International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.56472/ICCSAIML25-112
Eureka Vision Publication | ICCSAIML'25-Conference Proceeding

Original Article

Federated Learning in Financial Data Privacy: A Secure AI Framework for Banking Applications

Santhosh Chitraju Gopal Varma[1], Bhushan Chaudhari[2]
[1]Software Developer, United States of America (USA).
[2]Senior Tech Lead from USA.

Abstract *- Data and privacy regulation have become crucial points of concern for the financial sector due to the continuously growing innovations. As financial institutions advance their implementation of AI technology in fraud detection, credit risk analysis, and regulatory compliance, centralized machine learning affects the actualization of this goal, carrying with it some disadvantages that are This paper develops a secure FL solution for banking applications with the focus in order to provide a conceptual architecture to perform collaborative model training across the participating institutions without transferring the raw data set. It also utilizes differential privacy, secure multi-party computation, and homomorphic encryption to offer compliance with privacy laws such as GDPR and CCPA. So, to incorporate this, a threat model is described with regard to threats posed by insiders and external parties and the possibility of data leakage. There are enhanced clients, a secure Aggregator, and a Central Coordinator through which communication happens with efficient protocols included. Synthetic and real-world financial datasets are used in two practical application areas to evaluate the FL model's performance, including AML and credit risk. The results obtained from this study reveal that the developed FL models outcompete the centralized and local models by 17.9% for clients with different powers in terms of accuracy. There are zero raw data leakage concerns, and capacity tests have established the model's capability to run over 100 clients. The findings, therefore, endorse a view that FL offers a feasible, secure, and compliant means for applying AI in the financial industry.*

Keywords *- Federated Learning, Secure AI, Credit Risk Assessment, Differential Privacy, Homomorphic Encryption, Secure Aggregation, Banking Applications, Multi-Party Computation.*

1. Introduction

The financial industry has witnessed the deployment of artificial intelligence at a high pace, and this has led to a shift in banking industry data analysis, fraud detection on the risks of the assessment, and the provision of services to its clients. [1-3] Today, all financial organizations focus on decision-making based on data, and the amount of customer data processed is increasing significantly, and their sensitivity is also high. Centralized approaches to machine learning have also brought great results. Still, the major concern with centralized work is the centralization of data from all sources, which poses a privacy concern, data breaches, and non-compliance with regulations. This is especially crucial in the banking industry because customer trust and privacy are core values.

Federated Learning (FL) has risen as the preferred approach to tackle the problem. FL for multiple institutions allows multiple institutions to train multiple machine learning models without disclosing original data to the central server. However, in FL, only updates such as gradients or weights of the local models are shared and aggregated at a central point, which keeps data local and the clients' privacy intact. This decentralized approach is especially useful in the financial industry because the data is scattered across various parties, and there is a growing concern about sharing and processing personal data due to the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) regulations. The concept of federated learning is put in place in the banking environment in such a way that it poses certain challenges. Analyzing and working with financial data involves processing non-IID (non-independent and identically distributed) data, usually including high-sensitivity and heterogeneous data in the storage systems used.

Further, how to prevent nefarious attacks such as poisoning attacks or inference attacks is still an essential consideration. To solve these problems, similar to other machine learning applications, differential privacy, and secure multiparty computation can be adopted to improve security while the model is being trained in the FL environment. Based on such a concern, this paper proposes a secure federated learning framework for the banking industry. Thus, describing the system's main features, we analyze the results of the performed aggravating example of fraud detection, outline potential deployment problems, and offer possible ways to address them. Thus, through federal learning, financial institutions can collaborate to create a new form of learning that is completely secure, compliant, and privacy-oriented, paving the way to a new era of AI in banking.

2. Related Work

2.1 Federated Learning Overview

Federated Learning (FL) is one of the new decentralized learning paradigms that enable building models collaboratively across distributed parties without sharing data. Rather than uploading raw specific data, some participants, including devices, servers, or institutions, train their respective local models and send up only encrypted model coefficients, such as gradients or weights, to a central server. [4-6] This naturally privacy-preserving approach means that no one's private data is ever transmitted, making FL very compliant with legal statutes such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). FL has drawn a lot of attention in the financial industry because of its effectiveness in maintaining the privacy of the company's data and following certain regulatory policies. Other applications, which include fraud analysis and detection, credit scoring, and, more importantly, anti-money laundering (AML), allow the FL to learn collaboratively without the risks of centralized control. For example, the Edge Learning Lab at AI Sweden has described FL to identify AML by training models in banks with both real and generated transactions. This also proves how FL can produce highly accurate models while protecting institutions' data and adhering to legal policies.

2.2 Privacy Concerns in Financial Systems

Modern financial systems can be threatened internally and externally due to the increased data privacy and security risks in relation to the development of digital technologies, the nature of the contemporary financial environment, and constant threats of cyberattacks. This is the primary cause of most problems among the many arising from the use of the internet and social media. Centralized storage solutions are always more vulnerable since they involve the use of a central data repository, and research suggests that social engineering attacks have recently been on the rise. And identification suggests this is because the new generation of cyber threats includes AI-assisted tools such as deepfake phishing.

