ISSN: 2959-6386 (Online), Vol. 2, Issue 2

Journal of Knowledge Learning and Science Technology

journal homepage: https://jklst.org/index.php/home

Federated Learning Architecture: Design, Implementation, and Challenges in Distributed AI Systems

Lavanya Shanmugam[1], Ravish Tillu[2], Manish Tomar[3]

[1]Tata Consultancy Services, USA
[2]RBC Capital Markets, USA
[3]Citibank, USA

Abstract

Federated learning has emerged as a promising paradigm in the domain of distributed artificial intelligence (AI) systems, enabling collaborative model training across decentralized devices while preserving data privacy. This paper presents a comprehensive exploration of federated learning architecture, encompassing its design principles, implementation strategies, and the key challenges encountered in distributed AI systems. We delve into the underlying mechanisms of federated learning, discussing its advantages in heterogeneous environments and its potential applications across various domains. Furthermore, we analyse the technical intricacies involved in deploying federated learning systems, including communication efficiency, model aggregation techniques, and security considerations. By synthesizing insights from recent research and practical implementations, this paper offers valuable guidance for researchers and practitioners seeking to leverage federated learning in the development of scalable and privacy-preserving AI solutions.

Keywords: Federated learning, Distributed AI systems, Architecture design, Implementation strategies, Privacy preservation, Communication efficiency, Model aggregation, Security challenges.

Article Information:
Article history: *Received:* 01/09/2023 *Accepted:* 10/09/2023 *Online:* 16/09/2023 *Published*: 16/09/2023
DOI: https://doi.org/10.60087/jklst.vol2.n2.p384

Correspondence Author: Lavanya Shanmugam

Introduction

The proliferation of data-generating devices and the increasing demand for privacy-preserving AI solutions have propelled federated learning to the forefront of distributed artificial intelligence (AI) research. Traditional centralized machine learning approaches often face significant challenges when dealing with sensitive data, as data aggregation raises concerns regarding privacy breaches and regulatory compliance. Federated learning offers a compelling alternative by enabling model training directly on decentralized devices while keeping raw data localized, thus addressing privacy concerns without sacrificing model performance.

This paper provides a comprehensive examination of federated learning architecture, encompassing its design, implementation, and challenges within distributed AI systems. We begin by elucidating the fundamental principles of federated learning, highlighting its advantages in heterogeneous environments where data sources are diverse and geographically distributed. By leveraging federated learning, organizations can harness insights from data silos without compromising data privacy, making it particularly appealing in sectors such as healthcare, finance, and IoT.

Furthermore, we delve into the intricacies of implementing federated learning systems, discussing various strategies for orchestrating model training across a network of edge devices, mobile phones, or IoT sensors. This includes exploring techniques for efficient communication, robust model aggregation, and adaptive learning algorithms tailored to the decentralized nature of federated environments.

Despite its promise, federated learning also presents several challenges, ranging from communication overheads and heterogeneity in device capabilities to security vulnerabilities and algorithmic biases. Throughout this paper, we aim to address these challenges and provide insights into mitigating their impact on the effectiveness and reliability of federated learning systems.

By synthesizing insights from recent research advancements and practical implementations, this paper serves as a valuable resource for researchers, developers, and practitioners interested in harnessing the power of federated learning to build scalable, privacy-preserving AI solutions in distributed environments.

objectives

1. To Explore Federated Learning Architecture: This paper aims to delve into the underlying principles and architectural components of federated learning, elucidating how decentralized model training can be orchestrated across a network of devices while preserving data privacy.

2.To Investigate Implementation Strategies: We seek to analyse various implementation strategies and techniques employed in federated learning systems, including communication optimization, model aggregation methods, and adaptive learning algorithms tailored to distributed environments.

3. To Address Challenges in Distributed AI Systems: This paper endeavours to identify and address key challenges encountered in federated learning and distributed AI systems, such as communication overheads, heterogeneity in device capabilities, security vulnerabilities, and algorithmic biases, offering insights into effective mitigation strategies.

Literature Review

Federated Learning (FL) architectures in distributed AI systems aim to preserve privacy by aggregating locally trained models without sharing raw data externally. Challenges include accountability, fairness, communication costs, and non-IID data distribution. To address these, researchers propose innovative solutions. For instance, a blockchain-based architecture enhances accountability and fairness [1] [2]. Additionally, in-cluster training and gradient scarification techniques improve communication efficiency and performance, especially with Non-IID data, as demonstrated by the FedOES approach [3]. Understanding Quantum Federated Learning (QFL) is also crucial, with ongoing research focusing on new frameworks, applications, and critical design factors [4] [5]. These advancements contribute to shaping more robust and efficient federated learning systems in distributed AI environments.

