

Resource-Constrained Federated Learning with Heterogeneous Labels and Models

Gautham Krishna Gudur
Global AI Accelerator, Ericsson
gautham.krishna.gudur@ericsson.com

Bala Shyamala Balaji
Indian Institute of Technology,
Madras
balashyamala@gmail.com

Perepu Satheesh Kumar
Ericsson Research
perepu.satheesh.kumar@ericsson.com

ABSTRACT

Various IoT applications demand resource-constrained machine learning mechanisms for different applications such as pervasive healthcare, activity monitoring, speech recognition, real-time computer vision, etc. This necessitates us to leverage information from multiple devices with few communication overheads. Federated Learning proves to be an extremely viable option for distributed and collaborative machine learning. Particularly, on-device federated learning is an active area of research, however, there are a variety of challenges in addressing statistical (non-IID data) and model heterogeneities. In addition, in this paper we explore a new challenge of interest – to handle *label heterogeneities* in federated learning. To this end, we propose a framework with simple α -weighted federated aggregation of scores which leverages overlapping information gain across labels, while saving bandwidth costs in the process. Empirical evaluation on Animals-10 dataset (with 4 labels for effective elucidation of results) indicates an average deterministic accuracy increase of at least $\sim 16.7\%$. We also demonstrate the on-device capabilities of our proposed framework by experimenting with federated learning and inference across different iterations on a Raspberry Pi 2, a single-board computing platform.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**; *Distributed computing methodologies*.

KEYWORDS

On-Device Federated Learning, Heterogeneous Labels, Heterogeneous Models, Transfer Learning

ACM Reference Format:

Gautham Krishna Gudur, Bala Shyamala Balaji, and Perepu Satheesh Kumar. 2020. Resource-Constrained Federated Learning with Heterogeneous Labels and Models. In *The 3rd International Workshop on Artificial Intelligence of Things (AIoT'20) at KDD 2020: The 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, August 22-27, 2020, San Diego, CA*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD'20 Workshops: AIoT'20, August 22-27, 2020, San Diego, CA

© 2020 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

Contemporary machine learning, particularly deep learning has led to major breakthroughs in various domains, such as computer vision, natural language processing, speech recognition, Internet of Things (IoT), etc. Particularly, on-device machine learning has spiked up a huge interest in the research community owing to the compute capabilities vested in resource-constrained devices like mobile and wearable devices. Sensor data from various IoT devices have a vast amount of incoming data which have massive potential to leverage such on-device machine learning techniques on-the-fly to transform them into meaningful information coupled with supervised, unsupervised and/or other learning mechanisms. With the ubiquitous proliferation of such personalized IoT devices, collaborative and distributed learning is now possible more than ever to help best utilize the information learnt from multiple devices.

However, such collaborative data sharing across devices might always not be feasible owing to privacy concerns from multiple participants. Moreover, users might not prefer nor have any interest in sending their sensitive data to a remote server/cloud, particularly in fields like health-care, defense, telecommunication, etc. With the advent of *Federated Learning (FL)* [17], [2], it is now possible to effectively train a global/centralized model without compromising on sensitive data of various users by enabling the transfer of model weights and updates from local devices to the cloud, instead of conventionally transferring the sensitive data to the cloud. A server has the role of coordinating between models, however most of the work is not performed by a central entity anymore but by a federation of users. The *Federated Averaging (FedAvg)* algorithm [17] was proposed by McMahan et al. which aggregates the model parameters of each client (local device) by combining local Stochastic Gradient Descent (SGD) through a server. Federated learning has been an active and challenging area of research in solving problems pertaining to secure communication protocols, device and statistical heterogeneities, and privacy preserving networks [12].

Federated Learning deals with various forms of heterogeneities like device, system, statistical heterogeneities, etc. [12], [19]. Particularly, FL in IoT and edge computing scenarios, statistical heterogeneities have gained much visibility as a research problem predominantly owing to the non-IID (non-independent and identically distributed) nature of the vast amounts of streaming real-world data incoming from distinct distributions across devices. This leads to challenges in personalized federation of devices, and necessitates us to address various heterogeneities in data and learning processes for effective model aggregation.

