big data and cognitive computing

MDPI

Article

Adaptive, Privacy-Enhanced Real-Time Fraud Detection in Banking Networks Through Federated Learning and VAE-QLSTM Fusion

Hanae Abbassi *, Saida El Mendili and Youssef Gahi

Engineering Sciences Laboratory, National School of Applied Sciences, Ibn Tofail University, Kenitra 14000, Morocco; elmendili.saida@uit.ac.ma (S.E.M.); gahi.youssef@uit.ac.ma (Y.G.)
* Correspondence: hanae.abbassi@uit.ac.ma

Abstract

Increased digital banking operations have brought about a surge in suspicious activities, necessitating heightened real-time fraud detection systems. Conversely, traditional static approaches encounter challenges in maintaining privacy while adapting to new fraudulent trends. In this paper, we provide a unique approach to tackling those challenges by integrating VAE-QLSTM with Federated Learning (FL) in a semi-decentralized architecture, maintaining privacy alongside adapting to emerging malicious behaviors. The suggested architecture builds on the adeptness of VAE-QLSTM to capture meaningful representations of transactions, serving in abnormality detection. On the other hand, QLSTM combines quantum computational capability with temporal sequence modeling, seeking to give a rapid and scalable method for real-time malignancy detection. The designed approach was set up through TensorFlow Federated on two real-world datasets—notably IEEE-CIS and European cardholders—outperforming current strategies in terms of accuracy and sensitivity, achieving 94.5% and 91.3%, respectively. This proves the potential of merging VAE-QLSTM with FL to address fraud detection difficulties, ensuring privacy and scalability in advanced banking networks.

Keywords: real-time fraud detection; VAE-QLSTM; federated learning; TensorFlow Federated

check for updates

Academic Editors: Salvador García López and Moulay A. Akhloufi

Received: 26 May 2025
Revised: 22 June 2025
Accepted: 3 July 2025
Published: 9 July 2025

Citation: Abbassi, H.; El Mendili, S.; Gahi, Y. Adaptive, Privacy-Enhanced Real-Time Fraud Detection in Banking Networks Through Federated Learning and VAE-QLSTM Fusion. *Big Data Cogn. Comput.* **2025**, *9*, 185. https://doi.org/10.3390/bdcc9070185

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The digital era paves a broad chain of e-services provided by FinTechs, helping individuals and ventures to direct their finances. Those e-services might involve mutual funds, payment utilities, and loans. Online transactional fraud detection is a critical issue in the banking sector and has been growing at a rapid pace. With no rigid verification and monitoring, it could lead to tens of dollars in expenses, impacting the FinTech's reputation entirely [1]. According to [2,3], the fraud prevention and identification market was valued at USD 43 billion in 2023, compared to USD 36.89 billion in 2022, and is predicted to expand to USD 182 billion by 2030, rising at a CAGR of 22% throughout the forecast time frame.

Historically, rule-based tactics were the primary tech adopted to deal with such criminal cyberattacks, showing an inability to keep up due to their dependence on static and rigid rules. On the other hand, the rapid enhancement of machine learning, namely, Random Forest, Support Vector Machine, etc., has introduced fresh solutions to identify phony acts. Despite this success achieved, their reliance on centralized infrastructures threatens privacy while limiting adaptation in the face of ever-changing suspicious practices [4–7].

This poses serious problems in terms of data confidentiality and adaptation to unknown trends, particularly in light of increasing regulations and cybersecurity risks. Handling these limitations, the application of deep learning architectures incorporating GAN, AE, RNN, and LSTMs, among others, has displayed considerable promise in transforming the method of detecting and battling criminal behavior through their adeptness to perform auto feature engineering and capture complex conducts [8].

While DL-based methods have established success in capturing the complex relationships within malicious acts, they are innately imperfect due to the computational limits of usual processing and their incapacity to proficiently tackle involved feature spaces. This drives the pioneering of quantum AI as an intriguing field showing great promise for improving the identification of fraudulent patterns, due to its ability to evolve, adapt, and acquire observations from hidden patterns in massive datasets [9]. Combining quantum AI with DL offers a unique opportunity for overcoming the constraints of convolutional fraud detection systems, enhancing its robustness and accuracy.

Presently, FL has been established as a promising technology for coping with the privacy concerns that are characteristic of centralized ML systems. FL extends a decentralized paradigm that allows multiple FinTechs to train ML models of fraud detection collaboratively without exposing their critical data, thereby maintaining data security and privacy while promoting collaboration between entities [10,11]. In that context, FL warranties conformity with laws, likewise refining the banking industry's power to dispute suspects while guarding vulnerable data. Yet, with all these pros, there are what are commonly known as restricted or skewed data issues that may inhibit the usefulness of the FL [12].

Tackling these challenges, this study exploits a hybrid architecture integrating the powers of federated learning, deep learning, and Quantum ML. The designed scheme applies FL to guarantee the continuity of critical information across all banks while facilitating the collaborative enhancement of the algorithms between multiple industries. It also uses VAE for feature learning and QLSTM to point out patterns in transaction sequences that occur over time. This hybrid approach seeks to yield a scalable, resilient, and quick method to spot malignancy immediately, opening the way for improvements in secure e-banking.

The key contributions of our research are threefold:

- First, we supply a revolutionary strategy that fuses the capabilities of federated learning with the detection power of VAE-QLSTM to create a secure, resilient, and accurate fraud detection framework, overcoming the limits of traditional centralized methods and federated learning individually.
- We also implement and analyze the suggested technique using real-world banking datasets, proving its efficacy in enhancing the accuracy and resilience of fraud detection systems.
- We provide an in-depth comparison of the suggested technique with state-of-the-art fraud detection architectures in FinTech, noting its strengths and weaknesses.

The rest of this paper is organized as follows: Section 2 provides a review of current banking fraud detection studies, and Section 3 highlights the research gaps and our motivations. In Section 4, we display our proposed architecture, and then we detail our experimental setup and results in Section 5. Section 6 discusses the pros and cons of the provided architecture by making comparisons with current studies. Finally, Section 7 concludes the research.

2. Banking Fraud Detection Methods

The identification of fraudulent transactions in banking and finance is becoming increasingly complicated, necessitating the development of more effective and efficient fraud detection technologies. Researchers have investigated the implementation of machine

learning, blockchain, federated learning, and GenAI approaches to increase fraud detection model accuracy and efficacy. In this section, we will evaluate the current literature on banking and financial fraud detection, emphasizing these investigations' noteworthy contributions, methodologies, and findings.

Several studies have expanded the use of supervised methods, specifically, one study inspected the performance of three predictive algorithms—logistic regression, random forest, and decision trees—in identifying duplicitous credit card transactions [13]. The findings indicated that the random forest model had the greatest accuracy, with a peak of 96% and an area under the curve of 98.9%. Furthermore, the study discovered that credit card issuers over the age of 60 became more prone to fraud, with an elevated suspicion rate between 22:00 and 4:00 GMT. Similarly, Rb and Kr in [14] have proposed an artificial neural network that achieves an accuracy of approximately 100% in detecting fraudulent transactions. The proposed study contrasted the effectiveness of a neural network against conventional machine learning algorithms, including k-nearest Neighbor and Support Vector Machine, employing data pre-processing, standardization, and under-sampling to circumvent the issues caused by imbalanced datasets in the credit card fraud dataset.

