

Federated Edge Learning: Design Issues and Challenges

Afaf Tak and Soumaya Cherkaoui

ABSTRACT

Federated Learning (FL) is a distributed machine learning technique, where each device contributes to the learning model by independently computing the gradient based on its local training data. It has recently become a hot research topic, as it promises several benefits related to data privacy and scalability. However, implementing FL at the network edge is challenging due to system and data heterogeneity and resource constraints. In this article, we examine the existing challenges and trade-offs in Federated Edge Learning (FEEL). The design of FEEL algorithms for resources-efficient learning raises several challenges. These challenges are essentially related to the multidisciplinary nature of the problem. As the data is the key component of the learning, this article advocates a new set of considerations for data characteristics in wireless scheduling algorithms in FEEL. Hence, we propose a general framework for data-aware scheduling as a guideline for future research directions. We also discuss the main axes and requirements for data evaluation and some exploitable techniques and metrics.

INTRODUCTION

The growing interest in intelligent services motivates the integration of artificial intelligence (AI) in Internet of Things (IoT) applications. The collection of large volumes from the different devices and sensors is necessary for training AI models. However, uploading massive data generated by connected devices to the cloud is usually impractical, mainly due to issues including privacy, network congestion, and latency. Federated Edge Learning (FEEL) [1] is a Machine Learning (ML) setting that utilizes edge computing to tackle these concerns. In contrast to centralized ML, Federated Learning (FL) [2] consists of training the model on the devices, with the orchestration of a central entity, where only the resultant model parameters are sent to the edge servers to be aggregated. FEEL refers to the use of FL at the edge of the network, which makes it a promising solution for privacy preserving ML.

An important design decision for a FEEL algorithm is whether to choose either asynchronous or synchronous aggregation. Recent works tend to promote synchronous training, where, for instance, synchronization among participating devices is required for update averaging [2] and privacy preservation [3]. However, there are many challenges upon using synchronous FL in edge environments.

To begin with, the heterogeneity of resources across different devices sparks new system challenges. For instance, significant delays can be caused by stragglers. Moreover, communication loads across devices limit the scalability of FL for large models. Participating devices communicate full model updates during every training iteration, which are of the same size as the trained model. For large models, such as deep neural networks, the model size can be in the range of giga bytes. As a result, if communication bandwidth is limited or communication is costly, FEEL can be deemed impractical or unfeasible, as communication overhead becomes a bottleneck for FEEL.

Furthermore, end devices have limited battery lives and varying available energy levels. As training ML models is a computation-heavy task, only devices that have enough energy can be solicited to participate. Furthermore, energy and computational constraints limit both the size of the models that can be trained on-device, and the number of local training iterations.

Additionally, as the data collected by the clients depends on their local environment and usage pattern, both the size and the distribution of the local datasets will typically vary between different clients. This non-Independently and Identically Distributed (non-IID) and unbalanced nature of data across the network imposes significant challenges linked to models' convergence.

Consequently, designing an efficient FEEL algorithm should take into account the limited and heterogeneous nature of the resources, alongside the non-IID and unbalanced aspect of the data distributions. In general, proposed FEEL algorithms target efficient selection of participant devices, optimization of the resource allocation and usage, or adequate updates' aggregation. However, it is hard to capture both the resources problems and the learning goal, as there is no direct relation between the model's loss function and resource optimization. A manageable approach found in current works is to focus on resource optimization with certain learning guarantees, such as maximizing the number of collected updates and maintaining the level of local accuracy [4]. Nonetheless, these guarantees are not sufficient, as a significant drop in accuracy is observed when data is non-IID and unbalanced. Therefore, we propose to lighten the effects of design trade-offs through the direct integration of the data properties in the device selection and resource optimization algorithms. In fact, data properties were at the heart of FL since its inception, but they have been largely overlooked in the design of

FEEL algorithms. Moreover, data diversity has long been premised on in active learning, where models can be trained using few labelled data samples if the highly diverse data is selectively added to the training set. Thus, data diversity should be considered in the design of FEEL algorithms, as we advocate in this article.