One important step in this direction is the ability of end-users to have the choice of architecting their own models, rather than being constrained by the pre-defined architecture mandated by the global

model for model aggregation. One effective way to circumvent this problem is by leveraging the concept of knowledge distillation [5], wherein the disparate local models distill their respective knowledge into various *student models* which has a common model architecture, thereby effectively incorporating model independence and heterogeneity. This was proposed by Li et al. in FedMD [11]. However, as much independence and heterogeneity in architecting the users' own models is ensured in their work, they do not guarantee *heterogeneity and independence in labels across devices*.

For a concrete example, a central intelligent system/service acts as the global cloud/server with multiple telecommunication operators being the clients in a typical federated learning scenario. However, most operators (clients) typically choose to have their own machine learning model architectures, and also have their own labels which are governed by few standardized and also localized company bodies. This necessitates seamless interaction of the local clients with the global cloud/server independent of model architectures and labels. Many such scenarios and settings with such heterogeneous labels and models exist in federated IoT/on-device settings, such as behaviour/health monitoring, activity tracking, keyword spotting, next-word prediction, etc. Few works address handling new labels in typical machine learning scenarios, however, to the best of our knowledge, there is no work which addresses this important problem of *label and model heterogeneities* in non-IID federated learning scenarios.

The main scientific contributions in this work are as follows:

- Enabling end-users to build and characterize their own preferred local architectures in a federated learning scenario, so that effective transfer learning happens between the global and local models.
- A framework to allow flexible heterogeneous selection of labels by showcasing scenarios with and without overlap across different user devices, thereby leveraging the information learnt across devices pertaining to those overlapped classes.
- Empirical demonstration of the framework's ability to handle different data distributions (statistical heterogeneities and non-IIDness) from various user devices.
- Demonstrating the feasibility of on-device personalized federated learning from incoming real-world data independent of users, capable of running on simple mobile and wearable devices.

2 RELATED WORK

Federated Learning has contributed vividly in enabling distributed and collective machine learning across various devices. Federated learning and differentially private machine learning have/soon will emerge to become the de facto mechanisms for dealing with sensitive data, and data protected by Intellectual Property rights, GDPR, etc [2]. Federated Learning was first introduced by McMahan et al. in [17], and new challenges and open problems to be solved, aggregation and optimization methods, [12], strategies [10], and multiple advancements [7] have been proposed and addressed in many interesting recent works.

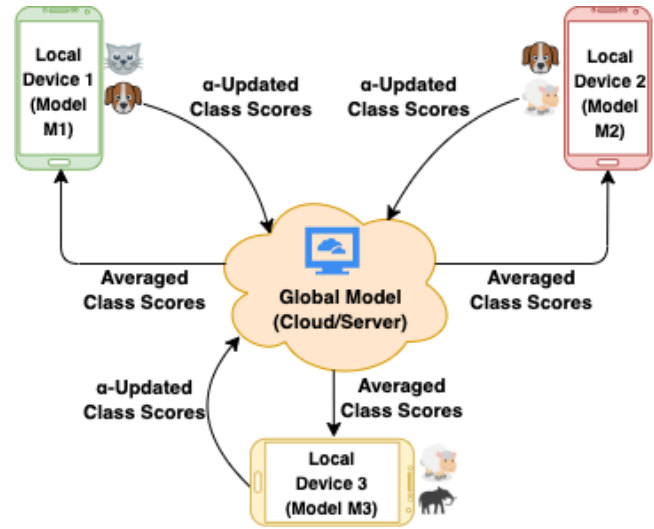


Figure 1: Overall Architecture of the proposed framework with *Weighted Local Update with α* . Each mobile device can consist of disparate set of local labels and models, and they interact with the global model (cloud/server). The models in each local device are first updated using weighted α score, the respective class scores are then aggregated in the global model, and the updated consensus is again distributed across local models.

Particularly for Federated Learning in IoT and pervasive (mobile/wearable/edge) devices, important problems and research directions are addressed FL on mobile and edge networks in this survey [16]. Federated Optimization for on-device applications is discussed in [9].