Background

In recent years, the field of machine learning (ML) and deep learning (DL) has experienced remarkable growth, largely fueled by the abundance of available data. However, many application domains lack centralized repositories of adequately labeled and complete data, such as medical image analysis for doctors' diagnoses. Creating such datasets is often time-consuming, labor-intensive, and reliant on domain expertise, leading to the formation of data silos within individual organizations. While some organizations manage to accumulate high-quality datasets, they are often small in scale, hindering the effectiveness of DL applications that require extensive, fully labeled data.

Traditionally, ML models were built using data gathered in centralized locations. However, concerns regarding data

ownership, confidentiality, user privacy, and evolving data management regulations like the General Data Protection Regulation (GDPR) necessitate private, secure, efficient, and fair distributed model training methodologies. Federated learning (FL) addresses these concerns by allowing separate model training on local data distributed across various devices or organizations. Local model updates are then aggregated to form a global model, ensuring that individual data remains private.

Initially introduced by researchers at Google for updating language models in keyboard systems, FL involves building a joint model using data from different sites without sharing raw data. The joint model is encrypted and shared among participants to maintain privacy, resulting in a performance approximation of an ideal model trained with centralized data. Despite potential accuracy losses due to added security measures, FL offers significant privacy and security benefits, often outweighing the drawbacks, especially in sensitive application domains.

FL architectures typically adhere to either the client-server or peer-to-peer model. In the client-server model, a coordinator aggregates model parameter using federated averaging (FedAvg). Each participating client receives an initial model, trains locally, and sends updates to the coordinator for aggregation. This process repeats until convergence. While the client-server architecture minimizes communication overhead, the peer-to-peer model enhances security by facilitating direct client communication. However, the peer-to-peer model requires more computational resources for encryption and decryption.

In summary, FL presents a promising approach to address data privacy concerns while enabling collaborative model training in distributed environments. Its adoption holds significant potential for improving model quality and security across various application domains.

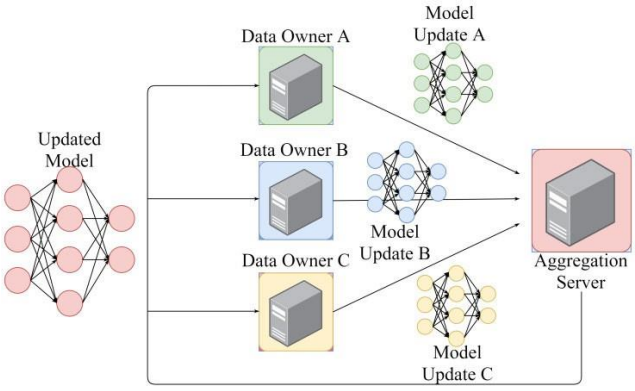

FIGURE 1. Client-server FL architecture.

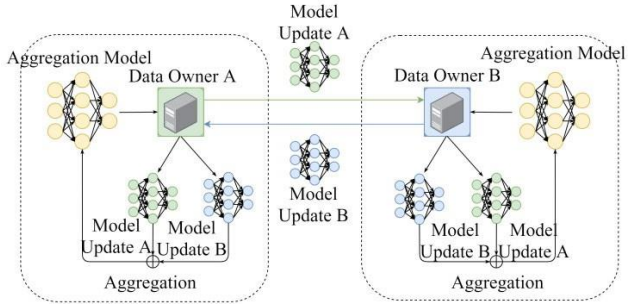

FIGURE 1. peer to peer FL architecture.

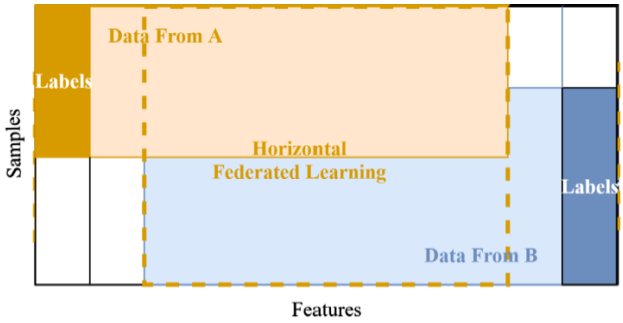

FIGURE 3. Horizontal FL architecture.

