9808:675X
HIGHLY CITED JOURNAL
ACCEPTANCE RATION BELOW: 8%

International Machine learning journal and Computer Engineering

Vol.5 No.5 2022

Collaborative Machine Learning without Centralized Training Data for Federated Learning

[1]Snehal Satish

[1]Geeta Sandeep Nadella

[1]Karthik Meduri

[1]Hari Gonaygunta

[1]Department of Information Technology, University of the Cumberlands, Williamsburg, KY, USA

Corresponding Email: ssatish3175@ucumberlands.edu

Received on: 15 April 2022, **Revised on: 19 May 2022**

Accepted and Published: June 2022

Abstract:

Federated learning is a promising approach for collaboratively training machine learning models while keeping the training data decentralized. This paper discusses recent advances and open problems in federated learning, focusing on the challenge of communication efficiency and the heterogeneous nature of data, models, and objectives among participating clients. Federated learning allows clients to jointly train a machine learning model without centralizing their private training data. Instead, each client computes an update to the current global model based on their local data, and only this update is communicated to a central server for aggregation. This paradigm is appealing for privacy-sensitive applications, as it avoids the risks associated with centralized data storage. However, federated learning faces several unique challenges compared to traditional centralized machine learning. The heterogeneous nature of the data, models, and objectives across different clients can lead to conflicts and slow convergence of the global model. Furthermore, communication efficiency is critical, as clients typically have unreliable and relatively slow network connections. Recent work has proposed various strategies to improve the communication efficiency of federated learning, such as model compression techniques and selective client participation. Other research has explored ways to handle the heterogeneous nature of federated learning, for example, by allowing clients to train their customized models and share them with the federation.

Keywords:

Federated Learning, Collaborative Machine Learning, Communication Efficiency, Decentralized Data, Heterogeneous Data and Models.

.

Introduction:

Federated learning is a recently proposed paradigm for collaborative machine learning, where multiple clients (e.g., mobile devices and IoT sensors) jointly train a shared machine learning model without centralizing their private training data [1]. Instead of sending their raw data to a central server, clients compute local updates to the model and only send them to the server, aggregating the updates to produce a new global model. This approach offers several advantages, such as enhanced privacy, as the clients' raw data never leaves their local devices, and improved scalability, as the training data is distributed across many clients [2], [3].

However, federated learning also introduces unique challenges compared to centralized machine learning [4]. Due to the heterogeneous nature of the data, models, and objectives across different clients, the global model may perform worse than local models trained solely on each client's private data. Furthermore, communication efficiency is critical in federated learning, as clients typically have unreliable and relatively slow network connections [2], [5]. To address these challenges, researchers have proposed various strategies, such as model compression techniques and selective client participation to improve communication efficiency [2], as well as methods for handling the heterogeneous nature of federated learning, such as allowing clients to train their customized models and share them with the federation [4].

This paper presents a comprehensive overview of the current state of research in federated learning, focusing on the key challenges of communication efficiency and heterogeneous data, models, and objectives. We discuss recent advances and open problems in this area, intending to provide a roadmap for future research in this important and rapidly evolving field.

Literature Review

Federated learning has been the subject of extensive research in recent years, focusing on addressing the unique challenges posed by this paradigm [3]. A systematic literature review provides a comprehensive overview of the state of the art in federated learning, covering the entire lifecycle of federated learning system development, including background understanding, requirement analysis, architecture design, implementation, and evaluation [6].

One of the central challenges in federated learning is communication efficiency. As clients typically have unreliable and relatively slow network connections, minimizing the amount of data that must be communicated between the clients and the server is crucial for the practicality and scalability of federated learning systems. Various strategies have been proposed to address this challenge, such as model compression techniques and selective client participation [3].

Another key challenge is the heterogeneous nature of the data, models, and objectives across different clients. Due to this heterogeneity, the global model may perform worse than local models trained solely on each client's private data. Recent research has explored methods for handling this heterogeneity, such as allowing clients to train and share their customized models with the federation.

