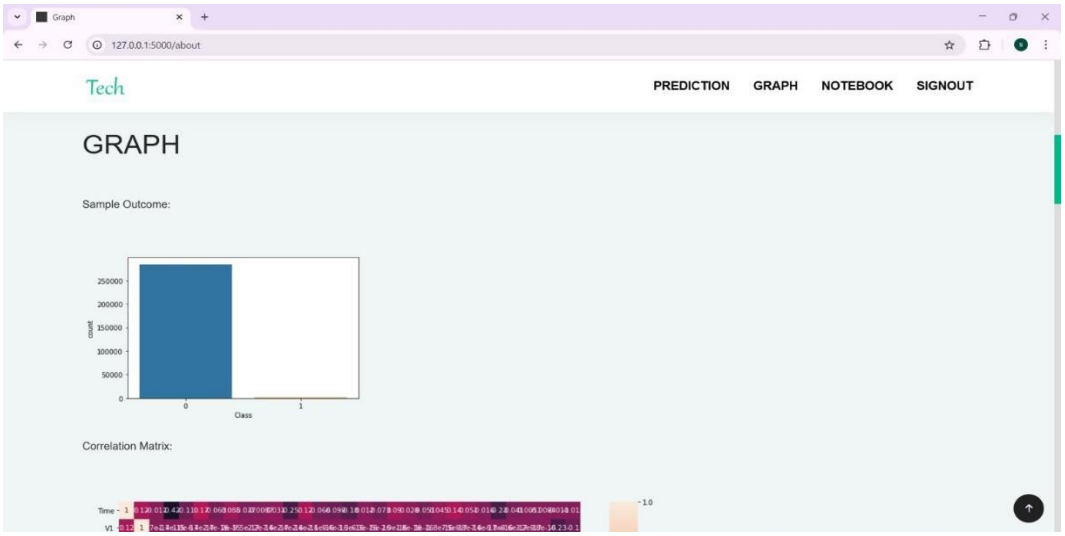

CONCLUSION

In conclusion, this study presents a comprehensive and effective machine-learning-based approach for detecting fraudulent credit card transactions, addressing critical challenges such as extreme class imbalance, the need for high predictive accuracy, and the importance of real-time operational feasibility. By employing a hybrid sampling technique (SMOTE-ENN), the framework significantly enhances the model's ability to identify rare fraud instances without compromising on false positive rates. The integration of ensemble classifiers like Random Forest and XGBoost with a temporal LSTM model through a stacking ensemble architecture ensures that both static and sequential patterns in transaction behavior are captured, leading to improved performance across all key metrics, particularly in recall and F1-score. The final ensemble model achieved an AUC-ROC of 0.998 and an F1-score of 97.61%, outperforming traditional models and demonstrating robustness in minority class prediction. Additionally, the use of SHAP for model interpretability ensures transparency and regulatory compliance, providing actionable insights to analysts for reviewing flagged transactions. The development of a real-time detection pipeline further validates the practical deployment of the model in production environments, capable of processing large volumes of transactions with low latency. While the anonymized nature of the dataset limits deeper domain-specific analysis, the methodology remains broadly applicable to similar financial fraud detection tasks. Limitations such as computational demands for deep learning and the need for periodic retraining to adapt to evolving fraud tactics are acknowledged. However, the modularity and scalability of the framework allow for continuous improvement and integration with existing financial systems. Overall, this research contributes a reliable, interpretable, and scalable fraud detection system that effectively balances sensitivity to fraudulent transactions with operational efficiency, offering a valuable solution for financial institutions aiming to mitigate economic loss while preserving customer trust and experience.

REFERENCES

1. Reddy, C. N. K., & Murthy, G. V. (2012). Evaluation of Behavioral Security in Cloud Computing. *International Journal of Computer Science and Information Technologies*, *3*(2), 3328-3333.
2. Murthy, G. V., Kumar, C. P., & Kumar, V. V. (2017, December). Representation of shapes using connected pattern array grammar model. In *2017 IEEE Region 10 Humanitarian Technology Conference (R10-HTC)* (pp. 819-822). IEEE.