In this context, Moreira et al. [15] have investigated the deployment of machine learning learners in predictive tests for financial fraud identification in banking systems, through a database of more than 6 million records to evaluate the potential of Logistic Regression, Naive Bayes, KNN, and Perceptron predictors with RUS, SMOTE, and ADASYN strategies. The study's findings revealed that Logistic Regression and KNN models outperformed Naive Bayes regarding accuracy and AUC metrics, especially on balanced training bases; the investigation points out the necessity of considering false negatives and positives, along with the need to pick the most appropriate model according to the company's operational regulations and community. A further study [16] has presented IFDTC4.5, an intuitionistic fuzzy logic-based decision tree that categorizes transactions as fraud, regular, or questionable depending on the information gain ratio, and has shown a great rate of detection and effectiveness in identifying online transactional fraud.

Moreover, a study by [17] tried to solve the important problem of finding fraud in real-time e-transactions by proposing an original framework that combines the benefits of unsupervised learners, such as autoencoders and extended isolation forests, with cutting-edge big data tools, notably, Spark streaming and Sparkling Water. This led to an improvement in accuracy of nearly 99% on two real-world transactional databases. On the other hand, Patil et al. [18] have introduced a graph-based machine learning model that detects fraudulent activity in various types of banking operations, such as credit card transactions, debit card transactions, and online banking transactions, using advanced methods for anomalies, behaviors, and patterns to analyze past transactions and user behavior. The investigations revealed a higher outcome.

Building on those advances, recent research has explored how blockchain technology might improve fraud detection in online transactions. In a thorough study, authors [19] have improved the efficacy of fraud detection in mobile transactions by integrating Bidirectional 3D Quasi-Recurrent Neural Networks with Proof-of-Voting consensus Blockchain technology (Bi-3DQRNN-PoV-FD-MT) on a Bitcoin transactions dataset. Demonstrating considerable gains in accuracy, surpassing existing methods by 12.09% in contrast to the K-Nearest Neighbor-Distributed Blockchain Consortium (KNN-DBC), 8.91% in contrast to the Decision Tree-Ethereum Blockchain-Enabled Smart Contract (DT-EBSC), and 6.92% in contrast to the Heterogeneous Graph Transformer Networks-Ethereum Smart Contract (HGTN-ESC). The experimental findings indicate that the Bi-3DQRNN-PoV-FD-MT framework represents a potential strategy for bettering fraud detection in mobile transactions. Additionally, the study by Raj et al. [20] has exploited the synergies of supervised machine

learning predictors with blockchain to enhance security and transparency in credit card fraud detection. The results showed that the XGBoost demonstrated exceptional performance, achieving an accuracy and a precision of 97% and 94%, respectively, with an AUC of 0.97. The GBM and RF learners also revealed remarkable performance. In addition, the proposed architecture underlines the potential profits of merging cutting-edge machine learning techniques with blockchain technology to meaningfully progress fraud detection systems, ultimately boosting user trust and the security of online transactions.

Federated learning has also gained popularity in fraud detection, as it empowers the expansion of more accurate models without compromising data privacy. In this context, [21] has offered an innovative approach for detecting fraudulent credit card transactions that relies on a hybrid algorithmic optimization-based deep learning strategy. The Jellyfish Namib Beetle Optimization Algorithm-SpinalNet (JNBO-SpinalNet) scheme was created to enhance fraud detection accuracy through the pre-processing of raw data with the quantile normalization, picking features with various distance measures, improving features with the bootstrapping method, and adjusting the SpinalNet with the JNBO model to boost the detection rate. The JNBO-SpinalNet model showed improved performance, with an accuracy of 89.10% and a tiny loss function of 13.16%. Exploring the same architecture, [22] laid out a credit card fraud detection architecture that uses federated learning using TensorFlow Federated and PyTorch. The research resolved the imbalance in transaction data by contrasting analyses of individual and hybrid resampling strategies. Simulations indicate that hybrid resampling increases machine learning effectiveness, with Random Forest attaining the best accuracy of 99.99%, ahead of Logistic Regression, K-Nearest Neighbors, Decision Tree, and Gaussian Naive Bayes at 94.61%, 99.96%, 99.98%, and 91.47%, correspondingly. Furthermore, PyTorch demonstrated greater prediction for federated learning than TensorFlow Federated, despite having a longer computing time, highlighting the promise of federated learning in improving detection while safeguarding data privacy.

Recently, GenAI has been explored as a promising approach to identifying suspicious activity. For instance, Generative Adversarial Networks (GANs) were deployed to generate artificial data, increasing the discriminative power of classifiers and boosting overall performance [23]. The suggested architecture was trained on the European cardholders 2013 dataset, which includes 30 features and exhibits considerable class imbalance, with less than 1% of fraudulent transactions. The results revealed that the model outperforms previous studies, with an AUC of 0.999. This research contributes to the development of new strategies for securing transactions and detecting credit card fraud, thus offering a potential solution to the class imbalance problem. A further study [24] has sought the adoption of GANs to produce synthetic data for credit card fraud detection, focusing on the impact of consumer distribution on the effectiveness of synthetic data augmentation. The results indicate that institutions with customers of higher credit quality are more likely to benefit from GAN-based augmentation, although banks offering cheaper credit may benefit asymmetrically due to variations in their consumption patterns. The study highlights the value of taking customer distribution into account when using synthetic data to improve model accuracy in unbalanced datasets.

In the upcoming section, we will highlight the gaps in the research and our motivations for resolving them.

3. Research Gaps and Motivations

Based on the review of the above fraud detection research, the subsequent weaknesses and open challenges have been pointed out.

Most of the discussed algorithms adopt conventional centralized fraud detection systems that rely heavily on collecting and storing critical client information, posing

data privacy and confidentiality risks [25]. Applying centralized ML approaches often overlooks atypical fraudulent acts, which result in enormous financial losses to banking and lending institutions. Furthermore, the lack of transparency and explicability of machine learning models applied to fraud detection makes it difficult to pinpoint the core reasons for suspects [26].

In addition, the lack of studies on fusion FL with quantum DL leaves a significant gap in the literature, particularly in the area of financial fraud detection. This gap is exacerbated by the scarcity of real data on financial fraud, which complicates the evaluation of the performance of fraud detection technologies [27]. On the other hand, the increasing scale of fraudulent practices and rapidly changing fraud trends require increasingly sophisticated and flexible fraud detection systems.

To cope with these gaps, our research leverages the power of DL and QDL models that enhance interpretability in complex and unbalanced fraud detection scenarios with privacy characteristics of FL. Primarily, by using FL, our architecture guarantees data privacy, permitting collaboratively trained models without breaching privacy rules, thus surmounting the hurdles of data exchange across banks. Secondly, to address the data imbalance problem, we apply the VAE owing to its capability to enhance model performance on atypical fraud patterns by generating additional samples that resemble the minority class [28,29]. It does not simply oversample data from the minority class, it utilizes latent space interpretations to capture complex interconnections to augment identification of fraud that other methodologies may overlook. The VAE also reduces feature noise and dimensionality, providing a more compact and informative representation for subsequent sequence modeling.

Additionally, QLSTM features quantum-enhanced temporal modeling to meet sophisticated sequential malignities. QLSTM leverages quantum-inspired mechanisms to better capture complex temporal patterns in transaction sequences, enabling the detection of subtle and evolving fraud behaviors that classical LSTM models might miss. Thus, our displayed scheme aims to outperform cutting-edge approaches, providing a scalable, secure, adaptable, and efficient solution in the landscape of online digital transaction fraud detection. In our study, we develop a new fraud detection architecture that exploits the strengths of federated learning and VAE-QLSTM. While FL allows multiple banks to collaboratively train a shared model without conveying their sensitive data, VAE-QLSTM's advanced pattern recognition capabilities enable spotting malicious transactions, seeking to improve the accuracy and robustness of fraud detection systems in banking networks.