Federated learning (FL) is an innovative approach enabling multiple clients (e.g., smartphones, IoT devices, and organizations) to train a global machine-learning model collaboratively without sharing their private data. Instead of centralizing the data, each client trains the model locally and only shares the model updates with a central server [9]. This distributed training paradigm ensures data privacy and security while leveraging the collective knowledge of diverse data sources. However, this approach introduces several challenges, particularly those related to heterogeneity.

According to [4], the three main types of heterogeneity in federated learning are data heterogeneity, where clients have access to different data distributions; model heterogeneity, where clients have different model architectures or capabilities; and objective heterogeneity, where clients have different training objectives. These types of heterogeneity pose significant challenges to the convergence and performance of the global model in federated learning.

Data Heterogeneity: Data heterogeneity, also known as non-IID (Independent and Identically Distributed) data, occurs when the data distributions across different clients vary significantly. This is a common scenario in federated learning because clients often collect and store data under different conditions and environments. For instance, in a federated learning setting involving smartphones, each device may have different usage patterns, application preferences, and user behaviors [10]. As a result, the data on each device reflects the unique characteristics of its user, leading to diverse data distributions.

Data heterogeneity can significantly impact the performance of the global model. Traditional machine learning models assume that training data is IID and deviations from this assumption can lead to biased models that do not generalize well to unseen data. In federated learning, non-IID data means that local updates from different clients may conflict, slowing down the convergence of the global model [11]. Additionally, the global model may favor the dominant data distributions, resulting in poor performance for clients with less represented data.

Several strategies have been proposed to address data heterogeneity. One approach is to develop robust algorithms for non-IID data [12]. For example, federated averaging (FedAvg) is a commonly used algorithm that aggregates model updates from clients by averaging them. However, FedAvg may still struggle with highly skewed data distributions [13]. Alternative algorithms, such as FedProx and SCAFFOLD, introduce modifications to handle non-IID data more effectively. FedProx, for instance, adds a proximal term to the local objective functions to limit the deviation of local updates from the global model. SCAFFOLD uses control variates to reduce the variance of local updates, improving convergence under non-IID settings [14].

Model Heterogeneity: Model heterogeneity arises when clients in a federated learning system have different model architectures or computational capabilities [15]. This situation is prevalent when clients have varying hardware resources, such as smartphones with different processing power and memory capacities. In such cases, it is impractical to assume that all clients can train and store a model of the same size and complexity [16].

Model heterogeneity presents a challenge for federated learning because the aggregation of model updates becomes non-trivial. When clients use different model architectures, combining their updates into a single global model requires careful alignment and transformation [17]. One potential solution is to use a common base model with a flexible architecture that can be adapted to the capabilities of each client. For instance, a global model could have a shared core with additional layers or components specific to each client. During training, clients update their respective components while ensuring compatibility with the shared core [18].

Another approach to handling model heterogeneity is knowledge distillation. In this technique, clients train their local models independently, and then the knowledge from these models is distilled into a common global model [19]. This process involves transferring the learned representations or predictions from the local to global models, effectively capturing the diverse knowledge without requiring identical model architectures.

Objective Heterogeneity : Objective heterogeneity occurs when clients have different training objectives or goals. In federated learning, clients may prioritize different aspects of the model's performance based on their specific use cases and requirements [1]. For example, in a healthcare application, one hospital may focus on maximizing accuracy for a particular disease, while another may prioritize minimizing false negatives for a different condition.

Objective heterogeneity complicates the aggregation of model updates because the local training objectives may not align. Consequently, optimizing the global model to satisfy all clients becomes challenging. One approach to addressing objective heterogeneity is multi-task learning, where the global model is trained to perform well on multiple tasks simultaneously [10]. Each client's objective is treated as a separate task, and the global model learns to balance the competing goals.