The main contributions of this article can be summarized as follows:

- We discuss the FEEL challenges imposed by the nature of the edge environment, from an algorithm design perspective. We review the challenges related to computational and communication capacities, as well as data properties, as they are at the core of the trade-offs in learning and resource optimization algorithms.
- We propose a general framework for incorporating data properties in FEEL, by providing a guideline for a thorough algorithm design, and criteria for the choice of diversity measures in both datasets and models.
- We present several possible measures and techniques to evaluate data and model diversity, which can be applied in different scenarios (e.g., classification, time series forecasting), in an effort to assist fellow researchers to further address FEEL challenges.

The remainder of this article is as follows. In the following section, we review the challenges found in designing FEEL algorithms, and we derive the main trade-offs. Then, we shed light on a new data-aware design direction for FEEL algorithms. Some possible techniques and methods to evaluate diversity are then detailed. In the final section, the conclusion and final remarks are presented.

DESIGN CHALLENGES: OVERVIEW

FEEL has several constraints related to the nature of the edge environment. In fact, FEEL involves the participation of heterogeneous devices that have different computation and communication capabilities, energy states, and dataset characteristics. Under device and data heterogeneity, in addition to resource constraints, participants selection and resource allocation have to be optimized for an efficient FEEL solution.

DESIGN CHALLENGES

The core challenges associated with solving the distributed optimization problem are twofold: resources and data. These challenges increase the FEEL setting complexity compared to similar problems, such as distributed learning in data centers.

Resources: The challenges related to the resources, namely computation, storage and communication, are mainly in terms of their heterogeneity and scarcity.

Heterogeneity of the Resources: Computation, storage and communication capabilities vary from one device to another. Devices may be equipped with different hardware (CPU and memory), network connectivity (e.g., 4G/5G, WiFi), and may differ in available power (battery level). The gap in computational resources creates challenges such as delays caused by stragglers. FEEL algorithms must therefore be adaptive to the heterogeneous hardware and be tolerant toward device drop-out and low or partial participation. A potential solution to the straggler problem is asynchronous

FEEL has several constraints related to the nature of the edge environment. In fact, FEEL involves the participation of heterogeneous devices that have different computation and communication capabilities, energy states, and dataset characteristics.

learning. However, the reliability of asynchronous FL and the model convergence in this setting are not always guaranteed. Thus, synchronous FL remains the preferred approach.

Limited Resources: In a contrast to the cloud, the computing and storage resources of the devices are very limited. Therefore, the models that can be trained on-device are relatively simpler and smaller than the models trained on the cloud. Furthermore, devices are frequently offline or unavailable either due to low battery levels, or because their resources are fully or partially used by other applications.

As for the communication resources, the available bandwidth is limited. It is therefore important to develop communication-efficient methods that make it possible to send compressed or partial model updates. To further reduce communication cost in FEEL settings, two potential directions are generally considered: reducing the total number of communication rounds until convergence [5], and reducing the size of the transmitted updates through compression and partial updates [2].

Data: In most cases, data distributions depend on the users' behavior [6]. As a result, the local datasets are massively distributed, statistically heterogeneous (i.e., non-IID and unbalanced), and highly redundant. Additionally, the raw generated data is often privacy-sensitive as it can reveal personal and confidential information.

Small and Widely Distributed Datasets: In FEEL scenarios, a large number of devices participate in the FL training with a small average number of data samples per client. Learning from small datasets makes local models prone to overfitting.

Non-IID: The training data on a given device is typically based on the usage of the device by a particular user, and hence any particular user's local dataset will not be representative of the population distribution. This data-generation paradigm fails to comply with the independent and identically distributed (IID) assumptions in distributed optimization, and thus adds complexity to the problem formulation and convergence analysis. The empirical evaluation of FEEL algorithms on non-IID data is usually performed on artificial partitions of MNIST or CIFAR-10, which do not provide a realistic model of a federated scenario.

Unbalance: Similar to the nature of the distributions, the size of the generated data depends on the user. Depending on users' use of the device, these may have varying amounts of local training data.

Redundancy: The unbalance of the data is also observed within the local datasets at a single device. In fact, IoT data is highly redundant. In sequential data (e.g., video surveillance, sensors data) for instance, only a subset of the data is informative or useful for the training.

