
Hybrid FL: Algorithms and Implementation

Xinwei Zhang
University of Minnesota
zhan6234@umn.edu

Wotao Yin
University of California, Los Angeles
wotaoyin@math.ucla.edu

Mingyi Hong
University of Minnesota
mhong@umn.edu

Tianyi Chen
Rensselaer Polytechnic Institute
chent18@rpi.edu

Abstract

Federated learning (FL) is a recently proposed distributed machine learning paradigm dealing with distributed and private data sets. Based on the data partition pattern, FL is often categorized into horizontal, vertical and hybrid settings. Despite the fact that many works have been developed for the first two approaches, the hybrid FL setting (which deals with partially overlapped feature space and sample space) remains less explored, though this setting is extremely important in practice. In this paper we first set up a new model-matching based problem formulation for hybrid FL, then propose an efficient algorithm which can collaboratively train the global and local models to deal with full and partial featured data. We conduct numerical experiments on the multi-view ModelNet40 data set to validate the performance of the proposed algorithm. To the best of our knowledge, this is the first formulation and algorithm developed for the hybrid FL.

1 Introduction

Federated learning (FL) is an emerging machine learning framework where multiple clients (e.g., mobile devices or organizations) collaboratively train a machine learning (ML) model [1]. FL specifically addresses the new challenges including the difficulty of synchronizing multiple clients, the heterogeneity of data, and the privacy and security of clients' data and, in some settings, also their local models. Due to these challenges, classic ML methods cannot be directly applied [2].

A popular FL setting that partitions data among clients is called *horizontal FL* (HFL). Each client has data of a different set of subjects, and the data of every client have the same set of features [3, 4]. Examples of such data include smartphone users' word-typing histories (from the same word dictionary), which are stored on individual devices and analyzed by the same features [1]. One can apply HFL to learn a model for word or sentence completion. In the setting known as *vertical FL* (VFL), it is features that are partitioned among clients, and all the clients share a common set of subjects [5, 6]. More features help build a more accurate model than using fewer features. For example, VFL can help an insurance company better predict someone's risk using not just this person's records at this company but also records from multiple other insurance businesses.

Once training is completed, HFL and VFL also have different prediction processes. In HFL, the jointly-trained model is typically shared among the clients, so each client performs predictions alone. In VFL, while a client can predict using its local model based on its local features, more accurate predictions are made when more clients work together and use their jointly learned model that take all the features available.

Motivation and main challenge. In practice, however, each client may possess only some subjects and some features. It is possible that no client has all features or all subjects. This is the case of financial institutions such as insurance providers, banks, and stock services providers, each of which may serve just a fraction of all customers and have only their partial records. This setting has been referred to as *hybrid FL* [3, 4], and it is the setting we focus in this paper. Both HFL and VFL are

special cases of hybrid FL. Compared to HFL and VFL, hybrid FL has its unique challenges. Some specific ones pertaining to algorithm design are:

1. **Client customized models.** In Hybrid FL, each client’s local data contain a subset of features, but in inference, it is possible that some clients may need to deal with data that has the full feature. So the server is often required to maintain a copy of the full feature.
2. **Limited data sharing.** In typical HFL, the clients do not share their local data or labels, but in VFL, the labels are either made available to the server [5] or stored in a designated client [6]. A Hybrid FL system needs to deal with both types of clients, so it is desirable that the training algorithm can operate without needing the server to have access to any data including the labels.
3. **Sample synchronization.** A typical issue with VFL (in which each client has some features of all the training samples) is that, when updating a given subset of features, all the clients need to draw the same (mini-batch of) samples; this problem is exacerbated in the hybrid FL system because not all the clients will have all the samples. An ideal algorithm shall work without any synchronization on the samples among the clients.

All the above points will directly translate to specific challenges when we design optimization algorithms, and it will become clear that none of the existing FL methods can meet all these requirements.

Related work on HFL. In HFL, a common algorithm is FedAvg [3], which adopts the computation-then-aggregation strategy, that is, the clients locally perform a few steps of model updates, and then the server aggregates the updated local models and averages them before sending the updated global model back. Beyond model communication, MIME [7] and SCAFFOLD [8] also send local gradients and other statistics to the server to achieve better convergence. Furthermore, PNFM [9] and FedMA [10] use a parameter matching based strategy in place of the model averaging step to get better global model performance, and they do not require the global model to have the same size as the local models. All HFL algorithms assume their data having the same size and format.