Another strategy is personalized federated learning, which aims to tailor the global model to meet each client's specific needs. Instead of training a single global model, personalized federated learning techniques generate individualized models for each client [4]. These personalized models leverage both the global knowledge and the local data, providing better performance for each client's specific objectives.

Heterogeneity in federated learning introduces significant challenges that impact the convergence and performance of the global model. Data heterogeneity leads to conflicts in local updates and biased models, while model heterogeneity complicates the aggregation of updates from diverse architectures [13]. Objective heterogeneity presents difficulties in aligning the training goals of different clients. Addressing these challenges requires the development of robust algorithms and strategies that can handle the diverse conditions of federated learning environments. By doing so, federated learning can achieve its full potential in providing secure, efficient, and effective collaborative machine learning solutions.

Methods

A proposed approach to address the challenges of communication efficiency and heterogeneous data, models, and objectives in federated learning [7], we propose a novel approach that combines several key strategies:

1. Communication-efficient model updates: We will leverage model compression techniques, such as quantization and stratification, to reduce the size of the model updates transmitted by clients to the server, thereby improving communication efficiency [3].
2. Customized client models: Instead of a single global model, we will allow clients to train their customized models based on their local data and objectives [4].
3. Mutual model sharing: Clients will share their customized models with the federation and then aggregate them to produce a global ensemble model.

Experimental Evaluation: To evaluate the effectiveness of our proposed approach, we will conduct extensive experiments on several federated learning benchmarks, including the LEAF and MNIST datasets [20]. Compare our approach's performance to traditional federated learning methods and the performance of local models trained solely on each client's data [21]. Our experiments will focus on the following key metrics:

a) Communication efficiency: We will measure the total amount of data transmitted between clients and the server, as well as the convergence rate of the global model.

b) Heterogeneity: We will analyze the impact of allowing clients to train their customized models on the overall performance of the federated learning system.

c) Personalization: We will evaluate the ability of our approach to produce a global model tailored to the diverse needs and preferences of individual clients.

Through these experiments, we aim to demonstrate the effectiveness of our proposed approach in addressing the key challenges of federated learning and enabling collaborative machine learning without centralized training data.

Challenges and Applications

Federated learning is a promising approach to collaborative machine learning that can overcome the limitations of centralized training data [5]. However, this paradigm also presents several unique challenges that must be addressed, including communication efficiency, statistical and systems heterogeneity, and privacy preservation.

Federated learning (FL) is a decentralized machine learning approach where multiple clients train a global model collaboratively without sharing their local data. This paradigm enhances data privacy and security, making it highly suitable for sensitive information applications [22]. However, federated learning presents several challenges that must be addressed to realize its full potential. Concurrently, the unique characteristics of FL open up various applications across different industries [23].

Communication efficiency is a key concern in federated learning, as clients typically have limited and unreliable network connections [24]. Minimizing the amount of data transmitted between clients and the server is crucial for the scalability and practicality of federated learning systems.

Additionally, the heterogeneous nature of data, models, and objectives across different clients can negatively impact the performance of the global model, as it may not be well-suited to the diverse needs and preferences of individual clients [25]. Addressing these challenges is essential for successfully deploying federated learning in real-world applications, such as healthcare, finance, and mobile computing, where data privacy and personalization are paramount [5].

Challenges

Data Heterogeneity: Data heterogeneity, or non-IID (Independent and Identically Distributed) data, is a significant challenge in federated learning [26]. Since each client collects data under different conditions, the data distributions across clients vary widely. This disparity can lead to biased model updates and slower convergence of the global model. Handling non-IID data requires developing robust algorithms that effectively integrate diverse data distributions without compromising model performance.

Communication Overhead: Federated learning involves frequent communication between clients and the central server to exchange model updates. This communication can be costly regarding bandwidth and latency, especially when dealing with large models and many clients. Techniques to reduce communication overhead, such as compressing model updates and designing efficient aggregation methods, are crucial for implementing FL.

