Dean&Francis

ISSN 2959-6157

Federated Learning Overview: Frameworks, Challenges, and Future Directions

Yuqing Shao

National University of Singapore, Singapore
E1561365@u.nus.edu

Abstract:

Federated Learning (FL) is a highly regarded distributed machine learning framework that aims to enable collaborative modeling among multiple parties without disclosing raw data. Compared to centralized learning, FL allows each participant to train models independently on-site and upload model update parameters rather than data to a central server, thereby improving model performance while effectively protecting data privacy. As artificial intelligence is increasingly applied in fields such as healthcare, finance, and industrial IoT, data privacy and compliance requirements are becoming increasingly stringent, highlighting the significant application potential and research value of federated learning. This paper systematically reviews the basic theories, core algorithms, and technical approaches of federated learning, focusing on research progress in areas such as communication efficiency, data heterogeneity, privacy protection, and trust mechanisms. The study also explored the development prospects and future directions of federated learning in the medical, financial, and precision computing fields. At the same time, we conducted a more in-depth analysis of the main challenges currently faced and analyzed potential future research directions in order to provide a reference for the continued development of this field.

Keywords: Federated Learning, Data Privacy, Communication Efficiency, Edge Computing.

1. Introduction

Although datasets keep getting bigger and used for multiple purposes, data privacy becomes a major concern, and the rigorous investigation in privacy-preserving machine learning frameworks is a pressing issue. This imperative is definitely on the last stage in sectors like healthcare, finance, and smart cities, as the amount of user data in these fields creates both possibilities and risks [1].

Today, instead of collecting data on a central server, which can result in data leaks, the federated learn-

ing methods process data at the local ends for the central server, which can in no way be a conduit for data violations and trust issues [2]. As research indicates, the grand objective has been to find a data security framework that preserves the delicacy of data without interrupting the efficiency of the system.

Federated learning (FL) became a significant possibility for these issues and allows now several clients to train an adequate model with only local data. Through distinct devices or institutions, these models are trained on rich data varieties, consequently broadening the generalizability of them, as they utilize intermediate results from other devices or institutions [3]. It assures the compliance with data governance laws, such as GDPR and HIPAA, this is one of the benefits, besides depth and richness of the data collected as various datasets get pooled together from different countries and regions.

McMahan, along with his co-authors, were the first ones to coin the notion of federated learning (FL), or neighborhood, in 2016. The critical assumption in its use is that participants can jointly learn, in the absence of raw data exchange, by constructing model parameters [4]. In this respect, it helps overcome the problems of data silos by removing the barriers of accessibility, and the data analytics and privacy policies become the secondary issues because it is not the actual data that is used. One of the hidden parts of the iceberg, and what breaches have really been taken care of through decentralization along with data silos and regulatory compliance, are data privacy issues.

2. Basic Principles

Federated learning is a branch of distributed machine learning that permits several clients (like IoT devices or research institutions) to collectively train the learning models while keeping the data on their local node [5]. Each client also has a model that it independently trains using local data and, consequently, sends model parameters or gradients to a central server, which also aggregates them to form a new world model. This process of iterations makes room for constant improvement of model performance. Weighting is a method commonly used to apply aggregation, where the weight assigned to each client's contribution is determined by factors like the volume of data they offer.

2.1 Horizontal Federated Learning

Horizontal federated learning is suitable for scenarios where clients share the same feature space but have different sample sets. For example, banks in different regions may record users' transaction behaviors with consistent field structures but distinct user populations. This type of federated learning is widely applied in fields such as finance and healthcare.

2.2 Vertical Federated Learning

Vertical federated learning is suitable for scenarios where client samples are consistent, but feature spaces are different [6]. A typical example is when two companies have the same users but different dimensions of information (e.g., one has purchase records, and the other has browsing records). Joint modeling is achieved by matching sample IDs.

2.3 Transfer Federated Learning

Federated learning is applicable in scenarios where clients have independent sample spaces and feature spaces [7]. In this context, knowledge transfer mechanisms are typically established through methods such as transfer learning and representation alignment to achieve joint modeling of homogeneous data. These approaches aim to bridge distributional gaps among clients, enabling more effective global model aggregation [8].

3. Key Challenges in Federated Learning

3.1 Communication Efficiency

Federated learning's communication efficiency is a crucial factor that determines the performance of the whole system. In most cases, the training process includes the constant model updates between clients and the central server from which the training takes place. Hence, the related communication cost can become the bottleneck very quickly, especially when it is a bandwidth-limited environment or when there are fully diverse devices. This problem has led to a variety of studies focusing on the approaches that have smaller communication frequency or size of data to maintain the integrity of the model. Techniques such as model quantization, sparsification, and local update aggregation have been widely explored to alleviate communication overhead [9].

The most common delivery model in federated learning is FedAvg, which performs a greater number of training steps on the client side before sending an update to the server. This way, more communication is less frequent. On the other hand, this technique is easy to use and rather innovative. However, it can be damaging to client data segregation issues. To do this, FedDyn inserts regularization terms that keep local training stable and equal while providing an opportunity to aggregate.

ISSN 2959-6157

Moreover, compression techniques have been found to be effective as well as algorithmic modifications. For instance, using STC (Sparse Ternary Compression), the size of the gradient updates after sparsification and quantisation can be achieved, thus resulting in the minimisation of the data streamed. Considering the same, FedPAQ can leverage periodic aggregation and quantized updates, which can help minimize communication requirements, leading to a higher level of accuracy and efficiency [10].

