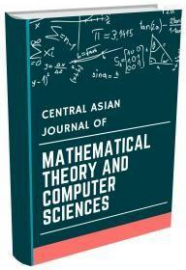

CENTRAL ASIAN JOURNAL OF MATHEMATICAL THEORY AND COMPUTER SCIENCES
https://cajmtcs.centralasianstudies.org/index.php/CAJMTCS
Volume: 07 Issue: 03 | July 2026 ISSN: 2660-5309

Article

An Intelligent Real-Time Credit Card Fraud Detection Framework Using Machine Learning

R. Sivakani[1], S. Tamil Selvam[2], N. Steevan D'Souza[3]

1. Department of Artificial Intelligence and Data Science, Dhaanish Ahmed College of Engineering, Chennai, Tamil Nadu, India
2. Department of Management Studies, Dhaanish School of Management, Dhaanish Ahmed College of Engineering, Chennai, Tamil Nadu, India
3. Department of Master of Business Administration, St. Aloysius Institute of Management and Information Technology, St. Aloysius (Deemed to Be University), Beeri, Mangaluru, Karnatka, India

* Correspondence: sivakani13@gmail.com

Citation: Sivakani R., Selvam S. T., D'Souza N. S. An Intelligent Real-Time Credit Card Fraud Detection Framework Using Machine Learning. Central Asian Journal of Mathematical Theory and Computer Sciences 2026, 7(3), 130-147.

Received: 18th Mar 2026
Revised: 13th Apr 2026
Accepted: 08th May 2026
Published: 30th Jun 2026

Abstract: In this report, we present the development and implementation of a Real-Time Credit Card Fraud Detection System using machine learning techniques. The exponential growth of electronic transactions demands robust and instantaneous fraud detection mechanisms. The main aim of the project was to construct a highly scalable and accurate system for real-time classification of transactions as legit or fraudulent. The methodology We had a very imbalanced transaction dataset. We did feature engineering to extract contextual variables like transaction history, a custom risk score. To tackle the class imbalance problem, ADASYN oversampling technique was applied. We chose the Light Gradient Boosting Machine (LGBM) model as it performed best for binary classification problems in both speed and accuracy. The model was then trained and deployed using a FastAPI web service with a PostgreSQL database for prediction logging and Redis caching to enhance performance and latency in production. The validation results showed high Recall and F1-scores, indicating that the system is effective in detecting fraudulent activities with low false negatives. -scores which proved that the system is capable of detecting fraudulent activities with low false negatives.

Keywords: Credit Card Fraud Detection, Machine Learning, Light Gradient Boosting Machine (LGBM), ADASYN Oversampling, Feature Engineering, FastAPI Web Service.

1. Introduction

Electronic payment systems are now an integral part of everyday life in the modern digital economy, enabling millions of financial transactions to be completed instantaneously throughout the world [58]. The rapid growth of online banking, mobile payment apps, e-commerce sites, contactless payment options and digital wallets has changed the way individuals and institutions complete financial transactions. While technological advancements have brought about increased convenience, efficiency and accessibility, they have also opened up new opportunities for cybercriminals to exploit vulnerabilities in digital payment ecosystems [44]. Credit card fraud is one of the most challenging and financially damaging forms of financial cybercrime faced by financial institutions, merchants, payment processors and consumers around the world. With electronic transactions growing exponentially every year, traditional fraud prevention

methods are no longer sufficient to cope with increasingly sophisticated fraudulent activities [72]. This has resulted in an increasing demand for smart automated real-time fraud detection systems that can accurately detect suspicious transactions with minimal processing delays.

Financial institutions process millions of transactions each day, each containing valuable information such as customer behaviour, spending habits, location of the transaction, amount of the purchase, merchant category, and temporal patterns. Hidden within these large datasets are subtle relationships and behavioural anomalies that are often markers of fraudulent activity [52]. The volume of transactions and the short time available to verify each payment prior to authorisation make it impossible to detect these anomalies manually. Traditional rule-based fraud detection systems rely on predefined thresholds and rules, designed by experts, like flagging unusually large purchases or transactions in geographically distant locations within a short time interval [64]. These systems offer a baseline of protection but are not very effective at detecting evolving fraud strategies and previously unseen attack patterns. Fraudsters are constantly changing their modus operandi in order to bypass static detection rules, therefore adaptive and intelligent detection mechanisms are needed to preserve the financial security [71].

Machine learning has become one of the most powerful technologies to solve this problem, because it allows computer systems to learn complex patterns directly from historical transaction data without the need to explicitly program every case of fraud [48]. Machine learning algorithms find statistical relationships between many transaction attributes, and improve their predictive ability by training on large datasets rather than using rules that are manually designed. Fraud detection is particularly well suited to supervised learning techniques that leverage prior labelled transactions to differentiate between legitimate financial behaviour and fraudulent activity [76]. The models can analyse hundreds of characteristics of each transaction at the same time, detecting subtle combinations of variables that a human analyst might not notice [57]. Thus, fraud detection systems based on machine learning can detect both known and unknown fraud patterns, while significantly reducing false alarms in contrast to traditional methods.

While machine learning offers many significant benefits, creating a successful fraud detection system poses several technical challenges [70]. Another big challenge is the large imbalance in the transaction datasets. In real-world financial systems, the vast majority of transactions are legitimate, and fraudulent transactions are typically less than one percent of all transactions recorded [51]. This imbalance means many traditional classification algorithms are biased toward predicting every transaction as legitimate because this results in very high overall accuracy. However, such models are not useful for practical fraud detection, because they do not identify the rare fraudulent cases that lead to large financial losses. Therefore, evaluation metrics such as Recall, Precision, F1-score and Area Under the Receiver Operating Characteristic Curve (AUC-ROC) are much more meaningful than overall accuracy in measuring fraud detection performance [63]. Recall is especially important among these metrics as it indicates the model's ability to detect real fraudulent transactions, thus reducing costly false negatives.

The handling of class imbalance requires the application of specific preprocessing techniques that can generate balanced training data while preserving the meaningful relationships in the original dataset [73]. More sophisticated oversampling methods, such as Adaptive Synthetic Sampling (ADASYN), generate synthetic samples of the minority class based on the local distribution of fraudulent observations [43]. Unlike traditional random oversampling, which merely replicates existing minority instances, ADASYN aims to create synthetic examples in regions where fraud detection is more difficult. The adaptivity of the approach allows the machine learning algorithms to learn more discriminative decision boundaries, and thus, improves the effectiveness of detecting fraudulent transactions while keeping false positive rates at a reasonable level [55]. The

use of such advanced pre-processing techniques greatly improves the robustness and reliability of fraud detection models deployed in real-world financial environments.

Another key aspect of fraud detection is effective feature engineering. The raw transaction data may not be sufficient in capturing the behavioural characteristics for accurate fraud prediction [47]. Therefore, it is possible to extract more contextual features to enrich the information that the machine learning model has access to. Examples include transaction frequency over time windows, average transaction amounts, customer spending trends, historical fraud occurrence, merchant risk levels, geographic movement patterns, account age, transaction velocity, and custom risk scores derived from multiple behavioural indicators [69]. These engineered features offer important insights into customer behaviour and allow the classification algorithm to better differentiate between normal purchasing patterns and suspicious activities [59]. Therefore, feature engineering is very important for improving the model accuracy and overall predictive performance.

Selecting appropriate machine learning algorithm is as important for high-performance fraud detection [65]. Among the various ensemble learning methods, the Light Gradient Boosting Machine (LightGBM) is widely recognised for its high efficiency, scalability, and predictive power [74]. LightGBM implements gradient boosting decision trees with histogram-based learning and leaf-wise tree growth strategies that significantly reduce the computational complexity and training time. LightGBM uses less memory than many traditional machine learning algorithms, supports parallel learning, handles high-dimensional datasets efficiently and provides better predictive performance on structured tabular data [53]. These characteristics make it especially suited for financial fraud detection applications where fast model training and low-latency inference are required.