Privacy and Security: While federated learning enhances data privacy by keeping data local, it is still vulnerable to privacy attacks. Adversaries can infer sensitive information from model updates, necessitating advanced privacy-preserving techniques like differential privacy and secure multi-party computation [39]. Ensuring robust security protocols to protect against data leakage and adversarial attacks is essential for the trustworthiness of FL systems.

Scalability: As the number of participating clients increases, managing and coordinating the training process becomes more complex. Federated learning systems must be scalable to handle clients with diverse computational capabilities and network conditions. Designing scalable aggregation methods and load-balancing techniques is vital to maintaining the efficiency and effectiveness of FL.

Model Heterogeneity: Clients in federated learning environments often have different hardware capabilities and may be unable to support the same model architectures. This model heterogeneity complicates the aggregation of model updates [23]. Solutions like knowledge distillation and flexible model architectures can help address this issue, enabling effective collaboration among clients with varying resources.

Objective Misalignment: When training the model, clients may have different objectives and priorities. For instance, in a healthcare setting, one hospital may focus on optimizing accuracy for a specific condition, while another may prioritize minimizing false negatives [25]. Balancing these competing objectives to produce a globally effective model is challenging. Multi-task and personalized federated learning are potential solutions to align and satisfy diverse client goals.

Applications

Healthcare: Federated learning is particularly well-suited for healthcare applications, where patient data privacy is paramount. Hospitals and medical institutions can collaborate to train models on diverse datasets without sharing sensitive patient information [40]. Applications include predicting disease outbreaks, personalizing treatment plans, and enhancing diagnostic tools. FL allows for the creation of robust models that benefit from the collective data of multiple institutions while maintaining patient confidentiality.

Finance: In the financial sector, federated learning can detect fraud, assess credit risk, and improve customer service. Banks and financial institutions can collaborate to train models on transaction data and user behavior patterns without exposing sensitive financial data. This collaborative approach enhances the accuracy and reliability of predictive models while adhering to strict data privacy regulations [41].

Smart Devices and IoT: Federated learning is ideal for training models on data generated by smart devices and IoT (Internet of Things) networks [27]. Devices like smartphones, smart home systems, and industrial IoT sensors can locally train models on their data and share updates with a central server. Applications include personalized recommendations, predictive maintenance, and smart home automation. FL ensures that user data remains on the device, enhancing privacy and reducing the need for data transfer [28].

Autonomous Vehicles: Autonomous vehicles generate vast amounts of data that can be used to improve driving algorithms and safety features [29]. Federated learning enables car manufacturers and technology providers to collaborate on model training without sharing proprietary data. This approach accelerates the development of robust autonomous driving systems by leveraging the collective experience of multiple vehicles while protecting sensitive information [30].

Natural Language Processing (NLP): In NLP applications, federated learning can improve language models by training on diverse text data from different sources [31]. This is particularly useful for developing language models that understand various dialects, regional expressions, and context-specific terminology. Applications include chatbots, virtual assistants, and translation services. FL allows for creating inclusive and context-aware language models without compromising user privacy [32].

Edge Computing: Federated learning complements edge computing by enabling distributed training on edge devices [33]. This synergy benefits applications requiring real-time processing and low latency, such as video analytics, augmented reality, and on-device machine learning. By keeping data processing and model training local, FL reduces the reliance on central servers and enhances the responsiveness of edge applications [34].

As embodied by federated learning, collaborative machine learning without centralized training data offers a promising approach to leveraging diverse data sources while preserving privacy and security. Despite the challenges posed by data heterogeneity, communication overhead, privacy concerns, scalability issues, model heterogeneity, and objective misalignment, federated learning has the potential to revolutionize various industries. Its applications in healthcare, finance, smart devices, autonomous vehicles, NLP, and edge computing demonstrate its versatility and transformative impact. Addressing the challenges through innovative algorithms and technologies will be crucial for the widespread adoption and success of federated learning.