By combining these components within a federated architecture, our approach ensures robustness against data sparsity, preserves privacy, and adapts effectively to dynamic fraud trends by modeling both latent feature distributions and sequential irregularities. This synergetic approach seeks to greatly enhance the accuracy and security of fraud detection systems in FinTech networks. Additionally, by utilizing deep learning for feature extraction and quantum-enhanced sequential modeling, our method enhances flexibility in adapting to changing fraud strategies. This integration intends to boost the identification of both established and emerging fraud trends in the fast-changing FinTech sector. The following section will detail our proposed semi-decentralized FL-VAE-QLSTM architecture.

4. Semi-Decentralized FL and VAE-QLSTM for Banking Fraud Detection

In order to address the above challenges, we establish a semi-decentralized FL architecture combined with VAE-QLSTM, which is specifically developed to improve real-time banking fraud detection. FL allows banking organizations to collaborate without exchanging raw data in this architecture. At the same time, VAE-QLSTM improves fraud detection

models by training them on local transaction data, guaranteeing privacy at all stages. Figure 1 displays our proposed architecture, which comprises the following three layers:

- Local Training: Each participating bank trains a VAE-QLSTM learner independently on its own set of transactions to maintain the privacy of the raw data.
- Cluster Aggregation: At this stage, the models trained locally are grouped into clusters. These are then aggregated to form a cluster-level model, which amplifies the detection of malicious transactions by mapping them across multiple entities.
- Global Aggregation: Finally, the outputs simulated from the cluster level are combined to build a Global Model for improved flexibility and resilience across the banking system. Once this model is created, it is then adopted to spot suspicious operations, leading to real-time decision making, and to continue to learn about new fraud tendencies.

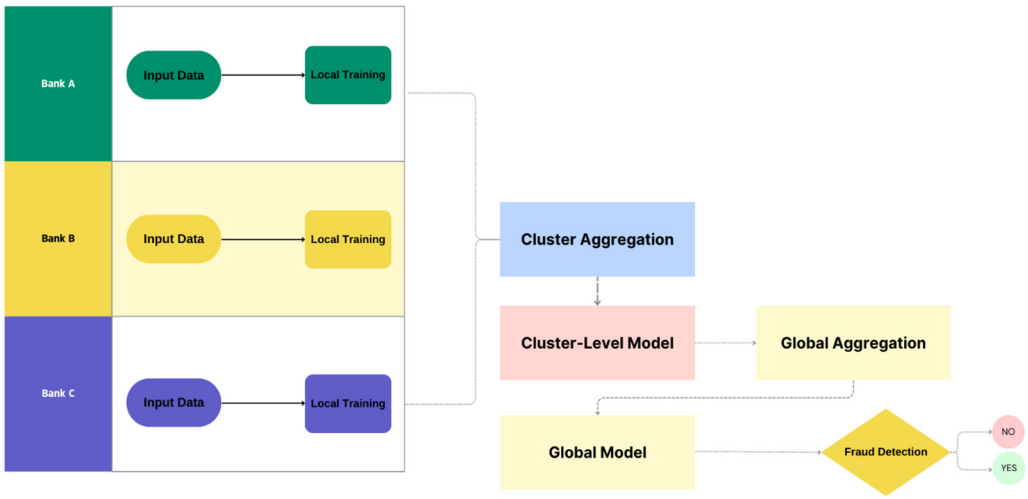


Figure 1. High-level architecture for federated fraud detection.

This layered scheme permits fraud detection algorithms to harness intelligence exchange among multiple entities while upholding data confidentiality. By incorporating federated learning with VAE-QLSTM, the framework significantly improves accuracy, adaptability, and privacy compliance, resulting in a more effective and safe fraud detection mechanism.

4.1. Local Training

During this local training phase, each participating bank or credit institution trains its VAE-QLSTM model independently using its own transaction database. This ensures the protection of critical consumer data, such as transaction history, user habits, and personal identifiers, within each institution while complying with financial legislation to reduce the risk of privacy breaches. The main objective is to have complete control over the data while contributing to the global fraud detection system. Each bank can identify fraud trends specific to its customers without exchanging raw data, thus ensuring high confidentiality standards. Through rigorous data analysis, banks limit the risk of security breaches and benefit from shared fraud detection information. Figure 2 illustrates how the training process is carried out individually within each bank.

The key stages of the local training process are the following:

- Each bank standardizes and optimizes its local transaction records in preparation for training the model. This involves standardizing the data and maintaining the consistency of features, such as transaction amounts and timestamps, to improve the model's ability to detect anomalies. At the same time, feature engineering converts raw data into useful characteristics such as time-based, amount-based, and behavioral

tendencies. In addition, additional features, such as transaction frequency and geographical patterns, are created to reveal hidden links in the data that static rule-based systems might overlook.

- Then, the Variational Autoencoder-Quantum Long Short-Term Memory (VAE-QLSTM) hybrid architecture will be adopted by each local bank to effectively identify anomalies in transactional data. The VAE module learns a low-dimensional conceptual representation of transactions, simulating a typical transaction distribution and identifying deviations as potential fraud. Meanwhile, the QLSTM module plays a crucial role in analyzing transaction sequences by identifying temporal patterns. This allows the model to spot unusual behaviors that might indicate a malicious act, which enhances computational efficiency and pattern recognition. This makes the system more adept at detecting both isolated anomalies and recurring fraudulent activities. By merging these powerful architectures, the VAE-QLSTM succeeds at detecting not only individual abnormalities but also complicated, time-varying fraud trends that traditional approaches may miss. This comprehensive technique ensures effective identification of both solitary and sequential fraudulent activity, hence improving overall fraud detection powers.
- After the training process, each bank retains its own data and only shares the model's configurations, such as weights and biases, through the aggregation layer to maintain privacy. This approach allows banks to collaborate on fraud detection while keeping critical info confidential. The revisions that are securely shared contribute to a federated learning framework, which helps identify suspects across various banks without compromising privacy. This federated paradigm encourages community knowledge for fraud detection while ensuring strong privacy and security requirements.

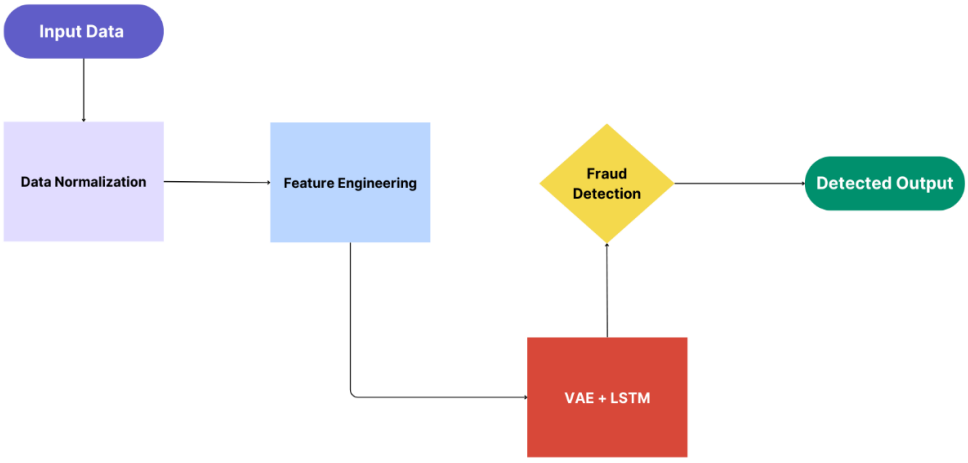


Figure 2. Local training stage for individual banks.

This stage leads to a comprehensive fraud detection system, tailored to each bank, paving the way for the next step, cluster aggregation, which combines findings from several institutions to improve fraud detection throughout the banking network.

4.2. Cluster Aggregation

After every single bank is done with local training, the upcoming stage is cluster aggregation, which strengthens the models from the banks in each cluster into a single model. This technique improves the model's capacity to detect fraud in a variety of transactions while maintaining data privacy. Figure 3 displays the cluster aggregation layer.

