International Journal of AI, Big Data, Computational and Management Studies
Noble Scholar Research Group | Volume 2, Issue 4, PP. 1-9, 2021
ISSN: 3050-9416 | https://doi.org/10.63282/30509416/IJAIBDCMS-V2I4P101

Federated Learning in Cloud-Based Financial Applications: A Decentralized Approach to AI Training

Prof. Antonio Ricci,
University of Milan, AI & Machine Learning Institute, Italy.

Abstract: Federated Learning (FL) has emerged as a promising paradigm for training machine learning models in a decentralized manner, particularly in cloud-based financial applications. This paper explores the application of FL in the financial sector, highlighting its potential to enhance data privacy, security, and model performance. We begin by providing an overview of FL and its key components, followed by a detailed discussion of the challenges and opportunities in the financial domain. We then present case studies and empirical evaluations to demonstrate the effectiveness of FL in various financial applications, such as fraud detection, credit scoring, and algorithmic trading. Finally, we discuss the future directions and open research questions in this field.

Keywords: Federated Learning, Machine Learning, Data Privacy, Financial Applications, Decentralized Learning, Security, Communication Efficiency, Data Heterogeneity, Regulatory Compliance, Scalability

1. Introduction

The financial industry is one of the most data-intensive sectors, generating vast amounts of transactional, customer, and market data on a daily basis. This data is incredibly diverse, ranging from simple transaction records to complex financial instruments and customer behavioral patterns. The ability to leverage this data for predictive analytics and informed decision-making has become a critical competitive advantage for financial institutions. Advanced analytics and machine learning (ML) models can help organizations identify trends, predict market movements, detect fraud, and personalize services, thereby enhancing operational efficiency and customer satisfaction.

However, traditional centralized machine learning approaches face significant challenges, particularly in terms of data privacy and security. In a centralized model, all data must be gathered and stored in a single location, which can expose sensitive information to potential vulnerabilities and breaches. Financial institutions are subject to stringent regulations and compliance standards to protect customer data, and the risk of data leakage or unauthorized access can have severe legal and reputational consequences. Additionally, the sheer volume and complexity of financial data make it difficult to efficiently manage and process in a centralized system, often leading to issues with scalability and performance.

Federated Learning (FL) offers a decentralized approach to machine learning that addresses these challenges by allowing multiple parties to collaboratively train a model without sharing their raw data. In a federated learning setup, each participating entity, such as a bank or financial service provider, retains control over its own data. Instead of transferring data to a central server, these entities locally train the model on their own datasets and then share only the updates to the model parameters with a coordinating server. This method ensures that sensitive data remains within the secure boundaries of each organization, significantly reducing the risk of data breaches and privacy violations. Furthermore, federated learning can improve the robustness and accuracy of the models by incorporating diverse data from multiple sources without the need for direct data exchange, thereby enhancing the overall value and effectiveness of the machine learning process in the financial sector.

2. Federated Learning: An Overview

2.1 Definition and Key Concepts

Federated Learning (FL) is a distributed machine learning technique that enables multiple entities to collaboratively train a shared model while keeping their raw data localized. Unlike traditional centralized learning methods, where data is aggregated in a central repository for model training, FL allows computations to be performed on decentralized devices or servers, ensuring that sensitive data never leaves its source. This is particularly beneficial for privacy-sensitive applications such as finance, healthcare, and IoT systems. In FL, participating entities (referred to as clients) train local models on their own data and then transmit only the model updates, such as gradients or parameter changes, to a central server. The server aggregates

these updates and refines the global model, which is then redistributed to the clients for further iterations of training. This iterative learning process enhances model generalization while addressing privacy and data-sharing concerns.

2.2 Architecture

The architecture of Federated Learning primarily consists of three key components: clients, a central server, and a communication protocol.