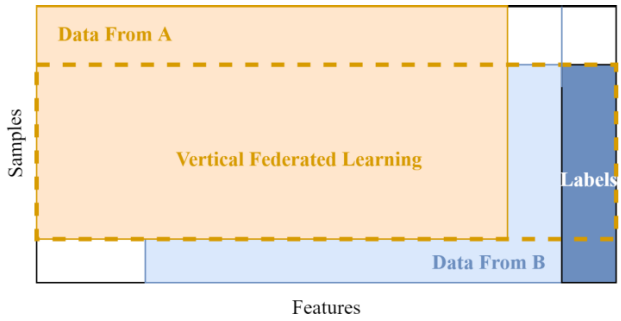

FIGURE 4. Vertical FL architecture

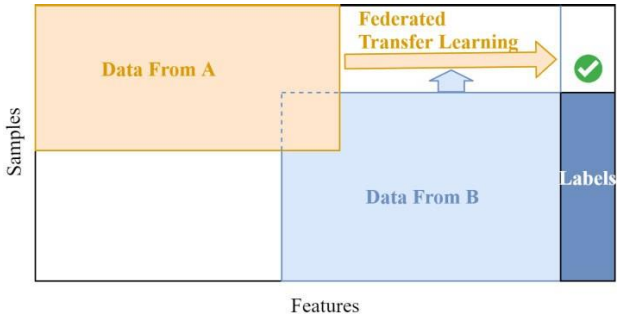

FL encompasses various approaches including horizontal FL (HFL) (Fig. 3), vertical FL (VFL) (Fig. 4), and federated transfer learning (FTL) (Fig. 5). In HFL, alignment occurs in data features across participants, while in VFL, alignment is in data samples rather than features. Both HFL and VFL may prove ineffective in highly heterogeneous data scenarios. In such cases, FTL offers an effective solution by transferring learned knowledge from a source domain to a target domain, inspired by transfer learning principles.

However, the discussed architectures and FL categories merely scratch the surface of FL research. Numerous research thrusts explore novel architectures, data partitioning schemes, and aggregation techniques. Efforts are also underway to address core challenges in FL such as privacy, security, communication costs, system and statistical heterogeneity, and personalization techniques. Each application area of FL presents unique challenges and considerations.

While considerable research has already been conducted in FL, various survey papers have summarized different focus areas. This study aims to review existing surveys covering various domains and focus areas in FL research. We delve into core challenges such as privacy, security, communication costs, system and statistical heterogeneity, architecture, and aggregation algorithm designs, each of which varies by domain and specific use case.

The motivation behind this paper lies in reviewing the current literature and summarizing state-of-the-art approaches developed to address these challenges. We also aim to identify gaps in existing FL surveys and fill them by surveying the latest developments in all aforementioned FL research areas. Our review offers a holistic examination of challenges, applications, and design factors, outlining promising future research directions.

We thoroughly investigate contemporary FL survey papers, classify FL research into broad categories of design aspects, challenges, and application areas, and discuss open issues and challenges in FL research. The remainder of this paper is organized as follows: Section II discusses related studies, Section III illustrates the taxonomy of survey papers, Section IV provides a discussion and analysis of topics under each category, Section V discusses open issues and challenges, and finally, Section VI concludes the paper.

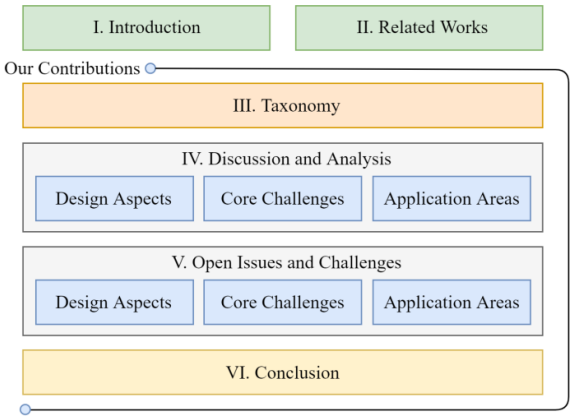

Related Work

In this section, we undertake an investigation and analysis of contemporary survey papers in the field. The reviewed papers, along with their summaries and focuses, are detailed in Table 1.

Li, Sahu et al. [3] examined the distinctions between federated learning (FL) and standard distributed machine learning (ML). They explored FL's unique characteristics, challenges, current methodologies, and future prospects. Although the paper did not concentrate on any specific domain, it addressed approaches tackling four core challenges: expensive communication, systems heterogeneity, statistical heterogeneity, and privacy/security. Local updating [1], [4] emerged as an approach to reduce the number of communication rounds. Conversely, compression schemes [5] aimed to diminish message sizes during communication rounds. Moreover, decentralized training [6], [7] alleviated the central server's communication burden. To tackle systems heterogeneity, asynchronous communication [8]–[10] was employed to minimize stragglers, while active sampling selected or influenced participating devices based on their system resources and overheads, with fault tolerance [11]–[16] disregarding failed devices through algorithmic redundancy. Addressing statistical heterogeneity, methods such as meta-learning and multitask learning were utilized to model heterogeneous data, adapting the selection between global and device-specific models, and leveraging transfer learning for personalization. Additionally, convergence guarantees for non-independent and identically distributed (non-IID) data [4], [10], [17], [18] were investigated. Finally, this survey encompassed secure multiparty computation (SMC) [19], [20], and differential privacy (DP) [21]–[24] approaches.

Authors in [25] directed their focus towards mobile edge networks, identifying core challenges including expensive communication, systems heterogeneity, and privacy/security. To address communication cost challenges, they

explored various approaches such as model compression [26], [27], importance-based updating for selective gradients [28], or local model updates [29], with a specific focus on edge and end computation. Methods to mitigate systems heterogeneity encompassed active sampling based on computation capabilities [33], data characteristics [34], and resource consumption [35] and allocation [36], [37].

