

An Aggregation-Free Federated Learning for Tackling Data Heterogeneity

Yuan Wang Huazhu Fu Renuga Kanagavelu Qingsong Wei Yong Liu Rick Siow Mong Goh
Institute of High Performance Computing (IHPC)
Agency for Science, Technology and Research (A*STAR), Singapore

Abstract

The performance of Federated Learning (FL) hinges on the effectiveness of utilizing knowledge from distributed datasets. Traditional FL methods adopt an aggregate-then-adapt framework, where clients update local models based on a global model aggregated by the server from the previous training round. This process can cause client drift, especially with significant cross-client data heterogeneity, impacting model performance and convergence of the FL algorithm. To address these challenges, we introduce FedAF, a novel aggregation-free FL algorithm. In this framework, clients collaboratively learn condensed data by leveraging peer knowledge, the server subsequently trains the global model using the condensed data and soft labels received from the clients. FedAF inherently avoids the issue of client drift, enhances the quality of condensed data amid notable data heterogeneity, and improves the global model performance. Extensive numerical studies on several popular benchmark datasets show FedAF surpasses various state-of-the-art FL algorithms in handling label-skew and feature-skew data heterogeneity, leading to superior global model accuracy and faster convergence.

1. Introduction

Federated Learning (FL) algorithms typically follow an iterative *aggregate-then-adapt* paradigm in which clients use their local private data to refine a global model provided by a central server. These locally updated models are subsequently returned to and aggregated by the server, where the global model is updated through averaging the local models according to predetermined rules [26]. Given the decentralized nature of client data in FL, substantial variations in the clients' data distribution are common. This scenario results in non-Independent and Identically Distributed (non-IID) data across clients, often referred to as data heterogeneity. Such heterogeneity presents considerable challenges to the convergence of FL algorithms, primarily due to the significant drift in local learning paths among clients [15, 30]. This phenomenon, known as client drift, can significantly

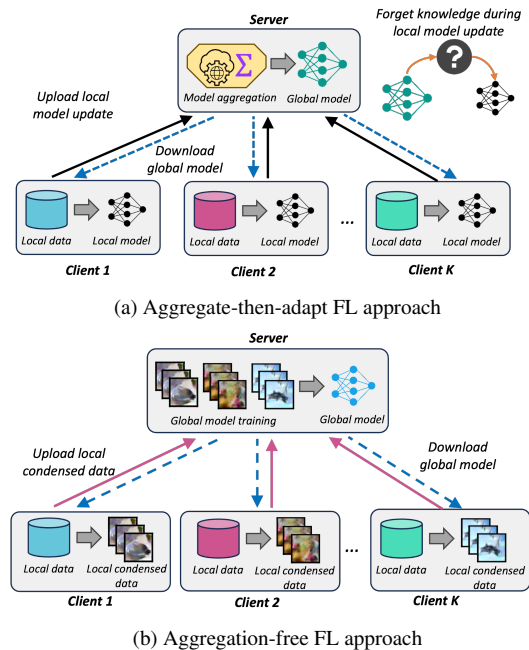


Figure 1. The conventional *aggregate-then-adapt* approach (a) is prone to client drift in data-heterogeneous scenarios, as clients update a downloaded global model and risk forgetting prior knowledge. In contrast, the aggregation-free paradigm (b) has the server train the global model directly using condensed synthetic data learned and shared by clients, which circumvents the client drift issue.

decrease the accuracy of the global model. [20, 39, 40].

Most existing efforts to address the challenge of data heterogeneity have focused either on modifying the local model training with additional regularization terms [1, 10, 15, 19, 22] or on utilizing alternative server-side aggregation or model update schemes [5, 13, 24, 29, 35, 41]. Nevertheless, these methods remain constrained by the conventional *aggregate-then-adapt* framework, as depicted in Figure 1a. In scenarios with strong cross-client non-IID data distribution, the fine-tuning of the global model over local data becomes susceptible to catastrophic forgetting. This means that clients tend to forget the knowledge learned in the global model and diverge away from the stationary point of global learning objective when they update the model in-

dividually [9, 14].