Privacy: The privacy-preserving aspect is an essential requirement in FL applications. The raw data generated on each device is protected by sharing only model updates instead of the raw

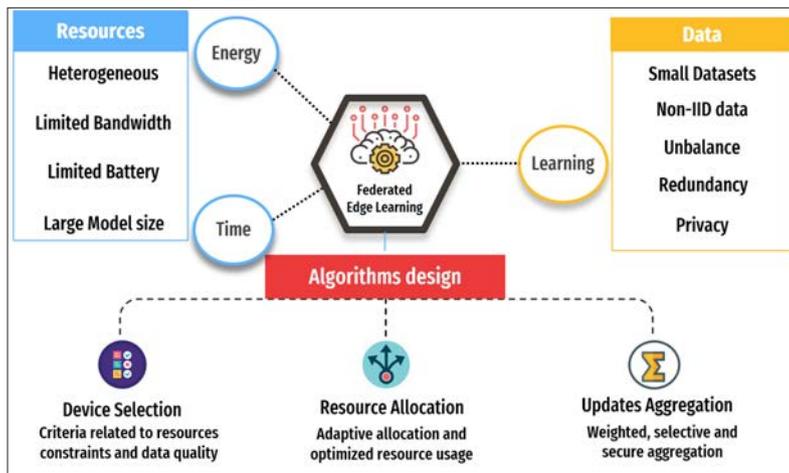


FIGURE 1. FEEL algorithms, challenges and optimization axes.

data. However, communicating model updates throughout the training process can still be reverse-engineered to reveal sensitive information, either by a third-party or a malicious central server.

DESIGN TRADE-OFFS

Several efforts were made to tackle the aforementioned challenges. However, FEEL is a multi-dimensional problem that brings about several trade-offs. As a result, algorithms designed to address one issue at a time are deemed impractical. Perhaps a tractable solution may be to combine several techniques when developing and deploying FEEL algorithms.

In general, an end-to-end FEEL solution should cover device selection, resource allocation, and update aggregation. In the following, we discuss major trade-offs that should be considered when designing solutions in the FEEL setting.

General FEEL Solution: Given the wide range of applications that can benefit from FEEL, there is no one-size-fits-all solution. However, in general, a FEEL solution needs to act on the following aspects:

Device Selection: Participant selection refers to the selection of devices to receive and train the model in each training round. Ideally, a set of participants is randomly selected by the server to participate. Then, the server has to aggregate parameter updates from all participants in the round before taking a weighted average of the models. However, due to the communication bottlenecks and the desire to tame the training latency, the device selection should be optimized in terms of resources [7] and data criteria.

Resource Allocation: Device selection should not be considered independently from resource allocation, especially computation and bandwidth. We refer to joint selection and resource allocation as a scheduling algorithm. Indeed, the number of scheduled devices is limited by the available bandwidth that can be allocated. Additionally, for an optimal learning round duration and energy consumption, both bandwidth and computation resources should be adapted based on the number of local iterations at each device, and the number of global iterations (i.e., learning rounds) [8]. Due to the fast-changing aspect of the FEEL environment, the computational complexity of

scheduling algorithms should be especially low. Therefore, the use of meta-heuristics and heuristics should be encouraged.

Updates Aggregation: This aspect of the solution design refers to how the updates are aggregated and how frequently they are aggregated. For instance, the frequency of the communication and aggregation can be reduced with more local computation [8], or reduced through selective communication of gradients [5]. For instance, FedAvg [9] is one of the most used methods in aggregation which uses weighted average Stochastic Gradient Descent updates, where the corresponding weights are decided by the volume of the training dataset. While FedAvg uses synchronous aggregation, in the FedAsync [10] algorithm, newly received local updates are weighted according to their staleness, where stale updates received from stragglers are weighted less based on how many rounds elapsed. It should also be noted that proposing new aggregation methods requires theoretical and empirical convergence analysis to guarantee that the learning loss function will converge to a global optimum. Update aggregation should also be communication-efficient [2, 5] and secure by the means of techniques such as differential privacy [3].

Optimization Axes: In addressing FEEL challenges, three optimization axes are often considered: time, energy and learning. In many cases, the FEEL algorithm can be viewed as a Pareto optimal problem [11]. The relation between the three axes and the challenges is illustrated in Fig. 1.

