


Federated Learning in Customer-Centric Applications: Balancing Privacy, Personalization and Performance

Xinxiang Gao ^a

Department of Computer Science, University of Toronto, Toronto, Canada

Keywords: Federated Learning, Data Heterogeneity, Model Convergence, Secure Collaboration.


Abstract: Federated Learning (FL) offers a practical way to meet the growing need for privacy-preserving machine learning, especially in customer-focused areas like finance, retail, and business collaboration. It allows models to be trained securely and in a decentralized way, without collecting sensitive data in one place. This paper examines how FL facilitates secure and decentralized model training while avoiding the centralization of sensitive data. It addresses key challenges such as data heterogeneity, model convergence, and privacy concerns, proposing methods like clustered FL, asynchronous updates, and attention-based mechanisms for improving model performance. The paper also discusses privacy vulnerabilities, such as gradient leakage, and explores solutions like differential privacy and secure aggregation. Although federated learning enhances data privacy and service personalization, it faces limitations, including increased computational complexity and communication overhead. Future research needs to prioritize enhancing privacy safeguards while ensuring model accuracy and scalability. This review can serve as a valuable reference for researchers looking to understand the advancements in this field.

1 INTRODUCTION

Machine Learning (ML) has established itself as a transformative technology in modern industry, as they enable businesses to utilize vast pools of data for automation, prediction, and personalization. Lu (2019) systematically analyzes Artificial Intelligence (AI) from fundamental mechanisms to practical applications and highlights the significant role of AI as a driver of industrial progress and the integration of emerging technologies. Machine learning effectiveness is often tied to the quantity of data, as demonstrated by researchers Kaplan et al. (2020), who show that larger datasets result in reduced loss in regarding neural networks. This drives organizations to gather vast amounts of user information for training and refining their models. Conventionally, data is centralized in cloud storage, where machine learning models are trained on aggregated user information. The centralized user information retrieval has clearly facilitated the improvement of the model performance but has also been the cause for serious questions on privacy, security, and compliance. Now, with rising consumer concerns for

data privacy and consequent laws by regions of the globe, companies are posed with the challenge of being innovative and yet responsible with their use of data.

Addressing such problems, federated learning is a promising approach by which ML models can be trained without centralizing the data, protecting the personal data from leaking. More specifically, federated learning combines elements of machine learning, distributed computing, and privacy-preserving methods, offering benefits such as reduced latency, lower communication overhead, and decreased power usage (Naik, 2024). For enterprises and customers alike, federated learning is a big leap in the trajectory of privacy-preserving technology. These companies can now deliver far better personalized services without compromising user data; the privacy and trust of the customer are also taken good care of. For example, Ahmed et al. (2022) proposes an innovative federated learning method for performing customer analysis and clustering based on transaction behavior, which addresses the privacy concerns associated with real-time streaming data, particularly in the context of retail sales. Since such data contains confidential information as well as

^a <https://orcid.org/0009-0009-1011-9148>

transactions that carry high intrinsic value, employing FL aims to take advantage of the data while preventing privacy leakage.

Federated learning has also emerged as a transformative solution in industries such as finance, retail, and enterprise collaboration by enabling secure, collaborative model training across multiple parties without sharing raw data. In the financial sector, federated learning enhances the evaluation of customers' credit and affordability by securely modeling their financial conditions while addressing the challenges of non-identically distributed data and adversarial risks, as highlighted by Long et al. (2020). In retail, as mentioned in Yang et al. (2019), it allows companies to deliver personalized recommendations while safeguarding sensitive customer information. Similarly, in cross-enterprise information exchange, federated learning facilitates the creation of collaborative ecosystems where enterprises can leverage shared insights without exposing proprietary or personal data. By ensuring data privacy, security, and regulatory compliance, federated learning is transforming the way businesses operate and interact, paving the way for more personalized and secure AI applications.

This paper explores how federated learning changes the game of customer-centric applications through the safe, efficient, and personalized experience that privacy requirements of today mandate. This review can serve as a valuable reference for researchers aiming to understand recent advancements in privacy-preserving machine learning through federated learning.

2 METHOD

2.1 Preliminaries of Federated Learning

Federated Learning, firstly proposed by McMahan et al. (2017), is a distributed machine learning approach where multiple clients like mobile devices collaboratively train a shared global model without exchanging their local data. The workflow is as follows: Instead of sending data to a central server, each client processes data locally, and only the model parameters gradients are sent to the central server. The server aggregates these updates to improve the global model, which is then redistributed to the clients for further training.




Figure 1: A Schema for FedAvg (Li et al., 2020)

The training process in Figure 1, as illustrated by Li et al. (2020), showed Federated Average (FedAvg), the baseline FL algorithm in numerous studies. In general, in the aggregation step, the server aggregates the updates using a weighted average approach. Due to its ability to keep data decentralized and not exposed to a central server, federated learning is well-suited for applications involving privacy-sensitive data.

