

Seamless Integration: Sampling Strategies in Federated Learning Systems

Tatjana Legler

University of Kaiserslautern-Landau (RPTU)
German Research Center for Artificial Intelligence (DFKI)
Kaiserslautern, Germany
0000-0002-7297-0845

Vinit Hegiste

University of Kaiserslautern-Landau (RPTU)
Kaiserslautern, Germany
0000-0001-6944-1988

Martin Ruskowski

University of Kaiserslautern-Landau (RPTU)
German Research Center for Artificial Intelligence (DFKI)
Kaiserslautern, Germany
0000-0002-6534-9057

Abstract—Federated Learning (FL) represents a paradigm shift in the field of machine learning, offering an approach for a decentralized training of models across a multitude of devices while maintaining the privacy of local data. However, the dynamic nature of FL systems, characterized by the ongoing incorporation of new clients with potentially diverse data distributions and computational capabilities, poses a significant challenge to the stability and efficiency of these distributed learning networks. The seamless integration of new clients is imperative to sustain and enhance the performance and robustness of FL systems. This paper looks into the complexities of integrating new clients into existing FL systems and explores how data heterogeneity and varying data distribution (not independent and identically distributed) among them can affect model training, system efficiency, scalability and stability. Despite these challenges, the integration of new clients into FL systems presents opportunities to enhance data diversity, improve learning performance, and leverage distributed computational power. In contrast to other fields of application such as the distributed optimization of word predictions on Gboard (where federated learning once originated), there are usually only a few clients in the production environment, which is why information from each new client becomes all the more valuable. This paper outlines strategies for effective client selection strategies and solutions for ensuring system scalability and stability. Using the example of images from optical quality inspection, it offers insights into practical approaches. In conclusion, this paper proposes that addressing the challenges presented by new client integration is crucial to the advancement and efficiency of distributed learning networks, thus paving the way for the adoption of Federated Learning in production environments.

Index Terms—Federated Learning, Decentralized Learning, Manufacturing Collaboration, Data Heterogeneity

I. INTRODUCTION

Federated Learning (FL) has emerged as a transformative approach in the field of machine learning, fundamentally changing how data is utilized and models are trained across distributed networks [1]. In contrast to traditional centralized machine learning paradigms, FL allows numerous devices to collaboratively train a shared model while retaining all training

data locally on the devices. This approach addresses significant concerns related to data privacy, security, and access rights [2]. The decentralized nature of FL not only reduces the risks associated with data centralization but also enables the use of data that was previously inaccessible or unusable due to privacy constraints.

The applications of machine learning in manufacturing span a wide range, from well-established uses such as fault diagnostics, process monitoring, and anomaly detection to more recent innovations like enhancing human-robot collaboration with the aid of Generative Pre-trained Transformers (GPTs) and optimizing construction designs [3]. The demand for efficient data processing, coupled with the sensitive and proprietary nature of manufacturing data, presents numerous challenges for machine learning applications [4]. These challenges are exacerbated in the dynamic environment of FL, particularly with the continuous integration of new clients. These new participants, characterized by potentially diverse data distributions and varying computational capabilities, introduce complexities that can impact the stability, efficiency, and scalability of FL systems. The heterogeneity in data, often non-independent and identically distributed (non-IID), along with the variance in computational resources among clients, necessitates innovative strategies to seamlessly incorporate these new participants into existing networks [5] [6].

The introduction of new clients within a FL system offers the potential to enhance the diversity of the data pool, thereby improving the robustness and performance of the collective learning process. Each new client contributes unique insights and perspectives, enriching the network and making the model more comprehensive and better suited for real-world applications. This is especially valuable in fields like optical quality inspection, where variations in data across different production lines can significantly influence the effectiveness of predictive models.

To fully harness the potential benefits of FL, it is crucial to

address the challenges related to the integration of new clients. This requires the use or development of strategies that enable effective onboarding and maintain the stability of the FL system. In this paper, we focus on client selection techniques as a foundation for future research into client integration, examining the complexities introduced by data heterogeneity and proposing strategies for efficient client integration.

For decades, various approaches, such as zero-defect manufacturing and lean Six Sigma methodologies, have been employed in production environments [7] [8]. These methods have successfully reduced defect rates to levels measurable in "parts per million" [9]. However, with the increasing variety of product variants, these traditional methods are reaching their limits [10]. To maintain quality assurance amidst this diversity, machine learning methods are increasingly being applied [11].