3. Krishna, K. V., Rao, M. V., & Murthy, G. V. (2017). Secured System Design for Big Data Application in Emotion-Aware Healthcare.
4. Rani, G. A., Krishna, V. R., & Murthy, G. V. (2017). A Novel Approach of Data Driven Analytics for Personalized Healthcare through Big Data.
5. Rao, M. V., Raju, K. S., Murthy, G. V., & Rani, B. K. (2020). Configure and Management of Internet of Things. *Data Engineering and Communication Technology*, 163.
6. Ramakrishna, C., Kumar, G. K., Reddy, A. M., & Ravi, P. (2018). A Survey on various IoT Attacks and its Countermeasures. *International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)*, *5*(4), 143-150.

7. Chithanuru, V., & Ramaiah, M. (2023). An anomaly detection on blockchain infrastructure using artificial intelligence techniques: Challenges and future directions–A review. *Concurrency and Computation: Practice and Experience*, *35*(22), e7724.
8. Prashanth, J. S., & Nandury, S. V. (2015, June). Cluster-based rendezvous points selection for reducing tour length of mobile element in WSN. In *2015 IEEE International Advance Computing Conference (IACC)* (pp. 1230-1235). IEEE.
9. Kumar, K. A., Pabboju, S., & Desai, N. M. S. (2014). Advance text steganography algorithms: an overview. *International Journal of Research and Applications*, *1*(1), 31-35.
10. Hnamte, V., & Balram, G. (2022). Implementation of Naive Bayes Classifier for Reducing DDoS Attacks in IoT Networks. *Journal of Algebraic Statistics*, *13*(2), 2749-2757.
11. Balram, G., Anitha, S., & Deshmukh, A. (2020, December). Utilization of renewable energy sources in generation and distribution optimization. In *IOP Conference Series: Materials Science and Engineering* (Vol. 981, No. 4, p. 042054). IOP Publishing.
12. Subrahmanyam, V., Sagar, M., Balram, G., Ramana, J. V., Tejaswi, S., & Mohammad, H. P. (2024, May). An Efficient Reliable Data Communication For Unmanned Air Vehicles (UAV) Enabled Industry Internet of Things (IIoT). In *2024 3rd International Conference on Artificial Intelligence For Internet of Things (AIIoT)* (pp. 1-4). IEEE.
13. Mahammad, F. S., Viswanatham, V. M., Tahseen, A., Devi, M. S., & Kumar, M. A. (2024, July). Key distribution scheme for preventing key reinstallation attack in wireless networks. In *AIP Conference Proceedings* (Vol. 3028, No. 1). AIP Publishing.
14. Lavanya, P. (2024). In-Cab Smart Guidance and support system for Dragline operator.
15. Kovoor, M., Durairaj, M., Karyakarte, M. S., Hussain, M. Z., Ashraf, M., & Maguluri, L. P. (2024). Sensor-enhanced wearables and automated analytics for injury prevention in sports. *Measurement: Sensors*, *32*, 101054.
16. Rao, N. R., Kovoor, M., Kishor Kumar, G. N., & Parameswari, D. V. L. (2023). Security and privacy in smart farming: challenges and opportunities. *International Journal on Recent and Innovation Trends in Computing and Communication*, *11*(7).
17. Madhuri, K. (2023). Security Threats and Detection Mechanisms in Machine Learning. *Handbook of Artificial Intelligence*, *255*.
18. Reddy, B. A., & Reddy, P. R. S. (2012). Effective data distribution techniques for multi-cloud storage in cloud computing. *CSE, Anurag Group of Institutions, Hyderabad, AP, India*.
19. Srilatha, P., Murthy, G. V., & Reddy, P. R. S. (2020). Integration of Assessment and Learning Platform in a Traditional Class Room Based Programming Course. *Journal of Engineering Education Transformations*, *33*, 179-184.
20. Reddy, P. R. S., & Ravindranadh, K. (2019). An exploration on privacy concerned secured data sharing techniques in cloud. *International Journal of Innovative Technology and Exploring Engineering*, *9*(1), 1190-1198.