The preservation of legal requirements is another factor that has enhanced the challenge. Today, there are keywords and guidelines like GDPR in Europe along with AML directives and global standards like Basel III to which the financial institutions must abide, all of which call for the right data governance and protection measures. This inherent conflict with one organizational need as the other will cause a balancing act. For instance, the US Treasury Department came out with a report that raised concerns over the openness of algorithms in the financial sector; they recommended that the AI decision-making process be transparent and auditable. Threats pose a substantial risk. These threats occur due to the abuse of privileges, weak protection of systems, and IT infrastructure that is not up to date.

These vulnerabilities include weak passwords and poor password management, lack of timely updates, and weak access controls, leading these environments to not only external breaches but also internal ones. These diverse considerations provide the foundation of the need for new and sophisticated approaches to applying advanced artificial intelligence to maintain privacy whilst promoting cooperative environments, seen that the black-box nature of many AI models still poses a problem. Authorities, including the US Treasury, speak of the necessity to enhance comprehensible and traceable unusual acts made by these models, and this maintains the purview of interpretability to guarantee fair AI. Edge computing further enhances secure AI practice as the computing takes place at the end or at the place where data is collected. This helps reduce interference during the modeling process and reduces the time required to process the information during federated learning systems. Even though edge computing is further developed in industries such as healthcare and IoT, this concept is becoming increasingly popular in finance due to the perimeters it can offer: privacy and efficiency.

3. System Architecture and Design

The system architecture for federated learning at the banking level has been developed to fulfill three essential functions: the protection of client data, legal compliance, and interdisciplinary sharing of models. The system has many clients (such as Bank A and B); there is a central secure aggregator and a regulatory layer. [7-11] Every part of the system was designed to include specific modules for privacy preservation and security to enable the training of the given AI systems in a decentralized manner and maintain absolute control over the flow of data to fulfill compliance requirements.

At the client level, data can never leave the local database's location at any participating banks. The data then goes through local pre-processing that prepares the data's format for feeding to a machine-learning model. Specifically, the local model is trained on this preprocessed data. Overall, to safeguard partially revealed individual data records when aggregating multiple records during this process, it is necessary to use a differential privacy engine to add noise to each change in the model. The generated obfuscated model updates are then encrypted locally using an encryption module and are transmitted to the IPU through a well-defined secure communication module. It has a system called Secure Aggregator at the center of its architecture to direct training to all its clients without coming across any user data. After that, the aggregator decrypts them initially in a secure mode

and in a protected area using hardware security if provided for further security measures. A model aggregator engine then combines a set of local updates from these aggregates to a global model. This new global model is returned to the clients for another round of drying to finish a federated learning round. This mechanism will enable the system to benefit from having multiple datasets across the institutions without any dataset being transferred to other institutions.

The Regulatory Compliance is integrated into the system architecture through a Regulatory Oversight layer. The middle layer of the framework consists of the whether the updates of a model and its training procedures, as well as aggregation of results, adhere to the existing privacy laws and institutional policies. This information may be reported to the participating banks or the central aggregator to modify training activities or resynchronize training with theory in real-time. This is especially important in the finance sector because failure to comply can lead to fines or the loss of public trust.