Several strategies exist to overcome data scarcity and expand the usable data pool. Generative AI methods, for example, can create synthetic data for multiple product variants and configurations, including training data for various error cases [12]. Additionally, breaking down existing data silos can increase both the volume and diversity of available data [13]. In the following, we explore FL as a potential approach to achieve this while preserving data privacy. Additionally, synthetic data and real images can be merged with FL, enhancing the robustness and applicability of machine learning models in production, as shown in [14].

II. STATE OF THE ART

The manufacturing industry is a data-intensive sector with numerous actors generating data from various sources, including sensors, production processes, and quality control systems [15]. However, the data in manufacturing are often sensitive and protected, complicating the sharing or collaborative training of machine learning models on a shared dataset [4]. FL offers a decentralized approach that aims to leverage a broad, heterogeneous database without the need to share raw data [2]. Unlike traditional centralized learning, where models are trained on large datasets housed on a central server [16] [17], FL only shares model parameters, such as network weights, rather than raw information like production data or images. This approach allows each participant to benefit from data-sovereign learning while reaping the advantages of collaborative learning, such as improved generalizability and stability. Additionally, by indirectly increasing the sample size, FL can provide access to a greater variety of features.

Beyond the basic concept of FL, McMahan et al. introduced a detailed method called Federated Averaging (FedAvg), in which a central server functions as an aggregator [18]. The FedAvg process includes the following steps:

- **Initialization:** A server initializes a global model and distributes it to all participating clients, initiating their local training.
- **Local Training:** Each selected client trains the model on its local data.
- **Aggregation:** After local training is completed on the device, the model updates (e.g., weights or weight changes)

are sent back to the central server. The server then aggregates these updates to create the global model, typically by calculating an arithmetic average, as done in FedAvg.

- **Update:** The global server sends the updated model to all clients.

This process, known as a communication round, is repeated until the model achieves the desired accuracy or meets a specific convergence or termination criterion (see Figure 1).

A. Client Selection Strategies

When FL and FedAvg were introduced, initially a *random selection* of clients was performed to participate in the next communication round [18]. The authors observed that incorporating additional clients beyond a certain threshold yields diminishing returns as the additional communication overhead exceeds the performance gains, but did not investigate further into client selection strategies. This simple approach works well when there is a sufficient number of clients with similar data distributions [20]. However, it often fails to account for device heterogeneity, which can lead to incomplete communication rounds, or the risk of over-sampling frequently provided data [21].

The available computation power, the network connection as well as a possible battery operation which can lead to phases of energy saving during which the clients are inaccessible, can be considered for *resource-aware selection* [22]. Using different hardware components can result in different calculation times. For instance, [23] focuses on reducing the influence of hardware in the selection process by neither waiting for the slowest clients (known as stragglers) every round nor completely discarding them, which would result in an imbalance in representation. To avoid a prolonged waiting phase in mobile edge computing as devices with varying resources communicate over a cellular network, [24] estimates how many clients can reasonably participate.

For *performance-aware selection* metrics such as accuracy or training loss can be used in addition to the consideration of resource constraints as mentioned above [25]. [26] demonstrates, that the choice of clients with a higher local loss can improve convergence of the model, as they have a bigger potential for improvement.

The term *Contribution-Based Selection* is used to categorize methods that use the effects of a client on the global model as a selection criterion for the next communication round [27]. [28] uses the Shapley value known from cooperative game theory to estimate the contribution of each client to the global model. A further development can be found in [29] resulting in a contribution estimation, which looks at the gradient space and data space for each client. [30] propose a flexible selection scheme that converges even when communication rounds are interrupted by clients leaving or joining.

[31] investigate weight selection strategies in federated learning, focusing on scenarios with few clients and repetitive product images in local datasets to mirror real-world industrial conditions. Their comprehensive evaluation across multiple




Fig. 1. Illustration of the federated learning process [19].

deep learning architectures reveals that Optimal Epoch Weight Selection (OEWS) consistently outperforms other strategies, significantly boosting model performance in industrial applications.

Clustering approaches can be applied in various contexts within FL. Clustering can be differentiated based on the input parameters used, such as training weights, known information like geolocation data, computed metrics such as accuracy, data distribution details, or resource availability. The objectives of clustering can also vary. Some methods aim to form clusters for further training as separate groups, thereby improving local accuracy. Others use clustering to select a diverse set of clients for the next training round, with the goal of developing a more robust global model.

[32] proposes a multi-criterial clustering approach, where the task owner can choose resource-based (such as CPU or RAM) or data-oriented (such as data quality or minimal size) requirements. They show that their DBSCAN-based selection of clients performs better than randomly selecting the same number of clients. One drawback with their approach on the one hand is that they rely on the correctness of sent information from clients regarding resources and data and that it has only been tested at MNIST.