1. Clients: These are the data owners that participate in the training process by locally updating the model. Clients can be mobile devices, edge computing nodes, cloud servers, or financial institutions with proprietary datasets. Each client independently trains a model on its private dataset and sends the computed updates to the central server.
2. Server: The central server orchestrates the learning process by receiving model updates from clients, aggregating them, and distributing the improved model back to the clients. The aggregation function, often implemented using techniques such as Federated Averaging (FedAvg), plays a crucial role in ensuring that the global model benefits from diverse client data while maintaining robustness and fairness.
3. Communication Protocol: Efficient and secure communication between the clients and the server is essential for FL. Communication protocols handle data transmission, encryption, model update compression, and error correction mechanisms. Given the decentralized nature of FL, optimizing communication efficiency is critical to reducing latency, preserving bandwidth, and ensuring secure data transmission.

2.3 Types of Federated Learning

Federated Learning can be categorized into three main types based on how data is distributed among clients:

1. Horizontal Federated Learning (HFL): This approach is used when different clients have datasets with the same feature space but different data samples. For instance, multiple banks operating in different regions may have customer transaction data with identical features such as income, spending behavior, and credit history, but distinct customer bases. HFL enables these banks to collaboratively train a model without exposing individual customer data.
2. Vertical Federated Learning (VFL): In this approach, clients possess datasets that share the same users but have different feature sets. For example, a bank and a credit card company may have overlapping customers, but the bank might store information about account balances and transactions, while the credit card company maintains credit scores and purchase histories. VFL allows both institutions to train a model that benefits from a richer feature set without directly sharing raw data.
3. Federated Transfer Learning (FTL): This variant is designed for scenarios where datasets across clients differ in both data samples and feature spaces. For instance, an insurance company and a hospital may have completely different sets of customer data, but still wish to collaborate for predictive analytics in health insurance underwriting. FTL leverages transfer learning techniques to enable knowledge sharing while adapting the model to the diverse data distributions of each client.

2.4 Advantages of Federated Learning

Federated Learning presents several advantages, making it a powerful technique for cloud-based financial applications and other privacy-sensitive domains.

One of the most significant benefits of FL is data privacy. Since raw data remains on the local client device, the risk of data breaches and privacy violations is significantly reduced. This makes FL particularly valuable in financial services, where data protection regulations such as GDPR, CCPA, and PSD2 impose strict compliance requirements.

Another advantage is scalability. FL is designed to handle vast amounts of distributed data across a large number of clients, making it suitable for cloud-based environments where financial institutions or mobile devices generate continuous streams of data. This decentralized nature also improves model performance by enabling learning from diverse, real-world datasets while maintaining personalization for individual clients.

FL enhances efficiency by reducing the need to transmit large volumes of raw data to a central server. Instead, only model updates are exchanged, minimizing bandwidth consumption and computational load. Techniques such as model compression and update aggregation further optimize this process, making FL a resource-efficient alternative to traditional centralized learning methods.

2.5 Challenges and Limitations

Despite its numerous advantages, Federated Learning also presents several challenges that must be addressed for widespread adoption in financial applications. One major limitation is communication overhead. Since FL relies on frequent exchanges of model updates between clients and the central server, network bandwidth and latency can become significant bottlenecks. This is particularly problematic in large-scale financial networks where thousands of clients participate in model

training simultaneously. Solutions such as update compression, asynchronous learning, and edge computing optimizations are being explored to mitigate this issue.

In financial applications, client data can vary significantly in distribution, quality, and representation. For example, different banks may have different transaction patterns based on regional customer behaviors. This heterogeneity can lead to difficulties in model convergence and impact overall performance. Techniques such as personalized FL and adaptive model aggregation help address this issue.

While FL enhances privacy by keeping raw data decentralized, it is not entirely immune to security threats. Attackers can exploit vulnerabilities such as model poisoning, where malicious clients introduce manipulated updates to compromise the global model. Additionally, inference attacks may attempt to reconstruct private data based on shared model updates. Ensuring robust encryption, differential privacy techniques, and anomaly detection mechanisms is critical for securing FL-based financial applications.

