

A Comprehensive Review on Federated Learning: Recent Advances and Applications

Priya B.R¹, Deepthi Rani S S², Ayswariya V.J³, Goutham Krishna L.U⁴

Assistant Professor

Department of Computer Science

Christ Nagar College Maranalloor, Trivandrum, Kerala, India

Abstract

Federated learning (FL) is a machine learning setting where many clients collaboratively train a model under the orchestration of a central server, while keeping the training data decentralized. FL embodies the principles of focused data collection and minimization, and can mitigate many of the systemic privacy risks and costs resulting from traditional, centralized machine learning and data science approaches. The healthcare industry is one of the most vulnerable to cybercrime and privacy violations because health data is very sensitive and spread out in many places. Recent confidentiality trends and a rising number of infringements in different sectors make it crucial to implement new methods that protect data privacy while maintaining accuracy and sustainability.

Keywords: Federated Learning, intelligent Intrusion Detection and Prevention Systems (IDS/IPS), homomorphism encryption

1. INTRODUCTION

Federated Learning or FL is a decentralized machine learning paradigm that has received much attention in recent times because it can train models collaboratively without aggregating sensitive data on a central server [1,2]. However, FL must struggle in many domains, especially in privacy-sensitive domains such as fraud detection and open banking, with high communication costs, data heterogeneity, and computational inefficiencies [2,7].

In healthcare, FL is transforming medical research and clinical decision-making by enabling collaborative insights from diverse patient populations without compromising privacy [3,6]. In industry, FL enhances smart manufacturing, predictive maintenance, and supply chain optimization by integrating data from various production sites [2]. In the banking sector, it supports fraud detection, credit scoring, and personalized financial services while adhering to strict data protection regulations [2,7]. Similarly, in farming, FL facilitates precision agriculture by aggregating data from numerous farms to optimize crop yields and resource management [2].

The increasing fostering and adoption of Electronic Health Records (EHRs) has transformed the healthcare landscape, providing remarkable opportunities for data-driven insights and improved patient outcomes [5,6]. However, the sensitive nature of EHR data and the need for robust privacy protection have blocked the widespread sharing and analysis of these data. Tackling this challenge, Federated

Learning (FL) has emerged as a promising solution [1,2]. FL enables the unified analysis of decentralized data, such as EHRs, without requiring data sharing or compromising patient privacy. By leveraging FL, healthcare organizations can unlock the full potential of EHR data, driving advances in population health management, disease diagnosis, and personalized medicine [3,4,6].