Multiple device/client and system heterogeneities, including client communication mechanisms and fair resource allocation, inherently making most of them optimization problems, and are addressed in various works [13], [4], [19], [14]. Personalized federated learning closely deals with optimizing the degree of personalization and contribution from various clients, thereby enabling effective aggregation as discussed in [3].

Recent works on convergence of Federated Averaging (FedAvg) algorithm on disparate data distributions – non-IID data, and creating a small subset of data globally shared between all edge devices are proposed in [20], [15] respectively. Mohri et al. propose Agnostic Federated Learning [18], which addresses about handling any target data distribution formed by a mixture of client distributions.

FedMD [11], which we believe to be our most closest work, deals with heterogeneities in model architectures, and addresses this problem using transfer learning and knowledge distillation [5], and also uses an initial public dataset across all labels (which can be accessed by any device during federated learning). On similar grounds, federated distillation and augmentation in non-IID data distributions are used in [6]. Current federated learning approaches predominantly handle same labels across all the users and do not provide the flexibility to handle unique labels. However, in many practical applications, having unique labels for each local client/model is a

Table 1: Model Architectures (filters/units in each layer indicated within parentheses), Labels and Number of Images per federated learning iteration across user devices. Note the disparate model architectures and labels across users.

	User_1	User_2	User_3	Global_User
Model Architecture	2-Layer CNN (16, 32) Softmax Activation	3-Layer CNN (16, 16, 32) ReLU Activation	2-Layer CNN (16, 32) ReLU Activation	-
Labels	{Cat, Dog}	{Dog, Sheep}	{Sheep, Elephant}	{Cat, Dog, Sheep, Elephant}
Images per iteration	{500, 500} = 1000	{500, 500} = 1000	{500, 500} = 1000	{500, 500, 500, 500} = 2000

very viable and common scenario owing to their dependencies and constrains on specific regions, demographics, etc. To the best of our knowledge, none of the works take into account, label and model heterogeneities.

The rest of the paper is organized as follows. Section 3 discusses the problem formulation of handling heterogeneous labels and models in on-device federated learning scenarios (section 3.1), and section 3.2 presents the overall proposed framework and the methods used to address these challenges. Systematic experimentation and evaluation of the framework across different users, devices, iterations, models, labels in a federated learning setting is showcased in section 4, while also proving feasibility of the same on resource-constrained devices (section 4.2). Finally, section 5 concludes the paper.

3 OUR APPROACH

In this section, we discuss in detail about the problem formulation of heterogeneity in labels and models, and our proposed framework to handle the same (showcased in Figure 1).

3.1 Problem Formulation

We assume the following scenario in federated learning. There are multiple local devices which can characterize different model architectures based on the end users. The incoming streaming real-world data in all the devices is non-IID in nature, wherein the distribution and data characteristics differ across devices. We hypothesize that the incoming data to the different devices also consist of heterogeneities in labels, with either unique or overlapped labels. We also have a public dataset with the label set consisting of all labels pre-defined – this can be accessed by any device anytime, and acts as an initial template of the data and labels that can stream through, over different iterations. We repurpose this public dataset as the test set also, so that consistency is maintained while testing. To make different FL iterations independent from the public dataset, we do not involve the public dataset during federated learning (training) in the local models. The research problem here is to create a unified framework to handle heterogeneous models, labels and data distributions (with non-IID nature) in a federated learning setting.

3.2 Proposed Framework

The proposed framework and methods to handle the heterogeneous labels and models in a federated learning setting is presented in Algorithm 1. There are three important steps in the proposed method.

Algorithm 1: Proposed Framework to handle heterogeneous labels and models in Federated Learning

Input - Public Data set $\mathcal{D}_0\{x_0, y_0\}$, Private datasets \mathcal{D}_m^i , Total users M , Total iterations I , LabelSet for each user l_m

Output - Trained Model scores f_G^I

Initialize - $f_G^0 = \mathbf{0}$ (Global Model Scores)

for $i = 1$ **to** I **do**

for $m = 1$ **to** M **do**

Build: Model \mathcal{D}_m^i and predict $f_{\mathcal{D}_m^i}(x_0)$

Local Update: $f_{\mathcal{D}_m^i}(x_0) = f_G^I(x_0^{l_m}) + \alpha f_{\mathcal{D}_m^i}(x_0)$, where $f_G^I(x_0^{l_m})$ are the global scores of only the set of labels l_m with the m^{th} user, and $\alpha = \frac{\text{len}(\mathcal{D}_m^i)}{\text{len}(\mathcal{D}_0)}$