Related work on VFL. In VFL, the features and thus models are separated on different clients [11, 12, 6, 5]. There are relatively few works on VFL. Federated Block Coordinate Descent (Fed-BCD) [6] uses a parallel BCD-like algorithm to optimize the local blocks and transmits essential information for the other clients to compute their local gradients. Vertical Asynchronous Federated Learning (VAFL) [5] assumes that the server holds the global inference model while local clients train the feature extractors that deal with the local features.

Our contributions. We summarize our main contributions as follows.

1. We propose a hybrid FL model that captures many key aspects of collaborative-learning scenarios, where neither the subject set nor the feature set is necessarily complete at a client. Such a formulation can be tailored to meet different requirements for specific hybrid FL models. To our knowledge, this is the first concrete hybrid FL model in the literature.
2. We develop a convergent hybrid FL algorithm that enables knowledge transfer among clients, which at the same time maintains data locality and improves communication efficiency (by removing the sample synchronization requirement).
3. We evaluate the performance of the hybrid FL algorithm on real data sets that learn a federated model with its achieved learning accuracy comparable to that learned in centralized settings.

2 Problem Formulation

We consider a total of N samples, written as $\{\mathbf{x}^n, y^n\}_{n=1}^N$, where each \mathbf{x}^n has D features $\mathbf{x}^n = [x[1]^n; \dots; x[D]^n]$. There are M clients, where each has some samples and their partial features. The features have indices $d = 1, \dots, D$, sample data have indices $n = 1, \dots, N$, and clients have indices $m = 1, \dots, M$. If the m th agent has the n th sample, we write $\mathbf{x}_m^n = [x[d_1]^n; \dots; x[d_m]^n]_{d_1, \dots, d_m \in \mathcal{D}_m}$, where \mathcal{D}_m denotes the set of features known to the m th agent.

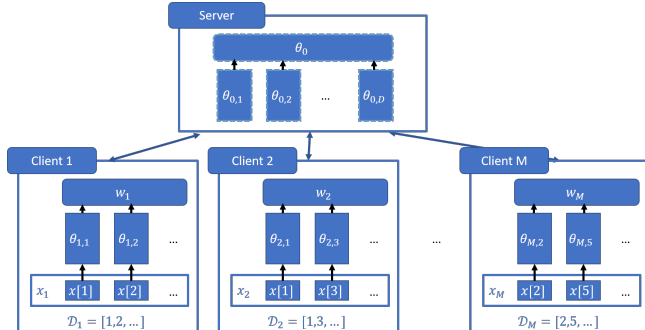


Figure 1: Hybrid FL system structure

A Generic Formulation. Consider a hybrid FL model consisting of an inference model and some feature extractors. For a given agent m , if it has enough samples that are related to feature d , then it will learn a *feature extractor* $\theta_{m,d}$, which is approximately the same as a global feature extractor $\theta_{0,d}$ located at the server, i.e., $\theta_{m,d} \approx \theta_{0,d}$. On the other hand, if it does not have enough sample containing feature d' , then it will not participate in learning $\theta_{m,d'}$. For a given agent m , it will process an input sample \mathbf{x}_m^n by going through an embedding $h_m(\{\theta_{m,d}\}_{d \in \mathcal{D}_m}; \mathbf{x}_m^n)$. The embedding vector then goes through the inference model w_m , also learned by agent m , and produces an output (label). During the aggregation step, the server may also create global copies of the aggregated models $\theta_0, \{\theta_{0,d}\}$. The setting discussed above is illustrated in Figure 1.

Use the above notation, let us first set up a high-level problem formulation as below:

$$\min_{\{\theta_{m,d}, w_m\}} \frac{1}{MN} \sum_{m=1}^M \sum_{n \in \mathcal{N}_m} \ell(h_m(\{\theta_{m,d}\}_{d \in \mathcal{D}_m}; \mathbf{x}_m^n); w_m; y^n) + r(\{\theta_{m,d}\}, w_m; \theta_0, \{\theta_{0,d}\}) \quad (1)$$

where $\ell(\cdot)$ measures the accuracy of using the embedding $h(\{\theta_{m,d}\}_{d \in \mathcal{D}_m}; \mathbf{x}_m^n)$ and the local model w_m to predict y^n , and $r(\cdot)$ is a generic regularizer that encodes the prior knowledge about the global and local models. Such a formulation is general enough to capture several special Hybrid FL settings. For example it is easy to see that the exiting VFL and HFL settings are both its special cases. We will demonstrate next how problem (1) can be customized to a practical hybrid FL problem.