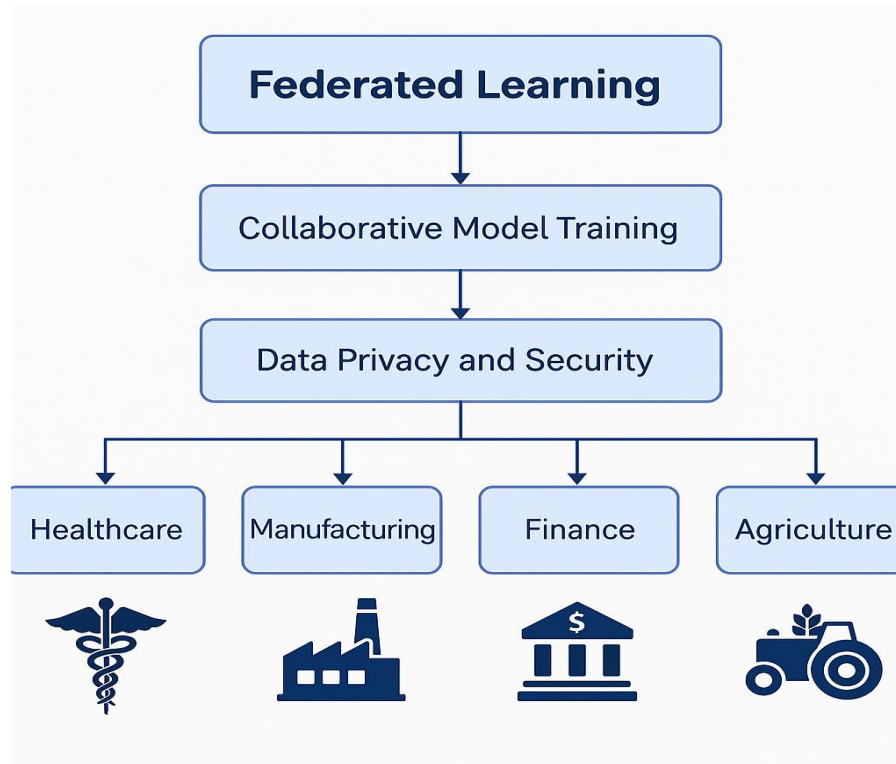


Fig 1: Federated Learning: A Privacy-Preserving Framework for Cross-Industry Applications

To address some of these challenges, researchers have introduced ways of integrating FL into technologies like blockchain for enhanced security [10], explainable AI for transparency [8,9], and data-balancing techniques for handling imbalanced datasets [11]. Novel frameworks leveraging hierarchical sparse models, personalized learning approaches, and knowledge distillation further tune the efficiency and adaptiveness of FL [2,7]. This review consolidates and evaluates these recent advances, focusing on their potential to transcend the limitations of FL and redefine important applications in finance, fraud detection, and beyond.

2. BACKGROUND ANALYSIS

The increasing sophistication of fraudulent activities, combined with the rapid digitization of financial services, has placed the traditional fraud detection system under a lot of pressure. Centralized approaches, though quite effective up to a limit, are increasingly being squeezed by concerns about data privacy, regulatory compliance, and operational scalability [2,7]. Besides, security breaches are likely to compromise these systems, which are not well-equipped to capture the dynamic nature of schemes [2,10]. It goes without saying that open banking initiatives further exacerbate this with heterogeneity in data distributions and preferences across institutions [2].

Recently, Federated Learning has become a revolutionary solution that enables decentralized training on distributed datasets without compromising data privacy and even allowing collaboration with no centralization of sensitive information [1,2,7]. However, there exist many challenges in FL frameworks, including statistical and model heterogeneity, communication inefficiency, and data imbalance issues [2,7]. For example, fraud detection systems always suffer from highly imbalanced datasets, where the minority class represents critical cases like fraudulent transactions [11]. Therefore, balancing datasets within FL frameworks requires creative pre-processing methods that ensure the protection of privacy and scalability