Time Optimization: Accelerating the learning time can be evaluated with different lenses: learning round duration and time until learning convergence. Due to the synchronous model aggregation of FEEL, the total duration of a round is determined by the slowest device among all the scheduled devices [12]. For this reason, more bandwidth should be allocated for transmission by stragglers and less for faster devices. This to some extent can equalize their total update time (computing plus communication time). Furthermore, to avoid squandering bandwidth on extremely slow devices, scheduling (i.e., joint selection and resource allocation) should exclude slowest devices by applying thresholds on their expected completion time, which can be inferred using their computing capacities and channel states. From a learning perspective, the learning latency is determined by the number of rounds until convergence. The optimization techniques centered on this aspect mainly focus either on the selective upload of updates, or on maximizing the participating devices in each round.

Energy Optimization: Optimizing energy consumption across the network is necessary to reduce the rate of drop-out devices because of battery drainage. In fact, training and transmission of large-scale models are energy consuming, while most edge and end devices have limited battery lives. Additionally, using the maximum capacity of the devices would likely make the users less willing to participate in the training. A design goal of a scheduling algorithm (i.e., joint selection and resource allocation) would be to allocate bandwidth based on the devices' channel states and battery levels. As a result, more bandwidth should be allocated to devices with weaker

channels or poorer power states, to maximize the collected updates [4].

Learning Optimization: In contrast to centralized learning, optimizing learning in the FEEL setting cannot be seen independently from time and energy optimization. However, capturing the optimization of time, energy and the learning goal in the same optimization problem is hard, because there is no direct relation between the objective function of the learning (i.e., the loss function) and the time and energy minimization goal. A manageable approach used is to minimize time and energy under a certain convergence speed guarantee. For instance, some works argue that the number of collected updates in each round is inversely proportional to the convergence speed, and therefore is used as a guarantee [4]. Indeed, multi-user diversity (i.e., collecting the maximum amount of updates) can yield a high convergence speed, especially in IID environments; however, there is a significant chance of choosing the same sets of devices repeatedly. To avoid this issue, a goal of the FEEL algorithm can be to maximize fairness in terms of the number of collected updates among devices [13]. The fairness measure maximizes the chance of more diverse data sources, thus achieving gradient diversity. Nonetheless, the number of collected updates in this setting might be low. Fairness is also considered in the aggregation by q-Fair FL (q-FFL) [14], which reweights the objective function in FedAvg to assign higher weights in the loss function to devices with higher loss. Another approach is to use data size priority, which maximizes the size of data used in the training, by using a probability of selection inversely proportional to the available dataset's size. In the background, these scheduling algorithms all share the same idea: if the size of the training data is large, then the training would converge faster. However, IoT data is highly redundant and inherently unbalanced. Thus, many of the proposed algorithms witness a drop in performance in non-IID and unbalanced experiments. Therefore, the data properties should be considered throughout the FEEL algorithm.

DATA-AWARE FEEL DESIGN: FUTURE DIRECTION

Even if FL was first proposed with data as a central aspect, it has been overlooked in the design of proposed FEEL scheduling algorithms. With the significant drop of accuracy of models trained with resource-aware FEEL algorithms in non-IID and unbalanced settings, it becomes clear that the data aspect should be considered. Henceforth, we propose a new possible data-aware end-to-end FEEL solution based on the diversity properties of the different datasets. In general, diversity consists of two aspects, namely, richness and uncertainty. Richness quantifies the size of the data, while the uncertainty quantifies the information contained in the data. In fact, it has been long proven in Active Learning that by choosing highly uncertain data samples, a model can be trained using fewer labelled data samples. This fact suggests that data uncertainty should be incorporated into the design of FEEL scheduling algorithms. Nonetheless, the uncertainty measures used in Active Learning targets individual samples from unlabeled data in a centralized setting, thus, these measures cannot be directly integrated in FEEL. In