The goal of this project is to create an efficient Real-Time Credit Card Fraud Detection System that is capable of detecting fraudulent transactions while fulfilling the high-performance needs of modern financial systems [50]. The project goes beyond predictive modelling and encompasses the full machine learning lifecycle including data preprocessing, feature engineering, addressing the class imbalance problem, model training and evaluation, as well as model deployment and monitoring in production [62]. The architecture of the developed solution has been designed to be production-ready, scalable, reliable and efficient under realistic operational conditions [77].

The system uses a very imbalanced credit card transaction dataset with historical financial records labelled as either fraudulent or legitimate [60]. Extensive data preprocessing procedures are applied to improve the quality of data and the meaningful behavioural features are derived which capture customer transaction history and contextual risk information. The ADASYN oversampling is then used to balance the training dataset to enhance the ability of the classifier in identifying the minority-class fraud cases [45]. The processed data is then used for training a LightGBM based binary classification model optimised for maximum Recall with high Precision and F1-score. Hyperparameter optimisation and cross-validation techniques improve generalisation performance and reduce overfitting, so that the trained model performs consistently on transactions that have not been seen before.

The trained model is then integrated into a FastAPI web application to allow for real-time deployment through RESTful endpoints that can accept transaction requests and provide fraud predictions in a matter of milliseconds [67]. FastAPI provides asynchronous request handling, automatic API documentation, and great runtime performance. This makes it a perfect framework for production machine learning services [54]. Every prediction request is processed efficiently and logged into a PostgreSQL relational database for purposes of transaction auditing, historical analysis, compliance reporting and continuous model monitoring [78]. Redis is also embedded as an in-memory caching layer to prevent redundant calculations, lower database access latency, and enhance the overall responsiveness of the system during heavy transaction load [66]. This architecture

demonstrates how modern machine learning based models can be successfully integrated into scalable cloud-ready applications that can support high volume financial transaction processing.

The main challenge this project aims to solve is to develop a robust fraud detection system that can operate in highly imbalanced data situations, while still satisfying rigorous real-time performance requirements. Fraud detection systems have to identify suspicious transactions before approving payments, so they have fractions of a second to make a prediction [49]. "If fraudulent transactions are not detected, it may lead to significant financial losses, regulatory penalties, customer dissatisfaction, and reputational harm. On the other hand, the false positive errors of valid transactions as fraudulent are a nuisance to customers, and can reduce the confidence in digital payment services [56]. Thus, the proposed system aims to efficiently balance the increased ability to detect fraud with the reduction of unnecessary interruptions to transactions.

The motivation for undertaking this project is not only academic learning but also, evidence of the increasing importance of cyber security in the modern financial systems [61]. Fraud schemes are becoming increasingly sophisticated, exploiting weaknesses in current payment infrastructures, and financial institutions are losing billions of dollars each year as a result. Intelligent systems for fraud detection directly contribute to the reduction of the losses and the building of consumer trust in digital financial services [68]. This project also offers valuable hands-on experience in implementing advanced machine learning algorithms, software engineering principles, API development, database integration, caching mechanisms, and production deployment strategies all in one end-to-end solution [75]. The system developed demonstrates how artificial intelligence can be effectively combined with scalable software architecture to build robust, efficient and reliable fraud detection platforms to tackle real-world financial security challenges [46]. This work concludes that the combination of feature engineering, adaptive data balancing techniques, LightGBM classification, FastAPI deployment, PostgreSQL logging, and Redis caching provides a complete and production-ready fraud detection framework that balances predictive accuracy with operational efficiency, making it a good fit for deployment in current digital payment environments where speed, scalability, and reliability are critical.

Review of Literature

The development of digital payment systems has significantly transformed the international financial environment by providing customers with the ability to make safe and easy transactions using online banking, mobile applications, credit cards, and digital wallets. However, the growing use of electronic transactions has also led to a corresponding rise in fraudulent activities, resulting in significant financial losses for banks, merchants and consumers [8]. Therefore, fraud detection has become one of the most important research and development areas in financial technology. Existing fraud detection systems have been developed over decades, starting with simple rule-based mechanisms and gradually incorporating statistical methods and machine learning algorithms [15]. However, despite these developments, many traditional systems are subject to several limitations that reduce their effectiveness in identifying sophisticated fraud patterns [33]. A good motivation for the development of a modern real-time fraud detection system, which should be able to provide high predictive accuracy with low processing latency, is these shortcomings.

Until recently, financial institutions tended to use mostly rule-based fraud detection systems [19]. These systems work by applying pre-set business rules based on expert knowledge and historical fraud trends. Common rules include flagging transactions above a certain amount, flagging multiple transactions in a short time period, flagging transactions from unusual geographic locations and monitoring repeated failed payment attempts [5]. These methods are relatively simple to implement and understand, but are

not very adaptable as they are completely dependent on manually constructed rules. The methods of attack continue to change constantly, making it ever more difficult to keep these systems running against fraudsters. Security analysts must constantly review transaction patterns, update detection rules and introduce new conditions to cater for emerging fraud techniques. This ongoing maintenance adds to operating costs and still does not guarantee full fraud detection.

Rule-based systems typically cannot detect new patterns of fraud [40]. As the rules are explicitly clarified, any transaction that does not breach the existing conditions will be treated as legitimate by default, regardless of any suspicious behavioural features. Contemporary financial fraud frequently manifests itself as subtle behavioural deviations, rather than obvious breaches of pre-established boundaries [12]. Fraud transactions are intentionally designed to look like legitimate customer activity, making it difficult to identify them using static rules. Consequently, many such frauds are discovered only after they have done significant financial damage [29]. This limitation suggests a need for intelligent systems that can learn complex behavioural relationships from historical transaction data, rather than just manually defined conditions.

Another major limitation of traditional fraud detection systems is the use of batch processing methods [27]. In the past, banking systems would gather transaction data for hours before analysing it for fraud. However, this approach reduced the computational requirement during transaction processing at the cost of significant delays in detecting fraudulent activities. In today's financial landscape, where electronic payments happen in seconds, delayed fraud detection is not an option [14]. If fraud is found through batch processing, the fraudulent transactions have already taken place and the stolen money is much more difficult to recover. Moreover, late detection diminishes customer confidence and increases the operational costs of fraud investigation and dispute resolution [34]. This requires real-time decision making in modern fraud detection systems.

Many legacy fraud detection platforms also suffer from performance limitations due to their technical infrastructure [18]. Previous systems were generally installed on general purpose hardware setups with less processing power than today's server architectures. Such systems were generally built around conventional processors, limited memory capacity, and operating systems that were intended primarily for general business applications, not high-performance analytical tasks [4]. While these platforms might provide enough reliability for routine transaction processing, they were often not capable of running computationally intensive algorithms for fraud detection under tight response time constraints [41]. As the volume of electronic transactions increased dramatically over time these infrastructure limitations became increasingly apparent, decreasing the scalability and responsiveness of existing fraud detection solutions.

Many traditional software architectures were built as monolithic applications, where all system components ran within a single execution environment [39]. These architectures often used sequential processing techniques that could not efficiently take advantage of multi-core processors or distributed computing resources [11]. As transaction volumes grew, these applications began to experience serious performance bottlenecks, as each transaction had to be evaluated sequentially against a host of predefined rules [20]. This inability to efficiently distribute computational workloads led to longer processing times and lower throughput, making it difficult for financial institutions to meet the increased demand for real-time fraud prevention.

The machine learning evolution has offered more advanced analytical methods for fraud detection [37]. But many early implementations suffered from a number of practical limitations. At first, machine learning models usually employed algorithms like logistic regression, decision trees, support vector machines or random forests and not much optimisation was done for real-time implementation [6]. While these algorithms showed improved predictive abilities compared with rule-based approaches, they often required

considerable computational resources during training and prediction [24]. In environments where large volumes of transactions are processed, longer inference times translated directly into longer authorisation delays, which lowered customer satisfaction and restricted the practical deployment of such models in production financial systems.

One of the challenges with early implementations of machine learning was the highly imbalanced nature of fraud datasets [35]. "Because fraudulent transactions are only a tiny percentage of total financial transactions, traditional classification algorithms are inherently biased towards predicting the majority legitimate class [3]. As a result, these models often had high overall accuracy, but missed a significant percentage of the actual fraudulent transactions. In practice, such behaviour is unacceptable in financial applications as failing to detect fraudulent activities may lead to serious financial and reputational damages [26]. Most of the earlier systems did not use sophisticated techniques to control class imbalance and had poor Recall values and high false negative rates.