2.2 Customers in Financial Evaluation

To evaluate the user's affordability and their credit, modeling customer's financial conditions becomes necessary. For such scenarios, confidential data needs to be protected, leading to a demand for federated learning. However, participants in the FL system having non-identically and independently distributed (non-IID) data, different participants varying feature spaces or model architectures, and the risk of adversarial attacks, such as poisoning local models, pose a threat to the model performance, as addressed by Long et al. (2020). Then, Long et al. (2022) introduced clustered federated learning, where participants are grouped by similar data distributions, each contributing to a unique global model. It contributes to personalized federated learning, where each participant has a unique model consisting of a shared part and personalized components, tailored to their specific data characteristics, which facilitates the modeling of the user.

Imteaj and Amini (2022) use federated learning to build a global model for assigning credit scoring by predicting consumers' financial conditions without requiring sensitive customer data to be shared with external entities. To address the potential delay caused by slower or weaker agents, they allow partial work contributions, ensuring that the overall model convergence is not significantly impacted. They

implement asynchronous FL to reduce times waiting for updates in the training process, allowing for more efficient updates as local models complete their training at different rates. The results show that their FL-based model achieves a similar F1-score to centralized approaches, even with a high level of data skewness. It also outperforms other state-of-the-art FL models, especially when dealing with resource-constrained agents, by achieving approximately 5-6% higher accuracy on average.

2.3 Customers in Retail Purchases

In the retail sector, understanding and predicting customer behavior is crucial for providing personalized services, enhancing customer satisfaction, and driving sales. However, the traditional approach of centralizing customer data for analysis poses significant challenges, especially in cases of security and resource management. FL provides a compelling framework to these obstacles as sensitive data is no longer required to be gathered for training. This is especially vital in retail, as transaction data often contains personal information that must be protected.

For addressing such issues, Ahmed et al. (2022) introduce a groundbreaking federated learning approach for analyzing and clustering customers according to their transaction patterns. More specifically, they elevate the Federated Learning Framework and use an attention-based model to generate low-dimensional embeddings from transaction data. The attention mechanism helps in focusing on relevant parts of the data, improving the quality of the embeddings. The master server consolidates the model updates from all edge devices involved in the process. Once the model converges, it is shared back with the edge devices or clients, and a Clustering-Based Dynamic Method (CBDM) is applied to the embedded data. CBDM is an efficient pattern mining method that clusters customers based on their purchase behavior. Semantic embedding is used to extract and cluster relevant patterns, leading to more meaningful and actionable insights. The proposed method achieved ROC values of 0.75 for random distribution and 0.70 for fixed distribution, indicating a good balance between true positive and false positive rates. This suggests that the method is effective in identifying meaningful clusters. Additionally, the approach significantly reduces communication costs and ensures privacy, making it a viable and efficient solution for retail applications.

2.4 Customers in Enterprise Information Exchange

In the modern business landscape, effective information exchange between different enterprises is crucial for creating a collaborative ecosystem that enhances customer experiences, drives innovation, and fosters mutual growth. More specifically, if model captures the collective knowledge from all participating enterprises, they are able to benefit from shared insights. Moreover, the collaborative approach leads to more robust and accurate recommendation models, as the system leverages a broader and more diverse dataset. However, the approach of sharing confidential data between enterprises poses significant threats to customer's privacy and security, thus fails to meet regulatory requirements for enterprises. Such cross-enterprise scenarios where require the protection of confidential data would be well-suited for the application of FL.

To tackle such issues, Li et al. (2023) proposed an innovative federated learning method for cross-enterprise recommendation systems. More specifically, they elevate Graph Neural Networks (GNNs) within FL framework to enhance the recommendation quality for users while preserving data privacy. The GNNs are used to model the complicated connections between users, items, and their interactions, which are essential for accurate and personalized recommendations. The federated learning framework ensures that each enterprise's data remains on-premises, minimizing the likelihood of data leakage and unexpected access. Secure communication protocols are employed to protect the model updates during transmission, ensuring that the entire process is secure and compliant with data protection regulations. The results suggest that, when compared to conventional centralized methods, the proposed approach not only theoretically safeguards privacy but also attains greater recommendation accuracy. More specifically, the integration of GNNs within the federated learning framework resulted in more accurate and relevant recommendations, demonstrated by enhanced precision, recall, and F1-score. Additionally, the federated learning method significantly lowered communication demands and computational expenses, rendering it an adaptable and effective option for cross-enterprise collaboration.

3 DISCUSSION

Despite federated learning presents the potential to addressing concerns around confidential data

management in the context of large-scale ML, several key challenges still need to be addressed for its widespread adoption across industries, especially in customer-centric applications. This section explores the major challenges, approaches proposed in the reviewed literature, and their limitations, focusing on the heterogeneity of data, model convergence, and privacy preservation.