TABLE 1. Summary table of survey papers and main focus.

Survey Paper	Summary	Main Focus
Li, Sahu, *et al.* [3]	Discusses the unique characteristics and challenges of FL, provides details of current approaches, outlines directions of future work.	Challenges
Lim *et al.* [25]	Highlights challenges of FL implementation and existing solutions and presents applications of FL for mobile edge network optimization.	Mobile edge networks
Briggs *et al.* [57]	Focusing on IoT, covers works related to FL challenges and privacy preserving methods, identify the strengths and weaknesses of different methods applied to FL, and outlines future directions.	IoT, privacy/security
Li, Wen, *et al.* [58]	Categorizes FL systems according to six different aspects to facilitate and guide the design of FL systems, provides case studies and future research opportunities.	FL systems
Li, Fan, *et al.* [59]	Illustrates the evolution of FL and reviews existing applications of FL in industrial engineering, mobile devices and healthcare.	Applications
Kurupathi, Maass [60]	Highlights existing privacy techniques and proposes applications of FL in industries.	Privacy/security, applications
Yang *et al.* [61]	Introduces a secure FL framework, which includes horizontal FL, vertical FL and federated transfer learning, and proposes building data networks among organizations based on federated mechanisms.	Architecture, applications
Xu *et al.* [62]	Provides a review for FL technologies mainly for biomedicine, and discusses the challenges, issues and potential of FL in healthcare.	Healthcare
Kulkarni *et al.* [63]	Highlights the need for personalization in FL and surveys research on the topic.	Personalization
Lyu *et al.* [64]	Introduces taxonomy of threat models and major attacks on FL, highlighting intuitions, techniques and assumptions adopted by different attacks and discusses future research directions.	Threat models and attack types
Aledhari *et al.* [65]	Provides a thorough summary of relevant protocols, platforms, challenges and real-life uses cases of FL.	Platforms, protocols, applications
Mothukuri *et al.* [66]	Provides a detailed study of security and privacy, and presents current approaches, challenges and future directions in FL.	Privacy/security

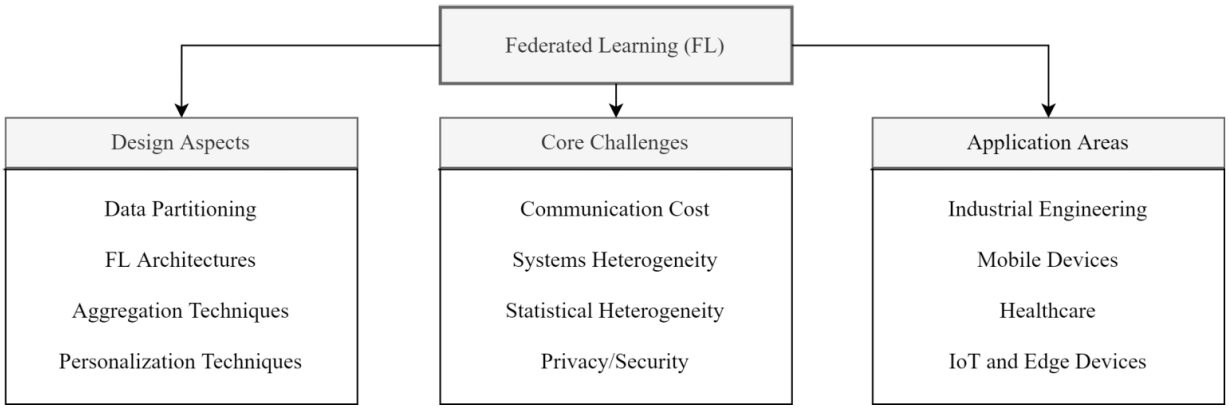

FIGURE 7. Classification of the reviewed survey papers.

Discussion and Analysis

In this section, we undertake a review and discussion of the design aspects, core challenges, and application areas to offer a comprehensive summary of the following subtopics: data partitioning, FL architectures, aggregation techniques, personalization techniques, communication cost, systems heterogeneity, statistical heterogeneity, privacy/security, and application areas.

Design aspects

Data partitioning

Data utilized for training in FL exhibits non-identical characteristics due to its distribution across various devices. The dataset's sample space comprises all instances, while the feature space encompasses different attributes. For instance, two hospitals may possess records of different patient sets (sample space) and diverse information about each patient in their Electronic Health Records (EHR) (feature space). FL systems (FLSs) categorically fall into horizontal FL (HFL), vertical FL (VFL), and hybrid FL based on how data are allocated across multiple devices in terms of sample and feature spaces.

Horizontal FL (HFL): This scenario occurs when datasets share the same feature space but differ in sample space. HFL is commonly employed in cross-device scenarios where individual users aim to enhance their model's performance on a task. It's characterized by horizontal partitioning, where local data overlap in the feature space, allowing each user to train their local models using a duplicate model architecture. For example, regional branches of an organization may possess different user groups but share the same feature spaces. Notably, research by McMahan et al. [1] falls within the horizontal partitioning paradigm, where individual users on the Android platform update model parameters locally, contributing to centralized model training. Additionally, to address finite labeled entities, hierarchical heterogeneous HFL frameworks have been proposed, allowing adaptation of each user multiple times as the target domain.