Feature engineering, in many traditional fraud detection systems, was also relatively restricted [16]. Prior implementations have largely depended on the use of raw transaction features such as transaction amount, merchant identifier, transaction location and timestamp without deriving more behavioural information from historical customer activities. Therefore, the input features available often did not capture meaningful behavioural patterns that distinguish legitimate customer spending habits from fraudulent activities [7]. Contextual information is now increasingly used in fraud detection; for example transaction velocity, customer spending frequency, average purchase values, geographical movement patterns, historical account behaviour and customised risk indicators [30]. Without these engineered features, earlier systems could not tap into the full predictive power of the transaction data that was available.

A further major weakness of traditional fraud detection environments were the deployment methods [23]. Many earlier machine learning systems were stand-alone analytical programs that had to be invoked manually or periodically in batch mode rather than being continuously available prediction services [10]. This architecture did not support integration into modern banking applications with instant responses to predictions through standardised web interfaces [38]. And often, previous deployment strategies did not have automated monitoring features, making it hard to continually assess model performance after deployment or to spot gradual declines in prediction accuracy from evolving fraud patterns.

Most legacy fraud detection systems also had limitations in data persistence and performance optimisation [21]. Transaction logs were frequently stored in basic file systems or old-style relational databases, without the benefit of optimised indexing or efficient query capabilities. Although these approaches supported the historical record keeping they were often inadequate to support real-time analytics or rapid auditing requirements [1]. Furthermore, many systems did not employ dedicated in-memory caching technologies that could reduce the repeated access to the database and accelerate the prediction responses. With the growth in volume of transactions, repeated disc operations caused unnecessary latencies that reduced system responsiveness and the overall efficiency of operations [36].

Today’s financial systems require infrastructure that can handle millions of concurrent transaction requests with response times measured in milliseconds. Traditional architectures were not typically designed to scale horizontally . This meant that it was difficult to scale fast growing workloads without significant hardware upgrades [13]. The lack of efficient asynchronous processing, intelligent caching mechanisms and lightweight application frameworks further limited the scalability of existing fraud detection platforms [17]. This led many organisations to face rising infrastructure costs, yet still failed to meet challenging performance objectives.

The limitations of current fraud detection systems highlight the need for a more intelligent, scalable and production-ready solution [31]. The proposed Real-Time Credit Card Fraud Detection System addresses these limitations through the use of state-of-the-art machine learning methods, optimised software architecture, and modern deployment technologies [25]. The proposed system instead relies on Light Gradient Boosting Machine (LightGBM), an advanced gradient boosting algorithm designed for high-performance classification tasks on structured data, instead of relying on just manually defined rules. LightGBM provides substantially faster training and prediction speed with higher classification accuracy, which makes it a good fit for real-time fraud detection applications [42].

In the data preprocessing, the proposed system applies the Adaptive Synthetic Sampling (ADASYN) technique to overcome the challenges of highly imbalanced transaction data [9]. ADASYN generates synthetic samples of fraud in the difficult-to-classify regions, so that the model can learn better boundaries, which improves the model's ability to identify fraudulent transactions significantly. This approach overcomes one of the major drawbacks of previous systems, achieving high Recall with a balanced trade-off between Precision and F1-score [28].

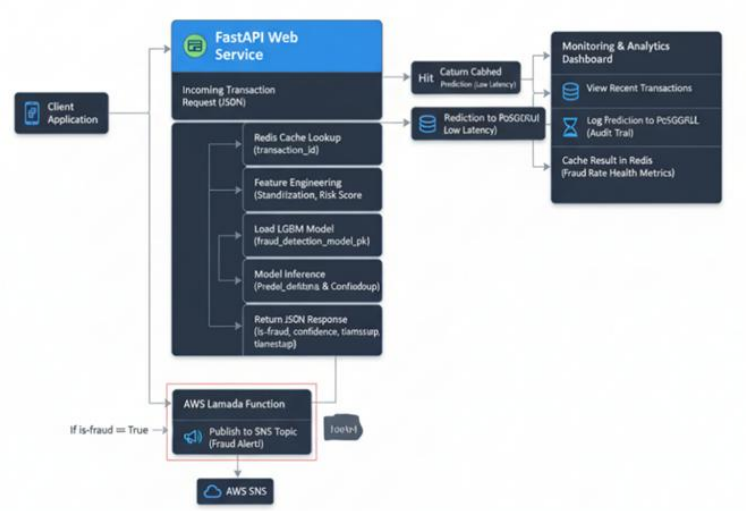

Figure 1. Block Diagram.

Another aspect that differentiates the proposed solution from the traditional implementations is the deployment architecture [22]. The trained ML model is deployed using the FastAPI framework which provides lightweight, asynchronous and high-performance RESTful web services to process prediction requests with low latency. The centralised database to securely store transaction records, prediction outcomes, and auditing information is PostgreSQL, which provides a way for comprehensive monitoring and future model improvement [2]. To accelerate repeated data access, lighten the database workload, and provide a stable response time even under the heavy transaction load, an in-memory cache system Redis is integrated. These technologies create a scalable production environment that can support modern financial applications where speed, reliability, security and continuous availability are critical [32]. The proposed system effectively tackles the shortcomings of traditional fraud detection systems by integrating intelligent feature engineering, advanced imbalance handling, efficient machine learning algorithms, optimised deployment architecture, database persistence, and high-speed caching, presenting an effective framework for accurate, low-latency, real-time credit card fraud detection in today's digital payment ecosystems.

2. Materials and Methods

The proposed methodology adheres to a structured and systematic machine learning pipeline to ensure the reproducibility, robustness, and high predictive performance over the fraud detection. The development starts with the collection of a historical dataset of credit card transactions [82]. The dataset contains legitimate and fraudulent transaction records. The raw financial data is usually noisy and contains limited predictive information. Hence, a lot of data preprocessing and feature engineering is done before training the model [101]. Transaction attributes of interest are cleaned, transformed and arranged in a format suitable for machine learning analysis. These additional contextual features include customer transaction history, transaction frequency, cumulative spending behaviour, custom risk scores, etc. and are generated to enhance the ability of the model to discriminate between legitimate and fraudulent activities [93]. Feature scaling techniques are used to standardise numerical features to ensure that the variables with different ranges contribute equally during model training and avoid bias towards the features with larger numerical values. The resulting dataset is then split into training and testing subsets while preserving the original distribution of fraud and non-fraud transactions.

A major challenge in fraud detection is the severe class imbalance in real-world financial datasets where fraudulent transactions are a very small proportion of the total number of transactions. Training ML models directly on such imbalanced data leads to a classifier biased toward the majority class, resulting in poor fraud detection performance, even though the overall accuracy is high [81]. To address this limitation, the Adaptive Synthetic Sampling (ADASYN) technique is used only on the training dataset. ADASYN intelligently creates synthetic samples of the minority class according to the distribution of the difficult-to-learn fraudulent cases. Unlike traditional oversampling methods that merely duplicate existing minority samples, ADASYN attempts to create new synthetic samples at the regions where the classifier needs more learning support [92]. This adaptive sampling strategy allows the model to learn better decision boundaries between fraudulent and legitimate transactions. It greatly improves the recall and decreases the false negative rate [103]. This means that the trained classifier becomes better at spotting transactions that are fraudulent, without too much false positives.

The prepared dataset is then used to train the Light Gradient Boosting Machine (LightGBM) model after data preprocessing and imbalance correction [100]. LightGBM is chosen due to its superior computational efficiency, scalability, and prediction performance on structured tabular data. LightGBM is a gradient boosting framework based on decision trees . It employs histogram-based learning and leaf-wise tree growth to enable faster training and lower memory consumption compared to many traditional ensemble learning algorithms [87]. The model is trained on the ADASYN-generated balanced dataset and can learn meaningful relationships between engineered features of transactions and the fraud labels. Hyperparameter tuning aims to optimise the most important parameters of the model such as the number of boosting iterations, learning rate, maximum tree depth, feature fraction, and minimum child samples. The key focus in fraud detection applications is to minimise false negatives [105]. Therefore, the model evaluation is mainly based on Recall, Precision, F1-score and Receiver Operating Characteristic, rather than on overall accuracy [94]. Cross validation techniques are used during training to improve generalisation performance and reduce the chance of overfitting so that the final model performs consistently on previously unseen transaction data.