Recent works on data condensation [4, 31, 32, 36–38] suggest an *aggregation-free* paradigm that has the potential to overcome the limitations mentioned above in FL with privacy preservation. As depicted in Figure 1b, in this new framework, each client first learns a compact set of synthetic data (i.e., the condensed data) for each class and then shares the learned condensed data with the server. The server then utilizes the received condensed data to directly update the global model. However, in the current research on the *aggregation-free* method [25, 34], two critical open challenges emerge: First, significant cross-client data heterogeneity can compromise the quality of locally condensed data, adversely affecting the global model training. Second, relying exclusively on condensed data for global model training can result in reduced convergence performance and robustness, particularly when the quality of the received condensed data is sub-optimal.

Motivated by the above research gap, this paper presents FedAF, a novel *aggregation-free* FL algorithm tailored to combat data heterogeneity. At the heart of our research is the question of how to optimally harness the inherent knowledge in each client’s original data to enhance both local data condensation and global model training. To achieve this, we first introduce a *collaborative data condensation* scheme. In this scheme, clients condense their local dataset by minimizing a loss function that integrates a standard distribution matching loss [37] with an additional regularization term based on Sliced Wasserstein Distance (SWD). This regularization aligns the local knowledge distribution with the broader distribution across other clients, granting individual clients a more comprehensive perspective for data condensation. Furthermore, to train the global model with superior performance, we incorporate a *local-global knowledge matching* scheme. This approach enables the server to utilize not only the condensed data shared by clients but also soft labels extracted from their data, thereby refining and stabilizing the training process. As a result, the global model retains more knowledge from earlier rounds, leading to enhanced overall convergence performance.

Extensive experiments demonstrate that FedAF can consistently deliver superior model performance and accelerated convergence speed, outperforming several state-of-the-art FL algorithms across various degrees of data heterogeneity. For instance, on CIFAR10, we achieve up to 25.44% improvement in accuracy and 80% improvement in convergence speed, compared with the state-of-the-art methods. In summary, our contributions are threefold as follows:

- We propose a novel *aggregation-free* FL algorithm, termed FedAF, to tackle the challenge of data heterogeneity. Unlike traditional approaches that aggregate local model gradients, FedAF updates the global model using client-condensed data, thereby effectively circumventing

client drift issues.

- We introduce a collaborative data condensation scheme to enhance the quality of condensed data. By employing a Sliced Wasserstein Distance-based regularization, this scheme allows each client to leverage the broader knowledge in the data of other clients, a feature not adequately explored in existing literature.
- We further present a local-global knowledge matching scheme which equips the server with soft labels extracted from client data for enhanced global insights. This scheme supplements the condensed data received from clients, thereby facilitating improved model accuracy and accelerating convergence speed.

2. Background and Related Works

FL Algorithms for Heterogeneous Data. The foundational FedAvg algorithm [26], widely used in FL, calculates the global model by averaging the local models from each client. Among its variants, FedAvgM [13] adds server-side Nesterov momentum to enhance the global model update. FedNova [29] normalizes aggregation weights based on the amount of local computation. FedBN [23] specifically excludes batch normalization layer parameters from global aggregation. FedProx [21] integrates a proximal term in local training loss to mitigate the issue of client drift, while SCAFFOLD [15] employs variance reduction and a control variate technique to directly address client drift. FedDyn [1] allows clients to update the regularization in their local training loss to align more closely with the global empirical loss. In contrast, FedDC [10] proposes learning a drift variable to actively mitigate discrepancies between local and global model parameters. MOON [19] uses model-contrastive regularization to foster similarity in feature representations between the global and local models. Meanwhile, FedDF [24] and FedBE [5] focus on knowledge distillation-based model fusion and Bayesian model ensemble, respectively, to transfer knowledge into the global model. To forego the need for an unlabelled transfer dataset, FedGen [41] enables the server to learn and share a generator model with clients, facilitating knowledge distillation from the global model to local models through generated feature representations.

