International Journal of Creative and Open Research in Engineering and Management
ISSN: **3108-1754 (Online)**
Volume 02 Issue 05 May-2026 | Impact Factor: 3.5

Federated Learning: A Comprehensive Review of Strategies, Challenges, and Applications

Sewa Khatter
Department of Electrical and Electronics Engineering Maharaja Surajmal Institute of Technology, New Delhi
sewa.khatter.2301@gmail.com

Abstract—Federated learning(FL) represents a subsection of machine learning that allows multiple clients to work collaboratively whilst not sharing the private raw data [2]. This review paper systematically goes over FL architectures, specifically cross-device and cross-silo deploy-ments [3]. We explain core optimization algorithms such as federated averaging (FedAvg) & its variants, and analyse their mathematical im-plementation and performance [4]. We also address the key challenges faced during FL such as data heterogeneity (non-IID data), commu-nication efficiency, system heterogeneity, fairness & bias, and system vulnerability to privacy breaches [5]. Additionally, the paper discusses the various applications of FL across healthcare, mobile edge computing and Internet of Things (IoT) [6]. Lastly, this paper outlines the future research trajectories so that FL models can be deployed at a global scale [7].

Index Terms—Federated Learning, Distributed Optimization, Cross-Device, Cross-Silo, Non-IID Data.

1 Introduction

Traditionally machine learning has always relied on cen-tralised data collection, where data from multiple users is collected on to a server for training of models. While this approach is effective, in settings such as healthcare & mobile systems, the data is highly sensitive and cannot be freely shared. The centralised data collection also raises strong concerns over privacy, data ownership and security. As the restrictions by regulatory environments get increasingly stricter over the movement of sensitive customer informa-tion, the traditional methods get legally and practically more constrained [6].

To bypass these challenges, the researchers have intro-duced an alternate framework that allows decentralised model training. In Federated Learning, instead of trans-ferring the data to a server and then training the model, the server distributes a global model directly to the client devices where the models get trained locally and only the updates are shared with the server [1]. This significantly reduces the privacy risks and allows the large scale training of decentralised datasets [5].

Despite the advantages, FL faces a few challenges. Data across various devices is often heterogeneous and non-IID, communication between the server and the device is limited and the client devices are not reliable. These factors make the training all more complex. This paper is a compre-hensive review of the foundational structure of these de-centralised systems, their core optimisation algorithms, the applications & technical challenges and the future research directions.

2 Federated Learning Architecture

Federated learning operates on the client-server architec-ture, meaning that a central server coordinates the training across multiple distributed devices. Each client device stores its data locally and the training process is done indepen-dently. The training process follows the below iterative process:

- The server initializes a global model.
- A subset of clients is selected.
- The model is sent to selected clients.
- Clients perform local training.
- Updates are sent back and aggregated.

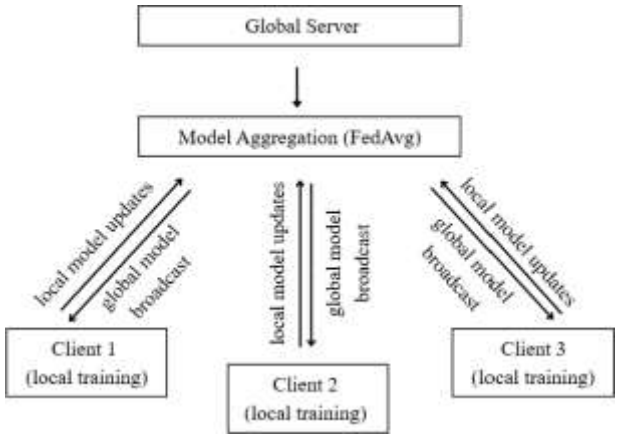

Fig. 1. Federated Learning workflow illustrating global model broadcast-ing, local client-side training, and server-side aggregation using FedAvg.

These iterations are continued until convergence is achieved. Due to this architecture the raw data never leaves

the client devices ensuring that there is no breach of privacy. Instead only the model updates are transmitted to the server which reduces the risk of data exposure.

The deployment of FL is generally bifurcated into two categories based on the scale and the reliability of participat-ing nodes. Cross-device FL involves millions of unreliable mobile or IoT devices that have constrained bandwidth and intermittent availability [6]. Whereas, Cross-silo FL consists of a small number of reliable and stable organisations that are spread geographically such as hospitals and financial institutions [1], [2].