[33] cluster the clients based on the label of the local faults of a computer room and use the Hamming distance to determine similarity. As the Hamming distance describes the number of symbols that distinguish two strings of equal length from each other [34]; this method is only useful if all possible error cases are predefined and displayed in a binary representation. This method is not suitable for recognizing similarities from only the underlying data or weights.

An initial Clustered Federated Learning (CFL) approach was introduced in [35], where clients are iteratively bipartitioned based on the cosine similarity of gradient changes. Hierarchical clustering offers the advantage of not requiring the number of clusters to be known in advance [36], and it also provides scalability and interpretability of the results. However, since multiple communication rounds are needed to fully separate all clients, recursive bipartitioning clustering demands significant computational and communication resources, which can limit its practical feasibility in scenarios with many participants. Additionally, because clients are randomly selected each round, the algorithm can be computationally inefficient and may even completely fail, further restricting its practical applicability [37].

[38] also employs hierarchical clustering to divide clients based on the similarity of gradient changes, but performs the clustering after a predetermined number of communication rounds. The procedure begins by training a global model, which is then fine-tuned on the private datasets of all clients to assess the difference between the global and local model parameters. This difference is used to create clusters and assign clients to them. Subsequent training is conducted independently within each cluster, resulting in the creation of multiple models. This approach is designed to accommodate a wider range of non-IID settings and allows for training on a subset of clients during each round of the FL model training.

To improve the efficiency of CFL, [39] proposes the Iterative Federated Clustering Algorithm (IFCA), which randomly generates cluster centroids and assigns clients to clusters in a way that minimizes the loss function. Although model accuracy with IFCA is significantly improved compared to

conventional FL, the communication costs remain unchanged, and the computational effort is higher, as all updates must be transmitted and clusters reformed in each round. Additionally, the probability of success is highly dependent on the initialization of the cluster centroids.

Cosine similarity followed by affinity propagation clustering has been shown to effectively cluster clients in image classification settings, as demonstrated by [40]. Their evaluation, conducted on the MNIST and CIFAR-10 datasets, included a verification procedure to ensure that the new global weights do not reduce accuracy below a predetermined threshold. If the accuracy drops below this threshold, the new weights are rejected, and training continues with the local weights.

Many purely academic approaches often yield effective solutions on benchmark datasets and demonstrate the convergence of their methods to address specific problems [41]. These problems are typically framed as closed systems, where the primary objective is to optimize specific metrics, such as achieving high accuracy or minimizing the number of communication rounds. However, this perspective frequently overlooks the dynamic nature of real-world systems, which can operate over extended periods. In practice, clients may join or leave the system at any time, or introduce entirely new data, adding layers of complexity that are not fully accounted for in controlled, static environments.

For our further considerations, we will disregard the hardware limitations of the clients, as it is assumed that companies have sufficient computing resources and stable network connections. In this context, optimizing model performance takes precedence over the duration of communication rounds. However, we will ensure that the required computational power remains reasonable for an industrial environment. Additionally, we assume a cross-silo problem definition, characterized by a relatively small number of clients [2], [42]. In this paper, we will focus on examining various client selection techniques to lay the groundwork for improving the integration of new clients in future work. This is particularly important because the careless integration of new clients into a stable system can lead to unwanted performance drops until the system stabilizes again with new clients. Our future work will evaluate how effectively these techniques can incorporate new clients, along with their data and data distributions, into the system.

III. METHODOLOGY

First, we examined how clustering approaches, which were primarily demonstrated on MNIST and minimal examples, can be transferred to real-world datasets and state-of-the-art neural networks, such as EfficientNet or DenseNet. Subsequently, we evaluated various clustering methods and selected the most effective one for further analysis. Initial testing was conducted on a somewhat simpler, proprietary dataset (see section III-A) to establish a baseline understanding, which was then validated on a subset of the more complex ImageNet dataset. This approach allowed us to assess the generalizability and robustness of the selected clustering method across different types of data and neural network architectures.




Fig. 2. Miniature Truck, captured by an inline quality inspection module. Only a cabin with one type of windshield (red) is pictured. In the following step, one trailer is added from a selection of options.

Even relatively small neural networks typically contain millions of trainable parameters. When fine-tuning on custom datasets, it is common practice to adapt only the final layers to the specific data and the required number of output classes. Consequently, analyzing only the final layer's activations is often sufficient for effective clustering, as it captures the most relevant features learned during fine-tuning. In this case, with a shape of $(n, 1280)$ corresponding to n classes in EfficientNet, this approach results in a small sample size with high dimensionality, which presents unique challenges for clustering.