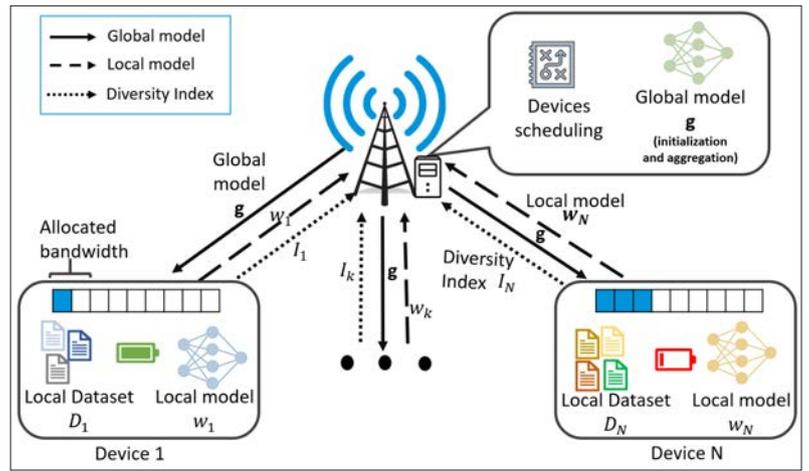


FIGURE 2. The proposed FEEL system model.

the FEEL setting, the updates' scheduling can be either before the training or after it, therefore the diversity measures should be selected depending on the time of scheduling. If the scheduling before the training is preferred, then the datasets' diversity is to be considered. Otherwise, if the scheduling is set after the training is over, the diversity to be considered is model diversity, as the diversity of the dataset can be reflected by the resulting model. In both cases, in addition to maximizing diversity through careful selection of participating devices, the scheduling algorithm can focus on minimizing the consumed resources in terms of completion time of FL and transmission energy of participating devices. For the pre-training scheduling, local computation energy can also be optimized. Furthermore, the scheduling problems' constraints are to be derived from the environment's properties concerning resources and data.

In this section, and to better illustrate the data-aware solutions, we consider the architecture illustrated in Fig. 2. The architecture is a cellular network composed of one base station (BS) equipped with a parameter server, and N devices that collaboratively train a shared model. In the following, we discuss different constraints related to the scheduling algorithms in this setting. Then, we present pre-training and post-training algorithm guidelines, where we detail the key criteria for the design of data-aware FEEL solutions, and we present some potential measures and methods to enable a variety of data-aware FEEL applications, which are summarized in Fig. 3.

SCHEDULING CONSTRAINTS

The scheduling algorithms must consider the following constraints that arise from the FEEL environment's properties.

Energy Consumption: Due to the limited energy level and the high computational requirements of training algorithms, it is necessary to evaluate a device's battery level before scheduling it for a training round. When first FL was proposed, the selected devices were limited to the ones plugged for charging. However, this criterion limits the number of devices that can be selected, leading to a slow convergence of the learning.

Radio Channel State: It is important to consider the radio channel state changes in the sched-

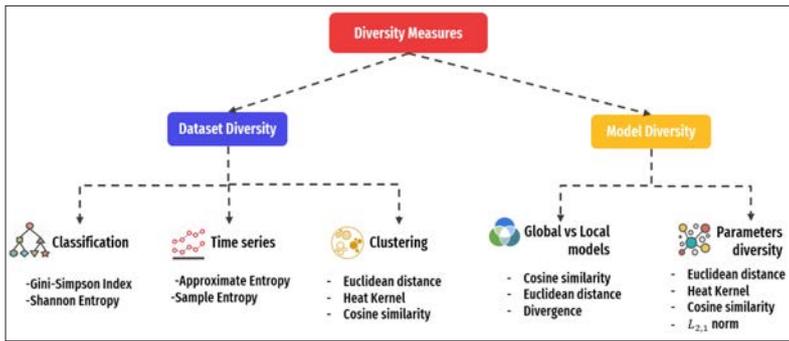


FIGURE 3. Diversity measures that can be used in pre-training and post-training scheduling.

uling. The quality of the communication is critical for both the device selection and resource allocation.

Expected Completion Time: The available computation resources, alongside data size, can be used to estimate the completion time of the device. Potential stragglers can be discarded even before the training process.

Number of Participants: A communication round cannot be considered valid unless a minimum number of updates is obtained. Therefore, a training round can be dropped if there are not enough devices to schedule.

Data Size: The available data in the device is smaller in size than a required minimum, it can be immediately discarded from the selection process. For instance, if the number of samples is less than the selected mini-batch size, the device should be excluded.