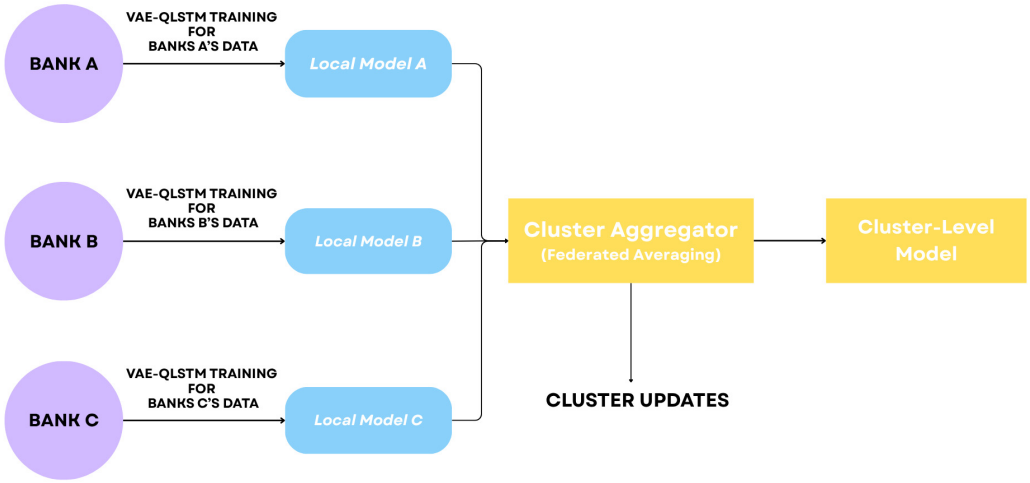


Figure 3. Cluster-level aggregation process.

The aggregation procedure goes as follows:

- Initially, each bank delivers updated models (the parameter alteration) to a local aggregator inside the cluster. This technique guarantees that only abstracted model details will be communicated, protecting the privacy of each bank's transaction information.
- Following that, the local aggregator acquires these updates from each bank in the cluster and uses Federated Averaging (FedAvg) to aggregate them. This method calculates the average of each bank's model parameters in order to create a single consolidated model that incorporates information from all participating institutions.
- The cluster-level model is then completed. This model exploits combined data and insights from all the banks in the cluster, enhancing its ability to identify complex fraud patterns that individual algorithms might overlook.
- Lastly, the cluster-level model will serve as input for the global aggregation layer. In this final stage, we combine the models from all clusters to form a comprehensive global model. This boosts our ability to identify suspect transactions throughout the entire network.

By balancing higher fraud detection with stringent data privacy, this strategy enables banks to leverage integrated insights without risking customers' safety. This latest model integrates models from all clusters.

4.3. Global Aggregation

The global aggregation layer joins cluster-level models from different clusters to form a single, global model. This stage tends to have the greatest degree of collaborative fraud detection, where intelligence across multiple banking businesses is synthesized to develop a more powerful and complete detection system. This layer is managed by a central server, which supervises the aggregation procedure and ensures each cluster's model parameters are integrated without jeopardizing the privacy of any particular bank's data. The global aggregation employs FL approaches, including FedAvg, to combine model parameters in a way that represents the different transaction patterns seen across all participating clusters. By synthesizing knowledge from numerous databases, the global model attains a greater level of expansion, successfully identifying fraud habits that might not be obvious at the local or cluster level.

Once consolidated, the global model is substantially more resilient and adaptive since it incorporates insights from a wide range of data sources, for instance, banks spanning different clusters, allowing it to detect sophisticated and previously unknown fraud patterns throughout the whole banking network. This capacity is crucial in fighting developing

fraud strategies since the model alters in reaction to real-world behavioral acts. According to operational needs, this global model may be installed locally at each bank to improve real-time fraud detection or may be centrally controlled. This versatility permits lenders to tailor fraud detection tactics to their infrastructures, guaranteeing flawless integration without impacting current processes. Figure 4 shows the global aggregation process.

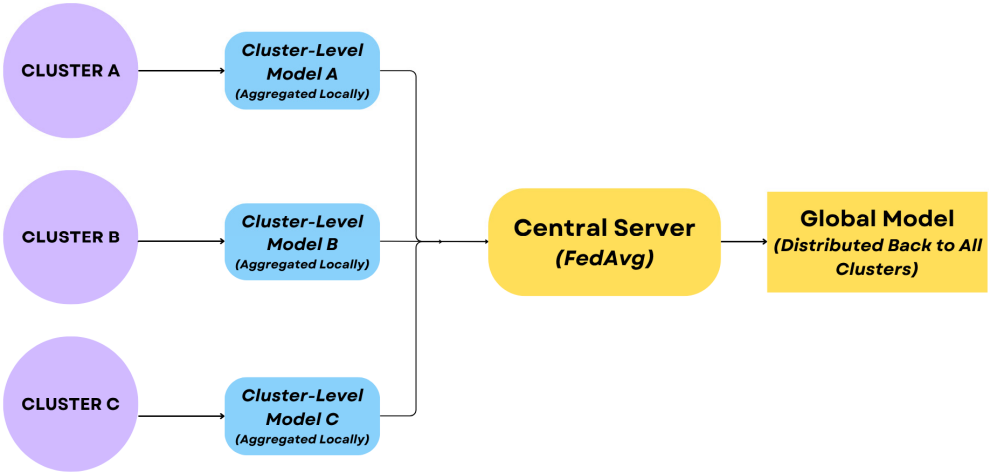


Figure 4. Global aggregation layer across clusters.

In the following section, we will deep dive into the experimental setup, including the preprocessing of the data used and the model deployment. This step involves the practical aspects of deploying the model, such as ensuring repeatability, as well as exhibiting real-world applications of our suggested framework

5. Experimentation and Deployment

This section describes the experimental and deployment procedures used to assess the suggested fraud detection architecture. Firstly, we present the databases adopted in our architecture. Then, we dive deep into the implementation of our VAE-QLST in each stage on the FL environment, followed by a depiction of the assessment criteria of our model.

5.1. Data Description

The availability of high-quality datasets is critical for developing and examining systems for detecting banking fraud. These datasets reflect real banking transactions exactly, permitting researchers to assess how effective their methods are at spotting suspicious operations. In this work, we leverage datasets that mirror real-world banking transactions to assess the efficacy of our proposed federated learning architecture for detecting banking fraud. We describe the datasets utilized in detail below.

European cardholder: This is a publicly accessible database for two days in September 2013 involving anonymized and aggregated transaction data. It comprises 284,807 transactions, 492 of which are fraudulent, and it serves as a critical and widely adopted benchmark in credit card fraud detection studies [30]. The following are the dataset's features:

- Time: the time frame of the operation in seconds.
- V1–V28: Transformed characteristics taken from transaction data, such as the geographical location and kind of transaction.
- Amount: The value of the transaction.
- Class: A binary label that indicates if the transaction is malicious (1) or lawful.

IEEE-CIS-Fraud-Detection dataset: This includes credit card transactions from a European bank. The dataset is made up of the following four files: train_transaction.csv,

train_identity.csv, test_transaction.csv, and test_identity.csv. We consider the transaction data as one dataset and the identity data as another dataset. Originating from a Kaggle competition, this dataset is widely recognized as a comprehensive benchmark for evaluating fraud detection models due to its detailed features and scale. By examining transaction and identity data as distinct datasets, we will assess the model's capacity for identifying suspicious transactions while assessing the model's ability to recognize high-risk cardholders [31].

These datasets provide a robust foundation for benchmarking and validating our proposed approach. Moreover, the federated learning framework coupled with the VAE-QLSTM model architecture supports continual learning and adaptability, positioning the system to respond effectively to evolving fraud patterns in dynamic financial environments. Table 1 presents the transaction and identity data's characteristics.