Vertical FL (VFL): This scenario arises when datasets across institutions share similar sample spaces but have dissimilar feature spaces. In VFL, participants possess homogeneous data but differ in feature space, necessitating privacy-preserving methods for model training. VFL aggregates distinct features while preserving privacy, ultimately constructing a model by combining data from multiple parties. Approaches such as linear regression have been proposed for vertically partitioned data, along with secure models like k-means, association rule mining, decision trees, and naive Bayesian classifiers. VFL systems typically perform entity alignment to combine common samples from different institutions, followed by encryption for collaborative model training.

In the subsequent sections, we continue to explore various aspects of federated learning, including architectures, challenges, and application domains.

Federated Transfer Learning (FTL)

FTL finds utility in situations where two datasets vary not only in sample but also in feature space. Initially proposed in [79], FTL enhances existing FL systems and extends beyond the scope of conventional FL algorithms. Its application has garnered significant attention across various industries, notably in healthcare [121]. FTL facilitates the sharing of diverse treatment and diagnosis information among hospitals to address a range of diseases. Transfer learning, a fundamental concept underlying FTL, aims to establish a common representation of features from disparate parties, requiring both parties to compute prediction results independently. Consequently, transfer learning techniques can be applied to the entire feature and sample space within a federated environment. To uphold privacy, FTL leverages encryption and approximation techniques, ensuring that sensitive data and models remain locally preserved [122]. Efforts to enhance FTL include Sharma et al.'s integration of secret sharing technology [123]. Furthermore, works by [124], [125] introduce the FedHealth model, which leverages FL to collect data from various institutions and deliver tailored healthcare services through transfer learning.

Each data partitioning paradigm presents its own set of advantages and limitations. For instance, two distinct clinics or hospitals can securely exchange data based on their respective needs in terms of instances or features required. While one clinic may possess millions of patient records with specialized information (e.g., oncology), another clinic, albeit newer, may have fewer records but diverse patient information. In such cases, the first clinic could benefit from VFL, while the second could leverage HFL. Ultimately, FTL enables healthcare providers to offer more personalized care by leveraging data from users' wearable devices for personal fitness.

Federated Learning (FL) architectures

Federated Learning (FL) architectures depict the integration of various components to establish an FL environment. Two prevalent architectures in FL are the client-server and peer-to-peer architectures.

Client-Server Architecture:
In the client-server architecture, depicted earlier in Fig. 1, a central server initiates a global model shared with clients for training on their local datasets. Post local training, trained models from involved clients are gathered by the server. The server then aggregates the models' parameters to construct a global model, disseminating it to all clients. This architecture, also known as centralized FL architecture, sees the server coordinating the continuous learning process. Unlike conventional client-server setups where the server hosts and trains a model on shared data, in FL, the server operates solely on locally received models from clients, synchronously or asynchronously. The primary advantage of this architecture lies in its reduced communication overhead. Google's development of Gboard for Android, a virtual keyboard, utilized this architecture. Presently, nearly all FL implementations adopt the client-server architecture.

Peer-to-Peer Architecture:
In the peer-to-peer architecture, illustrated in Fig. 2, the concept of a central server for model aggregations, as seen in the client-server architecture, is absent. Instead, algorithms ensure security and reliability in the absence of a central server. Each participant in the FL environment possesses its own model, refining it by leveraging information obtained from neighboring participants. In this adopted peer-to-peer topology, a protocol, facilitated by a central authority, guides the network during training rounds. This architecture offers enhanced security as participating clients communicate directly without a third-party coordinator. However, it necessitates increased computation for message encryption and decryption.

The aggregation algorithm is pivotal in federated learning (FL)

The aggregation algorithm is pivotal in federated learning (FL), dictating how the global model is amalgamated from local model updates contributed by all participating clients during each training round. Particularly crucial in horizontal FL (HFL) scenarios employing a centralized architecture, various popular aggregation algorithms are compared in Table 3 and summarized below.

FedAvg Algorithm: Proposed by Google, FedAvg operates on a stochastic gradient descent (SGD) optimization algorithm, ideally suited for HFL with a client-server architecture. Here, the server initializes the training process by disseminating global model parameters to a randomly selected group of clients. These clients then execute multiple epochs of SGD on their local datasets, training the global model, and subsequently share their locally trained models with the server. The server then computes the weighted average of all local models to generate a new global model. Despite its widespread success, FedAvg encounters convergence issues in certain settings, attributed to factors such as client drifting and the absence of an adaptive learning rate.

Scaffold: This algorithm tackles client drifting issues by employing a variance reduction technique in its local update. It estimates the update direction of both the server model and each client and measures client drifting using the difference, which is then utilized for local updates. This strategy effectively addresses client heterogeneity and reduces communication rounds during model convergence.