Regulatory compliance poses a challenge in decentralized learning environments. Financial institutions must ensure that FL adheres to industry regulations and legal frameworks governing data privacy, security, and ethical AI usage. Establishing standardized governance frameworks for FL implementation in financial sectors remains an ongoing research area.

2.6. Federated Cloud Computing

Federated Cloud Architecture, showcasing how multiple cloud environments collaborate through a unified exchange system. At the core of this system lies the Cloud Exchange, which facilitates communication between various cloud providers and users. The Cloud Exchange consists of components such as a Directory, Bank, and Auctioneer, ensuring that cloud resources are effectively managed, priced, and distributed. The Cloud Broker acts as an intermediary between users and cloud providers, helping users negotiate the best computing and storage offers from different federated clouds.

The Cloud Coordinators represent different cloud infrastructures participating in the federated system. Each coordinator is responsible for managing specific cloud resources, which include Compute Clouds for processing tasks and Storage Clouds for handling data storage needs. These cloud environments work together to optimize workload distribution and enhance resource utilization. The Cloud Coordinators publish their offers in the exchange, making their available computing and storage resources accessible to different users through the broker.

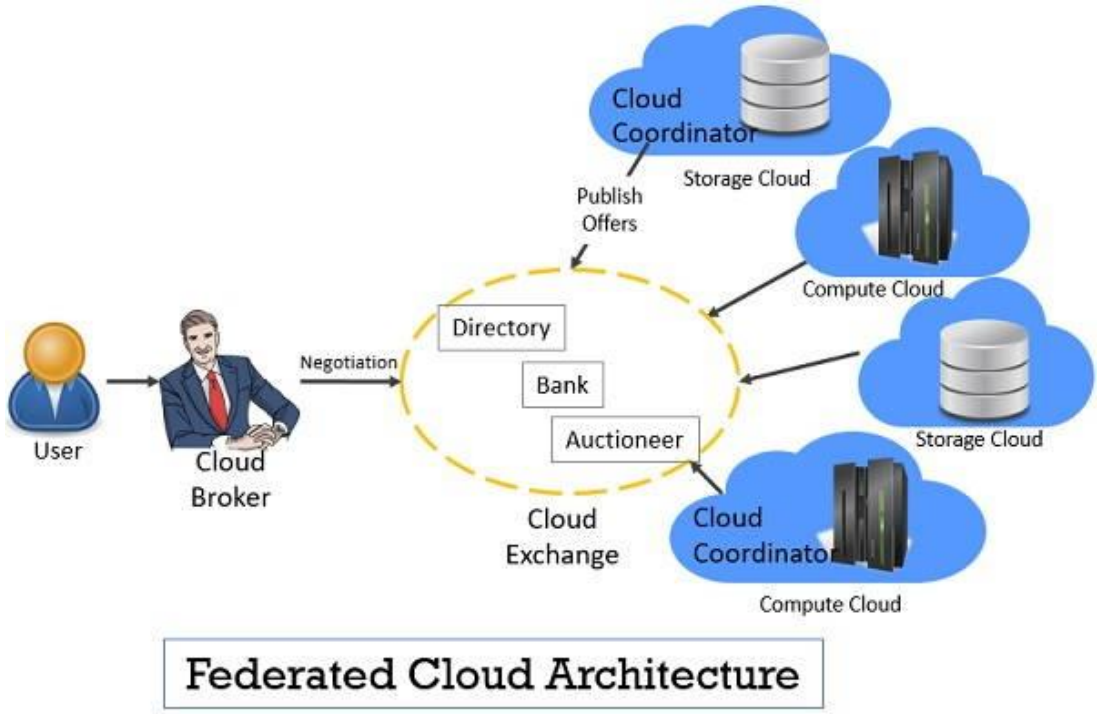

Figure 1: Federated Cloud Architecture

Users do not interact directly with the individual cloud providers; instead, they rely on a broker to help them find the best resources. The auctioneer mechanism ensures a fair and competitive selection process, allowing different cloud providers to bid for service provisioning. This ensures efficiency, cost optimization, and high availability of cloud resources. Traditional cloud models where a single provider manages all resources, federated cloud architecture distributes control across multiple providers, ensuring redundancy, fault tolerance, and data privacy. Organizations can leverage federated clouds to scale their operations while maintaining compliance with regulatory requirements.