Table 1. Summary of data attributes used in the study.

Characteristic	Transaction Data	Identity Data
Number of transactions	1,097,231 (590,540 training + 506,691 testing)	288,466 (144,233 training + 144,233 testing)
Class balance	0.17% (fraud) vs. 99.83% (normal)	
Features	TransactionID: unique identifier for each transaction TransactionDT: the timestamp TransactionAmt: the volume of the transaction ProductCD: the product code card1 to card6: attributes conveyed to the card applied for the transaction addr1 to addr2: attributes conveyed to the cardholder's address dist1 to dist2: features related to the distance between the cardholder's address and the operation spot P_emaildomain to R_emaildomain: attributes related to the email domain of the cardholder V1 to V339: anonymized attributes obtained from the transaction data	id: unique identifier for each cardholder DeviceType: device type utilized by the cardholder DeviceInfo: info about the device used id_01 to id_38: anonymized attributes obtained from the identity data DeviceInfo2 to DeviceInfo4: further information about the device used id_39 to id_41: further anonymized attributes obtained from the identity data

5.2. Model Implementation

The experimental setup was conducted on both local hardware and cloud resources. through an NVIDIA RTX 3070 16G GPU and an Intel Core i7 processor on the local workstation, adding to the flexibility of Google Colab Python 3.8, TensorFlow 2.12 with its built-in Keras API, and TensorFlow Federated version 0.53.0 were the primary tools used for seamless model training and FL workflows. Furthermore, we utilized a classical quantum simulator on the local GPU to train the QLSTM, ensuring fine execution without the need for a straight upsurge to physical quantum hardware.

5.2.1. Data Preprocessing

Feature engineering is necessary to convert the raw data from the European cardholder and IEEE-CIS datasets into relevant, organized inputs for the model. Each dataset is thoroughly processed to extract essential characteristics while maintaining data integrity and enabling successful fraud identification.

For the European cardholder database, our feature engineering focuses on the following:

- Transaction Time Trends: To spot suspicious recurring trends in malicious behavior, we produced attributes relating to transaction time, including the hour of the day, day of the week, and the last operation's time.
- Event Amount Scaling: Event amounts were normalized to reduce the influence of high or low values, enabling the model to emphasize relative transaction volumes instead of absolute amounts.
- Feature Interactions: We defined relationships between transaction amount and time to detect abnormalities in high-value transactions that occurred at atypical hours.

For the IEEE-CIS transaction dataset, the following specific feature engineering tactics were employed:

- Transaction timestamps were utilized to create attributes such as the time until the latest transaction and the time of day the transaction happened, which assisted in detecting temporal trends suggestive of fraudulent behavior.
- Card and Address Features: We merged card-related attributes (cards 1–6) with address data (addr1, addr2) to create novel features that allowed us to discover variations in the usage of cards across various geographical areas.
- Distance-Based Features: The dist1 and dist2 factors were used to calculate distances between the transaction sight and the cardholder's address, which helped identify geographical irregularities that may signal acts of fraud.

For the IEEE-CIS identity dataset, the applied feature engineering strategies are the following:

- To identify odd user habits, we developed features based on DeviceType and DeviceInfo. It is especially valuable for discovering trends in how a given device is utilized across many locations. For instance, if a user makes a transaction on a mobile device in a location but suddenly attempts a transaction from another area on a desktop system, this might trigger red flags.
- We engineered features to assess irregularities across the IP addresses for transactions and the attributed email domains. Through monitoring these outliers, we can spot high-risk users, as unhabitual conduct in IP addresses or email domains can suggest suspect or hacked accounts.

In complement to our rule-based feature engineering, we incorporated VAE to improve attribute representation and identify oddities. Although handcrafted features gather stated fraud sequences, VAE finds hidden walls within transactions, encoding complicated connections that may not be clear via manual selection. Particularly, VAE-generated inserts assist in improving the following:

- Transaction Time Patterns: By recording latent temporal changes and irregular interactions in transaction cycles, VAE improves anomaly identification for time-based behaviors.
- Event Amount Scaling: VAE detects small dispersion adjustments in transaction amounts, even when absolute numbers look normal.
- Geographical and Behavioral Trends: By encoding user mobility and device-based gestures within a latent area, VAE promotes the detection of fraud that standard intended attributes might ignore.

Considering the significant class imbalance evident in both databases, especially with fraud occurrences of less than 0.2%, we deployed the Variational Autoencoder (VAE) not only for feature representation but also as a means for decreasing the imbalance. The VAE incorporates latent structures of minority fraud patterns and implicitly augments the feature space, allowing the model to detect unusual fraudulent behaviors without resorting to naïve oversampling.

By concentrating on these high-impact attributes, we aimed to create a robust profile of real user behavior that the model might then use to identify fraud. Once those traits were retrieved, we began training our models. Following this, we will go into depth about the model training procedure, outlining the stages and techniques used to enhance and assess our proposed solution.

5.2.2. Model Training

Upon the data engineering stage, the model we offered relies on a semi-decentralized architecture involving three banking institutions. Each bank, represented by its database, develops a local model separately before collaborating with a federated learning process that gathers those models in two steps: local aggregation and global aggregation. This method keeps confidential information decentralized while leveraging common knowledge to enhance fraud detection skills.

During the training phase, we developed three distinct VAE-QLSTM models using the European cardholder dataset, the IEEE transaction dataset, and the IEEE identity dataset. The VAE served for feature encoding, capturing key characteristics before feeding them into the QLSTM architecture. QLSTM was selected for its superior ability to model intricate time-based relationships. Its quantum-inspired gates improve memory, enabling it to better spot irregular or subtle fraud trends that unfold over time. This makes it ideal for capturing both immediate and delayed fraudulent behaviors, a key advantage over conventional LSTMs. This is particularly useful for detecting evolving or delayed fraudulent activities as QLSTM handles both short- and long-term data dependencies. Before feeding data to QLSTM, we use VAE. The VAE preprocesses the data by denoising it and creating more expressive latent features, which helps QLSTM concentrate on significant transactional changes.

We adopted the Adam optimizer with a learning rate set to 0.001, and the loss function employed was binary cross-entropy. These hyperparameters were selected after empirical testing and tuning on the validation set, balancing convergence speed and generalization performance. The learning rate of 0.001 was found optimal for maintaining training stability across FL clients. Binary cross-entropy was chosen due to the binary classification nature of fraud detection.

Each model copes with 10 training epochs with a batch size and a validation split of 32 and 0.2, respectively. To prevent overfitting, we incorporated early stopping with patience of three epochs. The VAE-QLSTM architecture comprised a QLSTM layer, refined through settings including 128 number of units, a tanh activation function, and a dropout and recurrent dropout of 0.2. followed by a dense layer with 64 units and a sigmoid output layer. This allows the downstream QLSTM to focus on high-level behavioral patterns rather than surface-level noise. Figure 5 displays the VAE-QLSTM model training.

After training local models, we utilized TFF's FedAvg algorithm for local aggregation. The aggregation process was set up using the tff.learning.build_federated_averaging_process [32] function, and client-side logic was defined using the tff.learning.Client class [33]. Throughout the process, a learning rate and a batch size of 0.001 and 32 were maintained. The combined model was then validated using the three predefined datasets to evaluate its efficacy.

Beyond the local aggregation, we took advantage of TFF's hierarchical FL tactic to carry out global aggregation. At this step, we have fused locally aggregated models from different clusters to gather a global model. We kept the same learning rate and batch size as the previous step, implementing early stopping criteria to avoid overfitting. The achieved global model was assessed on the validation set to examine its overall performance, showcasing the effectiveness of the semi-decentralized approach in fraud detection.

