International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805016

Privacy Preserving AI in Financial Sector-Balancing Utility, Security and Compliance

Neha Tyagi

Senior Vice President, Bank of New York (BNY), USA

ABSTRACT: The fast implementation of the use of Artificial Intelligence (AI) in the financial industry has enhanced the efficiency of operations, risk evaluation, detection of fraud, and customized services to customers to a considerable extent. Nevertheless, the wide application of sensitive financial and personal data causes serious issues connected with data privacy, data security, and regulatory compliance. The paper explores Privacy Preserving Artificial Intelligence (PPAI) approaches that would allow financial institutions to utilize AI-based insights without violating confidential information, as well as legal regulations, including the GDPR, RBI data localization policy, and new regulations in the direction of artificial intelligence globally. The suggested methodology will consist of the comparative analysis of the solutions presenting privacy protection: federated learning, differential privacy, secure multi-party computation, and homomorphic encryption, and apply them to core financial applications, such as credit scoring, transaction monitoring, and fraud detection. An integrated framework between federated learning and differential privacy is developed to combine the utility of the model and the privacy guarantees. Benchmark financial datasets are used to evaluate the performance of an experimental evaluation based on the accuracy, the risk of data leakage, computational overhead, and regulatory alignment. Findings show that the hybrid PPAI structure attains close central performance of the model where the predictive accuracy is reduced by less than 3 percent and yet the data exposure is reduced to a minimum and the privacy laws are complied with. The results indicate that a combination of privacy-preserving methods with a well-planned approach will be a good balance of utility, security, and compliance that does not affect business goals. This paper finds that Privacy Preserving AI is not a regulatory requirement but a strategic facilitator of trust, resilience, and competitive edge in the new financial ecosystem.

KEYWORDS: Privacy Preserving AI, Financial Technology, Federated Learning, Differential Privacy, Regulatory Compliance, Data Security

I. INTRODUCTION

The financial industry is not a newcomer to the use of advanced computer technologies as a way of training efficiency and accuracy and increasing competitiveness. Over the last several years, Artificial Intelligence (AI) and Machine Learning (ML) have become the defining forces, redefining the essential financial processes, including credit risk evaluation, fraud detection, algorithmic trading, customer relationship management, anti-money laundering (AML), and regulatory reporting [1]. AI-based systems can allow financial institutions to operate on large amounts of structured and unstructured data, identify patterns in the data that had been missed, and provide real-time and data-driven decisions, which were once not possible to arrive at with the traditional rule-based systems. Consequently, AI can now be considered a key driver of digital transformation in the banking sector, insurance sector, capital markets, and fintech ecosystems [2].

In spite of the above benefits, the increased use of AI has accelerated the debate on the issues of data privacy, security, and regulatory adherence. Financial information is one of the most sensitive groups of personal data that include transaction history, income, credit history, investment portfolio and identity attributes. Any unauthorized access, misuse or leakage of such data can result in the loss of huge sums of money, identity theft, loss of customer confidence, and legal fines. The high-profile data breaches and the rising cases of cyberattacks have also highlighted the vulnerability of centralized systems of data-driven AI and have raised the pressing need to implement security-conscious solutions [3].

Governmental agencies everywhere throughout the globe have resorted to such threats by developing hard data protection and management systems. Such policies as the regulations on personal data protection (GDPR) in the European Union, the data protection regulations in the United States (California Consumer Privacy Act) make it essential to control data collection, processing, storage, and sharing tightly. The principles of minimizing data and restricting its purpose, transparency, and the consent of users are prioritized in the rules that do not always coincide with the data-driven character of the existing models of AI. This leaves the financial institutions in a tough position of trying to maximize the potential of AI and also fulfill the regulatory needs and patient privacy at the same time [4].