After a satisfactory model performance has been achieved, the trained LightGBM classifier is serialised and saved as a persistent model file for deployment in the production environment. Such a trained model with persistence does not need to re-train when the application restarts [80]. This greatly reduces the computational overhead and allows

immediate prediction. The serialisation methods store the trained model parameters and the preprocessing pipeline, so that the incoming transaction data can be processed consistently during the inference [104]. The saved model file is the main prediction engine used by the deployment framework and can be easily updated when a new trained version is available [91]. This provides an efficient model versioning approach, simplifying deployment and enabling continuous improvement through future retraining without affecting production services.

The deployment methodology has been designed to meet the demanding requirements of modern financial applications such as high availability, low response latency, scalability and operational reliability [102]. The deployed solution is built on the FastAPI framework as the core application server because of its lightweight architecture, asynchronous request handling features and excellent runtime performance. FastAPI provides a performant platform to serve ML models as RESTful web services that can respond to prediction requests in milliseconds [88]. During the application initialisation, the FastAPI service sets up all required resources, which includes loading configuration parameters, creating database connections, initialising caching services, and registering application endpoints. The implementation of pydantic data models provides strict request and response schemas to ensure full validation of incoming transaction information prior to prediction processing [98]. This validation mechanism ensures that invalid data doesn't get into the prediction pipeline, increasing the overall reliability and security of the deployed application.

In order to reduce prediction latency, instead of reading the serialised LightGBM model from the storage for each prediction request, it is loaded into the memory at the application startup. By keeping the trained model in memory, unnecessary disc access operation is avoided and the inference time is dramatically reduced so that transaction classifications are almost instantaneous [86]. When an incoming transaction request is received by the application, the features are first pre-processed and transformed using the same procedures used during model training. This processed transaction data is then fed as input into the in-memory LightGBM model, which in turn produces a fraud prediction with a probability score. Such architecture allows the system to react quickly even when there is a load of a large number of concurrent transaction requests.

In the deployment architecture, Redis is added as an in-memory cache solution to improve the application performance and reduce unnecessary computation overhead. A unique transaction identifier is assigned to each new transaction, and this is used as the cache key. Before it makes a machine learning prediction, the application checks if a prediction result for this transaction is already in the Redis cache [83]. In the event of a hit, the cached prediction is returned to the client immediately without calling the machine learning model and this can reduce the processing time and save computational resources drastically. If there is no cache entry, then the transaction is processed as normal, the prediction result is generated and the new response is stored in Redis for retrieval in the future [95]. This caching approach improves throughput and enables the system to maintain low latency during high transaction volumes (Figure 2).

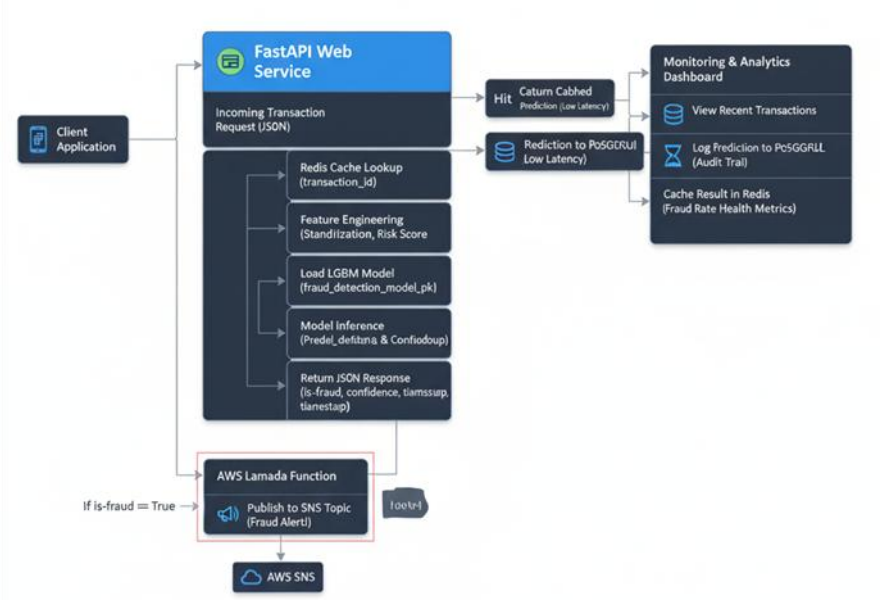

Figure 2. ER Diagram for Python Code.

Another important part of the deployment architecture is comprehensive transaction logging [90]. All transactions that have been processed, regardless of their classification result, are saved to a PostgreSQL relational database. The stored information includes the transaction identifier, input features, prediction probability, classification result, processing timestamp, and any other metadata necessary for auditing and regulatory compliance [79]. Database logging offers a permanent history of all prediction activity, enabling financial institutions to investigate suspicious transactions, assess model behaviour and conduct detailed forensic analysis when necessary. Furthermore, the gathered prediction logs allow for constant monitoring of the model's performance, allowing data scientists to identify concept drift, retrain the classifier with new transaction data and keep high fraud detection accuracy over time.

The proposed system, to demonstrate the integration of cloud-based event-driven services, includes an alert notification mechanism using Amazon Web Services Simple Notification Service (AWS SNS). Whenever the deployed machine learning model determines that a transaction might be fraudulent, an event is generated that triggers a serverless AWS Lambda function. This function will send an alert message to an SNS topic, which will then send notifications to the subscribers of the topic via email or other supported protocols [84]. The automated alerting system allows for instant alerts on suspicious activities to be automatically sent to fraud analysts and financial administrators, facilitating quick investigation and timely intervention before major financial loss is experienced [96]. This serverless implementation simplifies operations and delivers highly scalable event processing capabilities.

Dedicated API endpoints facilitate operational monitoring and system performance analysis, providing real-time statistical information related to transaction processing activities [99]. These monitoring services pull data directly from the PostgreSQL database to produce performance metrics such as total transactions processed, fraud detection rates, prediction distributions, hourly fraud occurrences, average processing time, and historical system trends. Administrators can use real-time operational statistics to constantly evaluate application performance, identify anomalous transaction patterns, monitor workload distribution, and measure overall system effectiveness [89]. These analytical capabilities enable data-driven decision making, proactive system health management, and contribute to the long-term robustness of the deployed fraud detection platform.

In summary, the proposed methodology combines state-of-the-art machine learning techniques with modern software engineering principles to create a production-ready fraud detection system capable of providing accurate, scalable, and real-time transaction classification [97]. The efficient end-to-end system, ready to be deployed in modern financial settings, consists of thorough data preprocessing, smart feature engineering, ADASYN class balancing, LightGBM classification, FastAPI deployment, Redis caching, PostgreSQL logging, AWS SNS alerting, and ongoing operational monitoring [85]. The integrated methodology guarantees a high predictive performance and fulfils the practical needs of modern digital payment systems such as speed, scalability, reliability, maintainability and continuous monitoring.

3. Results and Discussion

The final LGBM model returned a set of performance measures on the validation data set that supports the system's effectiveness, especially in the detection of the minority class [108]. Using was crucial to shift the model's performance in favour of Recall over raw Accuracy (Table 1).

Table 1. Interpretation of Key Performance Metrics.

Metric	Discussion Point
Precision	The higher the precision, the fewer false alarms there will be [107]. This is important to maintain customer confidence and reduce operational costs from investigating non-existent fraud.
Recall	The most important success factor is the high recall score, which shows that the system correctly detects all actual fraud cases, which significantly reduces financial losses.
F1-Score	The higher the F1-score, the better the balance between precision and recall. So overall, the model is quite reliable in its classifications.
Accuracy	Although this score is extremely high, it is mainly due to the large class imbalance and does not reflect the real success in this area.