Global Update: Update label wise,

$$f_G^{i+1} = \sum_{m=1}^M \beta_m f_{\mathcal{D}_m^i}(x_0), \text{ where,}$$

$$\beta = \begin{cases} 1 & \text{If labels are unique} \\ \text{acc}(f_{\mathcal{D}_m^i}(x_0)) & \text{If labels overlap} \end{cases}$$

end for

end for

Build: In this step, we build the model on the incoming data we have in each local user, i.e., local private data for the specific iteration. The users can choose their own model architecture which suits best for the data present in the iteration.

Local Update: In this step, we update the averaged global model scores (on public data) for the i^{th} iteration on the local private data. For the first iteration, we do not have any global scores and we initialize the scores to be zero in this case. For the rest of iterations, we have global averaged scores which we can use to update the local model scores according to Algorithm 1. In the local update, we propose a *weighted α -update*, where α is the ratio between the size of current private dataset and the size of public dataset, and governs the contributions of the new and the old models across different federated learning iterations.

Global Update: In this step, we first train the local model on the respective private datasets for that FL iteration. Further, we evaluate (test) this trained model on the public data, thereby obtaining the model scores on public data. We then perform the same operation across all the users and average them using the β parameter, where

β governs the weightage given to overlapping labels across users using test accuracies of the corresponding labels on public data (as given in Algorithm 1). This module gives the global averaged scores.

4 EXPERIMENTS AND RESULTS

We simulate a Federated Learning scenario with multiple iterations of small chunks of incremental data incoming (details in Table 1), across three different users to test our approach, and assume that the images arrive in real-time in the users' devices. We use the *Animals-10 dataset* [1], which consists of ten different animal categories (labels) extracted from Google Images. Now, we discuss the settings for label and model heterogeneities in our experiment.

Label Heterogeneities: In our experiment, we consider only four labels – $\{Cat, Dog, Sheep, Elephant\}$ from the dataset as shown in Table 1. Also, we include the number of images considered per user per iteration (500 images per iteration). The labels in each local user can either be unique (present only in that single user) or overlapped across users (present in more one user). We split the four labels into three pairs of two labels each, for convenience of showcasing the advantage of overlapping labels in experimentation. We also create a non-IID environment across different federated learning iterations wherein, the image data across different iterations are split with disparities in both labels and distributions of data (*Statistical Heterogeneities*).

Table 2: Details of Model Architectures (filters/units in each layer indicated within parentheses) changed across federated learning iterations and users.

Iteration	New Model Architecture
User_1 Iteration_10	3-Layer ANN (16, 16, 32) ReLU Activation
User_1 Iteration_14	1-Layer CNN (16) Softmax Activation
User_2 Iteration_6	3-Layer CNN (16, 16, 32) Softmax activation
User_3 Iteration_5	4-Layer CNN (8, 16, 16, 32) Softmax activation

Model Heterogeneities: We choose three different model architectures for the three different local user device. This is clearly elucidated in Table 1. To truly showcase near-real-time heterogeneity and model independence, we change the model across and within various FL iterations as shown in Table 2.

Initially, we divide the images across different users according to the four labels. Each image is re-scaled to size 128×128 . We create a Public Dataset - D_0 with 2000 images, with 500 images corresponding to each label. Next, we sample 500 images in every iterations as per the labels of the user and we use them for our problem (as shown in Table 1). In total, we ran 15 iterations in this whole

federated learning experiment, with each iteration running with early stopping (max 5 epochs). We track the loss using categorical cross-entropy loss function for multi-class classification, and use the Adam optimizer [8] to optimize the classification loss. We simulate all our experiments – both federated learning and inference on a *Raspberry Pi 2*.