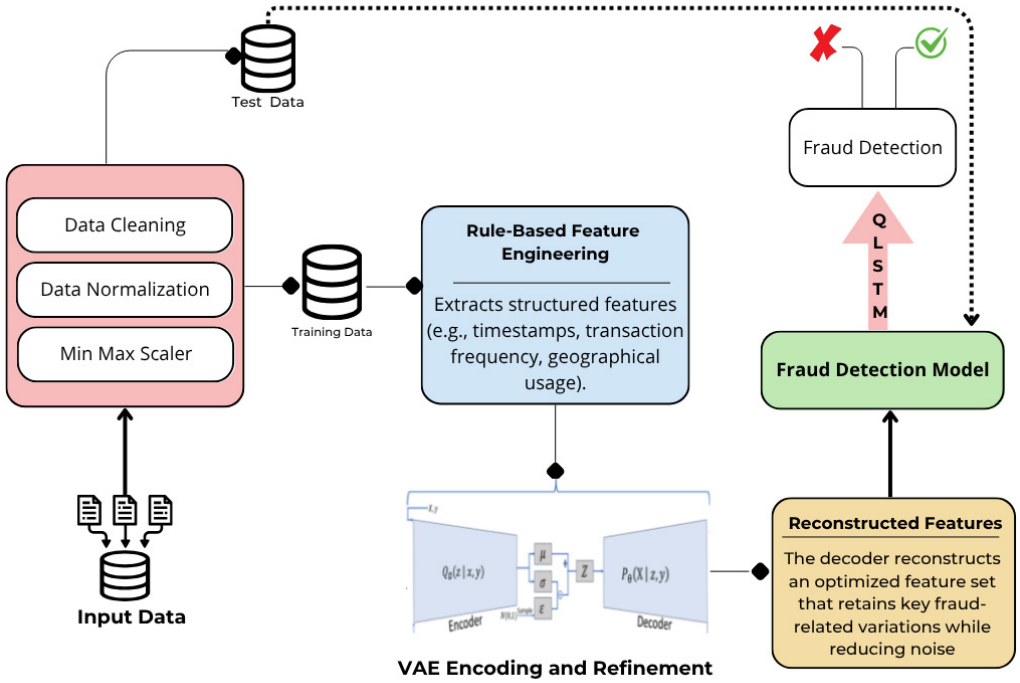


Figure 5. VAE-QLSTM-based fraud detection framework.

Moreover, the architectural design was structured with real-time applicability in mind. The VAE-QLSTM model is computationally efficient during inference when executed on GPU-enabled environments such as the RTX 3070. Moreover, the federated aggregation steps are implemented using TensorFlow Federated's lightweight protocols, which support timely model synchronization across distributed clients. While no empirical latency benchmarks were recorded, the setup demonstrates potential for integration in time-sensitive fraud detection scenarios.

5.3. Assessment Criteria

We have evaluated the efficiency of our suggested semi-decentralized architecture for identifying fraud, adopting three cutting-edge performance metrics, and examined the model at different stages of local (training-aggregation) and global aggregation. These metrics offer a thorough insight into the model's ability to detect fraudulent activities accurately.

Accuracy: The overall accuracy of the model is calculated as the percent of properly identified occurrences out of all analyzed examples. In other words, the accuracy can evaluate how often a fraud detection model is correctly predicting all scenarios. However, when the data is complexly skewed, accuracy is not sufficient to assess the model's performance. Consequently, we consider other metrics, namely the detection rate, which can be referred to as sensitivity or recall. It assesses the model's efficacy in picking up real positive fraud cases. A high detection rate implies that the model successfully spots potential malicious transactions. Similarly, AUC-ROC corresponds to the Area Under the Receiver Operating Characteristic Curve, which may more accurately convey the efficacy of the fraud detection learner on imbalanced distributions. AUC-ROC values vary between 0 and 1, with higher values suggesting more discriminating potential.

Those criteria were thoroughly evaluated at each level of the training process of our architecture, from the local model to the aggregation procedures, allowing us to evaluate the improvement of our fraud detection scheme. In the upcoming section, we will discuss the findings of our experiments and compare them to the state-of-the-art studies.

6. Results and Analysis

In this section, we will cover the simulation outcomes obtained using our private datasets, divided into five sets to train the VAE-QLSTM, assuring a complete evaluation. The VAE-QLSTM's performance is evaluated by looking at each model's accuracy, detection rate, and AUC-ROC value as summarized in Table 2. The analysis findings reveal the enhancements made by both the local and global aggregation procedures.

Table 2. Experimental results across multiple datasets.

Model	Accuracy	Detection Rate	AUC-ROC
European Cardholder Model	0.912	0.853	0.945
IEEE Transaction Model	0.923	0.875	0.962
IEEE Identity Model	0.905	0.838	0.938
Local Aggregated Model	0.933	0.893	0.973
Global Model	0.945	0.913	0.985

The findings demonstrate the efficacy of our semi-decentralized federated learning system. Amongst the different local models, the IEEE Transaction Model beat the others, with a detection rate of 0.875 and an AUC-ROC of 0.962, indicating high detection skills. The European cardholder model and the IEEE Identity Model both fared well, although they lagged.

Following local aggregation, efficiency increased, with the local aggregated model achieving an AUC-ROC of 0.973 and an accuracy of 0.933. This advancement was reinforced following global aggregation, leading to the global model achieving the greatest performance (AUC-ROC of 0.985 and detection rate of 0.913).

To further assess the effectiveness of our model, we contrasted the outcome metrics of our semi-decentralized federated learning system to those of related studies on the IEEE-CIS and European cardholder datasets. These comparisons provide an overview of how our approach performs compared to existing strategies that use both federated learning and classic machine learning models.

As shown in Table 3, our FL-VAE-QLSTM models outperformed other approaches across all metrics. The FL-GNN model from [32] achieved an AUC-ROC of 75.28% and a detection rate of 61.02%, which is significantly lower than the performance of our FL-VAE-QLSTM models applied to the IEEE Transaction dataset (AUC-ROC of 96.2%) and IEEE Identity dataset (AUC-ROC of 93.8%). Additionally, the SVM model from [33] and the FL-CNN model from [34] demonstrated weaker detection capabilities, with detection rates of 27.12% and 42.37%, respectively. These results underscore the superior performance of our approach, particularly in improving detection rates and AUC-ROC for fraud detection.

Table 3. Comparison of the proposed model with state-of-the-art methods on the IEEE-CIS dataset.

Ref	Model	Accuracy	Detection Rate	AUC-ROC
[34] 2024	FL-GNN	81.60%	61.02%	75.28%
[35] 2023	SVM	76.89%	27.12%	61.60%
[36] 2019	FL-CNN	77.36%	42.37%	66.61%
-	FL-VAE-QLSTM on Transaction	92.3%	87.5%	96.2%
-	FL-VAE-QLSTM on Identity	90.5%	83.8%	93.8%

Regarding the European cardholder dataset in Table 4, our FL-VAE-QLSTM model also achieved superior performance, with an accuracy of 91.2%, a detection rate of 85.3%, and an AUC-ROC of 94.5%. Comparatively, the JNBO-SpinalNet model in [19] achieved an accuracy of 89.10%, while the Adaboost + LGBM model from [35] attained a detection rate of only 64% and an AUC-ROC of 82%. On the other side, the NAG model mentioned in [36] performed well in terms of accuracy with 86.08%, but it fell short on providing metrics, namely, sensitivity or AUC-ROC, which makes it difficult for us to obtain direct comparisons.

Table 4. Comparison of the proposed model with state-of-the-art methods on the European cardholder dataset.