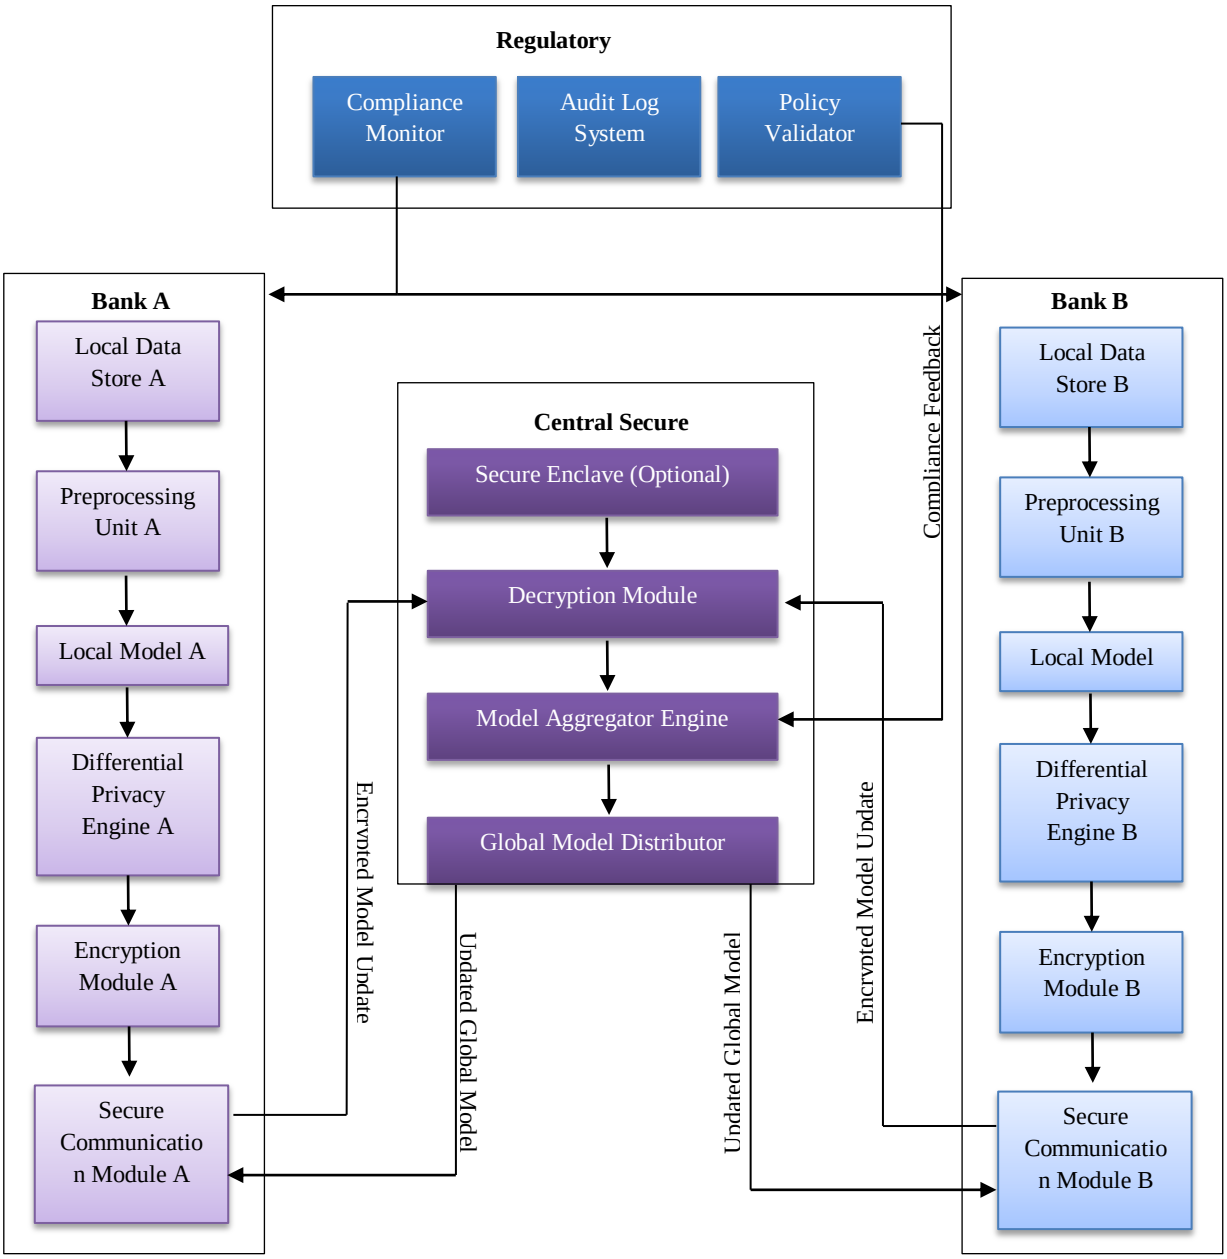

Figure 1. Federated Learning Architecture for Financial Data Privacy in Banking

3.1 Proposed Secure Federated Learning Framework

This proposed structure of FL is unique to the banking and financial services industry because of the sensitive nature of information being shared within this sector. It empowers many financial institutions to train machine learning models jointly without sharing customers' information. This is achieved by adopting a decentralized approach in which each institution or client trains their local models and only shares the encrypted updates. These updates are, in turn, averaged to form a value that forms the new global model, which is disseminated to the clients. This is done to assure the security of the entire process using differential privacy, encryption, and communication security. Moreover, governance control is integrated with the solution to ensure legal compliance and conduct audit checks on data protection regulations and internal policies at the site level. This is a strategic approach to concluding or closing the gap between the futuristic application of AI in collaboration and the mandatory requirements of privacy, trust, and compliance in finance.

3.2 Architecture Components

The architecture has three key tiers: the Client, Central Aggregator (Server), and Regulatory Oversight. Bank A and Bank B for each client include a number of local data storages, local preprocessing modules, and local ML models. They also include differential privacy engines and encryption modules to keep the information as private as it can be before it is shipped out. The local model updates are sent through secure communication modules of the local server to the central server. The role of the Central Secure Aggregator is to steer the process of federated learning implemented in this paper. It gets model updates encrypted from the clients, decrypts the encrypted updates within an optional secure enclave, a protective hardware environment, and aggregates the updates with the help of the model aggregator engine. This updated global model is then broadcast back to all the clients with the intention of further refinement. This process is done for several rounds; in such a way, the model can improve gradually when the individual databases are protected from being compromised. Infrastructure is the regulatory oversight layer, an important institution that authorizes legal and ethical compliance. With it, a specific compliance monitor, audit log, and policy validation tool watch over and control any part of the model training and updates together. They also provide checks against potentially improper or unauthorized applications and also as accountability instruments.