PRE-TRAINING SCHEDULING: DATASET DIVERSITY

The pre-training scheduling that we propose uses dataset diversity to choose devices that will conduct the training and send the updates. Scheduling the devices before the training makes it possible to eliminate potential stragglers, and adapt the number of epochs based on the battery levels available at the participating devices.

Scheduling Algorithm: In this algorithm, the global model is initialized by the BS. Afterwards, the following steps are repeated until the model converges or a maximum of rounds is attained:

Step 1: At the beginning of each training round, the devices send their diversity indicators and battery levels to the server.

Step 2: Based on the received information, alongside the evaluated channel state indicator, the server schedules a subset of devices and sends them the current global model.

Step 3: Each device in the subset uses its local data to train the model.

Step 4: The updated models are sent to the server to be aggregated.

Step 5: The PS aggregates the updates and creates the new model.

Datasets Diversity Measures: In the pre-training scheduling, dataset diversity will serve essentially as a lead for device selection, where it should prioritize devices that have potentially informative datasets with less redundancy, to speed up the learning process. While the richness of datasets can be easily quantified through the total number of samples, the uncertainty of the dataset depends strongly on the application.

For supervised learning, the uncertainty can be evaluated through the evenness of the dataset (i.e., the degree of balance between the classes in classification problems), which can be calculated through entropy measures. For sequence data, the uncertainty is reflected by the regularity of the series. Moreover, for unsupervised learning, local dissimilarity between pseudo-classes or randomly sampled data points can be considered. Furthermore, it is essential to consider privacy as a component of the used index. Sending the number of samples from each class for instance is a violation of the privacy principle of FEEL. In the following, we introduce some potential methods to evaluate dataset diversity.

Diversity Measures for Classification: The measures of diversity have long been used in Active Learning. In fact, uncertainty is used to choose the samples that should be labeled as this task is costly. However, in FL, the client selection does not concern independent samples; instead, the diversity should be evaluated at the level of the entire dataset. Moreover, in the premise of supervised FL, the labels are already known, which gives the possibility to use more informed measures. For instance, Shannon Entropy or Gini-Simpson index are suitable measures for datasets' uncertainty in classification problems. Shannon Entropy and Gini-Simpson index both favor IID partitions, where the maximum for both indexes is obtained for balanced distributions and the dataset with a single class has the minimum possible value. The Shannon Entropy quantifies the uncertainty (entropy or degree of surprise) of a prediction. It was first proposed to quantify the information content in strings of text. The underlying idea is that when a text contains more different letters, with almost equal proportional abundances, it will be more difficult to correctly predict which letter will be the next one in the string. However, Shannon Entropy is not defined for the case of classes with no representative samples. Therefore, it may not be practical in scenarios with high unbalance. Another possible measure is the Gini-Simpson index. The Simpson index measures the probability that two samples taken at random from the dataset of interest are from the same class. The Gini-Simpson index is its transformation $1 - \lambda$, which represents the probability that the two samples belong to different classes. Nonetheless, if the number of classes is large, the distinction using this index will be hard.

Diversity Measures for Time-Series Forecasting: In time series problems, other methods can be used, such as Approximate Entropy (ApEn) and Sample Entropy (SampEn). In sequential data, statistical measures such as the mean and the variance are not enough to illustrate the regularity, as they are influenced by system noise. ApEn is proposed to quantify the amount of regularity and the unpredictability of time-series data. It is based on the comparison between values of data in successive vectors, by quantifying how many data points vary more than a defined threshold. SampEn was proposed as a modification of ApEn. It is used for assessing the complexity of time-series data, with the advantage of being independent from the length of the vectors.

Diversity Measures for Clustering Tasks: For clustering tasks, a similarity measure between data

points from a randomly sampled subset should be considered. The measure can be distance based (e.g., Euclidean distance, Heat Kernel) or angular based (e.g., cosine similarity). A higher value is obtained if most of the data points in the sample are dissimilar, and thus the dataset should be considered as more diverse. It should be noted that angular based measures are invariant to scale, translation, rotation, and orientation, which makes them suitable for a wide range of applications, particularly multivariate datasets.

POST-TRAINING SCHEDULING: MODEL DIVERSITY

The post-training setting uses model diversity to choose devices that will send the updates. The model diversity is evaluated on two different aspects:

- By comparing the dissimilarity between the local model's parameters and the previous global model's parameters.
- By comparing the diversity within the model's parameters.