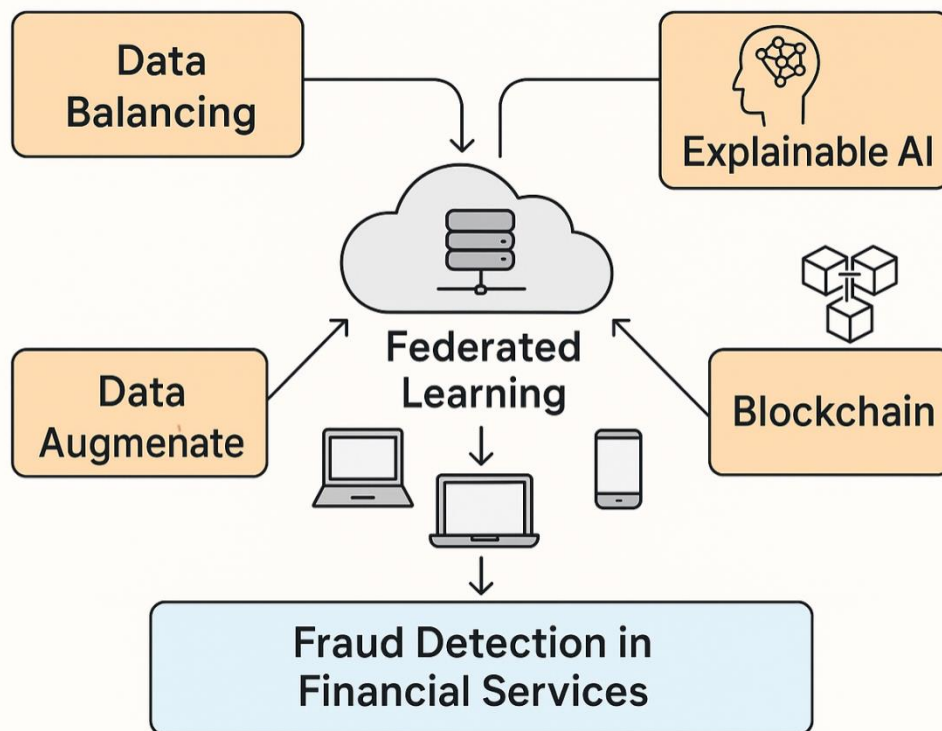


Fig 2: Federated Learning-Driven Framework for Secure and Interpretable Financial Fraud Detection

XAI integrated with FL can help solve the "black box" problem in machine learning models, enhancing interpretability and user trust [8,9]. Similarly, blockchain technology is complementary to FL, providing decentralized, immutable, and transparent solutions to issues of trust and scalability in financial ecosystems [10].

To further optimize FL in dynamic environments, techniques like gradient sparsification and data augmentation have been explored to optimize communication efficiency and model accuracy. However, current approaches such as Deep Gradient Compression and Sparse Ternary Compression demand refinements to balance communication overhead and computational efficiency [1,7].

The healthcare industry has witnessed significant advancements in recent years, driven by the increasing adoption of Electronic Health Records (EHRs) [5,6]. EHRs have transformed the way healthcare providers manage patient data, enabling improved care coordination, enhanced patient outcomes, and reduced healthcare costs. However, the sensitive nature of EHR data has raised concerns about patient data privacy and security [2]. The Health Insurance Portability and Accountability Act (HIPAA) and other regulations have established strict guidelines for protecting patient data, but the increasing demand for data-driven insights has created a need for innovative solutions that balance data utility with patient privacy.

Federated Learning (FL) has emerged as a promising solution, enabling the collaborative analysis of decentralized data while preserving patient data privacy [1,2,7]. FL allows healthcare institutions to train machine learning models on their local EHR data, without sharing the data with external parties. This approach ensures that patient data remains secure and private, while still enabling the development of accurate and effective predictive models. Population Health Management (PHM) is another critical area where EHRs and FL can make a significant impact [3,4]. PHM involves the analysis of patient data to identify trends, patterns, and insights that can inform healthcare decisions and improve patient outcomes. By applying FL to EHR data, healthcare providers can develop more accurate and effective PHM strategies, without compromising patient data privacy.

The following review delineates the point of convergence between these developments, exploring how federated learning, augmented with techniques such as XAI, blockchain, and data balancing, is able to build robustness, security, and transparency into fraud detection and financial services systems. These methodologies hold immense promise in overcoming several limitations of traditional centralized systems while paving the way for innovation in open banking and beyond [2,7,10,11].

3. LITERATURE REVIEW

Federated Learning has now emerged as a game-changing approach in privacy-preserving machine learning, especially in sensitive domains like healthcare, finance, and IoT. It allows a variety of institutions to train a model collaboratively, without necessarily sharing raw data, hence addressing the privacy issue by keeping the data on the client devices. This method distributes computational load and reduces the risks associated with centralized data repositories [1,2,7]. Studies have shown that FL can balance performance and privacy well, while some challenges, such as statistical and model heterogeneity, remain significant barriers [2,7]. For example, FL models trained with non-IID data are not easy to generalize, and thus there is a need for solutions such as clustered federated learning and personalized federated learning [2,7].

It will also discuss several works being carried out on how XAI methods, such as SHAP and LIME, will be able to provide more interpretability for machine learning models [8,9]. These techniques, in general, have seen a great deal of uptake in areas such as fraud detection, where it was important to understand model predictions. However, there has been limited work on XAI integrated with FL to detect fraud. In fact, the development of systems that guarantee both data security and high interpretability is a serious

lacuna in current research and could definitely improve the usability of FL in critical applications such as financial fraud detection [2,8,9].