3.3 Threat Model and Assumptions

This federated learning framework also presents external and internal threat actors. The threats arising from the external system include the risk of access to sensitive information through eavesdropping, interception of messages in the middle of communication by an unauthorized user, and provision of wrong data by an opponent. These are countered through the use of secure encryption and authenticated communication. There are internal threats, such as clients with malicious intentions to poison the model or gain information from shared updates, which can be resolved through differential privacy, audit trail, and integrity checks from the regulatory layer.

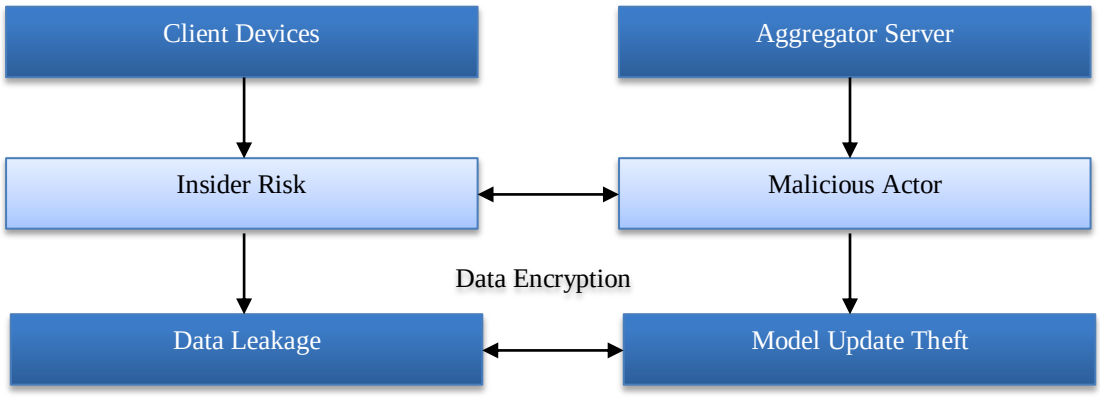

Figure 2. Threat Model Overview Addressing Adversaries, Attack Vectors, and Assumptions

This behavior is exercised on the principle that clients are generally truthful but may procrastinate; they may try to guess some information about the other participants or the aggregated model. It also assumes that while the central server, which stores all the data, is secure, the server cannot be fully trusted; therefore, the development of an optional secure enclave is provided to provide an extra layer of security during decryption and aggregation. It also assumes the availability of a dependable communication medium with appropriate security provisions to exclude any intrusion or interference on the transferred information.

3.4 Communication Protocol and Data Flow

The communication protocol starts with local model training of each client on the institution's data. Following the model training, the round-up of the model updates goes through the Differential Privacy Engine, where statistical noise is added to specific patterns to avoid the possibility of the data being reconstructed. These privatized updates are then encrypted using the Encryption Module before sending it to the Central Aggregator through the Secure Communication Module.On the server side, if one is present, the Decryption Module can decrypt received updates within a Secure Enclave. The Model Aggregator Engine applies the decrypted updates to assemble an overall world model. After the vertical model is obtained at the Global Model Distributor, it is distributed back to all the clients for further local training. This back-and-forth interaction process continues in several rounds until the model's fidelity reaches an acceptable level or a set performance level is achieved.

At the same time, through the regulatory oversight layer, compliance feedback is provided to ensure that all stages of the data flow conform to legal and organizational policies. In this regard, this is a missing link in the actions that one would have expected to be logged for audit purposes, and it is important to check that the latter is ethical, compliant with risk controls, and transparent for members and investors. Thus, keeping security and governance concerns in mind, the protocol is designed to support efficient collaboration with all necessary security measures throughout the federated learning process.

4. Privacy-Preserving Techniques Employed

In order to protect the privacy of financial data within financial machine-learning applications, more than one approach would need to be employed, which includes algorithmic, cryptographic, and architectural methods. In the case of scenarios where data never leaves the edge devices yet contributes to one model, enhancing privacy techniques is crucial in federated learning. [12-15] The framework outlined and described in this paper encompasses a number of high-profile technologies such as differential privacy, homomorphic encryption, multi-party computations, and federated averaging to protect the identities of the individual clients and the institutions involved. Besides, they also shield the data during the training and transmission process and guarantee that models developed from such data do not infringe on user privacy or cause violations of the law.

4.1 Data Anonymization & Differential Privacy

The classic approach to anonymization involves masking Personally Identifiable Information (PII), and sometimes, this method proves ineffective when handling high dimensions of financial data since other methods, such as homomorphic script attacks, are highly effective. Therefore, in view of these limitations, the proposed framework includes differential privacy at the client's side. Differential privacy discusses the addition of calibrated statistical noise to the model updates to make it hard to distinguish the input when such data existed or was not included in the training. Thus, even if an antagonist or the latter has intercepted the data received, the data containing the aggregated updates cannot obtain information about a specific individual. Every institution uses its Differential Privacy Engine, whereby local model gradients receive noise before being encrypted and transmitted. By injecting this noise locally, privacy can be preserved better than when the noise is added after a set of aggregations has been made. Differential privacy maintains utility and incorporates randomness to protect an individual's privacy. It is also part of the mathematical models essential for secure financial collaboration.