A Feature Matching Based Hybrid FL Formulation. Specifically, we assume that both features and labels are *private*, and they are not shared with other clients nor with the server. Further, assume that the feature extractors for the same feature are approximately consensual, that is, we require $\theta_{m,d} \approx \theta_{0,d}$ for every agent $m \in [M]$ who updates the d -th feature extractor. In this case, we denote the concatenates $\{\theta_{m,d}, d \in \mathcal{D}_m\}$ as θ_m , and denote the local data set as $\{\mathbf{x}_m^n, y^n\}_{n \in \mathcal{N}_m}$. Then the objective function of (1) can be separated into a sum of the following local objectives:

$$\min_{\{\theta_{m,d}, w_m\}} \frac{1}{N} \sum_{n \in \mathcal{N}_m} \ell(h_m(\theta_m; \mathbf{x}_m^n); w_m; y^n) + r_m(\theta_m, w_m; \theta_0, \{\theta_{0,d}\}). \quad (2)$$

Here $r_m(\cdot)$ indicates the local regularizer for client m that enforces the consensus among the local feature extractors θ_m and regularizes the difference of the local inference model w_m . Our main design effort will be devoted to finding the proper regularizer $r_m(\cdot)$, which has the following desired features: 1) It helps enforce the consensus of $\theta_{m,d}$ and $\theta_{0,d}$; 2) It facilitates the learning of a global inference model θ_0 from the local inference models $\{w_m\}_{m=1}^M$. Since these two tasks are relatively separable, then it is natural to expression $r_m(\cdot)$ as :

$$r_m(\theta_m, w_m; \theta_0, \{\theta_{0,d}\}) = \sum_{d \in \mathcal{D}_m} r_1(\{\theta_{m,d}\}, \{\theta_{0,d}\}) + r_2(\{w_m\}, \theta_0). \quad (3)$$

Notice that we can use any reasonable distance function to construct $r_1(\cdot)$ since $\theta_{m,d}$ and $\theta_{0,d}$ have the same dimension. However, it is not straightforward to construct $r_2(\cdot)$, for the following reasons. First, the dimension of θ_0 is much larger than each individual w_m since θ_0 is the inference model which takes all the features extractors as inputs. Second, it is not easy to identify the relation between different w_m 's parameters and combine them to yield a global θ_0 .

To deal with these challenges, we adopt a *matched averaging* idea proposed by [9, 10], which is used in the HFL setting to dynamically match the neurons of local models to build a global model. More specifically, it is assumed that the global θ_0 and local w_m are related through a linear mapping $w_m \approx \Pi_m \theta_0$, where such a mapping Π_m should be dynamically optimized. One special case of Π_m is a matching matrix, which contains only one non-zero entry at each row. By using such a matching matrix, it ensures that each neuron in the global model is a linear combination of a set of most closely related neurons in the local models. While the idea of model matching has been explored in the recent works [9, 10], here we design a special matching strategy specifically designed for hybrid FL. Hence, the main contribution of this work is the overall problem formulation, and the model matching is just one integral component of it.

By using the model matching strategy, the two regularizers in (1) can be expressed as following:

$$r_1(\{\theta_{m,d}\}, \{\theta_{0,d}\}) = \text{dist}(\theta_{m,d}, \theta_{0,d}), \quad r_2(\{w_m\}, \theta_0) = \text{dist}(w_m, \Pi_m \theta_0), \quad (4)$$