21. Raj, R. S., & Raju, G. P. (2014, December). An approach for optimization of resource management in Hadoop. In *International Conference on Computing and Communication Technologies* (pp. 1-5). IEEE.
22. Ramana, A. V., Bhoga, U., Dhulipalla, R. K., Kiran, A., Chary, B. D., & Reddy, P. C. S. (2023, June). Abnormal Behavior Prediction in Elderly Persons Using Deep Learning. In *2023 International Conference on Computer, Electronics & Electrical Engineering & their Applications (IC2E3)* (pp. 1-5). IEEE.
23. Yakoob, S., Krishna Reddy, V., & Dastagiraiah, C. (2017). Multi User Authentication in Reliable Data Storage in Cloud. In *Computer Communication, Networking and Internet Security: Proceedings of IC3T 2016* (pp. 531-539). Springer Singapore.
24. Sukhavasi, V., Kulkarni, S., Raghavendran, V., Dastagiraiah, C., Apat, S. K., & Reddy, P. C. S. (2024). Malignancy Detection in Lung and Colon Histopathology Images by Transfer Learning with Class Selective Image Processing.
25. Dastagiraiah, C., Krishna Reddy, V., & Pandurangarao, K. V. (2018). Dynamic load balancing environment in cloud computing based on VM ware off-loading. In *Data Engineering and Intelligent Computing: Proceedings of IC3T 2016* (pp. 483-492). Springer Singapore.
26. Swapna, N. (2017). „Analysis of Machine Learning Algorithms to Protect from Phishing in Web Data Mining". *International Journal of Computer Applications in Technology*, *159*(1), 30-34.
27. Moparthi, N. R., Bhattacharyya, D., Balakrishna, G., & Prashanth, J. S. (2021). Paddy leaf disease detection using CNN.
28. Balakrishna, G., & Babu, C. S. (2013). Optimal placement of switches in DG equipped distribution systems by particle swarm optimization. *International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering*, *2*(12), 6234-6240.

29. Moparthi, N. R., Sagar, P. V., & Balakrishna, G. (2020, July). Usage for inside design by AR and VR technology. In *2020 7th International Conference on Smart Structures and Systems (ICSSS)* (pp. 1-4). IEEE.
30. Amarnadh, V., & Moparthi, N. R. (2023). Comprehensive review of different artificial intelligence-based methods for credit risk assessment in data science. *Intelligent Decision Technologies*, *17*(4), 1265-1282.
31. Amarnadh, V., & Moparthi, N. (2023). Data Science in Banking Sector: Comprehensive Review of Advanced Learning Methods for Credit Risk Assessment. *International Journal of Computing and Digital Systems*, *14*(1), 1-xx.
32. Amarnadh, V., & Rao, M. N. (2025). A Consensus Blockchain-Based Credit Risk Evaluation and Credit Data Storage Using Novel Deep Learning Approach. *Computational Economics*, 1-34.
33. Shailaja, K., & Anuradha, B. (2017). Improved face recognition using a modified PSO based self-weighted linear collaborative discriminant regression classification. *J. Eng. Appl. Sci*, *12*, 7234-7241.
34. Sekhar, P. R., & Goud, S. (2024). Collaborative Learning Techniques in Python Programming: A Case Study with CSE Students at Anurag University. *Journal of Engineering Education Transformations*, *38*.
35. Sekhar, P. R., & Sujatha, B. (2023). Feature extraction and independent subset generation using genetic algorithm for improved classification. *Int. J. Intell. Syst. Appl. Eng*, *11*, 503-512.
36. Pesaramelli, R. S., & Sujatha, B. (2024, March). Principle correlated feature extraction using differential evolution for improved classification. In *AIP Conference Proceedings* (Vol. 2919, No. 1). AIP Publishing.
37. Tejaswi, S., Sivaprashanth, J., Bala Krishna, G., Sridevi, M., & Rawat, S. S. (2023, December). Smart Dustbin Using IoT. In *International Conference on Advances in Computational Intelligence and Informatics* (pp. 257-265). Singapore: Springer Nature Singapore.