A. Dataset

Our Dataset consists of cabins of miniature trucks, where the parts can be 3D printed, milled, or bought separately. See Figure 2 for two exemplary images from the quality inspection module. As the product is assembled on demand in lot size one, many combinations are possible. Previous publications proved that one advantage of FL is the higher generalizability [43], therefore allowing us to take many pictures away from the production line, where a higher variability of image date is easier to obtain and the risk of interfering with the production process is avoided. For testing purposes, we split the data set so that each client has only one color and one type of windshield, in addition to the class without windshield for every client. The distribution can be obtained from Table I: Classes 1 to 4 are evenly distributed at approximately 12%, whereas class 0 is the predominant category, comprising nearly 50% of the total distribution. In contrast, class 5 is underrepresented, accounting for only about 5%. For some tests, not all clients take part in the federation from the beginning and simulate newcomers that join with new data.

B. Usability Study on Clustering

A pretrained EfficientNet B4 [44] was used as the frozen backbone, with one fully connected layer replaced to match the number of classes. The training was done with SGD as the optimizer with a learning rate of 0.001, 200 communication rounds and 5 local epochs. Using Pytorch, the preliminary examinations for clustering were carried out with our imple-

TABLE I
DISTRIBUTION OF THE TRAINING DATA TO THE CLASSES

Class	0	1	2	3	4	5
CID						
1	565	659				
2	565		574			
3	565			625		
4	565				627	
5	499	499				
6	499		504			
7	499			511		
8	499				499	
9	528	565				
10	528		532			
11	528			559		
12	528				506	
13	510	525				
14	510		530			
15	510			488		
16	510				500	
17	565	607				
18	565		633			
19	565			566		
20	565				599	
21	547					580
22	544					480
Sum	11759	2855	2773	2749	2731	1060
Percent	49,15%	11,93%	11,59%	11,49%	11,41%	4,43%




Fig. 3. Comparison of clustering methods on test accuracy across 100 communication rounds.

mentation to exclude any other influences (e.g. calculating the global model manually based on client's weight files).

To determine which clustering method is most appropriate, three methods that are well-suited for a large number of data points and clusters were selected: Agglomerative clustering [45], HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise) [46] and BIRCH (Balanced Iterative Reducing and Clustering using Hierarchies) [47]. Figure 3 presents sample results from a larger set of tests that were conducted. The optimal parameters, such as the distance metric, for each method were determined through grid search on the weights from previous runs in which all clients participated. We observe that with agglomerative

clustering, accuracy increases rapidly, particularly during the first 10 rounds, and then remains stable while continuing to improve gradually. All other methods exhibit significantly higher fluctuations and, except for BIRCH ($n=5$), even achieve lower accuracy than agglomerative clustering. It should be noted that BIRCH occasionally even clearly surpasses it. The high fluctuation in BIRCH's performance could be due to its sensitivity to parameters and the high dimensionality combined with a small sample size. BIRCH is generally designed for larger datasets, and when applied to a small dataset with high dimensional features, it may struggle to form stable clusters, leading to the observed variability in test accuracy. HDBSCAN consistently performed worse in comparison.

Agglomerative clustering performs best in our scenario due to its ability to capture complex hierarchical structures in relatively small, but high-dimensional dataset. Its flexibility in distance metrics, robustness to small sample sizes, and stability in high-dimensional spaces make it well-suited to effectively clustering the final layer activations of your neural network, leading to more accurate and stable results compared to other methods like BIRCH or HDBSCAN.

A more detailed evaluation was carried out using saved weights from two different scenarios: one in which all clients participated in every communication round, and another where the experiment began with 8 clients (representing all classes but limited to red and blue colors), gradually adding new clients after an initial training phase. Clusters were recalculated after each communication round. The following observations were made during the experiments: New clients were immediately recognized as newcomers, allowing for precautions to be taken to prevent any adverse effects or using advanced approaches to integrate them effectively. Within a maximum of three communication rounds, each consisting of five local epochs, new clients were correctly assigned to the appropriate cluster. It was also evident that cabin types 1 and 4 were easily distinguishable, while types 2 and 3 were often grouped together into the same cluster or formed mixed clusters. Agglomerative clustering with complete linkage yielded good results, as did the BIRCH algorithm. However, the resulting clusters sometimes differed; for instance, BIRCH occasionally produced clusters more strongly correlated with cabin colors. In subsequent communication rounds, the clusters realigned with windshield types. Additionally, one group exhibited a higher loss than the others, with no further improvement. Upon closer inspection of the client IDs, it became clear that the group with the higher loss consisted entirely of clients that trained with type 4 windshields. An analysis of the weights confirmed this clustering. Depending on the use case and objectives, it may be beneficial either to develop multiple models tailored to these clusters or to design the client selection process in a way that disrupts this separation, such as by consistently selecting a client from each cluster.