Algorithm 1 Hybrid Federated Matching Algorithm

Input: $\{w_m^0\}, \{\theta_d^0\}, \eta, T, Q, P$
 Initialize: $\theta_{m,d}^0 = \theta_d^0, m = 1, \dots, M$
for $t = 0, \dots, T - 1$ **do**
 // Local Updates:
 for $i = 1, \dots, M$ **in parallel do**
 $\theta_m^{t,Q}, w_m^{t,Q} = \arg \min_{\theta_m, w_m} f_m(\theta_m, w_m)$
 end for
 // Global Aggregation of Feature Extractors:
 $\theta_d^{t+1} = \frac{1}{|\{m:d \in \mathcal{D}_m\}|} \sum_{m:d \in \mathcal{D}_m} \theta_{d,m}^{t,Q}, \text{ for } d = 1, \dots, D$
 // Global Matching Updates:
 for $p = 0, \dots, P - 1$ **do**
 Randomly select a client index m'
 Update matching pattern: $\Pi_{m'} = \arg \min_{\Pi_{m'}} f_0(\theta_0^{t,p}, \{\Pi_m\}; \{w_m^{t,Q}\})$
 Update global inference model: $\theta_0^{t,p+1} = \arg \min_{\theta_0} f_0(\theta_0, \{\Pi_m\}; \{w_m^{t,Q}\})$
 end for
 Distribute the aggregated models: $w_m^{t+1,0} = \Pi_m \theta_0^{t,P}, \theta_{m,d}^{t+1,0} = \theta_d^{t+1}$
end for

where $\text{dist}(\cdot)$ measures the distance between the models. It is important to note that the matching patterns Π_m 's have to be optimized as well. This leads to the following Hybrid FL problem:

$$\begin{aligned}
 \min_{\{\theta_m, w_m, \Pi_m\}, \theta_0, \{\theta_{0,d}\}} & \frac{1}{MN} \sum_{m=1}^M \sum_{n \in \mathcal{N}_m} \ell(h(\theta_m; \mathbf{x}_m^n); w_m; y_m^n) \\
 & + \mu \sum_{m=1}^M \text{dist}(\Pi_m \theta_0, w_m) + \sum_{m=1}^M \sum_{d \in \mathcal{D}_m} \text{dist}(\theta_{m,d}, \theta_{0,d}) \\
 \text{s.t. } & \Pi_m \mathbf{1} = \mathbf{1}, \Pi_m \geq 0, m = 1, \dots, M.
 \end{aligned} \tag{5}$$

In this formulation, we jointly minimize the classification loss and the consensus loss and μ is used to balance between the two losses. When μ is very small, the clients focus on training the local models which can be very different from each other. When μ is large, the emphasis is put on learning an accurate global model by integrating local information.

3 The Proposed Algorithm

We propose an algorithm for solving (5). The problem contains parameter blocks $\{\theta_{m,d}\}$ and $\{w_m\}$, the global parameters $\theta_{0,d}$, and the global parameters θ_0 and $\{\Pi_m\}$, so we can update each of them given the others. The problem related to the local parameters $\theta_m, w_m, m = 1, \dots, M$, is:

$$\min_{\theta_m, w_m} f_m(\theta_m, w_m) = \sum_{n \in \mathcal{N}_m} \ell(h(\theta_m; \mathbf{x}_m^n); w_m; y_m^n) + \mu \text{dist}(\Pi_m \theta_0, w_m) + \sum_{d \in \mathcal{D}_m} \text{dist}(\theta_{m,d}, \theta_{0,d}); \tag{6}$$

the problem related to global feature extractors $\theta_{0,d}$'s is:

$$\min_{\theta_{0,d}} \sum_{m=1}^M \sum_{d \in \mathcal{D}_m} \text{dist}(\theta_{m,d}, \theta_{0,d}); \tag{7}$$

and the third block related to the global inference model θ_0, Π_m is:

$$\min_{\theta_0, \{\Pi_m\}} \sum_{m=1}^M \text{dist}(\Pi_m \theta_0, w_m), \text{ s.t. } \Pi_m \mathbf{1} = \mathbf{1}, \Pi_m \geq 0, m = 1, \dots, M. \tag{8}$$

In view of its block structure, we propose a block coordinate descent type algorithm called Hybrid Federated Matched Averaging (HFMA) to solve this problem, which is summarized in Algorithm 1.

In each iteration, the local clients first fix Π_m and θ_0 , optimize the local objective function (6) with gradient-based local solvers such as SGD. The form of this local problem is similar to FedProx

[4] so we call it HFMA-Prox. Then the server aggregates the local feature extractors and optimize (7). For some common choice of distance function, such as the square of Euclidean norm, the problem has a closed-form solution. Finally, the server aggregates the local inference models and optimizes the global model matching problem (8). This subproblem can be optimized by another iterative procedure: a) randomly pick an index m and apply the Hungarian matching algorithm to find Π_m fixing the other $P_{i_{m'}}$'s; b) update θ_0 by fixed $\{\Pi_m\}$. After few rounds of update, we obtain the matched global model and the corresponding matching pattern. In practice, the matching is performed layer by layer. Dummy neurons with zero weights are padded to match the size of different models. Due to space limitation, we have to skip some technical details.