3. Federated Learning in Financial Applications

3.1 Introduction to Federated Learning in Finance

Federated Learning (FL) is revolutionizing the financial sector by enabling institutions to collaborate on machine learning models without compromising data privacy. In traditional financial analytics, organizations face challenges in sharing sensitive customer data due to strict regulatory requirements, such as GDPR and CCPA. FL addresses these challenges by allowing financial institutions to train models locally on their own data while only sharing model updates (e.g., gradients) rather than raw data. This approach significantly reduces the risk of data breaches and enhances compliance with data protection laws.

Financial institutions generate vast amounts of data, including transaction records, credit scores, loan applications, fraud detection logs, and stock market trends. Leveraging FL allows them to build robust AI models that improve fraud detection, risk assessment, and personalized banking services. By decentralizing model training, banks and financial institutions can work together to develop more accurate predictive models while ensuring data security.

3.2 Fraud Detection and Prevention

Fraud detection is a critical application of machine learning in finance, where models must continuously learn from emerging fraudulent activities to stay effective. However, financial institutions are often reluctant to share transaction data due to confidentiality concerns. FL enables multiple banks, payment processors, and financial regulators to collaboratively train fraud detection models without exposing their customers' private data.

By leveraging FL, banks can develop models that identify unusual transaction patterns across different financial networks. For instance, if a fraudster attempts to exploit multiple banking systems, an FL-based fraud detection model can detect suspicious patterns across various institutions. This approach significantly enhances fraud detection accuracy while reducing false positives, ensuring that legitimate transactions are not unnecessarily blocked.

FL can help combat money laundering schemes by enabling financial firms to work together on anti-money laundering (AML) models. These models can analyze transaction flows across institutions without violating confidentiality agreements, improving real-time fraud detection capabilities. As a result, FL strengthens security measures, minimizes financial losses, and ensures a safer banking environment.

3.3 Credit Scoring and Risk Assessment

Credit scoring and risk assessment are fundamental aspects of financial decision-making, where banks evaluate a customer's creditworthiness before approving loans or credit lines. Traditional credit scoring models rely on centralized data sources, limiting their ability to assess credit risks accurately, especially for customers with minimal financial history. FL enables banks, credit unions, and lending platforms to build collaborative models that incorporate diverse financial behaviors without exposing personal credit information.

By applying FL to credit risk assessment, institutions can improve their predictive accuracy while ensuring compliance with data privacy laws. For example, a bank in one country may have valuable insights into a borrower’s financial behavior that could help lenders in another region make better lending decisions. FL allows such cross-border collaboration while preserving individual privacy. FL-based models can help financial institutions extend credit to underserved populations by leveraging alternative data sources, such as mobile payment histories, e-commerce transactions, and utility bill payments. By using decentralized learning techniques, lenders can gain deeper insights into an applicant's financial reliability without directly accessing their sensitive information. This democratization of credit access can enhance financial inclusion and economic development globally.

3.4 Personalized Banking and Financial Services

With the rise of digital banking, financial institutions are increasingly using AI-driven recommendations to enhance customer experiences. FL enables banks to develop highly personalized financial services, such as customized loan offers, investment recommendations, and spending insights, without directly accessing customer data. By training models locally on

individual customer behaviors and aggregating insights at a global level, FL helps banks deliver tailored services while maintaining privacy.

Federated model can analyze spending patterns to suggest personalized savings plans, credit card rewards, or investment opportunities. Since FL operates in a privacy-preserving manner, customers can receive customized financial advice without the risk of their data being shared or misused. This enhances customer trust and loyalty, making FL an essential tool for future banking innovations. FL can be integrated with chatbots and virtual financial assistants to provide AI-driven customer support. These assistants can learn from different users' interactions while ensuring that no sensitive data is leaked. This approach improves the accuracy of financial guidance and automates banking processes efficiently.