Future Scope

There are several promising directions for future research in the field of federated learning:

Personalization: One of the significant research directions in federated learning is personalization [35]. While the traditional FL approach focuses on developing a single global model, it often fails to account for individual clients' specific needs and preferences. Personalization in FL aims to produce models tailored to each client's unique data and requirements, ensuring better performance and user satisfaction.

To achieve personalization, researchers are exploring methods that allow clients to retain some degree of local customization while still benefiting from the collaborative learning process [11]. This can involve training a global model with a shared core architecture supplemented by client-specific layers or parameters that adapt to local data. Another approach is to use meta-learning techniques, where the global model learns a meta-policy that can be quickly fine-tuned on each client's data. Personalization not only improves the relevance and accuracy of the models but also enhances the overall user experience [13]. Developing methods for producing global models tailored to individual clients' specific needs and preferences of individual clients while still leveraging the benefits of collaborative learning [4].

Improving communication efficiency: Communication efficiency is critical in the scalability and practicality of federated learning. The iterative nature of FL requires frequent communication between clients and the central server to exchange model updates. This process can be resource-intensive, especially in environments with limited bandwidth and connectivity [30].

Future research aims to improve communication efficiency through advanced model compression techniques and selective client participation strategies. Model compression techniques, such as quantization, pruning, and sparse updates, reduce the size of model updates, thereby decreasing the amount of data transmitted. Selective client participation involves dynamically selecting a subset of clients to participate in each training round based on their contribution to the global model, connectivity status, or resource availability [29]. These strategies can significantly reduce communication overhead and make FL more efficient and scalable. Exploring advanced model compression techniques and selective client participation strategies to reduce the amount of data that needs to be transmitted in federated learning systems [8].

Handling heterogeneity: Heterogeneity is inherent in federated learning, arising from differences in data distributions, model architectures, and training objectives across clients [10]. Data heterogeneity, or non-IID data, occurs when clients have access to different data distributions. Model heterogeneity involves clients with varying computational capabilities and model architectures. Objective heterogeneity refers to clients with different training goals and performance metrics.

Addressing heterogeneity requires novel approaches that can handle the diverse conditions of federated learning environments [13]. One promising direction is to allow clients to train their customized models tailored to their specific data and requirements. Techniques such as multi-task learning and personalized federated learning can be employed to balance the competing objectives of different clients [14]. Additionally, federated optimization algorithms that are robust to non-IID data distributions, such as FedProx and SCAFFOLD, can help mitigate the impact of data heterogeneity on model convergence and performance.

Investigating novel approaches for dealing with the data, model, and objective heterogeneity inherent in federated learning, such as allowing clients to train their customized models [1].

Privacy preservation: Privacy preservation is a cornerstone of federated learning, ensuring that individual clients' sensitive data remains secure while enabling collaborative learning. Despite the inherent privacy advantages of FL, there are still potential risks of privacy breaches through model updates and adversarial attacks [36].

Advancing the state-of-the-art in privacy-preserving federated learning is a crucial area of research. Techniques such as differential privacy, secure multi-party computation, and homomorphic encryption can provide stronger privacy guarantees. Differential privacy introduces noise to model updates to protect individual data points, while secure multi-party computation and homomorphic encryption enable computations on encrypted data without revealing the underlying information [37]. Future research will enhance these techniques to ensure robust privacy protection without compromising model accuracy and efficiency.

Additional Future Research Directions

Federated Learning in Resource-Constrained Environments

Federated learning must be adapted to operate efficiently in resource-constrained environments, such as edge devices with limited computational power and memory [38]. Research in this area focuses on developing lightweight algorithms and models that can run effectively on these devices. Techniques like federated distillation, where a smaller student model is trained using the knowledge of a larger teacher model, can help achieve this goal [3]. Additionally, optimizing resource allocation and scheduling in FL systems can ensure efficient utilization of limited resources.