The conventional methods used to develop AI mostly involve the principle of centralized data aggregation, so that all the data of various sources are combined into one repository where they are trained and inferred. Although this paradigm is effective in terms of performance, it has extremely high privacy and security risks. The centralized storage enlarges the attack surface of cyber threats and complicates in fulfilling the requirements of data sovereignty. Furthermore, by default centralized models can store sensitive data in memory, causing it to leak out via model inversion or membership inference attacks. Such restrictions have brought increased attention to the alternative AI paradigms, which are built to incorporate privacy protection into the model lifecycle [5].

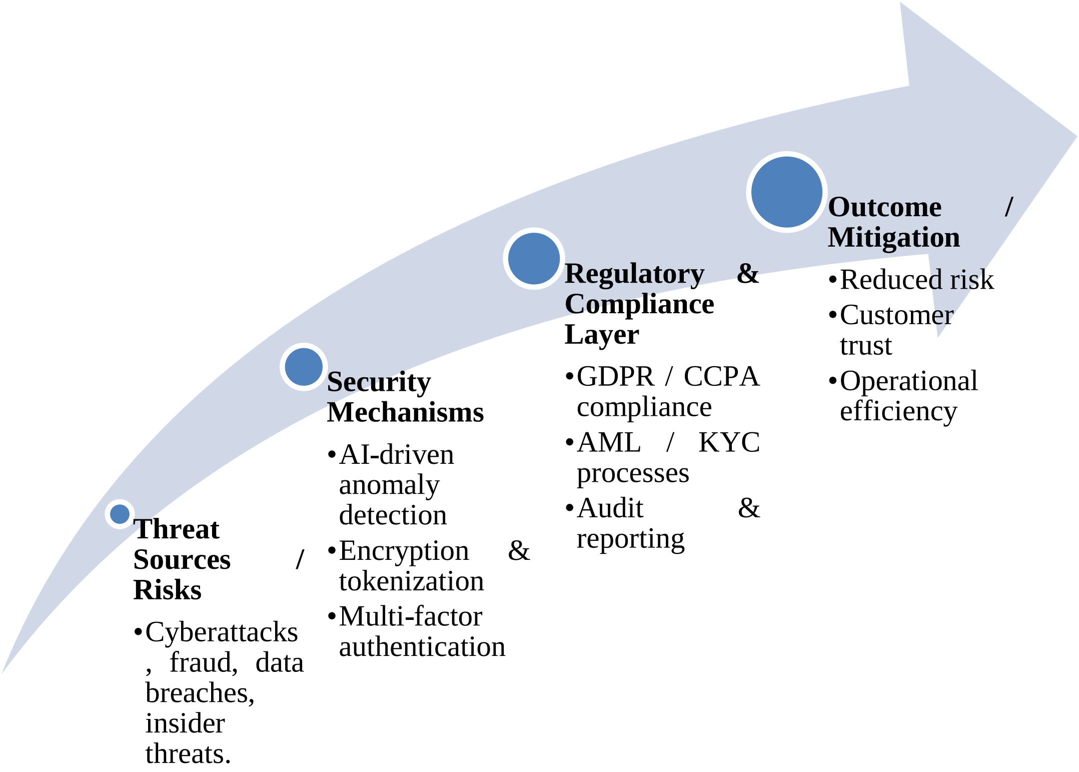

Figure 1: Cybersecurity & Regulatory Compliance in Fintech

Privacy Preserving Artificial Intelligence (PPAI) is one of the possible solutions to such issues. PPAI describes a set of methods and systems that are set up to facilitate the training and deployment of AI models without the disclosure of raw sensitive data. Subsequently, instead of viewing privacy as a secondary consideration, PPAI incorporates privacy guarantees into data processing, learning algorithms and system architectures. In this sphere, federated learning, differential privacy, secure multi-party computation (SMPC) and homomorphic encryption are some of the key methods with multiple trade-offs among privacy, accuracy, computational complexity, and scalability.

The applicability of PPAI is especially high in the financial sector as the area is pressured by the competition to innovate and the regulatory pressure. To illustrate, fraud detection systems can utilize learning patterns between several institutions, but the direct data sharing is frequently not allowed by the confidentiality and competition regulations. In the same vein, credit scoring models need various datasets to minimize bias and enhance generalization, however, the availability of customer level financial data across borders creates difficulties in compliance. PPAI methods facilitate collaborative intelligence through the provision of the ability to train models and draw inferences on collaborating institutions without disclosing any underlying sensitive information to support innovation, maintain trust, and ensure compliance.