C. Evaluation of Client Selection Strategies

In order to compare different strategies, the clients as listed in Table II were selected for training and the data of

the excluded clients were used for testing, whereby random images were selected to create a similar distribution of classes, e.g. to avoid over-representing class 0. The total number of clients was 22, 16 clients correspond to the setup shown in Table II. In order to simulate the complete absence of a class, the last client with class 5 was omitted from the training in some tests, resulting in 15 clients. The configuration for each run was: EfficientNet with SGD as an optimizer with learning rate 0.001, 200 communication rounds with 5 local epochs each. As a baseline, random sampling of approximately twenty and fifty percent (rounding up to the next integer number) of all participating clients was done. Between a sampling rate of 50% and 20% of all 22 clients, there is no noticeable difference, supporting the findings of [18] that saturation occurs at a certain point, and additional client participation no longer results in improvement.

TABLE II

OVERVIEW OF CLIENT SELECTION FOR SAMPLING TESTS, INDICATING THE CLASSES PRESENT IN EACH CLIENT'S DATASET. EACH 'X' REPRESENTS A CLASS THE CLIENT HAS, WHILE EMPTY CELLS INDICATE MISSING CLASSES. MISSING CLASSES WITHIN EACH COLOR-CODED GROUP ARE HIGHLIGHTED.

Client ID	Color	0	1	2	3	4	5
1	blue	x	x				
2	blue	x			x		
3	blue	x				x	
4	red	x	x				
5	red	x			x		
6	red	x				x	
7	green	x	x				
8	green	x			x		
9	green	x				x	
10	yellow	x		x			
11	yellow	x			x		
12	yellow	x				x	
13	orange	x	x				
14	orange	x			x		
15	orange	x				x	
16	silver	x					x
Sum		16	4	3	4	4	1

In comparison to random selection, we evaluate metric-based methods (specifically, highest training loss) and cluster-based methods (focussing on agglomerative clustering). For the initialization of both methods, each client is selected once at the beginning. Figure 4 shows the impact the methods have on client participation. A review of the selection based on the training loss reveals that client 16, with the underrepresented class, is selected with greater frequency. However, client 1, is selected with a similarly high frequency, although class 1 is represented at an average level.

Figure 5 illustrates the test accuracy over 200 communication rounds for four client selection methods: Random selection with 5 clients participating each round, loss-based selection of the 5 clients with the highest training loss, and agglomerative clustering with $k=3$ and $k=5$, corresponding to the number of clusters to be identified. The moving average with windows size 5 is used for clearer visualization. Agglomerative clustering with $k=5$ consistently outperforms the other meth-

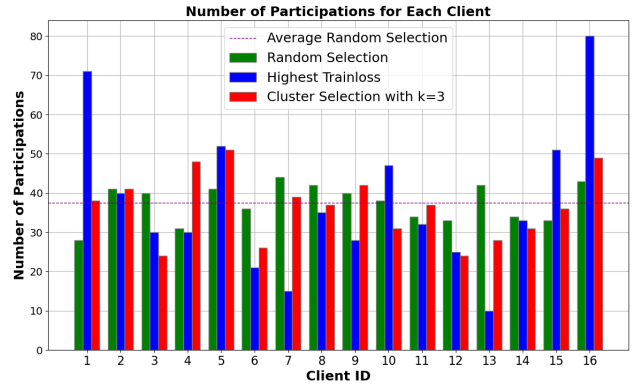


Fig. 4. Comparison of client participation frequencies across different selection methods: The chart displays the number of participations for each client using three selection strategies: Random selection, highest training loss and cluster selection with $k=3$.




Fig. 5. The plot shows the moving average of test accuracy across 200 communication rounds for three client selection methods: Random selection, training loss-based selection, and agglomerative clustering, each with their corresponding parameters.

ods, showing the highest and most stable accuracy throughout. Initially, it demonstrates a rapid increase in accuracy, maintaining superior performance across all phases. Agglomerative clustering with $k=3$ and random selection yield similar results, with $k=3$ exhibiting more fluctuations but eventually aligning with random selection. The training loss method, despite early instability, eventually achieves comparable accuracy, though with more fluctuations. Overall, clustering with $k=5$ offers the greatest accuracy and stability, while non-randomized methods provide better and more controlled integration of new clients.