System Performance Discussion The development of the real-time prediction service has provided an accurate and responsive solution.

Real-Time and Scalability: The system could meet the low-latency requirement thanks to the use of FastAPI as the web framework and the pre-loaded model and caching layer [106]. The cache is especially useful when there are multiple retries or repeated queries on the same transaction, significantly alleviating the burden on the model inference engine.

The addition of a database for persistent logging provides the system with a better auditability. The endpoint offers operators the ability to monitor system performance and fraud trends in real time, and visualise key metrics such as fraud rate by hour to identify time windows most vulnerable. This feature is essential to ongoing operational improvement and the identification of new fraud vectors [109]. The project is API-centric, but the structured response and fast inference time improve the 'user experience' for the integrating application (e.g. a banking application). Post-transaction review and system debugging are facilitated by detailed logging, leading to a robust and maintainable solution.

4. Conclusion

The development of the Real-Time Credit Card Fraud Detection System was successful in achieving the primary objective of designing and implementing an intelligent, scalable and production-ready solution that is able to detect fraudulent credit card transactions with high accuracy and low latency. The system was equipped with sophisticated data pre-processing, feature engineering, ADASYN over-sampling to overcome the severe class imbalance and LightGBM algorithm. It was able to learn complex transaction patterns and demonstrated promising results in fraud detection. The main goal was to optimise Recall so as to reduce false negatives, thereby detecting a large proportion of fraudulent transactions before the occurrence of financial losses. Built on FastAPI, PostgreSQL, and Redis, the deployment architecture added Fast API responses, secure transaction logging, efficient caching, and scalable performance for real-world financial environments. Our framework shows that by combining state-of-the-art machine learning techniques with strong software engineering practices, we can build a reliable end-to-end fraud detection solution that can support modern digital payment systems.

The proposed system shows good predictive performance and operational efficiency, but there are many chances for future improvements. Implementing a complete MLOps pipeline would automate the collection of data, retraining of the model, validation and deployment, and ensure that the model continually adapts to changing fraud patterns and minimises model drift. The system's capacity to capture sequential customer behaviour could be taken to a higher level by the implementation of sophisticated deep learning techniques like Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, or transaction embedding models. With a real-time monitoring dashboard, a fraud analyst would have intuitive visualisations of prediction statistics, fraud trends, and overall system health. Furthermore, the application of ensemble learning strategies, which combines LightGBM with other models such as XGBoost and Logistic Regression, can improve the accuracy, robustness, and confidence of the predictions, thus increasing the effectiveness of the system for large-scale deployment in modern financial institutions.

REFERENCES

[1] F. Al Shaibi, I. Ghassan, R. Salem, N. Yas, and A. Marks, "How to prove AI error cases," Scientific Culture, vol. 11, no. 3, pp. 36–47, 2025.

[2] G. C. Vegineni, "Sustainable tech development: Integrating AI, XR, and CI/CD pipelines for a greener future," in Exploring the Impact of Extended Reality (XR) Technologies on Promoting Environmental Sustainability, S. K. Gupta, N. Maurya, D. N. Le, and T. Mzili, Eds., vol. 38. Cham, Switzerland: Springer, 2025.

[3] M. M. Reddy Chinthala and M. Kalloji, "Policy-Oriented Zero Trust Microsegmentation for East-West Traffic Governance in Hybrid Cloud Architectures," 2025 6th International Conference on Smart Electronics and Communication (ICOSEC). IEEE, pp. 1330–1335, Sep. 24, 2025.

[4] H. Apuri and C. Yepuri, "Design of multi-agent autonomous workflow systems using agentic AI frameworks," in Proc. IEEE Int. Conf. Advances in Urban Computing (ICAUC), 2026.

[5] H. Apuri, M. Aurangabadkar, S. Goel, M. M. R. Chinthala, and C. Yepuri, "Advancing infrastructure-as-code resilience through generative AI agents for predictive remediation and autonomous security enforcement," Int. J. Engineering and Advanced Technology (IJEAT), vol. 15, no. 4, Mar. 2026.

[6] H. Apuri et al., "Self-healing infrastructure: Autonomous LLM agents for real-time remediation of configuration drift and security misconfigurations in IaC deployments," Int. J. Innovative Technology and Exploring Engineering (IJITEE), vol. 15, no. 4, Mar. 2026.

[7] H. Pandian, "AI-based capacity forecasting models for elastic cloud and hybrid enterprise systems," Journal of Information Systems Engineering and Management, vol. 10, no. 63s, pp. 1648–1658, Dec. 2025.

[8] H. Pandian, "AI-driven predictive performance bottleneck detection in mission-critical financial systems," Journal of Computational Analysis and Applications, vol. 30, no. 2, pp. 1019–1033, Mar. 2022.

[9] H. Pandian, "Architectural optimization techniques for high-volume batch processing in Hadoop ecosystems," International Journal of Intelligent Systems and Applications in Engineering, vol. 8, no. 4, pp. 428–439, Dec. 2020.
[10] H. Pandian, "Embedding performance engineering into CI/CD pipelines for regulated financial systems," Journal of Computational Analysis and Applications, vol. 33, no. 8, pp. 7892–7910, Dec. 2024.
[11] V. T. Manne, "An experimental comparison of enclave token vaults and HSMs for real-time card tokenization," in 2026 9th International Conference on Computational Intelligence in Data Science (ICCIDS), 2026, pp. 1–6.
[12] S. B. Venkata, "Intent-to-AVM: MCP-orchestrated AVM terraform stacks for regulated azure environments," in 2026 9th International Conference on Computational Intelligence in Data Science (ICCIDS), 2026, pp. 1–6.
[13] S. B. Venkata, "Runbook Mesh: MCP-Orchestrated Terraform and Ansible Co-Execution on Azure," in 2026 Second International Conference on Intelligent Systems for Communication, IoT and Security (ICISCoIS), 2026, pp. 1–7.
[14] S. B. Venkata, "Explainable repair-time intelligence for programmable hearing aids using digital repair twins," in 2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC), 2026, pp. 529–534.
[15] D. Antar and H. Elsayed, "Enhancing university students' Arabic language skills through social media: A case study of TikTok," Theory and Practice in Language Studies, vol. 15, no. 8, pp. 2683–2692, 2025.
[16] D. Antar, "The use of drama in developing the skill of speaking in Standard Arabic among third grade Arabic-speaking students," Theory and Practice in Language Studies, vol. 13, no. 7, pp. 1601–1613, 2023.
[17] D. Antar, "The effectiveness of using ChatGPT-4 in creative writing in Arabic: Poetry and short story as a model," Information Sciences Letters, vol. 12, no. 12, pp. 2445–2459, 2023.
[18] D. Antar, "Manifestations of color in the poetry of Arab crow poets," Journal of the Faculty of Arts, Helwan University, vol. 41, no. 1, pp. 7–26, 2016.
[19] K. Sharma and B. Goswami, "Rental market of pump-sets in the central and western parts of Nepal plains," Asia-Pacific Journal of Rural Development, vol. 30, no. 1–2, pp. 226–243, 2020.
[20] K. Sharma, "Mechanization without ownership: Market structure and pricing in Nepal's Terai," CABI Agriculture and Bioscience, vol. 7, no. 1, p. 0020, Feb. 2026.
[21] K. Sharma, P. Basnet, and K. R. Bhatt, "Social media discussion and short-horizon stock returns: Evidence from a retail coordination episode," Digital Finance, vol. 8, no. 1, pp. 1–17, 2026.
[22] K. Sharma and R. N. Shrestha, "The migration paradox: Why remittances fail to stimulate agricultural investment in Nepal's Terai plains," Economics Bulletin, vol. 45, no. 4, pp. 1649–1657, Dec. 2025.
[23] B. Kumar, "Blockchain-based authentication model for education data storage," International Journal of Industrial and Systems Engineering, vol. 51, no. 4, pp. 498–515, 2025.
[24] V. Ramalingam, B. Kumar, S. K. Gupta, D. M. Alsekait, and D. S. AbdElminaam, "A hybrid federated learning framework with generative AI for privacy-preserving and sustainable security in IoT-enabled smart environments," Scientific Reports, vol. 16, Art. no. 3071, 2026.
[25] B. Kumar, W. Shamas, J. Sandeep, and D. Albalushi, "Developing an advanced cybersecurity framework and blueprint: A contemporary approach to counter hacking through reverse engineering techniques," in Smart Cyber Physical Systems, Proceedings of ICSCPS 2024, Singapore: Springer, 2024, pp. 27–37.
[26] T. P. Krishna Kumar, "Shift of entertainment media to OTT platform – The paradigm shift of youth impacting business world," Turkish Online Journal of Qualitative Inquiry, vol. 11, no. 2, pp. 415–421, 2020.
[27] J. Muhamed Rafi, J. Moosa, T. P. Krishna Kumar, and K. Deepak, "Crude oil price influence on the performance of selected stocks from different sectors – An empirical analysis," Journal of Survey in Fisheries Science, vol. 10, Special Issue 3, pp. 1893–1902, 2023.
[28] S. K. Saravanan, R. Krishnamoorthy, T. P. Krishna Kumar, R. Narayana Rao, D. Udaya Suriya Rajkumar, and R. Thiagarajan, "IoT alert reflexion of forbidden deforestation regions with drone observation," IEEE Xplore, vol. 18, no. 5, pp. 201–211, 2023.
[29] T. P. Krishna Kumar, P. B. Acharjee, P. D. Sawant, P. Dabaria, and A. S. Mohideen, "The impact of using Facebook on consumer buying behaviour online," Journal of Chemical Health Risks, vol. 12, no. S4, pp. 744–752, 2023.
[30] S. Manochandar, T. P. Krishna Kumar, P. B. Acharjee, P. D. Sawant, P. Dabaria, and A. S. Mohideen, "The impact of using Facebook on consumer buying behaviour online," European Chemical Bulletin, vol. 12, SI 4, pp. 744–752, 2023.
[31] T. P. Krishna Kumar, P. Malhotra, B. Madhukumar, M. A. Raj, R. A. Isaac, and D. Balasubramanian, "Exploring the factors influencing the effectiveness of digital marketing in changing environment: A theoretical and