Responsible and ethical AI is also another significant factor that leads to the application of privacy-saving AI in the business world. Besides the legal compliance, financial institutions are also experiencing increased demand in their accountability, fairness and transparency in automated decision making. The concept of responsible AI must rest on the pillar of privacy since intrusive data habits may enhance prejudice, discrimination, and social injustices. This decreases the unnecessary data disclosure and responsible data utilization presupposed by the PPAI precondition the ethicality of AI regulation and heightens the confidence of individuals in automated financial systems.

However, the implementation of privacy-saving AI is not that straightforward. The vast majority of privacy-enhancing strategies have a computational cost (or communication cost) or reduced model accuracy. As a case in point, predictive performance may be impaired as a result of the introduction of noise in cases where disorderly privacy is utilized, and the homomorphic encryption may lead to a significant reduction in performance. Although federated learning can be applied in the decentralization of data, such issues as the heterogeneity of the system, the efficiency of the communication process, and adversarial updates are also possible. It implies that the privacy protection concerns must be struck with the balance between utility, security, and compliance needs of financial institutions without compromising the performance of the AI systems.

This is critical more so in high stakes financial systems where inaccuracies as little as small scale can bring about significant financial losses or ripple effects of risk. Therefore the necessity to introduce hybrid and adaptive privacy saving models that will effectively trade off various methods to achieve optimal trade-offs is increasing. These frameworks are aimed at sustaining a close-to-centralized level of performance of models and at providing high levels of privacy assurance and regulation compatibility.

In that regard, the current work is aimed at discussing the Privacy Preserving AI in the financial industry, as well as the importance of finding a balance between the utility of the model and the security of the data along with the level of compliance with regulations. The paper discusses the drawbacks of the traditional AI methods, evaluates the state-of-the-art privacy-preserving methods, and discusses their relevance to the real-life scenario in the financial context. Through a systematic evaluation of performance, privacy risk, and compliance factors, this piece of work aims to offer an organized insight on how financial institutions may embrace PPAI as a strategic enabling factor, and not a regulation liability.

Finally, there is a paradigm shift in financial innovation in the integration of privacy-preserving AI. Since intelligence that is data-driven becomes the key element to competitive advantage, organizations that actively identify privacy within their AI systems will be in a better place to navigate regulatory murk, risk reduction, and development of long-term customer and stakeholder trust. This overview establishes the base of the further discussion of the methodologies, experimental analysis, and practical consequences of Privacy Preserving AI in the contemporary financial ecosystem.

II. RELATED WORK- CYBERSECURITY, PRIVACY, AND AI IN THE FINTECH SECTOR

We witness a significant change in the financial technology (fintech) sector in the past ten years owing to the effects of digitalization, artificial intelligence (AI), cloud-based computing and blockchain technologies. Attached to these innovations, the issue of cybersecurity and data privacy have also gained a new perspective as the financial information can be quite sensitive, and cyber threats are becoming more complex. The reviewed literature also gives a detailed overview of the existing strategies, structures, and technology that are being employed to secure the financial system and balance the regulatory compliance with efficiency in their operations.

Olaiya et al. [1] provide an elaborate discussion of the cybersecurity policy in the industry of fintech and state that there is a need to adopt multi-layered defense measures in order to protect financial assets and client information. They note that the number of cyber-attacks (phishing, ransomware, identity theft) on fintech systems is slowly increasing, and technicologic protection against such attacks, such as intrusion-detection systems, encryption protocols, and AI-assisted anomaly-detection. According to the authors, the advanced cybersecurity frameworks should be included in the operational processes to guarantee the trust and minimize the financial losses. Matters of technological innovation and security governance are relevant to their production, making it a key backbone in the threat scenario and its complexity to the fintech organizations.