To ensure that cluster-based client selection is effective beyond our dataset, we conducted further experiments using a subset of ImageNet. This subset includes three superclasses: animals, tools, and vehicles, each containing multiple subclasses (for instance, various classes for cats, dogs, and birds under the animals' superclass). Throughout all runs, the mapping of clients to datasets remained consistent, as did the




Fig. 6. The plot shows the moving average of test accuracy on ImageNet across 100 communication rounds for three client selection methods: Random selection, training loss-based selection, and agglomerative clustering, each with their corresponding parameters.

randomization seed, such as when adding noise to the initially loaded ImageNet weights. The results, shown in Figure 6, demonstrate that clustering strategies with any amount of clusters outperform the random and training loss-based strategies, particularly during the first 25 communication rounds. This analysis was conducted using the same hyperparameters as before, highlighting the effectiveness of clustering approaches in improving model accuracy, especially in the early stages of training, but also proving that ImageNet is more complex and needs more fine-tuning to achieve a higher accuracy. Note, that clustering with $k=3$ outperforms random selection with $n=5$, meaning that in each communication round, two fewer clients are required to compute their updates, leading to a more efficient path to achieving high accuracy. It is important to highlight that the images within each superclass were re-clustered based on visual similarity using t-distributed stochastic neighbor embedding, leading to a non-IID distribution.

D. Computational Overhead

Finally, we measured the client selection time on a local computer, as consistent hardware configurations cannot be guaranteed throughout all tests when using a high-performance cluster. When selecting 10 out of 16 clients for training—already a higher ratio compared to our other setups—the results indicate that agglomerative clustering is only slightly slower than random selection, with selection times of 0.0049 seconds and 0.0027 seconds, respectively. However, HDBSCAN requires significantly more time, taking 0.0444 seconds. Despite this, the communication time for transmitting weights in a network will far exceed these client selection times, especially in a cross-silo setup with a relatively low number of clients making the additional computation time for clustering methods less impactful in the overall process.

IV. CONCLUSION AND OUTLOOK

In this study, we examined various client selection methods within the context of federated learning to address the chal-

lenges of integrating new clients with diverse and potentially non-IID data distributions. Our analysis demonstrated that clustering techniques, particularly agglomerative clustering, consistently outperformed random and loss-based client selection methods, offering both higher and more stable accuracy across communication rounds. This improvement is especially pronounced in the early stages of training, where efficient client selection plays a crucial role in accelerating the learning process while maintaining high model performance. The success of clustering-based methods underscores their potential in managing the complexities of data heterogeneity and enhancing the overall efficiency of FL systems.

Future research will focus on further refining clustering-based client selection strategies to better accommodate the dynamic nature of real-world FL environments, where clients may join or leave the system unpredictably. Additionally, the consideration of complex datasets like ImageNet, which has only been touched on in this study, is explored in greater depth in ongoing research and will be crucial to ensure their applicability across various domains. Another promising direction is the development of hybrid approaches that combine the strengths of different selection methods to optimize both accuracy and computational efficiency. As FL continues to evolve, addressing these challenges will be key to its successful adoption in production environments, particularly in industries where data privacy and diversity are paramount.

ACKNOWLEDGMENT

This work was funded by the Carl Zeiss Stiftung, Germany under the Sustainable Embedded AI project (P2021-02-009).

REFERENCES

- [1] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," *Artificial Intelligence and Statistics*, pp. 1273–1282, 2017.
- [2] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. Nitin Bhagoji, and e. Bonawitz, "Advances and open problems in federated learning," *FNT in Machine Learning (Foundations and Trends in Machine Learning)*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [3] I. K. Nti, A. F. Adekoya, B. A. Weyori, and O. Nyarko-Boateng, "Applications of artificial intelligence in engineering and manufacturing: a systematic review," *Journal of Intelligent Manufacturing*, vol. 33, no. 6, pp. 1581–1601, 2022.
- [4] A. Dogan and D. Birant, "Machine learning and data mining in manufacturing," *Expert Systems with Applications*, vol. 166, p. 114060, 2021.
- [5] M. Arafteh, A. Hammoud, H. Otrok, A. Mourad, C. Talhi, and Z. Dziong, "Independent and identically distributed (iid) data assessment in federated learning," in *GLOBECOM 2022 - 2022 IEEE Global Communications Conference*. IEEE, 2022, pp. 293–298.
- [6] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in *Proceedings of Machine Learning and Systems*, I. Dhillon, D. Papailiopoulos, and V. Sze, Eds., vol. 2, 2020, pp. 429–450.
- [7] O. Myklebust, "Zero defect manufacturing: A product and plant oriented lifecycle approach," *Procedia CIRP*, vol. 12, pp. 246–251, 2013.
- [8] M. K. Hassan, "Applying lean six sigma for waste reduction in a manufacturing environment," *American Journal of Industrial Engineering*, vol. 1, no. 2, pp. 28–35, 2013.
- [9] T. Calvin, "Quality control techniques for zero defects," *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*, vol. 6, no. 3, pp. 323–328, 1983.