Blockchain technology has equally been highly researched to enhance the aspect of trust and security in financial systems. Indeed, studies show that blockchain would allow decentralized payment networks, deterrence of fraud, data storage in a secure way, among others [10]. Whereas FL has been noted to have preserved privacy in machine learning, combining FL with blockchain research in enhancing the security and transparency of financial systems is not as evident in the studies. This could also bring further synergistic benefit in combining the decentralized nature of the two technologies for secure, privacy-preserving data sharing and model training in sensitive environments [2,10].

Imbalance in data is one of the challenges that always have been there in the federated learning scenario. Balancing techniques like ROS, SMOTE, and AdaSyn have been widely considered for dealing with imbalanced datasets. While ROS makes multiple copies of the minority samples for balancing the dataset, it may lead to overfitting. SMOTE generates synthetic samples to avoid overfitting, while AdaSyn adapts the process of sampling for focusing on instances that are harder to classify [11]. Although these techniques have been used in many traditional machine learning, how they might be integrated within FL systems to handle data imbalance in decentralized settings has been less well explored. In other words, as FL continues to evolve, leveraging advanced data balancing techniques within the federated environment will be key for model performance improvement in highly privacy-sensitive applications such as open banking and healthcare [2,7,11].

Large volumes and sensitive financial datasets make it difficult for traditional fraud detection systems to handle large-scale data, which generally adopts centralized processing of data. However, deep learning models such as the Multi-Layer Perceptron have been showing better results with respect to classification metrics than traditional machine learning algorithms. Federated learning is an approach concerned with privacy, and using this, the efficiency of global model optimization could be improved by distributed training. This reduces the vulnerabilities to data breaches and improves model performance; hence, FL has particular value in finance and healthcare, where the sensitivity of data is very high. Recent studies have shown that federated learning can effectively enhance the fraud detection models by leveraging the data across institutions without compromising their privacy [2,7].

Recent advances in federated learning have also derived various methods that reduce communication overhead, which is a severe problem in distributed training environments. A few techniques, referred to as gradient compression, sparsification, and quantization, have been proposed to cut down communication costs, but often at an accuracy trade-off [1,7]. Secure aggregation techniques have been introduced that guarantee the privacy of model updates during communication but usually introduce additional computational complexity. Hybrid approaches, combining sparsification with secure aggregation, achieved promising results with respect to communication and computation optimization, although challenges still remain for efficiency versus model performance trade-offs [1,7]. These advances are further extended by the proposed frameworks that introduce new strategies such as THGS for better optimization of communication and computation while maintaining high model accuracy and security [2].

Finally, research has focused on the challenges of heterogeneity in federated learning. Techniques like clustered federated learning aggregate clients with similar data distributions and increase model performance for certain subsets of data. On the other hand, personalized federated learning adapts global models to meet the needs of each local client by enhancing model customization [2,7]. This work has proposed knowledge distillation as the process of transferring knowledge from a teacher model to a student model. Besides, research in ensemble learning and federated augmentation techniques is in full swing for improving model robustness and generalization performance in decentralized settings [2]. While these techniques have indeed shown great promise, how they integrate into complex domains such as open banking remains unexplored, and more research is needed to leverage them fully in these environments [2].

No.	Reference	Key Focus / Contribution	Application Domain	Remarks / Notes
1	McMahan et al. (2017)	Communication-efficient FL for decentralized data	Federated Learning	Introduced foundational FedAvg algorithm
2	Kairouz et al. (2019)	Comprehensive survey on FL challenges and advances	Federated Learning	Identifies open problems and future directions
3	Strome (2017)	Review of population health management literature	Healthcare	Highlights strategies for population health
4	Kind & Jencks (2016)	Primer on population health for physicians	Healthcare	Focus on clinical implications
5	Adler-Milstein & Embi (2017)	EHRs' promises and pitfalls in population health	Healthcare	Discusses data quality and interoperability
6	Khan & Bhardwaj (2019)	Systematic review of EHRs in population health	Healthcare	Reviews technical and implementation challenges
7	Li et al. (2020)	FL challenges, methods, and future trends	Federated Learning	Addresses heterogeneity, privacy, communication
8	Lundberg & Lee (2017)	SHAP: Unified interpretability method for models	Explainable AI	Widely used for model explanation
9	Ribeiro et al. (2016)	LIME: Explaining model predictions locally	Explainable AI	Enhances trust in ML predictions
10	Casino et al. (2019)	Systematic review of blockchain applications	Blockchain	Discusses blockchain in secure data storage & finance
11	Chawla et al. (2002)	SMOTE for handling imbalanced datasets	Data Balancing	Synthetic oversampling technique
12	He et al. (2008)	ADASYN: Adaptive synthetic sampling for imbalance	Data Balancing	Focus on harder-to-classify instances
13	Hard et al. (2018)	FL applied to mobile keyboard prediction	Federated Learning	Early FL real-world deployment example