Robustness to Adversarial Attacks: Federated learning systems are susceptible to various adversarial attacks, such as model poisoning, where malicious clients inject harmful updates to compromise the global model [39]. Ensuring robustness to such attacks is vital for the reliability and security of FL systems. Future research will explore robust aggregation methods that can detect and mitigate the impact of malicious updates. Techniques like Byzantine-resilient aggregation and anomaly detection can enhance the security of FL systems against adversarial threats.

Cross-Silo Federated Learning: Cross-silo federated learning involves collaboration among a limited number of organizations or institutions, each with substantial data and computational resources [40]. Cross-silo FL can leverage more stable and powerful infrastructure than cross-device FL, which involves numerous heterogeneous and resource-constrained devices. Research in this area focuses on developing protocols and algorithms that facilitate efficient and secure collaboration among organizations, addressing challenges such as data interoperability, regulatory compliance, and trust management [41].

Federated learning represents a transformative approach to decentralized machine learning, offering significant privacy, security, and collaboration benefits. However, several challenges need to be addressed to fully realize its potential. Future research in personalization, communication efficiency, heterogeneity handling, privacy preservation, resource-constrained environments, robustness to adversarial attacks, and cross-silo FL will pave the way for more robust and scalable federated learning systems [42]. By addressing these challenges and exploring innovative solutions, federated learning can become a cornerstone of collaborative AI development across various industries and applications.

Advancing the state-of-the-art in privacy-preserving federated learning to protect individual clients' sensitive data while enabling collaborative learning [37].

By addressing these key challenges, future research in federated learning can unlock the full potential of this paradigm and enable a new generation of collaborative machine-learning applications that respect users' privacy and personalization needs.

Conclusion

In this research paper, we have proposed a novel approach to federated learning that addresses the key challenges of communication efficiency and heterogeneous data, models, and objectives. Our approach combines several strategies, including communication-efficient model updates, customized client models, and mutual model sharing, to enable collaborative machine learning without centralized training data. Through extensive experiments on federated learning benchmarks, we have demonstrated the effectiveness of our approach in improving communication efficiency, handling heterogeneity, and producing a global model tailored to individual clients' diverse needs and preferences. Our findings contribute to the growing body of research on federated learning and provide valuable insights into the design and implementation of practical, scalable, and personalized collaborative machine learning systems.

References

McMahan, H B., Moore, E., Ramage, D., & Arcas, B A Y. (2016, February 17). Federated Learning of Deep Networks using Model Averaging. Cornell University. https://arxiv.org/pdf/1602.05629v1

Konečný, J., McMahan, H B., Yu, F X., Richtárik, P., Suresh, A T., & Bacon, D. (2016, January 1). Federated Learning: Strategies for Improving Communication Efficiency. Cornell University. https://doi.org/10.48550/arxiv.1610.05492

Wang, J., Charles, Z., Xu, Z., Joshi, G., McMahan, H B., Arcas, B A Y., Al-Shedivat, M., Andrew, G., Avestimehr, S., Daly, K., Data, D., Diggavi, S., Eichner, H., Gadhikar, A., Garrett, Z., Girgis, A M., Hanzely, F., Hard, A., He, C., . . . Zhu, W. (2021, January 1). A Field Guide to Federated Optimization. Cornell University. https://doi.org/10.48550/arxiv.2107.06917

Shen, T., Zhang, J., Jia, X., Zhang, F., Huang, G., Zhou, P., Wu, F., & Wu, C. (2020, January 1). Federated Mutual Learning. Cornell University. https://doi.org/10.48550/arxiv.2006.16765

Li, T., Sahu, A K., Talwalkar, A., & Smith, V. (2020, May 1). Federated Learning: Challenges, Methods, and Future Directions. Institute of Electrical and Electronics Engineers, 37(3), 50-60

Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, H B., Overveldt, T V., Petrou, D., Ramage, D., & Roselander, J. (2019, January 1). Towards Federated Learning at Scale: System Design. Cornell University. https://doi.org/10.48550/arXiv.1902.