- [10] W. L. Berry and M. C. Cooper, "Manufacturing flexibility: methods for measuring the impact of product variety on performance in process industries," *Journal of Operations Management*, vol. 17, no. 2, pp. 163–178, 1999.
- [11] R. Rai, M. K. Tiwari, D. Ivanov, and A. Dolgui, "Machine learning in manufacturing and industry 4.0 applications," *International Journal of Production Research*, vol. 59, no. 16, pp. 4773–4778, 2021.
- [12] V. Buggineni, C. Chen, and J. Camelio, "Enhancing manufacturing operations with synthetic data: a systematic framework for data generation, accuracy, and utility," *Frontiers in Manufacturing Technology*, vol. 4, 2024.
- [13] Z. Adams and A. Nelson, "Breaking down silos: Integrating big data across organizational functions," *International Journal of Network and Communication Research*, vol. 7, no. 1, pp. 81–89, 2023.
- [14] V. Hegiste, S. Walunj, J. Antony, T. Legler, and M. Ruskowski, "Enhancing object detection with hybrid dataset in manufacturing environments: Comparing federated learning to conventional techniques," *1st International Conference on Innovative Engineering Sciences and Technological Research (ICIESTR-2024)*, 2024.
- [15] L. Liu, J. Zhang, S. H. Song, and K. B. Letaief, "Client-edge-cloud hierarchical federated learning," in *ICC 2020 - 2020 IEEE International Conference on Communications (ICC)*. IEEE, 2020, pp. 1–6.
- [16] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, M. A. Ranzato, A. Senior, P. Tucker, K. Yang, Q. Le, and A. Ng, "Large scale distributed deep networks," in *Advances in Neural Information Processing Systems*, F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger, Eds., vol. 25. Curran Associates, Inc, 2012.
- [17] F. N. Iandola, M. W. Moskewicz, K. Ashraf, and K. Keutzer, "Firecaffe: Near-linear acceleration of deep neural network training on compute clusters," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2016.
- [18] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Singh and J. Zhu, Eds., vol. 54. Fort Lauderdale, FL, USA: PMLR, 2017, pp. 1273–1282.
- [19] T. Legler, V. Hegiste, and M. Ruskowski, "Mapping of newcomer clients in federated learning based on activation strength," *32nd International Conference Flexible Automation and Intelligent Manufacturing*, 2023.
- [20] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of fedavg on non-iid data," in *International Conference on Learning Representations*, 2019.
- [21] Y. Jee Cho, S. Gupta, G. Joshi, and O. Yagan, "Bandit-based communication-efficient client selection strategies for federated learning," in *2020 54th Asilomar Conference on Signals, Systems, and Computers*. Place of publication not identified: IEEE, 2020, pp. 1066–1069.
- [22] F. Maciel, A. M. de Souza, L. F. Bittencourt, and L. A. Villas, "Resource aware client selection for federated learning in iot scenarios," in *19th International Conference on Distributed Computing in Smart Systems and the Internet of Things*. IEEE, 2023, pp. 1–8.
- [23] A. Tahir, Y. Chen, and P. Nilayam, "Fedsss: Federated learning with smart selection of clients," 11.07.2022.
- [24] T. Nishio and R. Yonetani, "Client selection for federated learning with heterogeneous resources in mobile edge," in *2019 IEEE International Conference on Communications (ICC)*. IEEE, 2019, pp. 1–7.
- [25] M. Panigrahi, S. Bharti, and A. Sharma, "A review on client selection models in federated learning," *WIREs Data Mining and Knowledge Discovery*, vol. 13, no. 6, 2023.
- [26] Y. J. Cho, J. Wang, and G. Joshi, "Client selection in federated learning: Convergence analysis and power-of-choice selection strategies," 03.10.2020.
- [27] W. Lin, Y. Xu, B. Liu, D. Li, T. Huang, and F. Shi, "Contribution-based federated learning client selection," *International Journal of Intelligent Systems*, vol. 37, no. 10, pp. 7235–7260, 2022.