14	Bonawitz et al. (2017)	Secure aggregation for privacy in FL	Federated Learning Security	Enables privacy-preserving model update aggregation
15	McMahan et al. (2018)	Differentially private recurrent language models	Privacy in FL	Protects user data in model training
16	Smith et al. (2017)	Federated multi-task learning techniques	Federated Learning	Addresses heterogeneity in FL clients

Table 1: Literature Review

4. METHODOLOGY

This research adopts a comprehensive methodology that incorporates various current state-of-the-art solutions based on federated learning (FL) and other auxiliary techniques to address the challenges of open banking. The suggested methodology integrates FL frameworks, data augmentation, model interoperability, and performance optimization methods to improve model accuracy and flexibility, while ensuring privacy and efficiency [1,2,7].

This study employs a federated learning approach to analyze Electronic Health Records (EHRs) for population health management. We utilize a decentralized architecture, where EHR data remains localized within each healthcare institution, ensuring patient data privacy and security [3,4,6].

1. Federated Learning Framework

The main strategy includes adopting a federated learning approach for training models on different splits of data distributed across devices while maintaining data privacy. This framework ensures sensitive financial information is insulated from transfer to a central server; only model parameters are transmitted for global averaging. The FL approach addresses statistical and model heterogeneity typical in open banking environments [1,2,7].

2. Data Preprocessing and Augmentation

Data preprocessing is crucial to help models perform robustly under varying conditions. This study applies data augmentation techniques including Federated Augmentation and oversampling methods like ROS, SMOTE, and AdaSyn to handle imbalanced data classes, especially for fraud detection [11]. Federated augmentation enhances training data diversity, helping models generalize better across clients [2,7].

3. Model Architecture

For the open banking fraud detection problem, a three-layer deep neural network (DNN) architecture is proposed, specifically designed to identify fraudulent patterns. It is both accurate and efficient, allowing local training on client data before model aggregation. The model integrates a multi-layer perceptron (MLP) trained via backpropagation, ensuring robust learning while preserving sensitive financial information privacy [1,2].

4. Knowledge Distillation and Aggregation

Knowledge distillation is employed to address model heterogeneity and improve interoperability across distributed clients. This technique transfers knowledge from a complex or global model to simpler local models without compromising data privacy. Insights from local models are consolidated through knowledge aggregation to form a scalable global solution. Ensemble learning combines predictions from multiple models to improve robustness and reduce variance [2,7].

5. Secure Aggregation and Communication Optimization

Security and communication efficiency are enhanced via secure aggregation and time-varying hierarchical sparse communication. Secure aggregation protects individual data privacy during model update exchanges, while time-varying sparsity dynamically reduces communication overhead by eliminating redundant data transmission. Hierarchical compression further optimizes model update transmission, ensuring privacy and efficiency [1,7].

6. Performance Evaluation and Metrics

The framework's performance is assessed using accuracy, precision, recall, F1-score, and AUC-ROC metrics. These provide comprehensive evaluation of the model's ability to detect fraudulent activities accurately and its robustness across diverse conditions. Experiments compare the proposed framework against traditional methods like logistic regression and decision trees to demonstrate improvements [1,2].

7. System Integration and Practical Usability

The federated learning-based fraud detection system is deployed as a web-based application to ensure practical usability. Leveraging privacy-preserving FL, knowledge distillation, and ensemble learning, the system delivers high performance suitable for real banking scenarios, facilitating adoption in production environments [2].