These methods, when taken together, reveal the ongoing effort to bring federated learning to practical deployment conditions where domain limitations and hardware variability are unavoidable. Nevertheless, this area requires an extensive research effort to make federated learning both scalable and feasible.

3.2 Data heterogeneity

In actual scenarios, clients within a federated system often have distinct data distributions (Non-IID), which can lead to increased local model bias and thus reduce the performance of the global model [6]. To alleviate this problem, FedProx introduces regularization terms into the objective function to reduce local model deviation. Personalized federated learning (e.g., PeFLL) allows clients to retain locally customized components on top of a shared model, enhancing adaptability; multi-task federated learning methods design differentiated model parameters for different clients based on task correlations, improving convergence.

3.3 Privacy and Security

Although FL structurally avoids the transmission of raw data, the model parameters or gradient information transmitted during the process may still be exploited by attackers. Attack types include gradient inversion, member inference, and backdoor attacks [3]. To enhance privacy protection, differential privacy mechanisms inject noise to increase uncertainty. Secure multi-party computation (MPC) enables multiple participants to perform joint computations without revealing inputs; and homomorphic encryption technology enables direct computation on encrypted data, effectively mitigating information leakage risks.

4. Typical Applications of Federated Learning

Federated learning has been applied in several industries, and the typical applications can be broadly categorized into the following three categories according to the distinctions between the data subjects and system architectures:

4.1 Medical and Health

In smart healthcare, hospitals, clinics, and other organizations are difficult to directly share private patient information data due to privacy compliance needs. Federated learning enables medical institutions to collaboratively train models for disease prediction, image diagnosis, etc. by building a joint modeling framework without sharing the original information data. For example, distinct hospitals can jointly train lung CT diagnostic models to achieve a win-win situation of model sharing and patient privacy protection.

4.2 Communication Efficiency

Financial institutions such as banks, insurance, and third-party credit bureaus usually hold distinct dimensions of users' information, and FL enables these organizations to integrate the information without disclosing sensitive data, which can be used to accomplish tasks such as credit assessment and fraud detection [8]. For example, banks and payment platforms can share modeling capabilities to predict users' credit risks, thus improving the accuracy of risk control.

4.3 Communication Efficiency

In the Internet of Things (IoT) and edge computing environments, terminal devices such as cell phones and computers have a large amount of local data. Federated learning can support devices to train models locally, e.g., for input method prediction, speech recognition, or smart home control, while guaranteeing that user data is not uploaded to the cloud, thus improving privacy and system response speed [5].

1. Research Trends and Future Outlook

Furthermore, federated learning has also gradually expanded to the fields of natural language processing, electric power system, intelligent manufacturing, etc., and has shown a widely application prospect. Its ability to facilitate collaborative model training without exposing raw data makes it especially suitable for domains with strict privacy or security requirements.

2. Conclusion

As a product of the integration of collaborative training and data privacy, federated learning provides a new paradigm for multi-data source modeling. By supporting local

modeling and cross-device collaboration, it effectively alleviates data silos and privacy compliance challenges, demonstrating significant theoretical value and application potential. However, current research still has many shortcomings, such as excessive communication and computing resource consumption, insufficient adaptability to heterogeneous data, and potential risks in terms of security and robustness. At the same time, some mainstream methods make many idealized assumptions in their experimental settings and lack an in-depth analysis of the challenges of deploying federated systems in real-world complex environments.

Future research should transcend the single goal of algorithmic precision and focus on a comprehensive balance of system performance, including personalized demand fulfillment, model convergence speed, energy efficiency management, and maintainability design. Additionally, more systematic evaluation frameworks and open-source benchmark tools will become critical enablers for collaborative innovation within the FL community. We believe that only by balancing practicality and scalability can federated learning truly become a key driver for privacy-preserving AI applications.

References

[1] McMahan H B, Moore E, Ramage D, et al. Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS). 2017.

[2] Yang Q, Liu Y, Chen T, et al. Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 2019, 10(2): 1–19.

[3] Li T, Sahu A K, Zaheer M, et al. Federated optimization in heterogeneous networks. Advances in Neural Information Processing Systems (NeurIPS). 2020.

[4] Acar D A, Zhao Y, Mattina M, et al. Federated learning based on dynamic regularization. International Conference on Learning Representations (ICLR). 2021.

[5] Sattler F, Wiedemann S, Müller K R, et al. Robust and communication-efficient federated learning from non-IID data. IEEE Transactions on Neural Networks and Learning Systems, 2020, 31(9): 3400–3413.

[6] Reisizadeh A, Mokhtari A, Hassani H, et al. FedPAQ: A communication-efficient federated learning method with periodic averaging and quantization. Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS). 2020.

[7] Liu J, Kang J, Xu Y, et al. Recent advances on federated learning: A systematic survey. arXiv:2301.01299.

[8] Lu X, Li Y, Liu X, et al. Privacy-preserving asynchronous federated learning for edge network computing. IEEE Access, 2020, 8: 48970–48981.

[9] Scott S, Yu L, Zhao S, et al. PeFLL: Personalized federated learning by learning to learn. International Conference on Learning Representations (ICLR). 2024.

[10] Carlini N, Tramer F, Wallace E, et al. The secret sharer: Measuring unintended neural network memorization & extracting secrets. arXiv:1802.08232.