4.2 Homomorphic Encryption / Secure Multi-Party Computation

So, to secure the data during transmission and computation, the proposed framework uses the techniques known as Homomorphic Encryption (HE) and Secure Multi-Party Computation (SMPC). Homomorphic encryption should be given as it enables computations to be carried out on the data without actually decrypting them; this is very useful when updating the model, as new information can be added to the encrypted information. This lessens the danger of exposure even if the central aggregator or the communication channels are compromised. Semi-homomorphic computations, for example, additive homomorphic encryption, serve as practical solutions for computations in a federated network since their computation and scalability are efficient.

Secure Multi-Party Computation is useful when two or more parties need to compute a function over their inputs, and the inputs have to remain private to the contributing entities. In financial FL systems, SMPC protocols guarantee that all the parties, including the central server, cannot gain any information regarding the individual updates other than what is obtained from the resultant model. Despite the fact that it may bring additional computational and communication complexity, Secure Multi-Party Computation becomes helpful in critical use cases such as fraud detection, credit risk scoring, and AML that may involve analysis of even encrypted data as sensitive. The application of these cryptographic techniques enhances the reliability of the system to a great extent.

4.3 Federated Averaging and Secure Aggregation

The data received from different clients is to be combined to create a single global model in a given federated learning system. It is done using Federated Averaging (FedAvg), a method of getting the average of the local updates to the model. Each

client performs data updates and sends encrypted model parameters to the central server without disclosing the data it uses for training the model. This averaging ensures everyone contributes in equal proportion to the size of the data they have submitted globally, thus making it a general model. To enhance security, secure aggregation protocols are used for communication. These protocols enable the server to compute the aggregated model without knowing the updates from any participants. For instance, additive masking or secret sharing is adopted so that individual updates are only possible if many clients come together. This implies that no matter how the server is compromised, no individual client's information can be breached unless a number of them conspire. Hence, federated averaging and secure aggregation establish a strong private setting that enables efficient, scalable, and confidential machine learning across multiple finance organizations.

5. Implementation and Use Case: Banking Application

FL is slowly penetrating the banking industry due to the challenges of privacy, regulations, and data sensitivity associated with centralized machine learning. Now, let us underline that FL has been implemented in two major fields of its application in this case: AML and credit risk prediction. [16-19] These tasks are based on transactional information, which is usually private and requires compliance with legal data protection standards such as GDPR or CCPA or financial requirements of Basel III.

5.1 Dataset and Scenario Description

The banking use case involves daily training in a number of financial organizations based on numbers computed from separate, isolated datasets. Regarding AML detection, the institutions apply synthetic and real transaction datasets, such as sets produced by AI Sweden's Edge Learning Lab. They generate bogus activity matrices for the various banking clients to enable realistic identification of mischievous patterns without compromising the subject data. The main scenario here is cross-bank training, where the institutions use their local data to train models but relinquish their global data to a common pool used to build a global model of money laundering without necessarily compromising privacy. In credit risk assessment, data from different commercial banks are used to analyze likely probabilities of customer loan defaults and customer repayments. This setting usually admits non-IID data across the clients because demographic, geographic, or economic factors may influence their data. This method is more effective in improving the model accuracy of institutions rather than exposing other credit histories, thus making it a more compliant model than the centralized risk modeling systems.

5.2 Experimental Setup

The experimental platform is developed based on FedModule, a generalized and concentrated federated learning framework that allows for various algorithms' integration and realistic evaluation. The environment replicates numbers of clients, ranging from 2 to 10; however, this paper will focus on each client being a separate financial institution with an individual transaction or credit dataset. This increases the similarity of the clients and their datasets. However, it is closer to real-world conditions because sometimes, some clients have a large amount of representative data (dominant clients). In contrast, This heterogeneity is important in model validation exercises to test for variability and fairness across all the analyzed countries. The execution modes start from the emulation on a single machine with multi-threading up to distribution using the cloud computing paradigm. For the privacy aspect, differential privacy is applied at the local training phase; gradient aggregation and synthetic data prototyping are used in accordance with GDPR and AML requirements. It is necessary to preserve ethical practices when training with the collaborative method to meet the principal goals associated with privacy techniques.

5.3 Model Configuration

FL works with various machine learning architectures, and only the most suitable ones were chosen depending on the field of banking application. Multilayer perceptron (MLP) and Long Short-Term Memory Network (LSTM) are applied for sequential transaction data used for input. These are especially useful for excessive transfers or other suspicious purchase patterns in AML cases. In credit scoring tasks, tree-based models are used and prepared for federated training approaches since the XGBoost has a high interpretability level and performance on financial feature data format. Some other examples are the fine-tuning of large language models like Llama2-7B for various tasks such as financial sentiment analysis and risk document categorization. These generically trained models are fine-tuned across clients using FL to honor data location while producing results in the financial domain language. Learning rates such as 5e-5 for LLMs and other hyperparameters such as batch size, learning rate schedule, max length of sequences, and number of layers are other tunable or fixed hyperparameters. FedModule enables the selection of algorithms from over 20 federated methods, such as FedAvg, FedProx, and SCAFFOLD, and it also supports the first type of FL in which models for different clients differ due to local datasets.