In addition to this, Oyewole et al. [2] are preoccupied with the regulatory perspective, the impact of the global data privacy regulations such as GDPR, CCPA, and the emerging regional regulations on the operations in the fintech. They point out that compliance is not merely a legal obligation but it is also a competitive edge because customers are currently demanding to receive transparency and control over their own personal data. The paper compares case studies of how fintech business altered data handling operations, data storage and data sharing operations to suit privacy regulations indicating the clashed considerations between the necessity to be innovative and the necessity to be regulated. The authors reach the conclusion that privacy-first strategy is active, enhances the level of customer confidence and reduces legal and financial risks.

Bhalla et al. [3] further this explanation by referring to the issue of consumer privacy in the framework of cybersecurity challenges. According to them, by the nature of fintech innovations (such as mobile banking, digital wallets, or peer-to-peer payment systems), the attack surface of possible breaches is enlarged. The paper highlights the importance of design based on security concerns, that is, encryption with access controls and real-time surveillance and both in conjunction with consumer education programs. The authors place stress on the holistic approach to cybersecurity by connecting technology and behavior of users to provide a combination of technical protection and policy and awareness campaigns.

Abikoye et al. [4] and Chennuri and Biyyala [5] study regulatory compliance and efficiency in the fintech setting. Abikoye et al. explore how fintech organizations can meet the requirements of financial regulations, the anti-money laundering (AML) regulations, and cybersecurity requirements without compromising efficiency. According to their results, automation, regulatory technology (RegTech), and AI-based monitoring tools can facilitate reporting and risk assessment and minimize the number of individuals working on them and the error rates. Similar studies are made by Chennuri and Biyyala, who examine the transformations taking place in banking and finance, such as digital onboarding, cloud banking, and AI-based credit scoring. Both articles point out that adoption of technology should be balanced with sound compliance guidelines to avoid operational and reputational risks.

Regarding technology, Mohammadi et al. [6] offer a review of the technology of federated learning (FL) used in AI to ensure privacy preservation. The paper discusses the means that can be used to train models in a decentralized manner without distributing raw data, which keeps the privacy intact but provides the possibility of collaborative AI solutions. The authors mention measures of assessing trade-offs between performance and privacy assurances in models, and it is possible to remark that FL can help mitigate the threat of the exposure of data in multi-institutional fintech partnership to a substantial degree. This is further elaborated upon by Hassan and Mohamed [7], who focus on FL usage in AI, IoT, healthcare, and finance, and illustrate the flexibility of the decentralized learning towards privacy-preserving analytics.

Chatterjee et al. [8] discuss how the concept of FL and blockchain technology could be developed to the advantage of the security of financial information. They indicate that they propose a hybrid approach where blockchain may ensure the integrity of both model updates and history of transactions, and FL may protect sensitive sets of data. It is a two-layered solution that addresses cybersecurity and privacy as a unit, which provides a scalable solution to the multi-party financial ecosystem. In federated learning, privacy and security issues are also consulted in Gosselin et al., [9], and they control the adversarial threats, the approaches of differential privacy, and the secure aggregation approaches. Their findings suggest that there is a need to develop FL systems that are resistant to malicious attacks and simultaneously remain useful.

Both Ahmad et al. [10] and Serneviciene and Kabasinskas [11] discuss the role of the explainable AI (XAI) in the detection and decision-making process in the sphere of financial fraud. According to Ahmad et al., models can be proposed that would integrate FL and XAI to leave privacy untouched and make the fraud detection processes transparent. The authors have concluded that responsible models are effective in ensuring regulatory compliance and trust that can allow the audit and regulatory community to understand and justify AI-driven decisions. Similarly, in a similar work, Serneviciene and Kabasinskas get a compendium of XAI literature in the field of finance, which may allow mitigating a risk exposure by employing a black-box AI model, e.g., biased credit scoring, fraud, and opaque risk assessment.