Depending on the statistical distribution of the data, these architectures are further divided into three categories:

- **Horizontal Federated Learning (HFL):** This method is employed when the clients share identical feature spaces but have entirely different sample spaces [2]. An example of this will be predictive texts on the mobile keyboard [5].
- **Vertical Federated Learning (VFL):** This technique is utilised when the participants share the same overlapping user base but possess different features, such as a bank and a retail company collaborating on mutual clients [2].
- **Transfer Federated Learning (TFL):** TFL gets de-ployed when the participants have minimal overlap in both sample spaces and feature spaces [3].

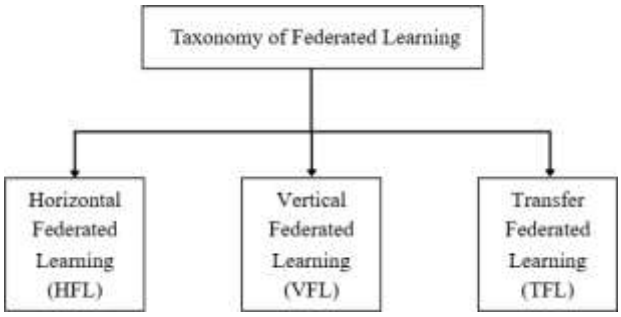

Fig. 2. Taxonomy of Federated Learning based on data partitioning and feature distribution across participating clients.

3 Core Algorithms

The core algorithm driving FL is Federated Averaging (Fe-dAvg). In a FedAvg round, the central server broadcasts the current global model weights to the available clients [5]. Each client k then executes E epochs of Stochastic Gradient Descent (SGD) on their local data set P_k of size n_k. Each client minimises its own local objective function using SGD, resulting in the local update:

$$w_{t+1}^{k} \leftarrow w^{t} - \eta g_k \tag{1}$$

Here η is the learning rate and g_k refers to the local gradient computed on the client k. After local training the updated model weights are transmitted back to the global server. The global server then computes the weighted sum of all the local models to obtain the updated global model:

$$w_{t+1} = \Sigma^{k}_{k=1} \frac{n_k}{n} w^{k}_{t+1} \tag{2}$$

where $n = \Sigma^{k}_{k=1} n_k$ represents the total number of train-ing samples across all participating clients [5]. The overall global objective function in federated learning is defined as:

$$f(w) = \Sigma^{k}_{k=1} \frac{n_k}{n} F_k(w) \tag{3}$$

where $F_k(w)$ denotes the local objective function of client k [5].

While FedAvg is extremely effective in reducing the communication frequency as compared to traditional meth-ods, it struggles mathematically when processing highly heterogenous, non-identically distributed (non-IID) data [4]. Theoretical analyses suggest that for highly non-IID set-tings, decaying learning rates may increase the chances of convergence [4]. Fixed learning rates can lead to instability or slower convergence in heterogeneous data distributions.

To mitigate these limitations, the researchers have come up with various variations of the aggregation algorithm. *FedProx* is a variation of FedAvg where proximal regular-ization term is introduced so that the local updates do not deviate far from the global model, ensuring convergence in highly heterogeneous data [2]. *SCAFFOLD* (Stochastic Controlled Averaging) employs control variates to correct the 'client drift' caused by non-IID data [4]. Adaptive op-timization strategies like *FedAdam*, *FedAdagrad*, and *FedYogi* apply server-level momentum and adaptive learning rates to match the communication efficiency of FedAvg whilst mimicking centralised adaptive algorithms [1].

4 Key Challenges

Despite its evident advantages, implementation of Feder-ated Learning at scale poses various complex problems regarding privacy, networking and statistical domain.

4.1 Privacy and Security

Although Federated Learning minimises raw data exposure, the local gradients transmitted can still leak sensitive infor-mation. Attackers can execute model inversion attacks to reconstruct the training dataset or membership inference attacks to figure out whether a specific data point was used during training [7]. Defense mechanisms mainly use Differential Privacy (DP) to protect the data [1]. DP works by clipping the update norms and adding Gaussian or Laplacian noise.

To keep the updates safe during transmission and ag-gregation, FL systems use Secure Multi-Party Computation (SMC) and Secure Aggregation protocols [3], [6]. These cryptographic methods are designed so that the server can only access aggregated updates while limiting visibility into individual client contributions. Federated Learning systems are also at risk of poisoning attacks [2]. Malicious clients can add corrupted data (data poisoning) or manipulate their model updates (model update poisoning) to add hidden flaws or make the model less accurate.