empirical investigation," Journal of Educational Administration: Theory and Practice, vol. 30, no. 4, pp. 7488–7493, 2024.

[32] T. P. Krishna Kumar, R. Suriakala, N. Shankar, and M. Deepak, "Global to local perspectives in succession planning of family business in unorganized sector," Journal of Educational Administration: Theory and Practice, vol. 30, no. 5, pp. 3056–3066, 2024.

[33] T. P. Krishna Kumar, S. Ramesh, D. R. Pallavi, M. Ramachandran, V. Saravanan, and M. Selvam, "Individual and group behavior-based customer product recommendation for designing information systems using SPSS statistics," in Springer Nature Switzerland, Jan. 10, 2025, pp. 292–299.

[34] S. Ramesh, T. P. Krishna Kumar, D. R. Pallavi, M. Ramachandran, K. Ramu, and S. Rajkumar, "Evidence on dividend preferences, attention, and IPO valuation of retail investors using SPSS statistics," in Springer Nature Switzerland, Jan. 10, 2025, pp. 651–659.

[35] T. P. Krishna Kumar, R. Suriakala, V. Seema, K. T. Vijayamol, and B. Tater, "The impact of artificial intelligence and machine learning on personalized marketing strategies," Journal of Chemical Health Risks, vol. 15, no. 2, pp. 1562–1576, 2025.

[36] B. Kumar, N. B. Najmusseher, P. K. Nizar Banu, and R. Dwivedi, "Epileptic seizure detection contribution in healthcare sustainability," in AI and IoT: Driving Business Success and Sustainability in the Digital Age, B. Awwad, Ed., Studies in Systems, Decision and Control, vol. 601, Cham, Switzerland: Springer, 2025, pp. 225–235.

[37] B. Kumar and O. Al Falhi, "Digital transformation through APIs," in Proc. Int. Conf. on Communication, Information Technology and Internet of Things (COM-IT-CON), 2022.

[38] K. Al Aufi and B. Kumar, "Security testing of Android application using Drozer," in Proc. Int. Conf. on Computational Sciences and Sustainable Technologies, Springer CCIS, 2024, pp. 8–18.

[39] B. Al Barwani, E. Al Maani, and B. Kumar, "IoT-enabled smart cities: A review of security frameworks, privacy, risks, and key technologies," in Proc. 1st Int. Conf. on Innovation in Information Technology and Business (ICIITB 2022), Advances in Computer Science Research, vol. 104, Springer, 2022, pp. 169–181.

[40] S. I. Sadigova et al., "Unveiling the dynamics of sociolinguistics, understanding language in social contexts, artificial intelligence effect," Forum for Linguistic Studies, vol. 7, no. 12, pp. 1–15, 2025.

[41] N. T. Safarli et al., "Foreign language proficiency among EFL students in higher education and usage of artificial intelligence," Forum for Linguistic Studies, vol. 7, no. 8, pp. 646–665, 2025.

[42] A. C. Samad et al., "Deep learning-driven smart pedagogy and assessment: A quality framework for outcome-based education in higher education," in Proc. 2025 Global Conf. Emerging Technology (GINOTECH), IEEE, Apr. 2025.

[43] N. K. Pasha et al., "The impact of AI tools on enhancing EFL learners' engagement in higher education," in Proc. 2025 Global Conf. Emerging Technology (GINOTECH), IEEE, Apr. 2025.

[44] A. Gullapelly et al., "AI-based pedagogical framework for improving engagement of EFL learners in higher education," in Proc. 2025 Int. Conf. Intelligent Communication Networks and Computational Techniques (ICICNCT), Sept. 5–6, 2025, IEEE.

[45] H. Pandian, "Performance engineering as a first-class cybersecurity control," in Proc. International Conference on Cyber Security, IoT, Data & Information Technology, vol. 459, pp. 1–12, Jun. 2024.

[46] H. Pandian, "Performance-driven development (PDD): A new software engineering paradigm," International Journal on Recent and Innovation Trends in Computing and Communication, vol. 9, no. 12, pp. 278–286, Dec. 2021.

[47] H. Pandian, "Quantifying business risk and financial loss from performance failures in enterprise systems," in Proc. International Conference on Computer Science, Engineering and Applications, vol. 978, pp. 160–170, Jul. 2021.

[48] H. Pandian, "Self-healing performance architectures for large-scale banking and payment platforms," International Journal of Intelligent Systems and Applications in Engineering, vol. 11, no. 11s, pp. 1019–1029, Dec. 2023.

[49] H. Pandian, "Unified observability framework for enterprise performance, capacity, and reliability," Journal of Electrical Systems, vol. 14, no. 4, pp. 96–106, Apr. 2018.

[50] H. Pandian, "Workload characterization models for distributed enterprise systems," International Journal on Recent and Innovation Trends in Computing and Communication, vol. 7, no. 12, 2019.

[51] M. M. R. Chinthala, H. Apuri, and K. Bitra, "Behaviour-Aware Hybrid Deep Networks for Detecting Zero-Day and Ransomware Threats," IJITEE, vol. 15, no. 5, pp. 1–10, Apr. 2026.

[52] D. Kodi and B. C. C. Marella, “Fraud Resilience: Innovating Enterprise Models for Risk Mitigation,” Journal of Information Systems Engineering and Management, vol. 10, no. 12s, pp. 683–695, Jan. 2025.

[53] V. R. Anumolu and B. C. C. Marella, “Maximizing ROI: The intersection of productivity, generative AI, and social equity,” in Advancing Social Equity Through Accessible Green Innovation. Hershey, PA, USA: IGI Global Scientific Publishing, 2025, pp. 373–386.

[54] B. C. C. Marella, “Streamlining Big Data Processing with Serverless Architectures for Efficient Analysis,” FMDB Transactions on Sustainable Intelligent Networks, vol. 1, no. 4, pp. 242–251, Dec. 14, 2024.