Finally, the article by Dalal and Roy [12] is focused on the total tradeoff between the security of cyberspace and privacy and the rights of individuals in the online financial services. They are more concerned with the ethical side of their activities, and in this context, they state that fintech companies must not disregard technological security and social effects. They show the need of the security models based on risks, privacy-by-design, and periodical compliance audits as the way to the protection of consumer data and its innovations. It aligns with the similarity present within the literature that goes through operational efficiency, technological advancement, and regulatory needs need to be hardened.

III. METHODOLOGY

The research will use a multi-layered, systematic approach in designing, implementing, and testing a Privacy Preserving Artificial Intelligence (PPAI) framework to be applied in the financial sector. The approach is aimed at attaining the best balance between the utility of the model, data security and compliance to the regulations, and it addresses the practical implementation limitations that are usually confronting financial institutions. The general procedure of the methodology will include data preparation, privacy-preserving technique selection, framework design, implementing the experiment, evaluation metrics, and compliance analysis.

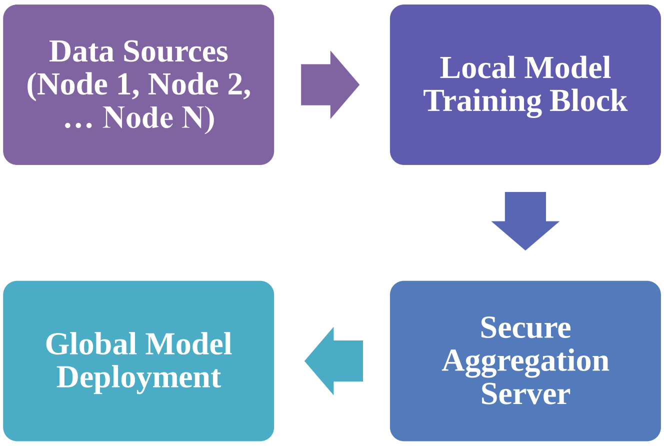
Data Sources (Node 1, Node 2, … Node N)
Local Model Training Block
Secure Aggregation Server
Global Model Deployment

Figure 2: Privacy-Preserving AI in Fintech – Federated Learning Architecture

1. Problem Definition and Use-Case Selection

The approach starts with the identification of representative financial application scenarios in which AI-based decision-making is effective and privacy-related. They are three fundamental applications that are selected, namely: fraud detection, credit risk assessment and transaction anomaly detection. These applications are selected because they are based on sensitive personal and transactional information, call-to-action needs, and high regulatory standards. The use cases pose unique privacy and performance issues and are, therefore, applicable in assessing the strengths and generalizability of privacy-sensitive AI practices.

2. Dataset Selection and Preprocessing

To guarantee the reproducibility and realistic assessment, publicly available financial benchmark datasets are used, but they are complemented by synthetically generated data to recreate inter-institutional data distribution. The data sets cover transaction based characteristics, behavioural characteristics, and risk factors that are usually employed in financial modelling. Personal data like identifiers of the customers are eliminated or anonymized in preprocessing.

Preprocessing of data includes feature normalization, missing value treatment, categorical encoding and class imbalance adjustment, and this is especially important in fraud detection processes. The data sets are decentralized, resembling the multiple financial organizations where the data is only stored locally. Such an arrangement is representative of real-life restrictions that are placed by localization of data and privacy policies.

3 Selection of Privacy-Preserving Techniques

The methodology systematically combines and analyzes four well-known privacy-sensitive AI methods to handle financial environment issues of data security and data compliance. Federated Learning (FL) provides a decentralized model training with multiple institutions that share a global model without sharing raw data and hence maintaining confidentiality. They are compared in terms of scalability, computational cost, privacy, and its compliance with regulatory requirements. In practice, experiments mainly involve FL and DP because of their efficiency and applicability to real-world fintech deployments, though SMPC and HE are viewed as complementary instruments to high-risk or highly sensitive operations that need to be assured of maximum privacy.