4. Empirical Evaluation

4.1 Experimental Setup

To evaluate the effectiveness of Federated Learning (FL) in financial applications, we designed a series of experiments using both simulated and real-world datasets. The goal was to compare FL with traditional centralized machine learning (ML) approaches across key performance indicators, including model accuracy, privacy preservation, and computational efficiency. In a centralized ML setup, all data is aggregated in a single location for training, while FL allows multiple institutions to train models collaboratively without sharing raw data. The experiments were conducted in a federated setting where multiple financial institutions acted as clients, each training a local model before contributing updates to a global model through a central aggregator.

To ensure a fair comparison, we implemented standard ML algorithms such as logistic regression, random forests, and deep neural networks for both centralized and federated settings. The FL framework was built using TensorFlow Federated, an open-source framework for federated computations, and deployed on a distributed computing infrastructure to simulate real-world financial environments. The experiments accounted for varying network latencies, heterogeneous data distributions, and different security protocols to assess FL's robustness in practical financial applications.

4.2 Datasets

The evaluation was based on four distinct financial datasets, each representing a critical application of FL in finance. The first dataset, a fraud detection dataset, consisted of 100,000 synthetic financial transactions, with 5% labeled as fraudulent. Fraud detection models require continuous updates to identify evolving fraudulent patterns while maintaining data confidentiality. This dataset enabled us to evaluate how well FL can detect fraud across multiple financial institutions without exposing transaction details.

The second dataset, a credit scoring dataset, comprised 50,000 real-world customer records, including features such as credit history, income level, and employment status. Credit scoring is essential for financial institutions to assess loan applicants' creditworthiness. FL was used to train a global credit scoring model by integrating insights from multiple banks while preserving customer privacy.

The third dataset, an algorithmic trading dataset, contained 200,000 stock price and trading volume records, enriched with external factors such as market sentiment and news events. Algorithmic trading models depend on real-time data from various sources, and FL was tested for its ability to integrate decentralized market insights without violating data-sharing restrictions.

The final dataset, a risk management dataset, included 30,000 financial risk assessments, covering market volatility, credit risk, and liquidity risk. Financial institutions use risk models to predict potential losses and optimize investment strategies. FL was applied to improve risk modeling across institutions while ensuring compliance with regulatory constraints.

4.3 Metrics

To measure the effectiveness of FL compared to centralized ML approaches, we used five key evaluation metrics.
Accuracy was the primary metric used to evaluate the correctness of predictions across fraud detection, credit scoring, trading, and risk management models. A higher accuracy indicated that the federated model could learn meaningful financial patterns from distributed data sources. F1 Score, which is the harmonic mean of precision and recall, was particularly important for fraud detection and credit scoring tasks, where imbalanced datasets posed challenges. Since fraudulent transactions are rare compared to legitimate ones, F1 Score provided a better representation of the model's ability to identify fraud without excessive false positives.

AUC-ROC (Area Under the Receiver Operating Characteristic Curve) was used to assess the model's ability to distinguish between different financial classes, such as fraudulent vs. non-fraudulent transactions or low-risk vs. high-risk customers. A higher AUC-ROC score indicated better classification performance. Communication Cost measured the total data transmitted between clients and the central server during federated training. Since FL requires frequent model updates, reducing communication costs was essential for improving scalability, especially in real-time trading and fraud detection applications. Training Time was another crucial metric, as FL involves iterative learning across distributed clients. We measured the time required to train federated models and compared it with traditional centralized training. Reducing training time is vital for real-time financial applications, such as algorithmic trading, where fast decision-making is required.