Adaptive Federated Optimization: Proposed by Google's research team, this approach introduces adaptability in server optimization, enhancing server optimization by incorporating adaptive learning rates informed by previous iterations. Clients minimize loss using local data over multiple training epochs, followed by server-based gradient-based optimization on the average of client model updates. While it incorporates adaptive learning rates in server optimization without increasing client storage or communication costs, it doesn't completely eliminate the effects of client heterogeneity. Nonetheless, it's particularly effective for moderate, naturally occurring heterogeneity, especially in cross-device settings.

Fed Boost: FedBoost is a communication-efficient FL algorithm based on ensemble learning techniques. It trains an ensemble of pretrained base predictors through FL, reducing both server-client and client-server communication

costs without relying on gradient or model compression. Besides communication efficiency, FedBoost offers computational speedups, convergence guarantees, privacy, and optimal solutions for density estimation tasks, with language modeling as a special case.

FedProx: Addressing the inherent challenges of FL, FedProx deals with system and statistical heterogeneity by offering a reparametrized and generalized version of FedAvg. It allows partial work tolerance based on resource availability, accepting variable amounts of local updates from resource-constrained devices for aggregation. Additionally, FedProx introduces a proximal term in a device's local solver objective to mitigate the impact of varying local updates.

FedMA: FedMA introduces FL into modern network architectures for deep learning. It performs matching and averaging layer-wise across convolutional layers, hidden states of long short-term memory networks, and fully connected layer neurons based on feature similarity to construct the shared global model at the server. FedMA exhibits robustness in handling client heterogeneity and outperforms FedProx and FedAvg within a few training rounds.

Open Issues and Challenges

Several open issues and challenges persist within the realm of federated learning (FL) [169]. Balancing accuracy, privacy, communication cost, and personalization levels is crucial when designing an FL system, with considerations often varying based on specific use cases or application areas. This section delves into some open issues concerning design aspects, core challenges, and application domains.

Design Aspects

1) Data Partitioning and FL Architectures:
Beyond the primary data partitioning schemes and FL architectures discussed herein, novel variations in FL architectures have emerged. For instance, PerFit a cloud-based platform, offers flexible selection of personalized FL approaches, catering to IoT applications. Another noteworthy architecture, FedHealth utilizes the Federated Transfer Learning (FTL) framework for wearable healthcare, facilitating the construction of personalized models for enhanced healthcare services. Future research endeavors could concentrate on developing FL architecture schemes tailored to meet the specific requirements of diverse industries and application domains.

2) Aggregation Techniques:
Facilitating developers in implementing FL solutions, toolkits featuring standardized and preconfigured aggregation algorithms suitable for distinct application areas and use cases could be invaluable. Analogous to AutoML solutions, such a toolkit for FL has the potential to lower the entry barrier for non-specialist developers.

3) Personalization Techniques:
Incorporating appropriate user and context features into the shared global model presents an alternative to device-specific personalization. For example, organizing filter orders in applications like Snapchat based on user features such as browsing history, age, gender, preferences, and usage patterns. Consequently, developing architectures adept at effectively accommodating such user and context features for various tasks remains an open challenge. Moreover, bridging the gap between the accuracy of personalized and global models, as observed in underscores the significance of personalization techniques as a vital research area in FL. Nonetheless, comprehensive metrics to evaluate the effectiveness of personalized approaches have yet to be formulated. While Wang et al. examined the conditions conducive to yielding desirable models through personalization, further research is warranted to develop holistic metrics for assessing personalized approach performance.

Core Challenges

1) Communication:
In federated learning (FL), a trade-off exists between communication costs and accuracy. Unlike conventional machine learning (ML) benchmarks, FL benchmarks typically lack a restriction criterion on communication. Introducing communication budgets as a constraint in communication-focused FL benchmarks could be beneficial. For instance, studies like and have explored one-shot or few-shot communication schemes in FL, aiming to

maximize performance within fixed rounds of communication. However, these methods require thorough evaluation and analysis in FL settings characterized by high network heterogeneity. In cross-device FL, asynchronous communication schemes are common, where only a few devices are active during each iteration. Investigating the consequences of such schemes, where device activity is event-driven, warrants further analysis.

2) Systems Heterogeneity:
Addressing systems heterogeneity in FL has been attempted through various algorithms [33], [35]. However, inconsistent wireless connectivity may lead to many participating devices dropping out during training. Future research could focus on designing FL algorithms that exhibit greater robustness, even when a significant number of devices experience connectivity issues. Recent efforts like the introduction of a proximal term in the optimization objective [4] have aimed to incorporate partial solutions from stragglers rather than entirely discarding them. Additionally, approaches like those proposed in [173] tackle device heterogeneity by selecting quantized models at different levels based on device-specific analyses conducted by the FL server.