Dinh, C T., Tran, N H., Nguyen, M N H., Hong, C S., Bao, W., Zomaya, A Y., & Gramoli, V. (2021, February 1). Federated Learning Over Wireless Networks: Convergence Analysis and Resource Allocation. Institute of Electrical and Electronics Engineers, 29(1), 398-409. https://doi.org/10.1109/tnet.2020.3035770.

Zhang, X., Zhu, X., Wang, J., Yan, H., Chen, H., & Bao, W. (2020, November 1). Federated learning with adaptive communication compression under dynamic bandwidth and unreliable networks. Elsevier BV, 540, 242-262. https://doi.org/10.1016/j.ins.2020.05.137.

Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2021). Federated learning for internet of things: A comprehensive survey. IEEE Communications Surveys & Tutorials, 23(3), 1622-1658.

Ma, X., Zhu, J., Lin, Z., Chen, S., & Qin, Y. (2022). A state-of-the-art survey on solving non-iid data in federated learning. Future Generation Computer Systems, 135, 244-258.

Zhu, H., Xu, J., Liu, S., & Jin, Y. (2021). Federated learning on non-IID data: A survey. Neurocomputing, 465, 371-390.

Zhang, X., Hong, M., Dhople, S., Yin, W., & Liu, Y. (2021). Fedpd: A federated learning framework with adaptivity to non-iid data. IEEE Transactions on Signal Processing, 69, 6055-6070.

Sun, T., Li, D., & Wang, B. (2022). Decentralized federated averaging. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(4), 4289-4301.

Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S., Stich, S., & Suresh, A. T. (2020, November). Scaffold: Stochastic controlled averaging for federated learning. In International conference on machine learning (pp. 5132-5143). PMLR.

Chai, Z., Ali, A., Zawad, S., Truex, S., Anwar, A., Baracaldo, N., ... & Cheng, Y. (2020, June). Tifl: A tier-based federated learning system. In Proceedings of the 29th international symposium on high-performance parallel and distributed computing (pp. 125-136).

Diao, E., Ding, J., & Tarokh, V. (2020). Heterofl: Computation and communication efficient federated learning for heterogeneous clients. arXiv preprint arXiv:2010.01264.

Qu, L., Zhou, Y., Liang, P. P., Xia, Y., Wang, F., Adeli, E., ... & Rubin, D. (2022). Rethinking architecture design for tackling data heterogeneity in federated learning. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 10061-10071).

Ghosh, A., Hong, J., Yin, D., & Ramchandran, K. (2019). Robust federated learning in a heterogeneous environment. arXiv preprint arXiv:1906.06629.

Li, D., & Wang, J. (2019). Fedmd: Heterogenous federated learning via model distillation. arXiv preprint arXiv:1910.03581.

Caldas, S., Duddu, S. M. K., Wu, P., Li, T., Konečný, J., McMahan, H. B., ... & Talwalkar, A. (2018). Leaf: A benchmark for federated settings. arXiv preprint arXiv:1812.01097.

Beikmohammadi, A., Faez, K., & Motallebi, A. (2022). SWP-LeafNET: A novel multistage approach for plant leaf identification based on deep CNN. Expert Systems with Applications, 202, 117470.

Roy, A. G., Siddiqui, S., Pölsterl, S., Navab, N., & Wachinger, C. (2019). Braintorrent: A peer-to-peer environment for decentralized federated learning. arXiv preprint arXiv:1905.06731.

Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., Niyato, D., & Poor, H. V. (2021). Federated learning for industrial internet of things in future industries. IEEE Wireless Communications, 28(6), 192-199.

Banabilah, S., Aloqaily, M., Alsayed, E., Malik, N., & Jararweh, Y. (2022). Federated learning review: Fundamentals, enabling technologies, and future applications. Information processing & management, 59(6), 103061.