[55] B. C. C. Marella, “AI and XR in Supply Chain: Revolutionizing Sustainable Practices for a Better Tomorrow,” in Exploring the Impact of Extended Reality (XR) Technologies on Promoting Environmental Sustainability. Cham, Switzerland: Springer Nature Switzerland, 2025, pp. 83–98.

[56] S. Venkatasubramanian, “G-MAPQR: GNN-assisted multi-agent predictive quality of service routing for high-mobility MANETs,” IETE Journal of Research, pp. 1–15, 2026.

[57] S. Venkatasubramanian, S. Raja, V. Sumanth, J. N. Dwivedi, J. Sathiaparkavi, S. Modak, and M. L. Kejela, “Fault diagnosis using data fusion with ensemble deep learning technique in IIoT,” Mathematical Problems in Engineering, vol. 2022, Art. no. 1682874, pp. 1–8, 2022.

[58] S. Raja, J. Logeshwaran, S. Venkatasubramanian, M. Jayalakshmi, N. Rajeswari, N. G. Olaiya, and W. D. Mammo, “OCHSA: Designing energy-efficient lifetime-aware leisure degree adaptive routing protocol with optimal cluster head selection for 5G communication network disaster management,” Scientific Programming, vol. 2022, Art. no. 5424356, pp. 1–11, 2022.

[59] S. Venkatasubramanian, J. N. Dwivedi, S. Raja, N. Rajeswari, J. Logeshwaran, and A. Praveen Kumar, “Prediction of Alzheimer’s disease using DHO-based pretrained CNN model,” Mathematical Problems in Engineering, vol. 2023, Art. no. 1110500, pp. 1–11, 2023.

[60] S. Venkatasubramanian, V. Mohan, A. Subasri, S. H. Prasath, A. Thenmozi, and M. A. D. Thirumanraj, “Decentralized IoT frameworks: Blockchain-enabled trust in smart ecosystems,” in Edge Computing and Applications, N. Ramalingam, Y. El Allioui, and S. Bhattacharyya, Eds., Cham, Switzerland: Springer, 2026.

[61] G. C. Vegineni, “Intelligent UI designs for state government applications: Fostering inclusion without AI and ML,” J. Adv. Develop. Res., vol. 13, no. 1, pp. 1–13, 2022.

[62] G. C. Vegineni, “Real-time performance optimization in modern UI applications,” in Proc. 3rd Int. Conf. Intelligent Cyber Physical Systems and Internet of Things (ICoICI), Coimbatore, India, 2025, pp. 528–533.

[63] G. C. Vegineni and V. R. Anumolu, “Sustainable supply chain practices: A teaching framework with data science,” in Integrating AI and Sustainability in Technical and Vocational Education and Training (TVET), A. S. Azar, S. K. Gupta, H. Taherdoost, and F. Alhamaty, Eds. USA: IGI Global Scientific Publishing, 2025, pp. 245–262.

[64] G. C. Vegineni, “Exploring anomalies in dark web activities for automated threat identification,” FMDB Trans. Sustainable Comput. Syst., vol. 2, no. 4, pp. 189–200, 2024.

[65] A. Alzaabi, N. Metawa, and A. Marks, "The role of digital transformation in enhancing government service delivery in the UAE: A literature review," International Review of Management and Marketing, vol. 15, no. 5, pp. 100–106, 2025.

[66] K. Bashar, M. Salama, M. Elhoseny, A. Marks, S. A. Mostafa, A. Mustapha, M. A. Mohammed, M. A. Mahmoud, B. A. S. Al-Rimy, and S. A. Razak, "An adaptive protection of flooding attacks model for complex network environments," Security and Communication Networks, vol. 2021, Art. no. 5515234, 2021.

[67] M. Elhoseny, N. Metawa, A. Marks, and Y. Damra, "Leveraging big data analytics to advance business intelligence: An innovative direction," presented at the 6th Innovation and Analytics Conference and Exhibition (IACE 2025), Kuala Lumpur, Malaysia, Oct. 2025.

[68] M. Elhosney and A. Marks, "A game-theoretic framework for optimizing pharmaceutical pricing negotiations between governments and private insurers," presented at the 25th International Symposium for Production Research (ISPR 2025), Antalya, Turkey, Oct. 2025.

[69] M. Elyat, Y. Al Bayati, N. A. Al Baloushi, M. Sarhan, A. Marks, K. Khudhair, and A. Allouzi, "Impact of intellectual property rights and technological factors on information security," International Journal on Technical and Physical Problems of Engineering (IJTPE), vol. 16, no. 3, pp. 234–241, 2024.

[70] N. B. Hernández, K. E. P. Izquierdo, D. M. V. Rodríguez, and A. Marks, "Information fusion for the training of public administrators: Design of a composite indicator through the integration of AHP and TOPSIS methods," International Journal of Neutrosophic Science, vol. 23, no. 4, pp. 405–414, 2024.

[71] M. Majdalaweih and A. Marks, "A reference framework for enterprise computing curriculum," in Recent Advances in Information Systems and Technologies, Á. Rocha, A. Correia, H. Adeli, L. P. Reis, and M. M. Costanzo, Eds., vol. 569, Advances in Intelligent Systems and Computing. Cham, Switzerland: Springer, 2017, pp. 451–460.

[72] M. Majdalaweih and A. Marks, "Conceptual framework for assessing information technology programs quality," in Recent Advances in Information Systems and Technologies, Á. Rocha, A. Correia, H. Adeli, L. P. Reis, and M. M. Costanzo, Eds., vol. 569, Advances in Intelligent Systems and Computing. Cham, Switzerland: Springer, 2017, pp. 461–470.

[73] Z. Mamadiyarov, M. Eshov, D. Otakhonova, S. Utegenova, K. Sabirov, A. Marks, and N. Bahodirov, "Does energy policy risk threaten energy sources diversification?" International Journal of Energy Economics and Policy, vol. 15, no. 5, pp. 730–740, 2025.

[74] S. Nagar, "An impact of performance of companies with sustainability goal on stock market movements," European Economic Letters, vol. 13, no. 3, pp. 417–420, 2023.

[75] S. Nagar and P. Mahajan, "An assessment of effectiveness of remote work mode in job internships," The Online Journal of Distance Education and E-Learning (TOJDEL), vol. 11, no. 2, p. 1154, 2023.

[76] S. Nagar, "Augmentation of insurance business in India: Role of commercial banks," in India Banking and Finance Report 2024, National Institute of Bank Management (NIBM), pp. 81–94, 2024.

[77] S. Nagar, "Cointegration of Indian stock market with global stock markets: An empirical analysis," European Economic Letters, vol. 14, no. 2, pp. 2457–2464, 2024.

[78] P. Jothilingam, "Towards autonomous commissioning: Integrating digital twins, artificial intelligence and smart sensors for next-generation process control systems," Certified Journal of International Research (CJIR), vol. 5, no. 1, pp. 1-8, Mar. 2025.

[79] P. Jothilingam, "Edge computing for industrial automation and control: Enabling real-time processing, scalable architectures and secure operations," Certified Journal of International Research (CJIR), vol. 5, no. 1, pp. 1–8, Mar. 2025.

[80] P. Jothilingam, "Advancing cybersecurity in industrial control systems: Frameworks, threat modeling, and resilience strategies," International Journal of Supportive Research (IJSR), vol. 2, no. 2, pp. 69–75, Jul. 2024.

[81] O. Alimbaeva, A. Joshi, G. Saritha, L. H. Alzubaidi, K. Senthamil Selvan, and A. Chaudhary, "Novel Materials for High-Performance Energy Storage Devices," in E3S Web of Conferences 13th International Conference on Power and Energy Systems (ICPES 2023) , Chengdu, China, 2024.

[82] B. Jayaprakash, D. Bordoloi, P. Mehta, A. Amudha, P. Marwaha, and K. S. Selvan, "A Network for Medical Segmentation of Images with Multiple Centers That Preserves Privacy," in Proceedings of the 2024 Global Conference on Communications and Information Technologies (GCCIT 2024), Bangalore, India, 2024.

[83] M. Vigenesh, M. Grover, K. Senthamilselvan, A. Singla, M. Chethan, and S. B. Patil, "Secure Data Aggregation in Wireless Sensor Ad Hoc Networks Using Homomorphic Encryption," in Proceedings of the 15th International C9onference on Computing, Communication and Networking Technologies (ICCCNT 2024), Himachal Pradesh, India, 2024.