In fact, choosing the local models that are divergent from the previous global model will possibly improve the representational ability of the global model directly, by aggregating updates that have potentially new information. Furthermore, if a dataset is highly unbalanced and limited in size, the model's parameters would be very similar. The redundancy within parameters negatively affects the model's representational ability. It is therefore necessary to prioritize updates with high diversity. In the following, we detail the post-training scheduling algorithm, then we present some possible measures for model diversity.

Scheduling Algorithm: Similar to pre-training scheduling, the global model is initialized by the BS. Afterwards, the following steps are repeated until the model converges or a maximum of rounds is attained.

Step 1: At the beginning of each training round, devices receive the current model.

Step 2: Each device in the subset uses its local data to train the model.

Step 3: The server sends an update request to the devices, to which each device responds by sending its model diversity index.

Step 4: Based on the received information, alongside the evaluated channel state indicator, the server schedules a subset of devices to upload their models. Then, the updated models are sent to the server to be aggregated.

Step 5: The PS aggregates the updates and creates the new global model.

Model Diversity Measures: While the richness aspect of the diversity is irrelevant in model diversity due to fixed model size among devices, the information contained in the models can be quantified through how the local models vary compared to the global model, and how the parameters within the same model repulse from each other. Some possible measures are as follows.

Local and Global Models' Dissimilarity: Choosing the local models that are divergent from the previous global model will possibly improve the representational ability of the global model directly [5]. Pairwise similarity measures such as cosine similarity and Euclidean distance can be used to evaluate the similarity of the new local param-

Federated Learning is a promising machine learning technique by virtue of its privacy-preserving aspect and ability to handle unbalanced and non-IID data. However, deploying FL based solutions at the edge of the network is subject to several challenges.

eters and the global parameters. Moreover, Divergence, a Bayesian method used to measure the difference between different data distributions, can be used to evaluate diversity of the learned model compared to the global model. Nonetheless, relying on models' dissimilarity might lead to collecting updates from outliers. It is thereby necessary to regulate these diversity measures through the use of thresholds in particular.

Parameters Dissimilarity: To evaluate the redundancy within the model's parameters, the same similarity measures used for clustering can also be applied to the parameters. Additionally, the $L_{2,1}$ norm can be used to obtain a group-wise sparse representation of the dissimilarity [15]. The internal L_1 norm encourages different parameters to be sparse, while the L_2 external norm is used to control the complexity of entire model.

CONCLUSION

Federated Learning is a promising machine learning technique by virtue of its privacy-preserving aspect and ability to handle unbalanced and non-IID data. However, deploying FL based solutions at the edge of the network is subject to several challenges. In fact, FEEL is a multi-disciplinary problem that requires optimization over both the resources and the data. Nonetheless, the data properties are overlooked in many parts of the proposed algorithms, despite being the essence of federated learning. Several FEEL design challenges and issues are introduced and discussed in terms of trade-offs. Furthermore, a new research direction is presented in an effort to incorporate the datasets' diversity properties into the design of FEEL algorithms. Our proposed method supposes that the data quality and veracity are guaranteed, which requires leveraging other techniques such as the blockchain as a trusted third party for data verification.

ACKNOWLEDGMENT

The authors would like to thank the Natural Sciences and Engineering Research Council of Canada for the financial support of this research.

REFERENCES

- [1] G. Zhu, Y. Wang, and K. Huang, "Broadband Analog Aggregation for Low-Latency Federated Edge Learning (Extended Version)," arXiv:1812.11494 [cs, math], Jan. 2019; available: <http://arxiv.org/abs/1812.11494>.
- [2] J. Konečný *et al.*, "Federated Learning: Strategies for Improving Communication Efficiency," arXiv:1610.05492 [cs], Oct. 2017; available: <http://arxiv.org/abs/1610.05492>.
- [3] K. Wei *et al.*, "Federated Learning with Differential Privacy: Algorithms and Performance Analysis," arXiv:1911.00222 [cs], Nov. 2019.
- [4] Q. Zeng *et al.*, "Energy-Efficient Radio Resource Allocation for Federated Edge Learning," arXiv:1907.06040 [cs, math], July 2019; available: <http://arxiv.org/abs/1907.06040>.
- [5] L. Wang, W. Wang, and B. Li, "CMFL: Mitigating Communication Overhead for Federated Learning," *Proc. IEEE ICDCS 2019*, Dallas, TX, USA, July 2019, pp. 954–64; available: <https://ieeexplore.ieee.org/document/8885054/>.
- [6] A. Taik and S. Cherkaoui, "Electrical Load Forecasting Using Edge Computing and Federated Learning," *Proc. ICC 2020 – 2020 IEEE Int'l. Conf. Communications (ICC)*, June 2020, pp. 1–6.