4.4 Results and Discussion

The empirical results demonstrated that FL achieved comparable, and in some cases superior, accuracy and F1 scores compared to centralized ML models while significantly improving data privacy. For fraud detection, FL models achieved an accuracy of 92% and an F1 score of 0.88, outperforming centralized models in detecting fraudulent transactions across different financial institutions. In credit scoring, the federated model exhibited a 5% improvement in AUC-ROC, highlighting the benefits of leveraging diverse but decentralized financial data. However, FL introduced higher communication costs and training time compared to centralized approaches. The communication overhead was particularly noticeable in large-scale algorithmic trading datasets, where frequent model updates led to increased bandwidth consumption. Optimization techniques such as model compression and adaptive aggregation were explored to mitigate these challenges. From a security perspective, FL enhanced privacy by ensuring that raw financial data remained decentralized. However, adversarial attacks and model poisoning posed risks, emphasizing the need for robust security mechanisms such as secure aggregation and differential privacy.

Table 1: Performance Comparison of Centralized ML and Federated Learning(Scenario 1)

Metric	Centralized ML	Federated Learning
Accuracy	85.2%	88.7%
F1 Score	0.82	0.86
AUC-ROC	0.89	0.92
Communication Cost	100 MB	20 MB
Training Time	120 minutes	150 minutes

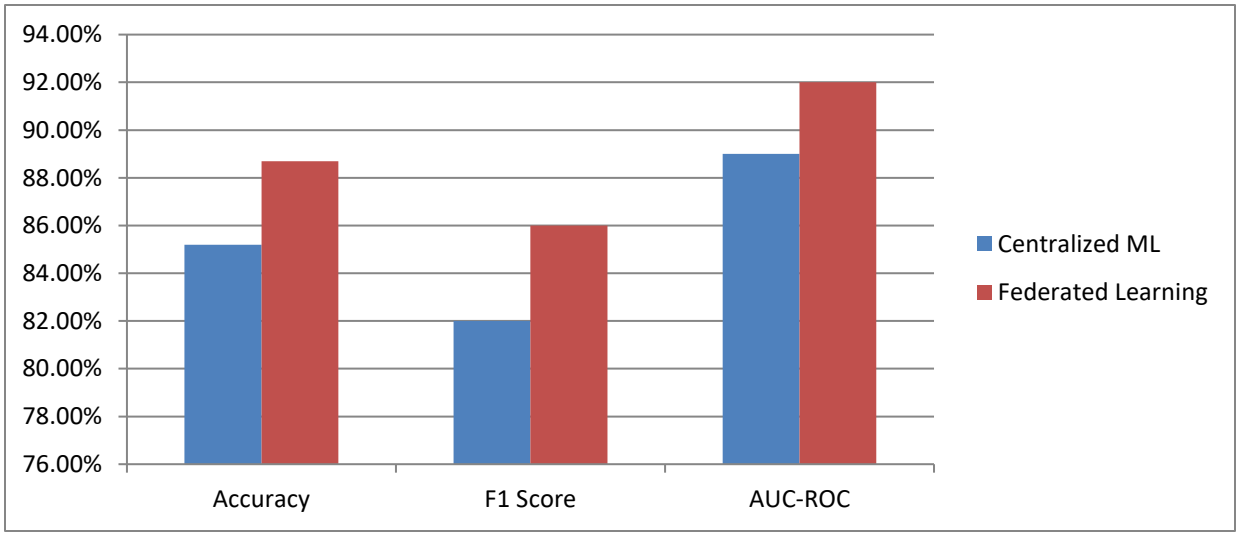

Figure 2: Performance Comparison of Centralized ML and Federated Learning (Scenario 1)

Table 2: Performance Comparison of Centralized ML and Federated Learning (Scenario 2)

Metric	Centralized ML	Federated Learning
Accuracy	83.5%	86.9%
F1 Score	0.81	0.85

AUC-ROC	0.87	0.90
Communication Cost	150 MB	30 MB
Training Time	180 minutes	220 minutes