4.2 Statistical Heterogeneity

Data generated by edge devices is intrinsically non-IID and unbalanced. A single user's dataset is not a repre-sentative of the global average. This heterogeneity causes

local model updates to drift toward client-specific optima, causing model drifts that damage the global aggregation process [4]. Theoretical analysis suggests that for the Fe-dAvg algorithm to converge on non-IID data at a rate of $O\ \frac{1}{T}$, the learning rate η must dynamically decay [4]. If we use a fixed learning rate, performance degradation may occur and it may not be able to find the optimal solution.

4.3 Communication Overhead

Communication bandwidth is the primary problem faced in cross-device Federated Learning. Mobile connections are often slow, costly and asymmetrical [6]. Transmitting model parameters for deep neural networks across thousands of rounds is too expensive. Researchers address it by im-plementing model compression techniques like gradient quantization and sparsification, which drastically reduce the payload size before transmission [5].

4.4 System-Induced Bias and Fairness

The operational constraints of FL can introduce bias. De-vices are only selected when they are idle, plugged in and connected to unmetered Wi-Fi so as to not disturb other users [6]. Due to this, devices from some specific socio-economic demographics, geographic region or time zones might be under-represented in the training pool. This results in a final model that performs unfairly across various user subsets.

5 Applications

FL is immensely useful in sectors where data privacy and confidentiality are big concerns. In the healthcare sector, the medical institutions employ cross-silo FL to train prediction models on electronic health records (EHR) and biomedical imaging without violating any patient confidentiality laws [7].

In mobile technologies, FL is utilised to improve user interface without centralising personal data. Keyboards use FL to train recurrent neural networks (RNNs) for next word prediction and out of vocabulary word learning [6]. This allows the model to learn details of a user's in-chat typing behavior. This training is done on the device so the personal data of the user stays private.

Financial institutes utilise FL for risk assessment and fraud detection. Banks can collaboratively train models to detect anomalies to identify any fraudulent transaction made whilst still keeping the client financial records private [2]. FL is also being employed in the advancement of the Internet of Things (IoT) and smart city infrastructure. This is achieved by training prediction models on decentralised sensor data to help optimise traffic management and envi-ronmental monitoring [3].

6 Open Challenges & Future Directions

As FL continues to mature, the research must focus on re-solving the existing structural limitations. The dependence on a central coordinating aggregator is the major point of failure and potential trust breaker. Integrating FL with blockchain technology or distributed ledgers, can lead to FL being completely decentralised, auditable and immutable record keeping for model updates [2].

FL also struggles with the 'straggler effect', where the model's updates are aggregated after all the selected devices are available thus resulting in the server waiting for the slowest devices. This issue can be tackled by developing an asynchronous FL optimisation algorithm for scaling FL to billions of highly volatile edge devices.

The shift towards Personalized Federated Learning (PFL) is necessary. Instead of forcing a single generalised global model onto the clients, PFL utilises meta-learning and multi-task learning techniques so that each client fine-tunes the global model to deeply fit its data distribution, turning statistical heterogeneity into an advantage [1].

7 Conclusion

Federated Learning is changing the way we approach artifi-cial intelligence. It separates machine learning optimization from centralized data storage. By training the models locally and securely aggregating the updates, FL implemented the principle of data minimisation.

While algorithms like FedAvg has moved from theory to production, to utilise FL's full potential the innovation must continue to balance communication efficiency, cryp-tographic privacy, and algorithmic fairness across highly heterogeneous datasets [4], [5]. The future of FL lies in personalized, fully decentralized architectures that adapt dynamically to the constraints of the edge [1].

References

[1] P. Kairouz et al., "Advances and Open Problems in Federated Learning," *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.

[2] B. Yurdem, M. Kuzlu, M. K. Gullu, F. O. Catak, and M. Tabassum, "Federated learning: Overview, strategies, applications, tools and future directions," *Heliyon*, vol. 10, e38137, 2024.

[3] J. Wen, Z. Zhang, Y. Lan, Z. Cui, J. Cai, and W. Zhang, "A survey on federated learning: challenges and applications," *International Journal of Machine Learning and Cybernetics*, vol. 14, pp. 513–535, 2023.

[4] X. Li, W. Yang, Z. Zhang, K. Huang, and S. Wang, "On the Convergence of FedAvg on Non-IID Data," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.

[5] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agu¨ era y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.

[6] K. Bonawitz et al., "Towards Federated Learning at Scale: System Design," in *Proceedings of the 2nd SysML Conference*, 2019.

[7] P. M. Mammen, "Federated Learning: Opportunities and Chal-lenges," *arXiv preprint arXiv:2101.05428*, 2021.