3) Statistical Heterogeneity:
Eichner et al. proposed a pluralistic solution to mitigate data heterogeneity, considering variations in device characteristics between day and night. Further research avenues could explore similar methods to address diurnal variations at more granular times of day or during different days of the week. For instance, in a federated network operating within a commercial neighborhood, data characteristics during weekdays might significantly differ from those on weekends. Investigating the effectiveness of pluralistic solutions in such scenarios warrants exploration. Additionally, further analysis could be conducted on block-cyclic data in nonconvex settings, employing methods such as parallel SGD, beyond the convex objectives and sequential SGD studied by [99].

4) Privacy/Security:
While device-specific privacy concerns have been extensively studied, finer privacy requirements at the sample level remain a promising research area. Techniques like the sample-specific privacy guarantee developed by Li et al. [174] trade off privacy for higher accuracy. Hybrid methods, addressing both sample- and device-level privacy requirements, could be explored. For instance, using sample-specific privacy for certain data subsets based on specific category levels or date ranges while employing device-specific privacy for the rest.

5) Ablation Analysis:
Evaluation in an FL system is often more complex than traditional ML and DL systems. A holistic industrial system necessitates considering various aspects such as privacy, accuracy/loss, communication rounds, and heterogeneity while building FL solutions. Developing a standard platform for facilitating comprehensive ablation analysis of different parts of an FL system is imperative for advancing research in this field.

Application Areas

Federated learning (FL) has predominantly found applications in supervised learning problems. However, future research endeavors can explore addressing challenges encountered when applying FL in domains requiring data exploration, unsupervised, semi-supervised, and reinforcement learning.

The challenges inherent in implementing FL solutions across various application domains have not been extensively explored. Existing studies have primarily focused on training FL models, overlooking domain-specific challenges. In addition to the core challenges delineated in this paper, it's imperative to consider issues pertinent to specific industries or application domains. For instance, domains like mobile edge networks may prioritize the emphasis on energy-efficient communication methods to a significant extent.

Conclusion

Federated Learning (FL) presents a collaborative approach for organizations to train prediction models without the need to share their sensitive data. This methodology has garnered significant interest from both industry and academia, offering solutions to challenges like data collection and privacy, particularly in sectors such as healthcare.

The escalating interest in FL prompted us to conduct a comprehensive review of contemporary survey papers on FL. We categorized FL into various topics encompassing design aspects, core challenges, and application domains. Through meticulous investigation and analysis of these survey papers, we provided an extensive overview of each FL topic. Lastly, we outlined potential avenues for future research.

This study is anticipated to serve as a valuable resource for upcoming researchers in FL and related fields, aiding them in delineating the scope of their work and fostering advancements in this burgeoning field.

References

[1] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. Y. (2016). Communication-efficient learning of deep networks from decentralized data. arXiv preprint arXiv:1602.05629. [Online]. Available: http://arxiv.org/abs/1602.05629

[2] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. Y. (2016). Federated learning of deep networks using model averaging. arXiv preprint arXiv:1602.05629. [Online]. Available: https://arxiv.org/abs/1602.05629

[3] 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. doi: 10.1109/MSP.2020.2975749.

[4] Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2018). Federated optimization in heterogeneous networks. arXiv preprint arXiv:1812.06127. [Online]. Available: http://arxiv.org/abs/1812.06127

[5] Tang, H., Gan, S., Zhang, C., Zhang, T., & Liu, J. (2018). Communication compression for decentralized training. arXiv preprint arXiv:1803.06443. [Online]. Available: http://arxiv.org/abs/1803.06443

[6] He, L., Bian, A., & Jaggi, M. (2018). COLA: Decentralized linear learning. arXiv preprint arXiv:1808.04883. [Online]. Available: http://arxiv.org/abs/1808.04883

[7] Lalitha, A., Wang, X., Kilinc, O., Lu, Y., Javidi, T., & Koushanfar, F. (2019). Decentralized Bayesian learning over graphs. arXiv preprint arXiv:1905.10466. [Online]. Available: http://arxiv.org/abs/1905.10466

[8] Dai, W., Kumar, A., Wei, J., Ho, Q., Gibson, G., & Xing, E. P. (2014). High-performance distributed ML at scale through parameter server consistency models. arXiv preprint arXiv:1410.8043. [Online]. Available: http://arxiv.org/abs/1410.8043

[9] Ho, Q., Cipar, J., Cui, H., Lee, S., Kim, J. K., Gibbons, P. B., Ganger, G., & Xing, E. P. (2013). More effective distributed ML via a stale synchronous parallel parameter server. In Proceedings of the Advances in Neural Information Processing Systems (pp. 1223–1231).

[10] Zinkevich, M. A., Weimer, M., Smola, A., & Li, L. (2010). Parallelized stochastic gradient descent. In Proceedings of the Neural Information Processing Systems (NIPS) (p. 4).