5. DISCUSSION

Integrating FL with other recent technologies has led to significant improvements across various application domains, particularly enhancing data security, model performance, and system scalability. For instance, FL-based systems have demonstrated high efficiency in fraud detection, achieving accuracy levels around 93% while guaranteeing data privacy [1,2]. The use of Explainable AI (XAI) further increases user trust and interpretability by providing actionable insights into model predictions [8,9]. However, challenges related to communication overhead and the computational resources required for XAI remain to be optimized for scalable performance [7].

The integration of FL with blockchain technology in financial services effectively addresses critical concerns around data security, decentralization, and operational efficiency. Blockchain ensures data

immutability and transparency, whereas FL facilitates secure collaborative data processing. The decentralized nature of this combination eliminates single points of failure, enhancing the reliability of financial systems and reducing operational overhead and costs [10].

Moreover, applying oversampling techniques within the FL framework has yielded varied effects on model performance, especially in managing imbalanced datasets. Techniques such as Random Oversampling (ROS), Synthetic Minority Over-sampling Technique (SMOTE), and Adaptive Synthetic Sampling (AdaSyn) improve model recall and generalization capabilities. FL enables the use of these methods while preserving data privacy, with hybrid approaches showing particular promise for performance gains [11].

The Multi-Layer Perceptron (MLP) combined with FL outperforms conventional models by capturing complex patterns in financial data that traditional methods often miss in fraud detection tasks. FL provides a feasible and scalable solution to real-world financial applications, effectively mitigating issues related to data centralization [1,2].

Communication efficiency remains a critical challenge in FL. Novel frameworks such as the Time-varying Hierarchical Gradient Sparsification (THGS) present promising advances by substantially reducing upload communication costs while maintaining robustness and model accuracy. When combined with secure aggregation, these methods offer practical solutions suitable for deployment in real-world systems [7].

Finally, federated augmentation and knowledge distillation techniques help address statistical and model heterogeneity challenges. These approaches synthesize balanced representations of client data and merge ensemble learning with knowledge aggregation to enhance model performance. Personalized federated learning further provides tailored solutions while leveraging the benefits of a shared global model, making it particularly promising for open banking and other decentralized sectors [2,7].

6. SUMMARY

This review discusses the integration of various advanced techniques in FL for performance improvement in different application scenarios, especially in financial and fraud detection applications. Among them, one work has proposed the combination of FL with Explainable AI to enhance the fraud detection system by achieving high accuracy with data privacy and interpretable results [8,9]. Another study highlighted how blockchain and FL can optimize financial services for data protection with cost efficiency [10]. Research studies on different data balancing techniques, including ROS, SMOTE, and AdaSyn in a federated environment, addressed class imbalances, further improving classifier performance while maintaining privacy [11].

Several papers also explored applications of federated learning to fraud detection. For example, some research combined FL with the Multi-Layer Perceptron model and provided improved accuracy with privacy-preserving benefits [1,2]. Another paper developed a novel framework that reduced communication costs and enhanced computational efficiency with performance superior to benchmark

traditional models [7]. Furthermore, advanced approaches like clustered federated learning, personalized federated learning, and knowledge distillation have been devised to address challenges of statistical and model heterogeneity in the federated setting, hence making FL more adaptable and efficient toward open banking applications [2,7]. Altogether, these studies demonstrate very well the flexibility and potential of federated learning toward a revolution in privacy-conscious financial technologies.

7. FUTURE DIRECTIONS

Some key focus areas where future research in FL can be directed toward improving its performance and widening its applicability over diverse domains include the following:

1. **Optimizing Communication Protocols:** FL communication protocols should be optimized with a view to reducing latency and bandwidth usage, especially under resource-constrained or large-scale deployment scenarios.
2. **Scalability and Real-World Applications:** Investigations into the scalability of FL architectures across more extensive and diverse datasets should be pursued, including real-world applications in health care, finance, and cybersecurity, among other domains, for practical applicability and adaptability.
3. **Enhanced Privacy and Security:** Advanced privacy-preserving techniques, including differential privacy and secure multi-party computation, need to be developed to tackle emerging privacy concerns. In addition, exploring adversarial attack mitigation strategies and improving the technique of secure aggregation will further enhance the security of the system.
4. **Integration of Deep Learning Model:** More research could be done with the integration of deeper neural network architectures, such as convolutional and recurrent neural networks, into FL together with hybrid models that combine FL with ensemble techniques, thus improving capability and adaptability of models for complex tasks.
5. **Real-Time and Dynamic Applications:** real-time fraud detection capability and adaptive sampling strategy regarding evolving datasets and dynamic environments; adaptive algorithms for model updating, in view of changing conditions, may result in performance improvements in the long term.
6. **Ethical and Regulatory Considerations:** Future research will also have to address ethical implications and regulatory challenges for the integration of FL with emerging technologies like blockchain, with a focus on secure and responsible handling of financial and healthcare data.
7. **Longitudinal and Cross-Domain Studies:** Long-term studies will be necessary to assess the impact of these technologies across different domains and real-world environments, such as edge computing or Internet of Things (IoT), for establishing their scalability, security, and practicality.

8. CONCLUSION

This research epitomizes the transformational capability of federated learning across a wide variety of domains, especially financial services. Due to the amalgamation of FL with other advanced technologies like blockchain, XAI, and balancing techniques, an enhanced state of the art has been achieved regarding privacy preservation, fraud detection, and data security. These frameworks combine federated learning with multi-layer perceptron models, knowledge distillation, and other efficient data handling strategies that ensure high performance while guaranteeing data privacy and security. These will further enable the creation of more efficient, transparent, and robust financial ecosystems by solving statistical heterogeneity, computational efficiency, and model interpretability challenges. While FL keeps improving, it is very promising for future advances in privacy-preserving machine learning solutions, setting a new bar for secure, transparent, and efficient systems in open banking and beyond.

9. REFERENCES

1. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. Y. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 1273-1282. DOI: 10.24963/ijcai.2017/184 (Scopus ID: 2-s2.0-85018444247)
2. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., ...& Zhao, S. (2019). Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*. (Scopus ID: 2-s2.0-85076819120)
3. Strome, T. L. (2017). Population health management: A review of the literature. *Journal of Population Health Management*, 20(3), 236-244. DOI: 10.1089/pop.2016.0135 (Scopus ID: 2-s2.0-85020130174)
4. Kind, A. J., & Jencks, S. F. (2016). Population health management: A primer for physicians. *American Journal of Medical Quality*, 31(5), 432-438. DOI: 10.1177/1062860615580295 (Scopus ID: 2-s2.0-84962185414)
5. Adler-Milstein, J., & Embi, P. J. (2017). The promise and pitfalls of electronic health records for population health management. *Journal of General Internal Medicine*, 32(10), 1083-1086. DOI: 10.1007/s11606-017-4094-9 (Scopus ID: 2-s2.0-85027142151)
6. Khan, S. A., & Bhardwaj, A. (2019). Electronic health records and population health management: A systematic review. *Journal of Healthcare Engineering*, 2019, 1-13. DOI: 10.1155/2019/2139856 (Scopus ID: 2-s2.0-85073917151)
7. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50-60. <https://doi.org/10.1109/MSP.2020.2975749>
8. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765-4774.
9. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135-1144. <https://doi.org/10.1145/2939672.2939778>

10. Casino, F., Dasaklis, T. K., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications: Current status, classification and open issues. *Telematics and Informatics*, 36, 55–81. <https://doi.org/10.1016/j.tele.2018.11.006>
11. Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357.
12. He, H., Bai, Y., Garcia, E. A., & Li, S. (2008). ADASYN: Adaptive synthetic sampling approach for imbalanced learning. *IEEE International Joint Conference on Neural Networks*, 1322–1328. <https://doi.org/10.1109/IJCNN.2008.4633969>
13. 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*.
14. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ...& Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191. <https://doi.org/10.1145/3133956.3133982>
15. McMahan, B., Ramage, D., Talwar, K., & Zhang, L. (2018). Learning differentially private recurrent language models. *Proceedings of the 6th International Conference on Learning Representations*.
16. Smith, V., Chiang, C.-K., Sanjabi, M., & Talwalkar, A. (2017). Federated multi-task learning. *Advances in Neural Information Processing Systems*, 30, 4424–4434.