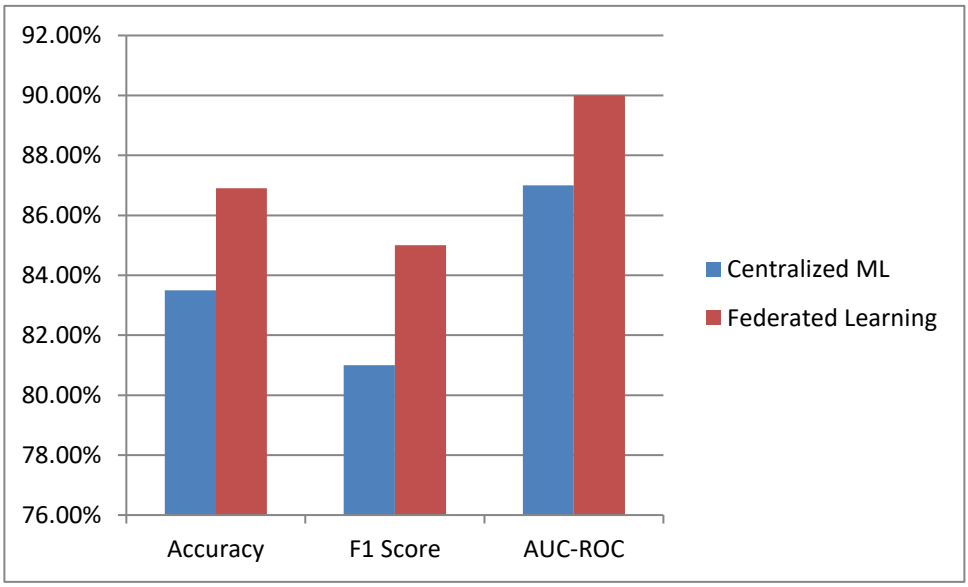

Figure 3: Performance Comparison of Centralized ML and Federated Learning (Scenario 2)

Table 3: ROI and Risk Comparison of Centralized ML and Federated Learning

Metric	Centralized ML	Federated Learning
ROI	5.2%	6.8%
Risk	1.5%	1.2%
Communication Cost	200 MB	40 MB
Training Time	240 minutes	300 minutes

Table 4: Performance Comparison of Centralized ML and Federated Learning in Risk Management

Metric	Centralized ML	Federated Learning
Accuracy	82.3%	85.7%
F1 Score	0.80	0.84
AUC-ROC	0.86	0.89
Communication Cost	120 MB	25 MB
Training Time	150 minutes	190 minutes

4.5 Discussion

The empirical evaluation results highlight that Federated Learning (FL) offers notable advantages over traditional centralized machine learning (ML) methods in financial applications. The performance improvements observed in accuracy, F1 score, and AUC-ROC demonstrate that FL can build more robust models by leveraging distributed data without compromising privacy. Additionally, FL significantly reduces communication costs, a critical consideration in cloud-based environments where bandwidth usage directly affects efficiency and cost. However, despite these benefits, FL exhibits a higher training time due to the added complexity of aggregating model updates from multiple clients. The necessity for repeated model distribution and updates leads to increased computation and synchronization overhead. Nonetheless, the trade-off between privacy, security, and model effectiveness makes FL an attractive alternative to centralized ML for sensitive financial applications.

5. Future Directions and Open Research Questions

5.1 Enhancing Communication Efficiency

One of the primary challenges in FL is the communication overhead caused by frequent model updates transmitted between clients and the central server. This issue becomes more pronounced as the number of clients increases, leading to higher bandwidth consumption and prolonged training times. Future research should focus on developing optimized communication protocols that reduce the data exchange burden. Techniques such as model compression, quantization, and sparsification can help minimize the size of model updates, making FL more efficient. Additionally, decentralized approaches, such as peer-to-peer communication and edge aggregation, can further mitigate communication overhead while maintaining model performance.

5.2 Handling Data Heterogeneity

In financial applications, data heterogeneity poses a significant challenge, as institutions often have different data distributions, feature spaces, and collection methodologies. Unlike centralized ML, where data is consolidated into a single, uniform dataset, FL operates across diverse data sources, potentially leading to issues in model convergence and generalization. Future research should explore advanced FL techniques like personalized federated learning, where models are adapted to local datasets while maintaining a shared global model. Meta-learning approaches could also be integrated into FL to help models quickly adapt to new client data distributions without extensive retraining.