4 Hybrid Privacy-Preserving Framework Design

A hybrid PPAI framework will be created so as to achieve a balance between privacy and utility. The model combines federated learning and differential privacy, and takes advantage of both methods. The financial entities involved in this architecture approach train local models using local datasets, and only exchange updates of the model with an aggregation server in the middle. The local model updates are at risk of differentiating privacy prior to transmission, such that contributions to individual data cannot be concluded.

The framework is based on a client-server design whereby the server does secure aggregation of noisy updates to create a global model. Privacy budgets are established with a lot of precision to establish an acceptable balance between accuracy and privacy. Secure aggregation using SMPC is optional, and it is used to further protect model updates in communication.

5 Model Architecture and Training Strategy

Financial applications of standard machine learning and deep learning models such as logistic regression, gradient boosting machines and neural networks are used. These models have been chosen due to their interpretability, performance and popular usage in the financial industry.

Training occurs in the form of iterative federated rounds in which local models are trained to run a specified number of epochs at which point the updates are shared. Adaptive learning rates and weighted aggregation is used to deal with data heterogeneity between institutions. Regularization methods are added to reduce overfitting and stabilize learning with noisy updates that have been added by differential privacy.

6 Security and Threat Model

The model spells out a clear threat model to assess privacy and security robustness. Potential attackers are honest-but-inquisitive servers, malicious clients trying to model poison, and external attackers trying to learn sensitive information about models through updating them. Among the risks of privacy leakage, there are membership inference and model inversion attacks.

To combat these attacks, the structure uses the noise injection, secure aggregation, and update validation systems. To test resilience in models, model performance and privacy resilience are tested under simulated attack conditions.

IV. RESULTS AND ANALYSIS

This part structure will be an in-depth examination of the experimental data in the assessment of the proposed Privacy Preserving AI (PPAI) framework in the financial industry. The outcomes are concerned with the realization of the effectiveness of the framework in balancing between the utility of the model, protection of privacy, security strength, and compliance with the regulator, as opposed to the conventional centralized AI models, as well as individual privacy-saving methods. Three typical financial applications were used in experimentation, namely detection of fraud, detection of credit risk, and detection of transaction anomaly. The first group of experiments compares the predictive accuracy of the suggested hybrid Federated Learning with Differential Privacy (FL+DP) framework with the centralized AI model and a federated learning model that does not provide any privacy guarantees. Accuracy, Precision, Recall, F1-score and AUC are used to measure performance.

Table 1: Performance Comparison Across Models

Model Approach	Accuracy (%)	Precision	Recall	F1-Score	AUC
Centralized AI (No Privacy)	94.8	93	92	92.5	0.96
Federated Learning (FL)	93.9	92	91	91.5	0.95
FL + Differential Privacy	92.1	90	89	89.5	0.93