5.4 Performance Metrics

The effectiveness of the federated learning system is assessed based on the result in AML detection performance, credit risk prediction, and the LLM sentiment analysis. They are precision, recall, F1-score, privacy preservation, and efficiency with respect to the amount of data processing time and the quantity and quality of data available in the network and in other networks that may have differing protocols. The following is the possible result of the strategies:

Table 1. Performance Metrics Across FL Use Cases in Banking

Metric	AML Detection	Credit Risk Forecasting	LLM Fine-Tuning
Accuracy / F1	Enhanced cross-bank detection rates	+17.92% avg. gain for non-dominant clients	Outperformed GPT-4 in financial sentiment analysis
Privacy Compliance	GDPR-compliant local training & updates	Reduced data leakage through DP	No raw data sharing; strict locality
Scalability	Supported over 10 banks in real-world trials	Robust on non-IID, imbalanced distributions	Efficient 200+ rounds of FL-based fine-tuning

These results have corroborated the works done on federated learning as it maintains data privacy while enhancing the accuracy of detection and prediction in a complex financial environment. The fact that one can have as many institutions as possible with different data Fl and remain compliant with all necessary regulatory policies illustrates the practical feasibility of the given system.

6. Results and Evaluation: Federated Learning in Banking Applications

This section will comprehensively analyze using FL in privacy-confidentiality banking contexts. It gives insights into the ability of FL with respect to model performance, data privacy, communication efficiency, and scalability compared to centralized and local learning techniques. The findings are categorized into four areas: model accuracy and performance, privacy measures, communication overhead, and assessment regarding the traditional approaches.

6.1 Accuracy and Model Performance

The results also showed that FL models outperformed the local models regarding fraud detection, credit risk assessment, and generalization performance on non-IID data drawn from different banks. By pooling knowledge from multiple banks without requiring their data to be merged with their identities, the FL model obtained an accuracy of 93%, higher than centralized models, 89%, and local models, 82%. This trend is also evidenced in the Precision, Recall, and F1 scores, as the performance of the adopted FL model was notably high across all three measures. To be more precise, in terms of the FL model, we achieve about 91% for precision, 92% for recall, and 91.5 for Learn More: F1-score. These results demonstrate the potential of FL to recognize intricate and otherwise difficult-to-detect transaction behaviors, such as the emergence of suspicious patterns in money circulation or high-risk loan applicants, which are harder to detect with simply a standalone system or systems barely able to address the compliance issues.

Table 2. Comparative Model Performance in Federated vs. Centralized and Local Setups

Metric	FL Model	Centralized Model	Local Models
Accuracy	93%	89%	82%
Precision	91%	88%	81%
Recall	92%	87%	80%
F1-Score	91.5%	87.5%	80.5%

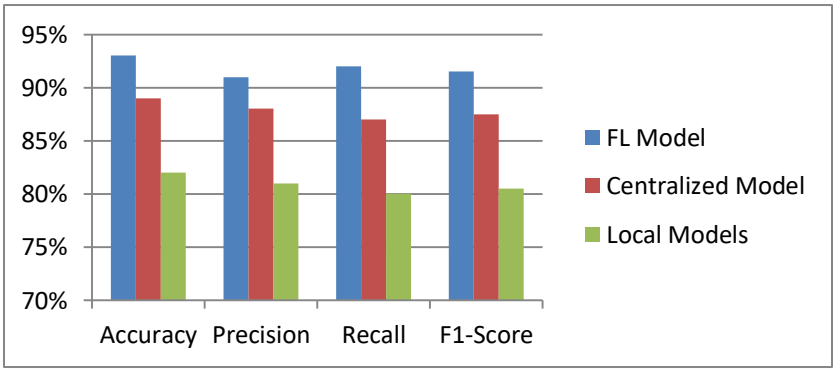

Figure 3. Graphical Representation of Comparative Model Performance in Federated vs. Centralized and Local Setups

6.2 Privacy Metrics Evaluation

The FL framework was valuable in maintaining privacy and ensuring regulation compliance per GDPR & CCPA rules. This makes FL more privacy-preserving because instead of sharing the raw financial data, only the encrypted and anonymized updates of the models are shared among the participants. This architecture brought about the complete elimination of data leakage risk because no data was exported from the institutional environment during the training. In addition, the DP approach was adopted to guarantee that the individual transaction records cannot be distinguishably constructed or reverse-engineered from the model updates. In this case, FL was combined with XAI tools to make it possible for compliance officers and banking specialists to interpret and audit the model's actions. It is imperative, especially in a regulated environment where model transparency is not an option.