38. Moreb, M., Mohammed, T. A., & Bayat, O. (2020). A novel software engineering approach toward using machine learning for improving the efficiency of health systems. *IEEE Access*, *8*, 23169-23178.
39. Ravi, P., Haritha, D., & Niranjan, P. (2018). A Survey: Computing Iceberg Queries. *International Journal of Engineering & Technology*, *7*(2.7), 791-793.
40. Madar, B., Kumar, G. K., & Ramakrishna, C. (2017). Captcha breaking using segmentation and morphological operations. *International Journal of Computer Applications*, *166*(4), 34-38.
41. Rani, M. S., & Geetavani, B. (2017, May). Design and analysis for improving reliability and accuracy of big-data based peripheral control through IoT. In *2017 International Conference on Trends in Electronics and Informatics (ICEI)* (pp. 749-753). IEEE.
42. Reddy, T., Prasad, T. S. D., Swetha, S., Nirmala, G., & Ram, P. (2018). A study on antiplatelets and anticoagulants utilisation in a tertiary care hospital. *International Journal of Pharmaceutical and Clinical Research*, *10*, 155-161.
43. Prasad, P. S., & Rao, S. K. M. (2017). HIASA: Hybrid improved artificial bee colony and simulated annealing based attack detection algorithm in mobile ad-hoc networks (MANETs). *Bonfring International Journal of Industrial Engineering and Management Science*, *7*(2), 01-12.
44. AC, R., Chowdary Kakarla, P., Simha PJ, V., & Mohan, N. (2022). Implementation of Tiny Machine Learning Models on Arduino 33–BLE for Gesture and Speech Recognition.
45. Subrahmanyam, V., Sagar, M., Balram, G., Ramana, J. V., Tejaswi, S., & Mohammad, H. P. (2024, May). An Efficient Reliable Data Communication For Unmanned Air Vehicles (UAV) Enabled Industry Internet of Things (IIoT). In *2024 3rd International Conference on Artificial Intelligence For Internet of Things (AIIoT)* (pp. 1-4). IEEE.
46. Nagaraj, P., Prasad, A. K., Narsimha, V. B., & Sujatha, B. (2022). Swine flu detection and location using machine learning techniques and GIS. *International Journal of Advanced Computer Science and Applications*, *13*(9).
47. Priyanka, J. H., & Parveen, N. (2024). DeepSkillNER: an automatic screening and ranking of resumes using hybrid deep learning and enhanced spectral clustering approach. *Multimedia Tools and Applications*, *83*(16), 47503-47530.
48. Sathish, S., Thangavel, K., & Boopathi, S. (2010). Performance analysis of DSR, AODV, FSR and ZRP routing protocols in MANET. *MES Journal of Technology and Management*, 57-61.
49. Siva Prasad, B. V. V., Mandapati, S., Kumar Ramasamy, L., Boddu, R., Reddy, P., & Suresh Kumar, B. (2023). Ensemble-based cryptography for soldiers' health monitoring using mobile ad hoc networks. *Automatika: časopis za automatiku, mjerenje, elektroniku, računarstvo i komunikacije*, *64*(3), 658-671.
50. Elechi, P., & Onu, K. E. (2022). Unmanned Aerial Vehicle Cellular Communication Operating in Non-terrestrial Networks. In *Unmanned Aerial Vehicle Cellular Communications* (pp. 225-251). Cham:

Springer International Publishing.
51. Prasad, B. V. V. S., Mandapati, S., Haritha, B., & Begum, M. J. (2020, August). Enhanced Security for the authentication of Digital Signature from the key generated by the CSTRNG method. In *2020 Third International Conference on Smart Systems and Inventive Technology (ICSSIT)* (pp. 1088-1093). IEEE.
52. Mukiri, R. R., Kumar, B. S., & Prasad, B. V. V. (2019, February). Effective Data Collaborative Strain Using RecTree Algorithm. In *Proceedings of International Conference on Sustainable Computing in Science, Technology and Management (SUSCOM), Amity University Rajasthan, Jaipur-India*.