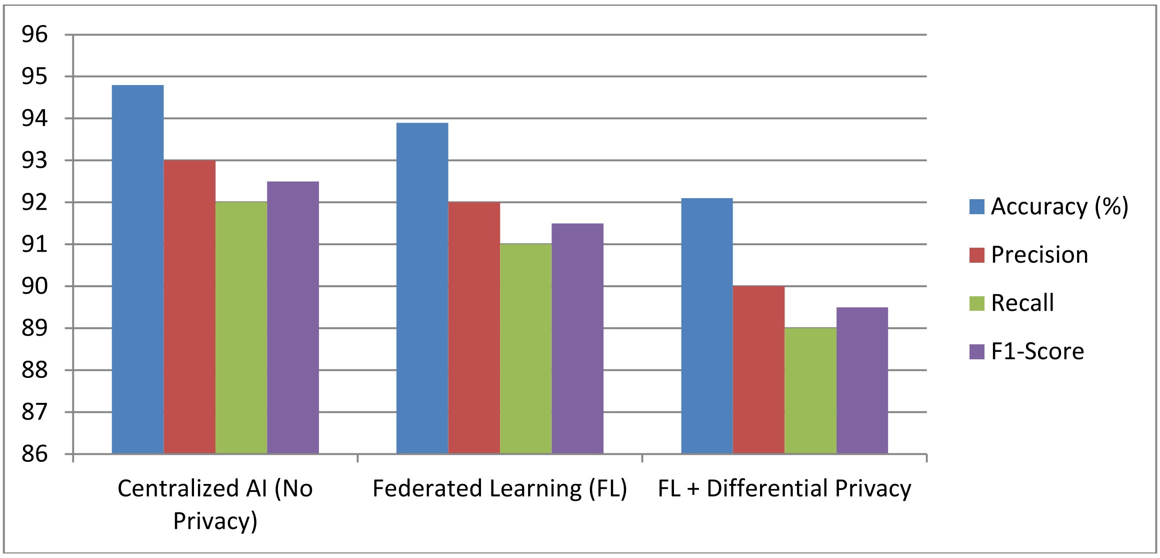

Figure 3: Performance Comparison- Across Different Approaches

The findings show that the centralized AI model has the biggest performance because of the free access to all the data. Nonetheless, this solution is highly dangerous with regard to privacy and compliance. The federated learning model shows a minimal decline in performance, which proves the fact that only the majority of the utility can be maintained by decentralized training. With the introduction of a differential privacy, there is a controlled decrease in performance (around 2-3%). This loss is warranted in practice in financial contexts, particularly compared with the high level of privacy that is obtained.

Table 2: Privacy and Security Evaluation

Model Approach	Attack Success Rate (%)	Data Exposure Risk
Centralized AI	68.5	High
Federated Learning	42.7	Medium
FL + Differential Privacy	18.9	Low

The centralized model also has a strong susceptibility to inference attacks, and attackers have identified training data as a member with more than two-thirds of success. Federated learning in itself is an effective solution to this risk because it does not imply sharing raw data, yet information is leaked through updating the model. The hybrid FL+DP model has the best privacy protection with attack success rate of less than 20%. This proves that even in collaborative learning, differential privacy is important in mitigating sensitive financial information. The robustness of security was considered by modeling the malicious client behavior, such as poisonous updates and unnatural gradient submissions. To ensure threat reduction, the proposed framework includes the noise injection and aggregation validation mechanisms. Findings indicate that decentralized AI systems are extremely vulnerable to single-point failures and attacks. Federated learning enhances resilience due to sharing training, but is susceptible to coordinated harmful updates. The hybrid framework shows an increased strength, because of the presence of differential privacy noise and secure aggregation, the impact of a single malicious party decreased. Whereas high noise may result in an unstable training process, privacy parameters can be carefully selected to keep the convergence stable. An important issue of privacy-preserving AI is the efficiency, especially in a large-scale financial implementation. The time of training, cost of communication, and cost of computation was measured in all the approaches.

Table 3: Efficiency and Overhead Comparison

Model Approach	Training Time (Relative)	Communication Overhead	Computational Cost
Centralized AI	1.0×	Low	Low
Federated Learning	1.3×	Medium	Medium
FL + Differential Privacy	1.5×	Medium–High	Medium–High

V. CONCLUSION AND FUTURE WORK

The given research is able to prove that Privacy Preserving Artificial Intelligence (PPAI) can be efficiently implemented in the financial industry to find the balance between model utility, privacy protection, security level, and regulatory compliance. The proposed hybrid Federated Learning with Differentiated Privacy (FL+DP) framework turned out to be a viable and trustworthy solution through extensive experiments in the field of fraud detecting, credit risk measurements, and anomaly detection in transactions as well. Although centralized AI models had the best predictive performance, they were very vulnerable to privacy and compliance risks. Federated learning appeased data-sharing anxieties without the official privacy guarantees. By contrast, the hybrid FL+DP system had almost centralized performance at relatively low cost: the accuracy decreased by only 2-3 percent, whereas the privacy loss and inference attack vulnerability were significantly lower. The findings also underline the fact that preservation of privacy does not necessarily negatively affect the effectiveness of AI when applied in a well thought-out hybrid form. The overhead that is presented by the framework does add some extra computation and communication overhead, but these costs are manageable in a contemporary distributed and cloud-based financial infrastructure. More to the point, the framework is compliant with regulatory demands, including the minimization of data, its localization, and auditing, which puts privacy-compliant AI as an enabler, instead of a compliance burden. The research can be further expanded in the future in a number of directions. To maximise the utility-privacy trade-off, first, there can be adaptive privacy mechanisms which dynamically adapt privacy budgets, depending on the sensitivity of risks and the criticality of use-cases. Second, more extensive implementations of secure multi-party computation and homomorphic encryption could be more effective in protecting systems in high-adversarial settings. Third, future research must consider the use of PPAI to justify and fair AI making it transparent and bias-free together with privacy. Lastly, deployment to actual pilots in various financial institutions would offer more in-depth information on scalability, interoperability, and long-term operational effects. Together, these instructions will enhance the purpose of Privacy Preserving AI as a platform of reliable, compliant, and sustainable financial intelligence systems.