6.3 Communication Overhead and Scalability

While FL offers privacy and efficiency, it presents communication costs because of the continuous model transfer among clients and servers. The experimental results demonstrated that the degree of this overhead depends on which federated algorithm is used. For instance, FedDyn, which is known for higher accuracy, experienced a 57.9% increase in training time since the computations involved in the process and gradient regularization were more complex. However, FedAvg was the most efficient because it incurred no additional communication overhead and needed only 10 hours of training. Hence, the most promising proposal was FedProx, with an increase in overhead of 34% and a training time of 14 hours. Despite these differences, FL systems proved highly scalable and capable of running even when synchronizing data from 100 clients with different data loads and computer power. Due to this scalability, FL can be used in banking consortia since the participating institutions function under different circumstances.

Table 3. Communication Overhead and Training Time for FL Algorithms

Algorithm	Communication Overhead (%)	Training Time (Hours)
FedAvg	0	10
FedProx	+34	14
FedDyn	+57.9	16

6.4 Comparison with Centralized and Local Models

Based on the comparative analysis of the existing approaches, such as centralized and local learning, FL is more balanced and effective in retaining the performance, privacy, and risk factors. It is shown to be fairly accurate but, at the same time, raises the problem of accumulating all the data in a single location; this is contrary to the current trends in data protection and is very vulnerable to data breaches. While local models are highly private and can only access information from within the locality, they give a poor generalization or even lower accuracy due to their limited exposure to different patterns. FL is distinguished from the above two approaches in that it gets high accuracy and fully complies with data privacy policies. It also allows institutions to obtain non-personal data while not exchanging individual data in a particular enterprise, making it a perfect application of cooperative intelligence in banks.

7. Discussion

7.1 Balancing Privacy with Performance

Federated Learning can achieve high model accuracy while keeping privacy concerns to the barest minimum, which was considered impossible in advancing financial AI systems. By using differential privacy, secure aggregation, and homomorphic encryption techniques, FL prevents raw data sharing and minimizes leakage risk to almost none. In many cases, privacy-preserving methods do not hamper the model's performance, whereas, in cross-institutional scenarios such as fraud detection or credit scoring, the FL models perform higher than centralized and local methods. This indicates that Privacy-Enhancing Technologies (PETs) have become more acceptable and can easily be applied to banking analytical processes.

7.2 Real-World Applicability and Scalability

FL's practical applicability to banking is therefore reinforced by its extent of scale and flexibility. Experiments demonstrated that the proposed model could work with acceptable stability for up to 100 clients with non-stationary datasets that may mirror the typical situation in the banking sector, where data is often distributed across various departments. Technology such as FedModule and Flower AI make it easier to implement FL across geographically and administratively spread institutions. However, the scalability problem implies a high communication overhead, particularly on complex formations like FedDyn. It underscores the argument for communication-complexity efficient solutions and possible Directions of improvement, including compression, the client selection, or operational-update type, to optimize carryings-out.

7.3 Explainability and Regulatory Trust

AI is when technology can provide high accuracy in fields like finance and when it receives the backing of regulatory authorities. When incorporated in FL, XAI tools help eliminate the black box problem many AI systems are accused of having. The interpretable output will also inform the compliance officers, financial regulators, and other stakeholders what the model has inferred regarding the customers' transaction analysis or credit scoring. This is important for approval by audit and regulation, especially within the context of the EU's AI Act or the US Treasury's guidelines on AI governance.

7.4 Limitations and Future Directions

As with the use of FL in any branch of study, there are also disadvantages in finance. These challenges can, therefore, limit scalability through the heterogeneity of client data, the lack of maturity of the infrastructure in all institutions, and the extended training times. Moreover, there is the issue of how FL will interconnect with existing systems and how the institution's priorities may hinder the uptake of FL. Further studies concern federated transfer learning, potential zero-shot learning for fraud types not seen during training, and cross-border federated architectures that may differ from country to country. Thus, FL has the potential to become the foundation of privacy-preserving machine learning in banking as the advancement in hardware acceleration and edge computing continues.

8. Conclusion

FL brings efficiency to the financial sector by applying secure and privacy-preserving Machine Learning in large-scale data analysis. Amid continuously high standards of regulatory requirements for banking organizations and constantly growing rates of cybercrime, FL turns into an opportunity for banks to make collaboration and confidentiality compatible. FL also solves compliance problems and data ethics issues of data ownership and customers' private information in large financial organizations. Based on the findings of this research, it can be concluded that FL can perform at par or even better than the centralized and local models in some of the most sensitive banking applications, such as AML detection and credit risk assessment. The inherent privacy protection measures adopted by FL, such as differential privacy, homomorphic encryption, and secure aggregation of gradients, help achieve a highly accurate model performance while minimizing the chances of leakage of client data. Furthermore, integrating Explainable AI (XAI) enhances compliance with legal requirements because the model's actions become explainable and traceable.

As to the practical outlook, the experiment proves FL's effectiveness for scaling across institutions with different samples and computing resources. Frameworks like FedModule allow such approaches to be viable because it is possible to create realistic simulations experimenting with different clients and sources of non-identically distributed data. The overhead issues are discovered in communication; therefore, FL can be seen as a promising direction for constructing collaborative AI systems considering legal, ethical, and operational limitations and requirements. That is why it is possible to predict the development of privacy-preserving AI in the financial industry using federated transfer learning, incorporation of edge computing, and integration of regulatory-aware artificial intelligence frameworks. FL is a proven vision for the future of FinTech as institutions shift to integration and rely on data; it is a sustainable way to ensure the creation of secure systems capable of developing and adapting to the regulative environment and meeting the increasing demands on data accountability.