Ref	Model	Accuracy	Detection Rate	AUC-ROC
[21] 2020	JNBO-SpinalNet	89.10%	-	-
[37] 2022	Adaboost + LGBM	-	64%	82%
[38] 2022	NAG	86.08%	-	-
[39] 2018	SA-GAN	-	79.37%	-
[40] 2020	LSTM	-	74.08%	87.02%
[40] 2020	GRU	-	72.08%	86.02%
[41] 2022	LSTM	-	76.1%	78%
[41] 2022	MLP	-	75.5%	81%
-	FL-VAE-QLSTM on European Cardholder	0.912	0.853	0.945

Furthermore, an expanded comparison encompassing current cutting-edge deep learning approaches reveals distinct performance landscapes. For instance, SA-GAN, LSTM, GRU, and MLP [37–39] exhibited promising outcomes in fraud detection. However, the classical LSTM [38,39] had inferior detection rates of 74.08% and 76.1%, respectively, with AUC-ROC scores of 87.02% and 78%. Compared to these, our FL-VAE-QLSTM model performs significantly better, highlighting the importance of quantum-enhanced sequence modeling intricate fraud patterns that traditional recurrent architectures may overlook. In contrast, our architecture demonstrates improved detection and accuracy, validating the advantages of federated learning for fraud detection in banking systems.

7. Discussion

Although experimental results have shown remarkable effectiveness, they are not without their challenges. These limitations have a direct impact on their implementation and performance. The datasets we relied on—notably, the European cardholder and IEEE-CIS databases—are widely used benchmarks in fraud detection research; however, due to strict data privacy regulations and confidentiality concerns in the financial sector, access to more recent and diverse real-world banking data is limited. Consequently, these datasets may not capture the full spectrum of fraud types and evolving tactics, which limits the model's applicability and generalizability to current and emerging fraud scenarios.

Moreover, scaling the proposed architecture to handle larger volumes of banking records and an increasing number of clients presents considerable computational challenges. This is particularly true for the quantum-inspired components, which require significant processing power, as well as for the federated learning aggregation process that must operate efficiently across geographically distributed and heterogeneous data sources. These factors could affect processing latency and the ability to maintain real-time fraud detection performance.

Furthermore, model interpretability remains a crucial concern, especially in the banking sector, where regulatory demands and stakeholder trust require transparent decision-making processes. Our architecture's complexity, which combines variational autoencoders with quantum-enhanced recurrent networks, makes it difficult to clearly explain predictions. Using explainable AI techniques and interpretable surrogate models will be vital to give actionable insights to fraud analysts and meet compliance standards.

Finally, ethical considerations like fairness and bias are essential when implementing federated learning systems in financial services. The diversity of client data and possible imbalances in fraud representation can unintentionally introduce bias, impacting the fairness of detection results among different customer groups. Future research should include fairness-aware learning algorithms and perform bias audits to ensure equitable fraud detection performance across various demographic and geographic segments.

In summary, although our work advances the integration of federated learning and quantum-inspired deep learning for fraud detection, it is essential to address practical, ethical, and scalability issues to ensure real-world adoption and impact.

8. Conclusions and Future Directions

Fraud detection is critical in digital banking, and the major challenge is the financial load of suspicions. As criminal conduct continues to evolve, there is a surge for a real-time fraud detection system that maintains data confidentiality and scalability. In this study, we presented a hybrid semi-decentralized FL centered on combining the VAE architecture for feature encoding with up-to-date QLSTM for time-series behavioral prediction. This scheme aims to balance speed and privacy concerns by adopting QML to supply observations of sequence data. However, the study confronts several hurdles, including dataset fairness, scalability, and model interpretability. However, practical limits remain, including data heterogeneity among FinTechs affecting the general model and delay in the aggregation algorithm slowing real-time fraud detection. Furthermore, QML simulations display significant hurdles such as hardware restrictions and model training difficulties.

We are interested in furthering our research into two areas. Firstly, it focuses on increasing FL agility with slope reduction and targeted updating clients to lower computational expenses. Secondly, we aim to explore explainable AI models to enhance the explainability of the fraud detection schemes in combination with cutting-edge Big Data engines for real-time fraud detection.

Author Contributions: Conceptualization, H.A.; methodology, H.A.; software, H.A.; validation, Y.G. and S.E.M.; formal analysis, H.A.; investigation, H.A., Y.G., and S.E.M.; resources, H.A., Y.G., and S.E.M.; data curation, H.A.; writing—original draft preparation, H.A.; writing—review and editing, S.E.M. and Y.G.; visualization, H.A.; supervision, S.E.M. and Y.G.; project administration, Y.G. and S.E.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data is contained within the article.

Conflicts of Interest: Authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. Faroukhi, A.Z.; El Alaoui, I.; Gahi, Y.; Amine, A. Big Data Value Chain: A Unified Approach for Integrated Data Quality and Security. In Proceedings of the 2020 IEEE 2nd International Conference on Electronics, Control, Optimization and Computer Science (ICECOCS), Kenitra, Morroco, 2–3 December 2020; pp. 1–8.
2. Leveraging Generative AI (GenAI) for Fraud Detection and Prevention. Available online: https://www.turing.com/resources/generative-ai-fraud-detection (accessed on 14 September 2024).