53. Balaraju, J., Raj, M. G., & Murthy, C. S. (2019). Fuzzy-FMEA risk evaluation approach for LHD machine–A case study. *Journal of Sustainable Mining*, *18*(4), 257-268.
54. Thirumoorthi, P., Deepika, S., & Yadaiah, N. (2014, March). Solar energy based dynamic sag compensator. In *2014 International Conference on Green Computing Communication and Electrical Engineering (ICGCCEE)* (pp. 1-6). IEEE.
55. Vinayasree, P., & Reddy, A. M. (2025). A Reliable and Secure Permissioned Blockchain-Assisted Data Transfer Mechanism in Healthcare-Based Cyber-Physical Systems. *Concurrency and Computation: Practice and Experience*, *37*(3), e8378.
56. Acharjee, P. B., Kumar, M., Krishna, G., Raminenei, K., Ibrahim, R. K., & Alazzam, M. B. (2023, May). Securing International Law Against Cyber Attacks through Blockchain Integration. In *2023 3rd International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE)* (pp. 2676-2681). IEEE.
57. Ramineni, K., Reddy, L. K. K., Ramana, T. V., & Rajesh, V. (2023, July). Classification of Skin Cancer Using Integrated Methodology. In *International Conference on Data Science and Applications* (pp. 105-118). Singapore: Springer Nature Singapore.
58. LAASSIRI, J., EL HAJJI, S. A. Ï. D., BOUHDADI, M., AOUDE, M. A., JAGADISH, H. P., LOHIT, M. K., ... & KHOLLADI, M. (2010). Specifying Behavioral Concepts by engineering language of RM-ODP. *Journal of Theoretical and Applied Information Technology*, *15*(1).
59. Prasad, D. V. R., & Mohanji, Y. K. V. (2021). FACE RECOGNITION-BASED LECTURE ATTENDANCE SYSTEM: A SURVEY PAPER. *Elementary Education Online*, *20*(4), 1245-1245.
60. Dasu, V. R. P., & Gujjari, B. (2015). Technology-Enhanced Learning Through ICT Tools Using Aakash Tablet. In *Proceedings of the International Conference on Transformations in Engineering Education: ICTIEE 2014* (pp. 203-216). Springer India.
61. Reddy, A. M., Reddy, K. S., Jayaram, M., Venkata Maha Lakshmi, N., Aluvalu, R., Mahesh, T. R., ... & Stalin Alex, D. (2022). An efficient multilevel thresholding scheme for heart image segmentation using a hybrid generalized adversarial network. *Journal of Sensors*, *2022*(1), 4093658.
62. Srinivasa Reddy, K., Suneela, B., Inthiyaz, S., Hasane Ahammad, S., Kumar, G. N. S., & Mallikarjuna Reddy, A. (2019). Texture filtration module under stabilization via random forest optimization methodology. *International Journal of Advanced Trends in Computer Science and Engineering*, *8*(3), 458-469.
63. Ramakrishna, C., Kumar, G. K., Reddy, A. M., & Ravi, P. (2018). A Survey on various IoT Attacks and its Countermeasures. *International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)*, *5*(4), 143-150.
64. Sirisha, G., & Reddy, A. M. (2018, September). Smart healthcare analysis and therapy for voice disorder using cloud and edge computing. In *2018 4th international conference on applied and theoretical computing and communication technology (iCATccT)* (pp. 103-106). IEEE.
65. Reddy, A. M., Yarlagadda, S., & Akkinen, H. (2021). An extensive analytical approach on human resources using random forest algorithm. *arXiv preprint arXiv:2105.07855*.
66. Kumar, G. N., Bhavanam, S. N., & Midasala, V. (2014). Image Hiding in a Video-based on DWT & LSB Algorithm. In *ICPVS Conference*.
67. Naveen Kumar, G. S., & Reddy, V. S. K. (2022). High performance algorithm for content-based video retrieval using multiple features. In *Intelligent Systems and Sustainable Computing: Proceedings of ICISSC 2021* (pp. 637-646). Singapore: Springer Nature Singapore.