[11] Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, H. B., Van Overveldt, T., Petrou, D., Ramage, D., & Roselander, J. (2019). Towards federated learning at scale: System design. arXiv preprint arXiv:1902.01046. [Online]. Available: http://arxiv.org/abs/1902.01046

[12]. Dixit, R. R., Schumaker, R. P., & Veronin, M. A. (2018). A Decision Tree Analysis of Opioid and Prescription Drug Interactions Leading to Death Using the FAERS Database. In *IIMA/ICITED Joint Conference 2018* (pp. 67-67). INTERNATIONAL INFORMATION MANAGEMENT ASSOCIATION.
https://doi.org/10.17613/1q3s-cc46

[13]. Veronin, M. A., Schumaker, R. P., Dixit, R. R., & Elath, H. (2019). Opioids and frequency counts in the US Food and Drug Administration Adverse Event Reporting System (FAERS) database: A quantitative view of the epidemic. *Drug, Healthcare and Patient Safety*, 65-70.
https://www.tandfonline.com/doi/full/10.2147/DHPS.S214771

[14]. Veronin, M. A., Schumaker, R. P., & Dixit, R. (2020). The irony of MedWatch and the FAERS database: an assessment of data input errors and potential consequences. *Journal of Pharmacy Technology*, *36*(4), 164-167.
https://doi.org/10.1177/8755122520928

[15]. Veronin, M. A., Schumaker, R. P., Dixit, R. R., Dhake, P., & Ogwo, M. (2020). A systematic approach to'cleaning'of drug name records data in the FAERS database: a case report. *International Journal of Big Data Management*, *1*(2), 105-118.
https://doi.org/10.1504/IJBDM.2020.112404

[16]. Schumaker, R. P., Veronin, M. A., & Dixit, R. R. (2022). Determining Mortality Likelihood of Opioid Drug Combinations using Decision Tree Analysis.
https://doi.org/10.21203/rs.3.rs-2340823/v1

[17]. Schumaker, R. P., Veronin, M. A., Dixit, R. R., Dhake, P., & Manson, D. (2017). Calculating a Severity Score of an Adverse Drug Event Using Machine Learning on the FAERS Database. In *IIMA/ICITED UWS Joint Conference* (pp. 20-30). INTERNATIONAL INFORMATION MANAGEMENT ASSOCIATION.

[18]. Dixit, R. R. (2018). Factors Influencing Healthtech Literacy: An Empirical Analysis of Socioeconomic, Demographic, Technological, and Health-Related Variables. *Applied Research*

in Artificial Intelligence and Cloud Computing, *1*(1), 23-37.

[19]. Dixit, R. R. (2022). Predicting Fetal Health using Cardiotocograms: A Machine Learning Approach. *Journal of Advanced Analytics in Healthcare Management*, *6*(1), 43-57. Retrieved from https://research.tensorgate.org/index.php/JAAHM/article/view/38

[20]. Dixit, R. R. (2021). Risk Assessment for Hospital Readmissions: Insights from Machine Learning Algorithms. *Sage Science Review of Applied Machine Learning*, *4*(2), 1-15. Retrieved from https://journals.sagescience.org/index.php/ssraml/article/view/68

[21]. Ravi, K. C., Dixit, R. R., Singh, S., Gopatoti, A., & Yadav, A. S. (2023, November). AI-Powered Pancreas Navigator: Delving into the Depths of Early Pancreatic Cancer Diagnosis using Advanced Deep Learning Techniques. In *2023 9th International Conference on Smart Structures and Systems (ICSSS)* (pp. 1-6). IEEE. https://doi.org/10.1109/ICSSS58085.2023.10407836

[22]. Khan, M. S., Dixit, R. R., Majumdar, A., Koti, V. M., Bhushan, S., & Yadav, V. (2023, November). Improving Multi-Organ Cancer Diagnosis through a Machine Learning Ensemble Approach. In *2023 7th International Conference on Electronics, Communication and Aerospace Technology (ICECA)* (pp. 1075-1082). IEEE. https://doi.org/10.1109/ICECA58529.2023.10394923

[23]. Sarker , M. (2023). Assessing the Integration of AI Technologies in Enhancing Patient Care Delivery in U.S. Hospitals. *Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online)*, *2*(2), 338-351. https://doi.org/10.60087/jklst.vol2.n2.p351

[24] Geyer, R. C., Klein, T., & Nabi, M. (2017). Differentially private federated learning: A client level perspective. arXiv preprint arXiv:1712.07557. [Online]. Available: http://arxiv.org/abs/1712.07557

[25] McMahan, H. B., Ramage, D., Talwar, K., & Zhang, L. (2018). Learning differentially private recurrent language models. arXiv preprint arXiv:1710.06963. [Online]. Available: https://arxiv.org/abs/1710.06963

[26] Lim, W. Y. B., Luong, N. C., Hoang, D. T., Jiao, Y., Liang, Y.-C., Yang, Q., Niyato, D., & Miao, C. (2020). Federated learning in mobile edge networks: A comprehensive survey. IEEE Communications Surveys & Tutorials, 22(3), 2031–2063. doi: 10.1109/COMST.2020.2986024.