- [28] Y. Qiao, M. S. Munir, A. Adhikary, A. D. Raha, and C. S. Hong, "Cdfed: Contribution-based dynamic federated learning for managing system and statistical heterogeneity," in *NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2023, pp. 1–5.
- [29] M. Jiang, H. R. Roth, W. Li, D. Yang, C. Zhao, V. Nath, D. Xu, Q. Dou, and Z. Xu, "Fair federated medical image segmentation via client contribution estimation," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Piscataway, NJ: IEEE, 2023, pp. 16 302–16 311.
- [30] Y. Ruan, X. Zhang, S.-C. Liang, and C. Joe-Wong, "Towards flexible device participation in federated learning," in *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Banerjee and K. Fukumizu, Eds., vol. 130. PMLR, 2021, pp. 3403–3411.
- [31] V. Hegiste, T. Legler, and M. Ruskowski, "Towards robust federated image classification: An empirical study of weight selection strategies in manufacturing." [Online]. Available: <http://arxiv.org/pdf/2408.10024v1>
- [32] S. Arisdakessian, O. A. Wahab, A. Mourad, and H. Otrok, "Towards instant clustering approach for federated learning client selection," in *2023 International Conference on Computing, Networking and Communications (ICNC)*. IEEE, 2023, pp. 409–413.
- [33] C. Li and H. Wu, "Fedcls:a federated learning client selection algorithm based on cluster label information," in *2022 IEEE 96th Vehicular Technology Conference:(VTC 2022-Fall)*. Piscataway, NJ: IEEE, 2022, pp. 1–5.
- [34] R. W. Hamming, *Coding and information theory*. Prentice-Hall, Inc, 1986.
- [35] F. Sattler, K.-R. Muller, and W. Samek, "Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 8, pp. 3710–3722, 2021.
- [36] S. C. Johnson, "Hierarchical clustering schemes," *Psychometrika*, vol. 32, no. 3, pp. 241–254, 1967.
- [37] M. Duan, D. Liu, X. Ji, R. Liu, L. Liang, X. Chen, and Y. Tan, "Fedgroup: Efficient federated learning via decomposed similarity-based clustering," in *2021 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking*. IEEE, 2021, pp. 228–237.
- [38] C. Briggs, Z. Fan, and P. Andras, "Federated learning with hierarchical clustering of local updates to improve training on non-iid data," in *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2020, pp. 1–9.
- [39] A. Ghosh, J. Chung, D. Yin, and K. Ramchandran, "An efficient framework for clustered federated learning," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc, 2020, pp. 19 586–19 597.
- [40] P. Tian, W. Liao, W. Yu, and E. Blasch, "Wssc: A weight-similarity-based client clustering approach for non-iid federated learning," *IEEE Internet of Things Journal*, vol. 9, no. 20, pp. 20 243–20 256, 2022.
- [41] D. J. Beutel, T. Topal, A. Mathur, X. Qiu, J. Fernandez-Marques, Y. Gao, L. Sani, K. H. Li, T. Parcollet, P. P. B. d. Gusmão, and N. D. Lane, "Flower: A friendly federated learning research framework." [Online]. Available: <http://arxiv.org/pdf/2007.14390v5>
- [42] Y. Huang, L. Chu, Z. Zhou, L. Wang, J. Liu, J. Pei, and Y. Zhang, "Personalized cross-silo federated learning on non-iid data," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 9, pp. 7865–7873, 2021.
- [43] V. Hegiste, T. Legler, and M. Ruskowski, "Application of federated machine learning in manufacturing," in *2022 International Conference on Industry 4.0 Technology (I4Tech)*. IEEE, 2022, pp. 1–8.
- [44] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 2019, pp. 6105–6114.
- [45] K. Chidananda Gowda and G. Krishna, "Agglomerative clustering using the concept of mutual nearest neighbourhood," *Pattern Recognition*, vol. 10, no. 2, pp. 105–112, 1978.
- [46] L. McInnes, J. Healy, and S. Astels, "hdbscan: Hierarchical density based clustering," *The Journal of Open Source Software*, vol. 2, no. 11, p. 205, 2017.
- [47] T. Zhang, R. Ramakrishnan, and M. Livny, "Birch: A new data clustering algorithm and its applications," *Data Mining and Knowledge Discovery*, vol. 1, no. 2, pp. 141–182, 1997.