68. Reddy, P. S., Kumar, G. N., Ritish, B., SaiSwetha, C., & Abhilash, K. B. (2013). Intelligent parking space detection system based on image segmentation. *Int J Sci Res Dev*, *1*(6), 1310-1312.
69. Naveen Kumar, G. S., Reddy, V. S. K., & Kumar, S. S. (2018). High-performance video retrieval based on spatio-temporal features. *Microelectronics, Electromagnetics and Telecommunications*, 433-441.
70. Kumar, G. N., & Reddy, M. A. BWT & LSB algorithm based hiding an image into a video. *IJESAT*, 170-174.
71. Lopez, S., Sarada, V., Praveen, R. V. S., Pandey, A., Khuntia, M., & Haralayya, D. B. (2024). Artificial intelligence challenges and role for sustainable education in india: Problems and prospects. *Sandeep Lopez, Vani Sarada, RVS Praveen, Anita Pandey, Monalisa Khuntia, Bhadrappa Haralayya (2024)*

Artificial Intelligence Challenges and Role for Sustainable Education in India: Problems and Prospects. Library Progress International, *44*(3), 18261-18271.

72. Yamuna, V., Praveen, R. V. S., Sathya, R., Dhivva, M., Lidiya, R., & Sowmiya, P. (2024, October). Integrating AI for Improved Brain Tumor Detection and Classification. In *2024 4th International Conference on Sustainable Expert Systems (ICSES)* (pp. 1603-1609). IEEE.
73. Kumar, N., Kurkute, S. L., Kalpana, V., Karuppannan, A., Praveen, R. V. S., & Mishra, S. (2024, August). Modelling and Evaluation of Li-ion Battery Performance Based on the Electric Vehicle Tiled Tests using Kalman Filter-GBDT Approach. In *2024 International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS)* (pp. 1-6). IEEE.
74. Sharma, S., Vij, S., Praveen, R. V. S., Srinivasan, S., Yadav, D. K., & VS, R. K. (2024, October). Stress Prediction in Higher Education Students Using Psychometric Assessments and AOA-CNN-XGBoost Models. In *2024 4th International Conference on Sustainable Expert Systems (ICSES)* (pp. 1631-1636). IEEE.
75. Anuprathibha, T., Praveen, R. V. S., Sukumar, P., Suganthi, G., & Ravichandran, T. (2024, October). Enhancing Fake Review Detection: A Hierarchical Graph Attention Network Approach Using Text and Ratings. In *2024 Global Conference on Communications and Information Technologies (GCCIT)* (pp. 1-5). IEEE.
76. Shinkar, A. R., Joshi, D., Praveen, R. V. S., Rajesh, Y., & Singh, D. (2024, December). Intelligent solar energy harvesting and management in IoT nodes using deep self-organizing maps. In *2024 International Conference on Emerging Research in Computational Science (ICERCS)* (pp. 1-6). IEEE.
77. Praveen, R. V. S., Hemavathi, U., Sathya, R., Siddiq, A. A., Sanjay, M. G., & Gowdish, S. (2024, October). AI Powered Plant Identification and Plant Disease Classification System. In *2024 4th International Conference on Sustainable Expert Systems (ICSES)* (pp. 1610-1616). IEEE.
78. Dhivya, R., Sagili, S. R., Praveen, R. V. S., VamsiLala, P. N. V., Sangeetha, A., & Suchithra, B. (2024, December). Predictive Modelling of Osteoporosis using Machine Learning Algorithms. In *2024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS)* (pp. 997-1002). IEEE.
79. Kemmannu, P. K., Praveen, R. V. S., Saravanan, B., Amshavalli, M., & Banupriya, V. (2024, December). Enhancing Sustainable Agriculture Through Smart Architecture: An Adaptive Neuro-Fuzzy Inference System with XGBoost Model. In *2024 International Conference on Sustainable Communication Networks and Application (ICSCNA)* (pp. 724-730). IEEE.
80. Praveen, R. V. S. (2024). *Data Engineering for Modern Applications*. Addition Publishing House.