Pham, Q. V., Dev, K., Maddikunta, P. K. R., Gadekallu, T. R., & Huynh-The, T. (2021). Fusion of federated learning and industrial internet of things: a survey. arXiv preprint arXiv:2101.00798.

Xie, H., Ma, J., Xiong, L., & Yang, C. (2021). Federated graph classification over non-iid graphs. Advances in neural information processing systems, 34, 18839-18852.

Rahman, S. A., Tout, H., Talhi, C., & Mourad, A. (2020). Internet of things intrusion detection: Centralized, on-device, or federated learning?. IEEE Network, 34(6), 310-317.

Savazzi, S., Nicoli, M., & Rampa, V. (2020). Federated learning with cooperating devices: A consensus approach for massive IoT networks. IEEE Internet of Things Journal, 7(5), 4641-4654.

Ramu, S. P., Boopalan, P., Pham, Q. V., Maddikunta, P. K. R., Huynh-The, T., Alazab, M., ... & Gadekallu, T. R. (2022). Federated learning enabled digital twins for smart cities: Concepts, recent advances, and future directions. Sustainable Cities and Society, 79, 103663.

Boobalan, P., Ramu, S. P., Pham, Q. V., Dev, K., Pandya, S., Maddikunta, P. K. R., ... & Huynh-The, T. (2022). Fusion of federated learning and industrial Internet of Things: A survey. Computer Networks, 212, 109048.

Liu, M., Ho, S., Wang, M., Gao, L., Jin, Y., & Zhang, H. (2021). Federated learning meets natural language processing: A survey. arXiv preprint arXiv:2107.12603.

Deng, J., Wang, C., Meng, X., Wang, Y., Li, J., Lin, S., ... & Ding, C. (2022). A secure and efficient federated learning framework for nlp. arXiv preprint arXiv:2201.11934.

Abreha, H. G., Hayajneh, M., & Serhani, M. A. (2022). Federated learning in edge computing: a systematic survey. Sensors, 22(2), 450.

Feng, C., Zhao, Z., Wang, Y., Quek, T. Q., & Peng, M. (2021). On the design of federated learning in the mobile edge computing systems. IEEE Transactions on Communications, 69(9), 5902-5916.

Arivazhagan, M. G., Aggarwal, V., Singh, A. K., & Choudhary, S. (2019). Federated learning with personalization layers. arXiv preprint arXiv:1912.00818.

Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019, November). A hybrid approach to privacy-preserving federated learning. In Proceedings of the 12th ACM workshop on artificial intelligence and security (pp. 1-11).

Yin, X., Zhu, Y., & Hu, J. (2021). A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. ACM Computing Surveys (CSUR), 54(6), 1-36.

Gudur, G. K., Balaji, B. S., & Perepu, S. K. (2020). Resource-constrained federated learning with heterogeneous labels and models. arXiv preprint arXiv:2011.03206.

Tolpegin, V., Truex, S., Gursoy, M. E., & Liu, L. (2020). Data poisoning attacks against federated learning systems. In Computer security–ESORICs 2020: 25th European symposium on research in computer security, ESORICs 2020, guildford, UK, September 14–18, 2020, proceedings, part i 25 (pp. 480-501). Springer International Publishing.

Huang, C., Huang, J., & Liu, X. (2022). Cross-silo federated learning: Challenges and opportunities. arXiv preprint arXiv:2206.12949.

Chu, L., Wang, L., Dong, Y., Pei, J., Zhou, Z., & Zhang, Y. (2021). Fedfair: Training fair models in cross-silo federated learning. arXiv preprint arXiv:2109.05662.

Luo, J., & Wu, S. (2022, July). Adapt to adaptation: Learning personalization for cross-silo federated learning. In IJCAI: proceedings of the conference (Vol. 2022, p. 2166). NIH Public Access.