[84] K. S. Selvan, S. Goyal, O. S. Kulkarni, K. Yuvaraj, and N. Vashisht, "Generative Adversarial Networks for 3D Scene Reconstruction," in Proceedings of the 15th International Conference on Computing, Communication and Networking Technologies (ICCCNT 2024), Himachal Pradesh, India, 2024.

[85] K. Senthamil Selvan, G. Aravind, C. Pavankumar, K. Bhopate, M. Ravshan, and T. S. Senthil Kumar, "Varicose Veins Treatment Using Automated Stockings," in Proceedings of the International Conference on Newer Engineering Concepts and Technology (ICONNECT 2023), Tiruchirappalli, India, 2023.

[86] M. Almusawi, A. P. Singh, Y. S. Bisht, K. Senthamil Selvan, J. D., and N. Esanmurodova, "Cryptography-Based Privacy-Preserving Data Analysis and Empowering Data Privacy through Secure Multi-Party Computation: Challenges and Solutions," in Proceedings of the 2023 International Conference for Technological Engineering and its Applications in Sustainable Development (ICTEASD 2023), Al-Najaf, Iraq, 2023.

[87] A. S. Priya, K. S. Selvan, P. Jeevananthan, D. Dhabliya, P. K. Parida and S. Pund, "An Analysis of Partition Tree Clustering Techniques for Automated Classification of Hyper Spectral Scans," 2023 International Conference on Power Energy, Environment & Intelligent Control (PEEIC), Greater Noida, India, 2023.

[88] G. Ahluwalia, S. K, D. Dhabliya, M. K. Singar, G. Ezhilarasan and V. Singh, "Assessing the Benefits of Data Mining for Predictive Analytics," 2023 International Conference on Emerging Research in Computational Science (ICERCS), Coimbatore, India, 2023.

[89] I. Nayak, G. Radha, R. M. M. Shareef, T. E, K. Senthamilselvan and J. V, "Enhancing Student Performance through Adaptive Knowledge Assessment: A Bayesian Optimization Approach," 2023 International Conference on Sustainable Communication Networks and Application (ICSCNA), Theni, India, 2023.

[90] M. C. Jobin Christ, D. Kalaiyarasi, J. G, K. Senthamilselvan, D. Kirubakaran and N. S. Gowri Ganesh, "PoBTx(Proof of Block and Transaction): An Efficient Consensus Algorithm for IoT Business Blockchain," 2023 Second International Conference on Augmented Intelligence and Sustainable Systems (ICAISS), Trichy, India, 2023.

[91] T. V. Hai, N. V. Quang, T. P. Thao, N. T. Linh, H. V. Quynh, and N. T. Binh, "Motivation and job satisfaction of healthcare workers in private hospitals in Vietnam," The Journal of V. N. Karazin Kharkiv National University: Series Medicine, vol. 33, no. 2(53), pp. 291–308, 2025.

[92] T. V. Hai and L. Thi, "Reasons for capital increase of securities brokers," Salud, Ciencia y Tecnología – Serie de Conferencias, vol. 3, Art. no. 1212, 2024.

[93] T. V. Hai, "Maintaining life insurance contracts of joint commercial banks with insurance links," in Proceedings of the Eighth International Conference on Information System Design and Intelligent Applications, Singapore: Springer Nature Singapore, 2025, pp. 131–141.

[94] T. V. Hai, "Factors affecting short-term solvency of securities companies: The case of Vietnam," Webology, vol. 18, Special Issue 5, pp. 393–404, 2021.

[95] P. P. Anand, U. K. Kanike, P. Paramasivan, S. S. Rajest, R. Regin, and S. S. Priscila, "Embracing Industry 5.0: Pioneering Next-Generation Technology for a Flourishing Human Experience and Societal Advancement," FMDB Transactions on Sustainable Social Sciences Letters, vol.1, no. 1, pp. 43–55, 2023.

[96] G. Gnanaguru, S. S. Priscila, M. Sakthivanitha, S. Radhakrishnan, S. S. Rajest, and S. Singh, "Thorough analysis of deep learning methods for diagnosis of COVID-19 CT images," in Advances in Medical Technologies and Clinical Practice, IGI Global, pp. 46–65, 2024.

[97] G. Gowthami and S. S. Priscila, "Tuna swarm optimisation-based feature selection and deep multimodal-sequential-hierarchical progressive network for network intrusion detection approach," Int. J. Crit. Comput.-based Syst., vol. 10, no. 4, pp. 355–374, 2023.

[98] A. J. Obaid, S. Suman Rajest, S. Silvia Priscila, T. Shynu, and S. A. Ettyem, "Dense convolution neural network for lung cancer classification and staging of the diseases using NSCLC images," in Proceedings of Data Analytics and Management, Singapore; Singapore: Springer Nature, pp. 361–372, 2023.

[99] S. S. Priscila and A. Jayanthiladevi, "A study on different hybrid deep learning approaches to forecast air pollution concentration of particulate matter," in 2023 9th International Conference on Advanced Computing and Communication Systems (ICACCS), Coimbatore, India, 2023.

[100] S. S. Priscila, S. S. Rajest, R. Regin, and T. Shynu, "Classification of Satellite Photographs Utilizing the K-Nearest Neighbor Algorithm," Central Asian Journal of Mathematical Theory and Computer Sciences, vol. 4, no. 6, pp. 53–71, 2023.

[101] S. S. Priscila and S. S. Rajest, "An Improvised Virtual Queue Algorithm to Manipulate the Congestion in High-Speed Network"," Central Asian Journal of Medical and Natural Science, vol. 3, no. 6, pp. 343–360, 2022.

[102] S. S. Priscila, S. S. Rajest, S. N. Tadiboina, R. Regin, and S. András, "Analysis of Machine Learning and Deep Learning Methods for Superstore Sales Prediction," FMDB Transactions on Sustainable Computer Letters, vol. 1, no. 1, pp. 1–11, 2023.

[103] R. Regin, Shynu, S. R. George, M. Bhattacharya, D. Datta, and S. S. Priscila, "Development of predictive model of diabetic using supervised machine learning classification algorithm of ensemble voting," Int. J. Bioinform. Res. Appl., vol. 19, no. 3, 2023.

[104] S. Silvia Priscila, S. Rajest, R. Regin, T. Shynu, and R. Steffi, "Classification of Satellite Photographs Utilizing the K-Nearest Neighbor Algorithm," Central Asian Journal of Mathematical Theory and Computer Sciences, vol. 4, no. 6, pp. 53–71, 2023.

[105] S. S. Rajest, S. Silvia Priscila, R. Regin, T. Shynu, and R. Steffi, "Application of Machine Learning to the Process of Crop Selection Based on Land Dataset," International Journal on Orange Technologies, vol. 5, no. 6, pp. 91–112, 2023.

[106] T. Shynu, A. J. Singh, B. Rajest, S. S. Regin, and R. Priscila, "Sustainable intelligent outbreak with self-directed learning system and feature extraction approach in technology," International Journal of Intelligent Engineering Informatics, vol. 10, no. 6, pp.484-503, 2022.

[107] S. S. Priscila, D. Celin Pappa, M. S. Banu, E. S. Soji, A. T. A. Christus, and V. S. Kumar, "Technological frontier on hybrid deep learning paradigm for global air quality intelligence," in Cross-Industry AI Applications, IGI Global, pp. 144–162, 2024.

[108] S. S. Priscila, E. S. Soji, N. Hossó, P. Paramasivan, and S. Suman Rajest, "Digital Realms and Mental Health: Examining the Influence of Online Learning Systems on Students," FMDB Transactions on Sustainable Techno Learning, vol. 1, no. 3, pp. 156–164, 2023.

[109] S. R. S. Steffi, R. Rajest, T. Shynu, and S. S. Priscila, "Analysis of an Interview Based on Emotion Detection Using Convolutional Neural Networks," Central Asian Journal of Theoretical and Applied Science, vol. 4, no. 6, pp. 78–102, 2023.