3. How Big Is the Fraud Detection and Prevention in 2032 | LinkedIn. Available online: https://www.linkedin.com/pulse/how-big-fraud-detection-prevention-2032-1osef/ (accessed on 14 September 2024).
4. Hernandez Aros, L.; Bustamante Molano, L.X.; Gutierrez-Portela, F.; Moreno Hernandez, J.J.; Rodríguez Barrero, M.S. Financial fraud detection through the application of machine learning techniques: A literature review. *Humanit. Soc. Sci. Commun.* **2024**, *11*, 1130. [CrossRef]
5. Gu, X.; Sabrina, F.; Fan, Z.; Sohail, S. A Review of Privacy Enhancement Methods for Federated Learning in Healthcare Systems. *Int. J. Environ. Res. Public Health* **2023**, *20*, 6539. [CrossRef]
6. Bin Sulaiman, R.; Schetinin, V.; Sant, P. Review of Machine Learning Approach on Credit Card Fraud Detection. *Hum. Cent. Intell. Syst.* **2022**, *2*, 55–68. [CrossRef]
7. Abbassi, H.; Berkaoui, A.; Elmendili, S.; Gahi, Y. End-to-End Real-time Architecture for Fraud Detection in Online Digital Transactions. *Int. J. Adv. Comput. Sci. Appl. (IJACSA)* **2023**, *14*, 749–757. [CrossRef]
8. Kumar, K.; Bhushan, B. Augmenting Cybersecurity and Fraud Detection Using Artificial Intelligence Advancements. In Proceedings of the 2023 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS), Greater Noida, India, 3–4 November 2023.
9. Singh, B. Generative Artificial Intelligence: Prospects for Banking Industry. *Int. J. Res. Eng. Sci. Manag.* **2024**, *7*, 83–86. [CrossRef]
10. Yang, A.; Ma, Z.; Zhang, C.; Han, Y.; Hu, Z.; Zhang, W.; Huang, X.; Wu, Y. Review on application progress of federated learning model and security hazard protection. *Digit. Commun. Netw.* **2023**, *9*, 146–158. [CrossRef]
11. Bhat, P.; Mahonara Pai, M.M.; Pai, R.M. Anomaly detection using Federated Learning: A Performance Based Parameter Aggregation Approach. In Proceedings of the 2023 3rd International Conference on Intelligent Technologies (CONIT), Hubli, India, 23–25 June 2023; pp. 1–6.
12. Wen, J.; Zhang, Z.; Lan, Y.; Cui, Z.; Cai, J.; Zhang, W. A survey on federated learning: Challenges and applications. *Int. J. Mach. Learn. Cyber.* **2023**, *14*, 513–535. [CrossRef]
13. Afriyie, J.K.; Tawiah, K.; Pels, W.A.; Addai-Henne, S.; Dwamena, H.A.; Owiredu, E.O.; Ayeh, S.A.; Eshun, J. A supervised machine learning algorithm for detecting and predicting fraud in credit card transactions. *Decis. Anal. J.* **2023**, *6*, 100163. [CrossRef]
14. Rb, A.; Kr, S.K. Credit card fraud detection using artificial neural network. *Glob. Transit. Proc.* **2021**, *2*, 35–41. [CrossRef]
15. Moreira, M.Â.L.; Junior, C.D.S.R.; de Lima Silva, D.F.; de Castro Junior, M.A.P.; de Araújo Costa, I.P.; Gomes, C.F.S.; dos Santos, M. Exploratory analysis and implementation of machine learning techniques for predictive assessment of fraud in banking systems. *Procedia Comput. Sci.* **2022**, *214*, 117–124. [CrossRef]
16. Askari, S.M.D.S.; Hussain, M.D.A. IFDTC4.5: Intuitionistic fuzzy logic based decision tree for E-transactional fraud detection. *J. Inf. Secur. Appl.* **2020**, *52*, 102469. [CrossRef]
17. Abbassi, H.; Mendili, S.E.; Gahi, Y. Real-Time Online Banking Fraud Detection Model by Unsupervised Learning Fusion. *HighTech Innov. J.* **2024**, *5*, 185–199. [CrossRef]
18. Patil, A.; Mahajan, S.; Menpara, J.; Wagle, S.; Pareek, P.; Kotecha, K. Enhancing fraud detection in banking by integration of graph databases with machine learning. *MethodsX* **2024**, *12*, 102683. [CrossRef] [PubMed]
19. Ranganatha, H.R.; Syed Mustafa, A. Enhancing fraud detection efficiency in mobile transactions through the integration of bidirectional 3d Quasi-Recurrent Neural network and blockchain technologies. *Expert Syst. Appl.* **2025**, *260*, 125179. [CrossRef]
20. Raj, A.T.; Shobana, J.; Nassa, V.K.; Painuly, S.; Savaram, M.; Sridevi, M. Enhancing Security for Online Transactions through Supervised Machine Learning and Block Chain Technology in Credit Card Fraud Detection. In Proceedings of the 2023 7th International Conference on I-SMAC (IoT in Social Mobile, Analytics and Cloud) (I-SMAC), Kirtipur, Nepal, 11–13 October 2023; pp. 241–248.
21. Reddy, V.V.K.; Reddy, R.V.K.; Munaga, M.S.K.; Karnam, B.; Maddila, S.K.; Kolli, C.S. Deep learning-based credit card fraud detection in federated learning. *Expert Syst. Appl.* **2024**, *255*, 124493. [CrossRef]
22. Abdul Salam, M.; Fouad, K.; Elbably, D.; Elsayed, S. Federated learning model for credit card fraud detection with data balancing techniques. *Neural Comput. Appl.* **2024**, *36*, 6231–6256. [CrossRef]
23. Almarshad, F.A.; Gashgari, G.A.; Alzahrani, A.I.A. Generative Adversarial Networks-Based Novel Approach for Fraud Detection for the European Cardholders 2013 Dataset. *IEEE Access* **2023**, *11*, 107348–107368. [CrossRef]
24. Langevin, A.; Cody, T.; Adams, S.; Beling, P. Generative adversarial networks for data augmentation and transfer in credit card fraud detection. *J. Oper. Res. Soc.* **2022**, *73*, 153–180.
25. Sha, X. Research on financial fraud algorithm based on federal learning and big data technology. In Proceedings of the 2024 IEEE 2nd International Conference on Sensors, Electronics and Computer Engineering (ICSECE), Jinzhou, China, 29–31 August 2024.
26. Huang, J.; Chai, J.; Cho, S. Deep learning in finance and banking: A literature review and classification. *Front. Bus. Res. China* **2020**, *14*, 13. [CrossRef]
27. Bockel-Rickermann, C.; Verdonck, T.; Verbeke, W. Fraud analytics: A decade of research: Organizing challenges and solutions in the field. *Expert Syst. Appl.* **2023**, *232*, 120605. [CrossRef]

28. Chen, W.; Yang, K.; Yu, Z.; Shi, Y.; Chen, C.P. A survey on imbalanced learning: Latest research, applications and future directions. *Artif. Intell. Rev.* **2024**, *57*, 137. [CrossRef]
29. Louis, L.D.; Dunton, A.; Saklecha, S.R.; Sivakumar, S.N.K.; Ahmed, A.S.; Sheth, S.; Chang, S.Y. Mitigating Risk in P2P Lending Network: Enhancing Predictions with GenAI and SMOTE. *J. Netw. Netw. Appl.* **2024**, *4*, 48–59.
30. Credit Card Fraud Detection. Available online: https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud (accessed on 26 September 2024).
31. IEEE-CIS Fraud Detection. Available online: https://kaggle.com/competitions/ieee-fraud-detection (accessed on 26 September 2024).
32. Alkhpor, H.K.; Alserhani, F.M. Collaborative Federated Learning-Based Model for Alert Correlation and Attack Scenario Recognition. *Electronics* **2023**, *12*, 4509. [CrossRef]
33. Marfo, W.; Tosh, D.; Moore, S. Network Anomaly Detection Using Federated Learning. In Proceedings of the MILCOM 2022—2022 IEEE Military Communications Conference (MILCOM), Rockville, MD, USA, 28 November–2 December 2022.
34. Tang, Y.; Liang, Y. Credit card fraud detection based on federated graph learning. *Expert Syst. Appl.* **2024**, *256*, 124979. [CrossRef]
35. Priyaradhikadevi, T.; Vanakovarayan, S.; Praveena, E.; Mathavan, V.; Prasanna, S.; Madhan, K. Credit Card Fraud Detection Using Machine Learning Based on Support Vector Machine. In Proceedings of the 2023 Eighth International Conference on Science Technology Engineering and Mathematics (ICONSTEM), Chennai, India, 6–7 April 2023; pp. 1–6.
36. Yang, W.; Zhang, Y.; Ye, K.; Li, L.; Xu, C.Z. FFD: A Federated Learning Based Method for Credit Card Fraud Detection. In Proceedings of the Big Data—bigData 2019: 8th International Congress, San Diego, CA, USA, June 25–30 2019; pp. 18–32.
37. Malik, E.F.; Khaw, K.W.; Belaton, B.; Wong, W.P.; Chew, X. Credit Card Fraud Detection Using a New Hybrid Machine Learning Architecture. *Mathematics* **2022**, *10*, 1480. [CrossRef]
38. Ghosh Dastidar, K.; Jurgovsky, J.; Siblini, W.; Granitzer, M. NAG: Neural feature aggregation framework for credit card fraud detection. *Knowl. Inf. Syst.* **2022**, *64*, 831–858. [CrossRef]
39. Chen, J.; Shen, Y.; Ali, R. Credit Card Fraud Detection Using Sparse Autoencoder and Generative Adversarial Network. In Proceedings of the 2018 IEEE 9th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON), Vancouver, BC, Canada, 1–3 November 2018.
40. Forough, J.; Momtazi, S. Ensemble of deep sequential models for credit card fraud detection. *Appl. Soft Comput.* **2021**, *99*, 106883. [CrossRef]
41. Esenogho, E.; Mienye, I.D.; Swart, T.G.; Aruleba, K.; Obaido, G. A Neural Network Ensemble with Feature Engineering for Improved Credit Card Fraud Detection. *IEEE Access* **2022**, *10*, 16400–16407. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.