We have the following convergence results about the proposed algorithm.

Claim 1. *Suppose that for each $m \in \{1, 2, \dots, M\}$, f_m has Lipschitz continuous gradient w.r.t. θ_m and w_m ; further assume that $\text{dist}(\cdot, \cdot)$ has Lipschitz gradient w.r.t. each of its argument. Suppose that θ_0 has a fixed dimension and its size is bounded. Then Algorithm 1 converges to a first-order stationary solution for problem (5).*

This claim can be proved by observing that the proposed algorithm can be viewed as the classical block-coordinate gradient descent (BCGD) algorithm, so classical result can apply [13, 14]; It can also be extended to the case where the θ_m and w_m steps are not solved to global min, but to some approximate stationary solution of (6). Due to space limitation, we omit the detailed proofs.

Remark 1. We highlight the merits of the proposed approach: 1) Unlike the typical VFL formulations [6, 5], our approach keeps labels at the clients. Thus, the local problems are fully separable. There is no sample synchronization need during local updates; 2) By utilizing the proposed merging technique, we can generate a global model at the server, which makes use of the full featured data. This makes the testing stage flexible: the clients can process data with either partial feature (by using its local parameters $(\{\theta_{m,d}\}, w_m)$), or the full data set (by requesting $(\{\theta_{0,d}\}, w_0)$ from the server).

Remark 2. Our formulation (5) and our proposed algorithm naturally reduces to existing FL models. For example, in the extreme case, all the local clients have all the features and the matching patterns are fixed, and the distance function is chosen as ℓ_2 -norm square (i.e., $\text{dist}(a, b) = \|a - b\|_2^2$), then the HFMA becomes FedProx if the local problem is (6).

Remark 3. In practice, we perform inexact minimization for the local problem (6). As an alternative, locally we can ignore the regularizer and optimize the following local problem for a few iterations:

$$\min_{\theta_m, w_m} f_m(\theta_m, w_m) = \sum_{n \in \mathcal{N}_m} \ell(h(\theta_m; \mathbf{x}_m^n); w_m; y_m^n). \quad (9)$$

We name this alternative as HFMA-Avg. Compare with HFMA-Prox, the gradient estimation is easier, and it requires much less memory for the local clients.

3.1 Experiment Settings

To evaluate the proposed algorithms, we conducted experiments using the ModelNet40 data set. This data set is used for multiview object classification. It has totally 40,000 samples from 40 classes. Each sample consists of 12 views from different angels, which are features, of an object. During the experiment, we use $M = 4, 8$ clients, and each client has data from only partial views in some of the classes. We use the convolutional layers of Resnet-34 for feature extraction for HFMA and use an MLP with one hidden layer for local inference. For comparison, we also trained the model using the entire data set with all the features, which we label ‘‘centralized’’ in the figures.






Figure 2: Illustration of how many clients (the numbers in boxes) possess the training data of each view in each class, in an experiment with $M = 8$ clients, $D = 12$ views, and 40 classes.




(a) The testing loss value of Centralized training, HFMA-Avg and HFMA-Prox ($\mu = 0.1$) with respect to the number of communication rounds.




(b) The testing accuracy of of Centralized training, HFMA-Avg and HFMA-Prox ($\mu = 0.1$) with respect to the number of communication rounds.

Figure 3: Algorithm performance on training neural networks on the ModelNet40 data set with total 4 views and each client trained with 3 views on 30,000 data points.



(a) The testing loss value of Centralized training, HFMA-Avg, HFMA-Prox ($\mu = 0.1$), HFMA-Prox ($\mu = 0.5$) w.r.t. the number of communication rounds.



(b) The testing accuracy of of Centralized training, HFMA-Avg, HFMA-Prox ($\mu = 0.1$), HFMA-Prox ($\mu = 0.5$) w.r.t. the number of communication rounds.

Figure 4: Algorithm performance on training neural networks on the ModelNet40 data set with total 12 views and each client trained with 6 views on 15,000 data points.