- [7] T. Nishio and R. Yonetani, "Client Selection for Federated Learning with Heterogeneous Resources in Mobile Edge," *Proc. ICC 2019*, May 2019, pp. 1–7; available: <http://arxiv.org/abs/1804.08333>.
- [8] S. Wang et al., "Adaptive Federated Learning in Resource Constrained Edge Computing Systems," *IEEE JSAC*, vol. 37, no. 6, June 2019, pp. 1205–21.
- [9] Y. Zhao et al., "Federated Learning with Non-IID Data," arXiv:1806.00582 [cs, stat], June 2018; available: <http://arxiv.org/abs/1806.00582>.
- [10] C. Xie, S. Koyejo, and I. Gupta, "Asynchronous Federated Optimization," arXiv:1903.03934 [cs], Sept. 2019; available: <http://arxiv.org/abs/1903.03934>.
- [11] J. Wang et al., "Thirty Years of Machine Learning: The Road to Pareto-Optimal Wireless Networks," *IEEE Commun. Surveys Tutorials*, 2020, pp. 1–1.
- [12] W. Shi, S. Zhou, and Z. Niu, "Device Scheduling with Fast Convergence for Wireless Federated Learning," arXiv:1911.00856 [cs, math], Nov. 2019; available: <http://arxiv.org/abs/1911.00856>.
- [13] H. H. Yang et al., "Age-Based Scheduling Policy for Federated Learning in Mobile Edge Networks," *Proc. ICASSP*, May 2020.
- [14] T. Li et al., "Fair Resource Allocation in Federated Learning," arXiv:1905.10497 [cs, stat], Feb. 2020; available: <http://arxiv.org/abs/1905.10497>.
- [15] Z. Gong et al., "Diversity in Machine Learning," *IEEE Access*, vol. 7, 2019, pp. 64 323–50.

BIOGRAPHIES

AFAF TAİK received the B.Eng. degree in software engineering from the Ecole Nationale Supérieure d'Informatique et d'Analyse des Systèmes, Rabat, Morocco in 2018, and a DESS in

applied sciences from the Université de Sherbrooke, Canada in 2018. She is currently pursuing her Ph.D. degree in the Department of Electrical and Computer and Electrical Engineering, Université de Sherbrooke, Canada. Her research interests include edge computing, machine learning, and the development of statistical and mathematical methods for distributed systems in wireless networks. She is a Student Member of IEEE.

SOU MAYA CHERKA OUI is a full professor in the Department of Electrical and Computer Engineering, Université de Sherbrooke, Canada, which she joined as a faculty member in 1999. Her research and teaching interests are in wireless networks. In particular, she works on next generation networks (5G and beyond), edge computing/network intelligence, and communication networks for verticals such as connected and autonomous vehicles, IoT, and Industrial IoT. Since 2005, she has been the Director of INTERLAB, a research group which conducts research funded both by government and industry. Before joining the University of Sherbrooke, she worked for industry as a project leader on projects targeted at the aerospace industry. Her work resulted in technology transfer to companies and to patented technology. She has published over 200 research papers in respected journals and conferences. She has been on the editorial board of several journals including *IEEE Journal on Selected Areas in Communications* (JSAC) and *IEEE Systems Journal*. Her work was awarded with recognitions and best paper awards including a best paper award at the IEEE Communications Society Flagship conference IEEE ICC in 2017. She is currently an IEEE ComSoc Distinguished Lecturer. She is a professional engineer in Canada, a senior IEEE Member, and she is serving as the Chair of the IEEE ComSoc IoT-Ad hoc and Sensor Networks Technical Committee since 2020.