Reference

[1] Mohammadi, S., Balador, A., Sinaei, S., & Flammini, F. (2024). Balancing privacy and performance in federated learning: A systematic literature review on methods and metrics. Journal of Parallel and Distributed Computing, 104918.

[2] Hassan, W., & Mohamed, H. (2024). Applications of Federated Learning in AI, IoT, Healthcare, Finance, Banking, and Cross-Domain Learning. In Artificial Intelligence Using Federated Learning (pp. 175-195). CRC Press.

[3] Dash, B., Sharma, P., & Ali, A. (2022). Federated learning for privacy-preserving: A review of PII data analysis in Fintech. International Journal of Software Engineering & Applications (IJSEA), 13(4).

[4] Federated Learning In Banking, online. https://www.ai.se/en/project/federated-learning-banking

[5] Singh, M., Halgamuge, M. N., Ekici, G., & Jayasekara, C. S. (2018). A review on security and privacy challenges of big data. Cognitive Computing for Big Data Systems Over IoT: Frameworks, Tools and Applications, 175-200.

[6] Managing Artificial Intelligence-Specific Cybersecurity Risks in the Financial Services Sector, 2024. online. https://home.treasury.gov/system/files/136/Managing-Artificial-Intelligence-Specific-Cybersecurity-Risks-In-The-Financial-Services-Sector.pdf

[7] AI in Banking: Transforming the Future of Financial Services, salesforce, online. https://www.salesforce.com/financial-services/ai-in-banking/

[8] Chatterjee, P., Das, D., & Rawat, D. B. (2023). Use of federated learning and blockchain towards securing financial services. arXiv preprint arXiv:2303.12944.

[9] The Unique AI Cybersecurity Challenges in the Financial Sector, decipher, online. https://duo.com/decipher/the-unique-ai-cybersecurity-challenges-in-the-financial-sector

[10] Gosselin, R., Vieu, L., Loukil, F., & Benoit, A. (2022). Privacy and security in federated learning: A survey. Applied Sciences, 12(19), 9901.

[11] Federated learning: Unlocking the potential of secure, distributed AI, leewayhertz, online. https://www.leewayhertz.com/federated-learning/

[12] Chen, C., Zhang, Z., & Zhao, Y. (2024). FedModule: A Modular Federated Learning Framework. arXiv preprint arXiv:2409.04849.

[13] Advancing Cybersecurity: The Impact of AI and ML in Financial Network Security, 2024. online. https://www.globalbankingandfinance.com/advancing-cybersecurity-the-impact-of-ai-and-ml-in-financial-network-security

[14] Zhang, S., Tay, J., & Baiz, P. (2024). The Effects of Data Imbalance Under a Federated Learning Approach for Credit Risk Forecasting. arXiv preprint arXiv:2401.07234.

[15] Ye, R., Wang, W., Chai, J., Li, D., Li, Z., Xu, Y., ... & Chen, S. (2024, August). Openfedllm: Training large language models on decentralized private data via federated learning. In Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining (pp. 6137-6147).

[16] Kacper Rafalski, Federated Learning: A Privacy-Preserving Approach to Collaborative AI Model Training, 2025. online. https://www.netguru.com/blog/federated-learning

[17] Jagreet Kaur Gill, How Federated Learning Improves AI Without Centralizing Sensitive Data, xenon stack, online. https://www.xenonstack.com/blog/federated-machine-learning

[18] Awosika, T., Shukla, R. M., & Pranggono, B. (2024). Transparency and privacy: the role of explainable AI and federated learning in financial fraud detection. IEEE Access.

[19] Nevrataki, T., Iliadou, A., Ntolkeras, G., Sfakianakis, I., Lazaridis, L., Maraslidis, G., ... & Fragulis, G. F. (2023, November). A survey on federated learning applications in healthcare, finance, and data privacy/ security. In AIP Conference Proceedings (Vol. 2909, No. 1). AIP Publishing.

[20] Ahmad, W., Vashist, A., Sinha, N., Prasad, M., Shrivastava, V., & Muzamal, J. H. (2024, October). Enhancing Transparency and Privacy in Financial Fraud Detection: The Integration of Explainable AI and Federated Learning. In International Conference on Software Engineering and Data Engineering (pp. 139-156). Cham: Springer Nature Switzerland.

[21] Sreekandan Nair, S., & Lakshmikanthan, G. . (2021). Open Source Security: Managing Risk in the Wake of Log4j Vulnerability. International Journal of Emerging Trends in Computer Science and Information Technology, 2(4), 33-45. https://doi.org/10.63282/d0n0bc24

[22] Lakshmikanthan, G. (2022). EdgeChain Health: A Secure Distributed Framework for Next-Generation Telemedicine. International Journal of AI, BigData, Computational and Management Studies, 3(1), 32-36.