4.1 Discussion on Results

Figure 2 represents the results across all three users on Animals Dataset. Also, from Table 3, we can clearly observe that the global updates – which represent the accuracies of the global updated model (and averaged across all users' labels in the i th iteration governed by β), are higher for all three users than the accuracies of their respective local updates. For instance, from Figure 2a, we can infer that the corresponding accuracies of labels $\{Cat, Dog\}$ (User 1 labels) after global updates in each iteration are deterministically higher than their respective local updates by an average of $\sim 17.4\%$ across all iterations with weighted α -update. Similarly for User 2, labels consisting of $\{Dog, Sheep\}$, we observe an average accuracy increase of $\sim 23.2\%$ from local updates to the global updates, while for User 3 labels consisting of $\{Sheep, Elephant\}$, we observe an average increase of $\sim 9.3\%$ from local updates to global updates.

Table 3: Average Accuracies (%) of Local and Global Updates, and their respective Accuracy increase with Weighted α -update

	Local_Update	Global_Update	Acc. Increase
User_1	63.66	81.02	17.36
User_2	74.3	97.47	23.17
User_3	68.72	78.02	9.3
Average	68.89	85.5	16.61

We would like to particularly point out that the overlap in labels significantly contributed to highest increase in accuracies, owing to the fact that information gain (weighted global update) happens only for overlapping labels. This is vividly visible in User 2 (Figure 2b, whose labels are $\{Dog, Sheep\}$), where in spite of an accuracy dip in local update at iteration 10, the global update at that iteration does not take a spike down which can be primarily attributed to the information gain from overlapped labels between User 1 and User 3 (in this case, *Dog* and *Sheep* respectively), thereby showcasing the robustness of overlapped label information gain in User 2. On the contrary, when we observe User 1 (Figure 2a), in spite of the accuracies of global updates being inherently better than local updates, when a dip in accuracies of local updates are observed at iteration 8 and 12, the accuracies of global updates at that iteration also spike down in a similar fashion. Similar trends of local and global accuracy trends like those observed in User 1 can also be observed in User 3 (Figure 2c). This clearly shows that when there are lesser overlapped labels (User 1 and User 3), the global model does not learn the label characteristics as much as when there are more overlapped labels – the global updates are more robust in spite of spikes and dips in local updates with overlapped labels (User 2), thereby leading to higher average increase in accuracies (as observed in Table 3).




Figure 2: Iterations vs Accuracies across all three users with Weighted Local Update with α . *Local_Update* signifies the accuracy of each local updated model (after i^{th} iteration) on Public Dataset. *Global_Update* signifies the accuracy of the corresponding global updated model (after i^{th} iteration) on Public Dataset.




Figure 3: Iterations vs Final Global Average Accuracies (%) with Weighted α -update

An overall average increase of $\sim 16.7\%$ deterministic accuracy increase (not relative) is observed with our proposed framework across all the different users, which can be calculated from the global model updates. The overall global model accuracies alone (which is different from global update accuracies observed in Figure 2) after each iteration are also elucidated in Figure 3. We propose that, with our current framework, communication (transfer) of just the model scores of respective labels between clients (local devices) and the central cloud is enough, without necessitating transfer of the whole model weights, which significantly reduces latency and memory overheads. We also point out that we tried local model distillation similar to the work done in FedMD [11] for handling label and model heterogeneities, however, we felt using only score consensus from the model without distillation yielded far lesser on-device computation times.

4.2 On-Device Performance

We observe the on-device performance of our proposed framework, which is experimented on a Raspberry Pi 2. We choose this single-board computing platform since it has similar hardware and software (HW/SW) specifications with that of predominant contemporary IoT/edge/mobile devices. The computation times taken

for execution of on-device federated learning and inference are reported in Table 4. This clearly shows the feasibility of our proposed system on embedded devices.

Table 4: Time taken for Execution

Process	Computation Time
Training time per epoch in an FL iteration (i)	~ 1.8 sec
Inference time	~ 15 ms

5 CONCLUSION

This paper presents a unified framework for flexibly handling heterogeneous labels and model architectures in federated learning in a non-IID fashion. By leveraging transfer learning along with simple scenario changes in the federated learning setting, we propose a framework with α -update aggregation in local models, and we are able to leverage the effectiveness of global model updates across all devices and obtain higher efficiencies. Moreover, overlapping labels are found to make our framework robust, and also helps in effective accuracy increase. We empirically showcase the successful feasibility of our framework on resource-constrained devices for federated learning/training and inference across different iterations on the Animals-10 dataset. We expect a good amount of research focus hereon in developing statistical, model and label based heterogeneities.