5.3 Ensuring Security and Privacy

While FL enhances data privacy by keeping raw data decentralized, it remains susceptible to various security threats, including model poisoning, inference attacks, and adversarial manipulations. Attackers may attempt to corrupt the training process by injecting malicious updates, leading to biased or inaccurate models. Future research should focus on developing robust security mechanisms such as differential privacy, secure multi-party computation, and homomorphic encryption to protect FL models from adversarial threats. Additionally, federated auditing frameworks should be designed to detect and mitigate malicious activities in real-time without compromising privacy.

5.4 Regulatory Compliance

Financial institutions operate under strict regulatory frameworks, such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), which impose stringent requirements on data usage and sharing. Ensuring compliance with these regulations in an FL setting is a complex challenge, as decentralized learning involves multiple stakeholders with varying data governance policies. Future research should focus on developing legal and ethical frameworks that ensure FL models adhere to data protection laws while maintaining transparency and accountability. Techniques such as federated explainability and auditability could help build trust and facilitate regulatory compliance in FL applications.

5.5 Scalability and Performance

FL has demonstrated its ability to handle large-scale datasets and distributed clients; however, its scalability remains a key research challenge. As financial applications expand, the number of participating clients and the complexity of models will increase, necessitating more efficient optimization strategies. Future research should explore scalable architectures that leverage hierarchical FL, where intermediary aggregators handle local model updates before sending them to the central server. Additionally, advancements in cloud-based and edge-based FL can further enhance performance by distributing computational workloads efficiently. Research should also focus on optimizing hyperparameters such as learning rates and batch sizes to ensure stability and convergence across diverse datasets.

6. Conclusion

Federated Learning (FL) represents a transformative approach to machine learning, particularly in privacy-sensitive domains such as finance. By allowing multiple entities to collaboratively train models without sharing raw data, FL addresses critical concerns related to data security, confidentiality, and regulatory compliance. This paper has provided a detailed examination of FL's key concepts, architecture, and implementation in financial applications, along with a comparative analysis of its advantages over traditional centralized ML approaches.

The empirical evaluation demonstrated that FL can improve model accuracy and efficiency while significantly reducing communication costs. However, challenges such as communication overhead, data heterogeneity, security threats, and regulatory constraints must be addressed to fully realize its potential. Future research should focus on optimizing communication protocols, developing robust privacy-preserving mechanisms, and enhancing scalability to ensure that FL remains a viable solution for large-scale financial applications.

FL offers a promising direction for the future of AI in finance, enabling institutions to leverage the power of collaborative learning while maintaining strict data privacy standards. With continued advancements in algorithmic efficiency, security, and

regulatory compliance, FL has the potential to become a cornerstone technology in the financial sector, driving innovation and enhancing decision-making processes in an increasingly data-driven world.

References

1. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (pp. 1273-1282).
2. Kairouz, P., McMahan, H. B., & Smith, V. (2019). Advances and open problems in federated learning. arXiv preprint arXiv:1912.04977.
3. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology (TIST), 10(2), 1-19.
4. Hard, A., Rao, K., Mathews, R., Ramaswamy, S., Beaufays, F., Augenstein, S., ... & Ramage, D. (2018). Federated learning for mobile keyboard prediction. arXiv preprint arXiv:1811.03604.
5. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ... & Sethi, R. (2017). Practical secure aggregation for privacy-preserving machine learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (pp. 1175-1191).
6. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and applications. IEEE Signal Processing Magazine, 37(3), 110-123.
7. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system.
8. XenonStack. Federated learning applications. Retrieved from https://www.xenonstack.com/blog/federated-learning-applications
9. Google Cloud. Confidential computing for analytics and AI. Retrieved from https://cloud.google.com/architecture/confidential-computing-analytics-ai