3.1 Data Heterogeneity

One of the major obstacles in federated learning is the variability of data among various edge devices. In practical applications, the data among various participants is usually naturally Non-identically Distributed (non-IID), which can pose a threat, if not addressed properly, to the robustness of the global model. More specifically, users may belong to different organizations with various multi-dimensional representations, while the same user might have different consumptive behaviors among multiple organizations. This is particularly relevant in customer-based applications, where each user's data distribution can differ based on region, usage patterns, and other variables.

To resolve this, Long et al. (2022) proposed clustered federated learning, that groups edge devices based on similar empirical distributions, allowing for the development of multiple global models tailored to each cluster. This approach not only improves the personalization of services but also handles the non-IID data issue. Similarly, Ahmed et al. (2022) addressed data heterogeneity in the retail sector by employing an attention-based mechanism to focus on relevant transaction data, enabling better clustering and analysis of customer behavior.

Moreover, to leverage non-IID data in FL, a key approach as well as a challenge is aligning the different vector spaces through federated transfer learning, as mentioned by Yang et al. (2019) Federated Transfer Learning, as introduced by Liu et al. (2020), handles differences in both samples and features by finding a common representation between different feature spaces and applying this to predict outcomes for samples with one-sided features.

Despite these advances, limitations remain. Clustered federated learning, for instance, requires accurate clustering of clients, which can be difficult in cases of subtle data heterogeneity or where client data distributions change dynamically. This method may also increase computational complexity and communication overhead by managing multiple models. Additionally, while attention-based mechanisms improve data focus, they may not fully resolve non-IID data issues, particularly in highly varied datasets, limiting the generalization ability of

the global model. Furthermore, While Federated Transfer Learning aims to establish a common representation across different feature spaces, its limitations are evident. Assessing the validity or optimality of this representation can be challenging before evaluating model results, as its effectiveness typically emerges only after training. Consequently, an inadequately determined representation can complicate outcome assessments and negatively impact results.

3.2 Model Convergence and Resource Constraints

Another challenge lies in the efficiency of model convergence, particularly in resource-constrained environments. As noted by Imteaj and Amini (2022), FL often involves participants with varying computational resources, which can lead to delays in the overall model convergence. These participants, referred to as stragglers, slow down the global model's updates, especially in real-time applications like financial evaluation.

To overcome this, asynchronous federated learning has been introduced, allowing participants to send updates as they complete training rather than waiting for all clients to finish (Xu et al., 2023). This method helps mitigate the impact of slower devices, ensuring faster convergence without significantly compromising accuracy. Moreover, the partial work contribution technique, where weaker agents contribute only to the extent of their capability, was found to be effective in maintaining model performance while reducing training delays.

However, the limitations of asynchronous updates and partial contributions must be considered. Asynchronous federated learning may lead to stale updates from slower clients, which can impact the master model accuracy and stability. In particular, model divergence risk increases as updates become less synchronized, especially in environments where device performance varies significantly. Partial contributions, while reducing delays, may inadvertently lower the contribution of certain clients, leading to a bias in the model toward clients with higher computational power.

These limitations are particularly critical in edge-based federated learning, common in sectors like retail, where data is gathered from devices with limited power. While asynchronous updates and communication optimization can help maintain real-time services, maintaining an equilibrium between accuracy and efficiency remains a challenge.

3.3 Privacy and Security Concerns

Federated learning, though dedicated to preserve privacy by keeping data from the server and only exchanging gradients as updates instead, still faces significant privacy and security challenges. The central concern lies in how gradients can inadvertently leak sensitive information. Attackers can exploit these gradients to reconstruct original data, as demonstrated by "deep leakage from gradients" technique from Zhu et al (2019), which shows how seemingly innocuous gradient information can be reverse engineered into sensitive training data. This undermines the foundational premise of FL, which states that even when data remains decentralized, the model update transfers can expose vulnerabilities.

Wei and Liu (2022) offer insights into how differentially private algorithms, though intended to protect against privacy breaches, can still fall prey to gradient leakage attacks when using fixed privacy parameters. The introduction of dynamic privacy parameters, which adjust noise based on the behavior of gradient updates, shows promise in enhancing privacy resilience while maintaining model accuracy. Despite this progress, however, fully mitigating these threats remains an ongoing challenge.

Further research into secure federated learning has explored various techniques to bolster privacy defenses. Approaches such as differential privacy by Zhu et al. (2019), secure aggregation by Kairouz et al. (2021), and homomorphic encryption by Phong et al. (2018) have been proposed to protect against not only gradient leakage but also other adversarial attacks like backdoor injections and data poisoning. However, many of these methods either reduce model accuracy or do not provide full protection against sophisticated attacks, leaving open problems for future research.