In the training phase, the local clients train with their partial data and partial features. In the testing phase, the local clients will test their local model with partial features on all the samples and we average over all clients to obtain the averaged local accuracy. The global accuracy, is computed using the matched global model ($\theta_0, \{\theta_{m,0}\}$) on all samples with full features.

During the training, we uniformly set the total communication rounds $T = 32$ and local update step $Q = 32$ with HFMA-Avg with mini-batch size 32 on local clients. The initial learning rate is set to be $\eta = 0.005$ and it decays by 0.2 for every 8 rounds of communication.

In the case when $M = 4$, we select four views ($D = 4$) as the full feature. Each client has 30 classes of data, and has $|\mathcal{D}_m| = 3$ views. Therefore, none of the clients has full sample nor full features and the data distribution is heterogeneous. Figure 3 shows the result when $|\mathcal{D}_m| = 3$. In this setting the local model as well as the global model trained with HFMA can behave well on the data set and even obtained higher accuracy than the centralized training. We believe this is due to the data heterogeneity among the clients which helps prevent the model from over-fitting.

In the more complicated case when $M = 8$, we set the total number of views to $D = 12$. Each client only has 15 classes and 6 views, and the way we divide data and features are illustrated in Figure 2. From the figure, we can see that 12.08% of the data have never been used by any of the local clients during training and the local data sets are more heterogeneous than the case when $M = 4$. In addition, we set the penalty weight $\mu = 0.1$ and $\mu = 0.5$ for HFMA-Prox to understand how this parameter affects the global and the local accuracy. Figure 4 shows the testing loss and accuracy of the algorithms. We can see that the federated trained models behave worse than the centralized trained model, which is predictable because the high data heterogeneity and the missing data. We can also observe that μ balance between the global and the local accuracy. When μ is small ($\mu = 0.1$), the local accuracy is high and the global accuracy is low, and when μ becomes larger, local accuracy drops while the global accuracy improves. This is intuitive since by using larger parameter μ , we put more emphasis on the global model integration.

References

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Proc. Intl. Conf. Artificial Intell. and Stat.*, Fort Lauderdale, FL, Apr. 2017, pp. 1273–1282.
- [2] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al., “Advances and open problems in federated learning,” *arXiv preprint:1912.04977*, Dec. 2019.
- [3] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.
- [4] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [5] T. Chen, X. Jin, Y. Sun, and W. Yin, “Vaf: a method of vertical asynchronous federated learning,” *arXiv preprint arXiv:2007.06081*, 2020.
- [6] Y. Liu, Y. Kang, X. Zhang, L. Li, Y. Cheng, T. Chen, M. Hong, and Q. Yang, “A communication efficient vertical federated learning framework,” *arXiv preprint arXiv:1912.11187*, 2019.
- [7] S. P. Karimireddy, M. Jaggi, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh, “Mime: Mimicking centralized stochastic algorithms in federated learning,” *arXiv preprint arXiv:2008.03606*, 2020.
- [8] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh, “Scaffold: Stochastic controlled averaging for on-device federated learning,” *arXiv preprint arXiv:1910.06378*, 2019.
- [9] M. Yurochkin, M. Agarwal, S. Ghosh, K. Greenewald, N. Hoang, and Y. Khazaeni, “Bayesian nonparametric federated learning of neural networks,” in *International Conference on Machine Learning*, 2019, pp. 7252–7261.
- [10] H. Wang, M. Yurochkin, Y. Sun, D. Papailiopoulos, and Y. Khazaeni, “Federated learning with matched averaging,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [11] S. Hardy, W. Henecka, H. Ivey-Law, R. Nock, G. Patrini, G. Smith, and B. Thorne, “Private federated learning on vertically partitioned data via entity resolution and additively homomorphic encryption,” *arXiv preprint arXiv:1711.10677*, 2017.
- [12] J. Ma, Q. Zhang, J. Lou, J. C. Ho, L. Xiong, and X. Jiang, “Privacy-preserving tensor factorization for collaborative health data analysis,” in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2019, pp. 1291–1300.
- [13] J. Zeng, T. T.-K. Lau, S. Lin, and Y. Yao, “Global convergence of block coordinate descent in deep learning,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 7313–7323.
- [14] M. Razaviyayn, M. Hong, and Z.-Q. Luo, “A unified convergence analysis of block successive minimization methods for nonsmooth optimization,” *SIAM Journal on Optimization*, vol. 23, no. 2, pp. 1126–1153, 2013.