REFERENCES

[1] O. P. Olaiya, T. O. Adesoga, A. Ojo, O. D. Olagunju, O. O. Ajayi, and Y. O. Adebayo, "Cybersecurity strategies in fintech: safeguarding financial data and assets," *GSC Advanced Research and Reviews*, vol. 20, no. 1, pp. 50–56, 2024.
[2] A. T. Oyewole, B. B. Oguejiofor, N. E. Eneh, C. U. Akpuokwe, and S. S. Bakare, "Data privacy laws and their impact on financial technology companies: a review," *Computer Science & IT Research Journal*, vol. 5, no. 3, pp. 628–650, Mar. 2024.
[3] S. Bhalla, C. M. Gupta, and P. Dewan, "Fintech Revolution: Navigating Consumer Privacy Concerns and Cybersecurity Challenges," in *E-banking, Fintech, & Financial Crimes: The Current Economic and Regulatory Landscape*, Cham: Springer Nature Switzerland, 2024, pp. 1–9.
[4] B. E. Abikoye, S. C. Umeorah, A. O. Adelaja, O. Ayodele, and Y. M. Ogunsuji, "Regulatory compliance and efficiency in financial technologies: Challenges and innovations," *World Journal of Advanced Research and Reviews*, vol. 23, no. 1, pp. 1830–1844, 2024.
[5] S. Chennuri and S. Biyyala, "The Fintech Revolution: Analyzing Key Innovations Reshaping the Future of Banking and Finance," *International Journal of Research in Computer Applications and Information Technology (IJRCAIT)*, vol. 7, no. 2, pp. 662–673, Nov. 2024.
[6] S. Mohammadi, A. Balador, S. Sinaei, and F. Flammini, "Balancing privacy and performance in federated learning: A systematic literature review on methods and metrics," *Journal of Parallel and Distributed Computing*, 104918, 2024.
[7] W. Hassan and H. Mohamed, "Applications of Federated Learning in AI, IoT, Healthcare, Finance, Banking, and Cross-Domain Learning," in *Artificial Intelligence Using Federated Learning*, CRC Press, 2024, pp. 175–195.
[8] P. Chatterjee, D. Das, and D. B. Rawat, "Use of federated learning and blockchain towards securing financial services," arXiv preprint arXiv:2303.12944, 2023.
[9] R. Gosselin, L. Vieu, F. Loukil, and A. Benoit, "Privacy and security in federated learning: A survey," *Applied Sciences*, vol. 12, no. 19, p. 9901, 2022.
[10] W. Ahmad, A. Vashist, N. Sinha, M. Prasad, V. Shrivastava, and J. H. Muzamal, "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*, Cham: Springer Nature Switzerland, 2024, pp. 139–156.
[11] J. Šernevičienė and A. Kabašinskas, "Explainable artificial intelligence (XAI) in finance: A systematic literature review," *Artificial Intelligence Review*, vol. 57, p. 216, 2024.
[12] A. Dalal and R. Roy, "Cybersecurity and privacy: Balancing security and individual rights in the digital age," *Journal of Basic Science and Engineering*, vol. 18, no. 1, Dec. 2021.