While federated learning allows organizations to build comprehensive models without sharing raw data, the evolving landscape of attacks continues to expose gaps in current defenses. More robust and performance-preserved encryption techniques as well as adaptive privacy mechanisms remain critical areas of research. Moreover, issues such as fairness and heterogeneity across clients contribute additional complexity to designing secure and efficient federated models (Bagdasaryan et al., 2020).

4 CONCLUSIONS

This study offers a comprehensive review of how federated learning researches faces and tackles crucial challenges in customer-centric applications, particularly in finance, retail, and cross-enterprise

collaboration. FL enables decentralized, privacy-preserved and collaborative model training, which is increasingly important in privacy-conscious industries. Through the studies reviewed, FL has shown its capacity to improve services like financial evaluation, personalized retail recommendations, and secure information sharing among enterprises, offering a balance between privacy preservation and machine learning effectiveness.

However, this study also highlights several challenges that limit the widespread deployment of FL. Data heterogeneity, where clients have varying and non-IID data, can negatively impact model performance, and methods like clustered federated learning attempt to mitigate this issue but add computational complexity. Additionally, achieving model convergence in environments with resource-constrained devices remains difficult, despite the introduction of techniques like asynchronous updates and partial contributions. Security risks, particularly gradient leakage, still pose threats to privacy, even in decentralized systems. Applying privacy-preserving methods along with FL offers partial solutions, but they often reduce model accuracy and introduce communication overhead. Future investigation is needed to refine methods and improve the scalability and robustness of federated learning systems.

REFERENCES

- Ahmed, U., Srivastava, G., & Lin, J. C.-W. 2022. Reliable customer analysis using federated learning and exploring deep-attention edge intelligence. *Future Generation Computer Systems*, 127, 70-79.
- Bagdasaryan E, Veit A, Hua Y, Estrin D, Shmatikov V 2020. How to backdoor federated learning. In: *International Conference on Artificial Intelligence and Statistics*. PMLR, pp 2938–2948.
- Imteaj, A., & Amini, M. H. 2022. Leveraging asynchronous federated learning to predict customers financial distress. *Intelligent Systems with Applications*, 14, 200064.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Yang, Q. 2021. *Advances and open problems in federated learning*. arXiv.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., ... & Amodei, D. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Li, L., Fan, Y., Tse, M., & Lin, K.-Y. 2020. A review of applications in federated learning. *Computers & Industrial Engineering*, 149, 106854.

- Li, Z., Bilal, M., Xu, X., Jiang, J., & Cui, Y. 2023. Federated learning-based cross-enterprise recommendation with graph neural networks. *IEEE Transactions on Industrial Informatics*, 19(1), 673–682.
- Liu, Y., Kang, Y., Xing, C., Chen, T., & Yang, Q. 2020. A secure federated transfer learning framework. *IEEE Intelligent Systems*, 35(4), 70–82.
- Long, G., Tan, Y., Jiang, J., & Zhang, C. 2020. Federated learning for open banking. In Q. Yang, L. Fan, & H. Yu (Eds.), *Federated learning: Privacy and incentive* (pp. 240–254). Springer International Publishing.
- Long, G., Xie, M., Shen, T., Zhou, T., Wang, X., & Jiang, J. 2022. Multi-center federated learning: Clients clustering for better personalization. *World Wide Web*, 26(1), 481–500.
- Lu, Y. 2019. Artificial intelligence: a survey on evolution, models, applications and future trends. *Journal of Management Analytics*, 6(1), 1–29.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., & Aguera y Arcas, B. 2017. Communication-efficient learning of deep networks from decentralized data. In A. Singh & J. Zhu (Eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (Vol. 54, pp. 1273–1282). PMLR.
- Naik, D., Naik, N. 2024. An Introduction to Federated Learning: Working, Types, Benefits and Limitations. In: Naik, N., Jenkins, P., Grace, P., Yang, L., Prajapat, S. (eds) *Advances in Computational Intelligence Systems*. UKCI 2023. *Advances in Intelligent Systems and Computing*, vol 1453. Springer, Cham.
- Phong, L. T., Aono, Y., Hayashi, T., Wang, L., & Moriai, S. 2018. Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions on Information Forensics and Security*, 13(5), 1333–1345.
- Wei, W., & Liu, L. 2022. Gradient leakage attack resilient deep learning. *IEEE Transactions on Information Forensics and Security*, 17, 303–316.
- Xu, C., Qu, Y., Xiang, Y., & Gao, L. 2023. Asynchronous federated learning on heterogeneous devices: A survey. *Computer Science Review*, 50, Article 100595.
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. 2019. Federated machine learning: Concept and applications. *arXiv*.
- Zhu, L., Liu, Z., & Han, S. 2019. Deep leakage from gradients. *arXiv*.