REFERENCES

- [1] ALESSIO, C. *Animals-10 Dataset*, 2019. <https://www.kaggle.com/alessiocorrado99/animals10>.
- [2] BONAWITZ, K., EICHNER, H., GRIESKAMP, W., HUBA, D., INGERMAN, A., IVANOV, V., KIDDON, C., KONECNY, J., MAZZOCCHI, S., MCMAHAN, H. B., OVERVELDT, T. V., PETROU, D., RAMAGE, D., AND ROSELANDER, J. Towards federated learning at scale: System design. In *SysML 2019* (2019).
- [3] DENG, Y., KAMANI, M. M., AND MAHDAVI, M. Adaptive personalized federated learning. *arXiv preprint arXiv:2003.13461* (2020).
- [4] GHOSH, A., HONG, J., YIN, D., AND RAMCHANDRAN, K. Robust federated learning in a heterogeneous environment. *arXiv preprint arXiv:1906.06629* (2019).
- [5] HINTON, G., VINYALS, O., AND DEAN, J. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop* (2015).
- [6] JEONG, E., OH, S., KIM, H., PARK, J., BENNIS, M., AND KIM, S.-L. Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data. *arXiv preprint arXiv:1811.11479* (2018).

- [7] KAIROUZ, P., MCMAHAN, H. B., AVENT, B., BELLET, A., BENNIS, M., BHAGOJI, A. N., BONAWITZ, K., CHARLES, Z., CORMODE, G., CUMMINGS, R., ET AL. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977* (2019).
- [8] KINGMA, D. P., AND BA, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [9] KONEČNÝ, J., MCMAHAN, H. B., RAMAGE, D., AND RICHTÁRIK, P. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527* (2016).
- [10] KONEČNÝ, J., MCMAHAN, H. B., YU, F. X., RICHTÁRIK, P., SURESH, A. T., AND BACON, D. Federated learning: Strategies for improving communication efficiency. In *NIPS Workshop on Private Multi-Party Machine Learning* (2016).
- [11] LI, D., AND WANG, J. Fedmd: Heterogenous federated learning via model distillation. *arXiv preprint arXiv:1910.03581* (2019).
- [12] LI, T., SAHU, A. K., TALWALKAR, A., AND SMITH, V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine* 37, 3 (2020), 50–60.
- [13] LI, T., SAHU, A. K., ZAHEER, M., SANJABI, M., TALWALKAR, A., AND SMITH, V. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems 2020* (2020), 429–450.
- [14] LI, T., SANJABI, M., BEIRAMI, A., AND SMITH, V. Fair resource allocation in federated learning. In *8th International Conference on Learning Representations (ICLR)* (2020).
- [15] LI, X., HUANG, K., YANG, W., WANG, S., AND ZHANG, Z. On the convergence of fedavg on non-iid data. In *8th International Conference on Learning Representations (ICLR)* (2020).
- [16] LIM, W. Y. B., LUONG, N. C., HOANG, D. T., JIAO, Y., LIANG, Y.-C., YANG, Q., NIYATO, D., AND MIAO, C. Federated learning in mobile edge networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials* (2020).
- [17] MCMAHAN, H. B., MOORE, E., RAMAGE, D., HAMPSON, S., AND ARCAS, B. A. Y. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (2017), vol. 54, pp. 1273–1282.
- [18] MOHRI, M., SIVEK, G., AND SURESH, A. T. Agnostic federated learning. In *Proceedings of the 36th International Conference on Machine Learning* (2019), vol. 97, pp. 4615–4625.
- [19] NISHIO, T., AND YONETANI, R. Client selection for federated learning with heterogeneous resources in mobile edge. In *IEEE International Conference on Communications (ICC)* (2019), pp. 1–7.
- [20] ZHAO, Y., LI, M., LAI, L., SUDA, N., CIVIN, D., AND CHANDRA, V. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582* (2018).