Data Condensation. Recent years have witnessed the rise of data condensation (or data distillation) techniques. These methods aim to compress a large training dataset into a significantly smaller set of synthetic data, enabling models trained on this condensed dataset to achieve performance comparable to those trained on the original dataset. For example, [32] explores a bi-level learning framework to learn the condensed data, allowing models trained on it to minimize the loss on the original data. [38] matches the gradients produced by training models on both the original and

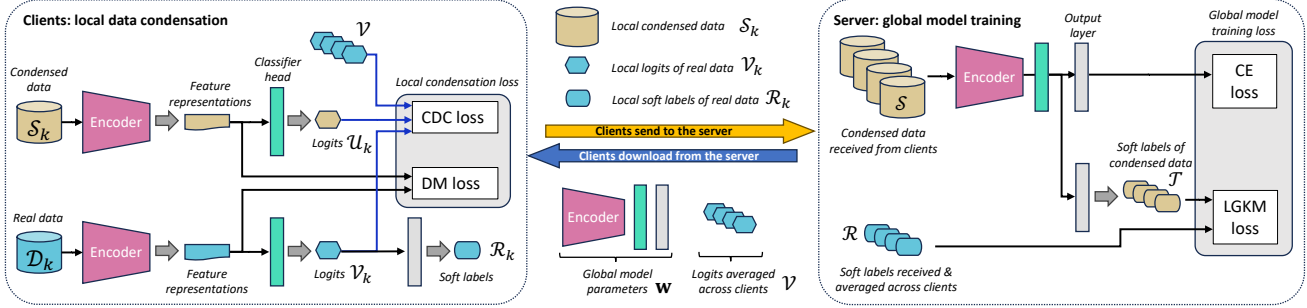


Figure 2. Overview of FedAF’s workflow. Left: Clients download the global model \mathbf{w} and the class-wise mean logits \mathcal{V} , averaged from \mathcal{V}_k at the server. They then update the condensed data \mathcal{S}_k using a combination of Distribution Matching (DM) loss and Collaborative Data Condensation (CDC) loss, with local real data \mathcal{D}_k and \mathcal{V} as inputs. Right: The server updates the global model \mathbf{w} by employing both cross-entropy loss and Local-Global Knowledge Matching (LGKM) loss. This utilizes both condensed data \mathcal{S}_k and soft labels \mathcal{R}_k received from each client $k \in \{1, 2, \dots, N\}$. The entire process iterates over a pre-defined number of communication rounds.

condensed data. This approach is further enhanced by [36], which applies differential Siamese augmentation to enable the learning of more informative condensed data. Differing from single-step gradient matching, [4] suggests matching multiple steps of training trajectories resulting from both the original and condensed data. To reduce the complexity of the expensive bi-level optimization used in earlier works, [37] introduces a distribution matching (DM) protocol. In this approach, condensed data is optimized to closely match the distribution of the original data in the latent feature space. Building upon DM, [31] further develops this concept by aligning layer-wise features between the real and condensed data.

Unlike Generative Adversarial Networks (GANs) [11], which focus on generating realistic-looking images, the objective of data condensation methods is to boost data efficiency by creating highly informative synthetic training samples. A recent study by [8] indicates that condensed data not only provides robust visual privacy, but also ensures that models trained on this data exhibit resistance to membership inference attacks.

The potential of condensed data has recently drawn interest in the FL community, prompting efforts to combine data condensation with FL within an aggregation-free framework. FedDM [34] implements a DM-based data condensation on the client side, with the server using condensed data from clients to approximate the original global training loss in FL. Employing the condensation method in [32] as the backbone, [25] introduces a dynamic weighting strategy for local data condensation and enhances global model training with pseudo data samples obtained from a conditional generator. However, both of these works do not adequately investigate how to utilize knowledge from other clients to improve the quality of local condensed data and the performance of the global model. While [25] allows the sharing of condensed data among clients, it relies on an assumption that all clients possess data of the same class, which may not be true under strong cross-client data heterogeneity.

3. Notations and Preliminaries

To facilitate a clearer understanding of our proposed method, we begin by introducing some essential notations and preliminaries.

Federated Learning. FL is a decentralized machine learning framework where K clients jointly learn a global model parameterized by \mathbf{w} without uploading local raw data $\mathcal{D} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K\}$, where $\mathcal{D}_k = \{(x_k^i, y_k^i)\}_{i=1}^{|\mathcal{D}_k|}$ is the local data owned by client k and $|\cdot|$ refers to the number of samples in \mathcal{D}_k . These clients are then set to learn the model parameter \mathbf{w} by solving the following problem collaboratively:

$$\arg \min_{\mathbf{w}} \mathcal{L}_{\text{glob}}(\mathbf{w}) = \sum_{i=k}^K p_k \mathcal{L}_k(\mathbf{w}, \mathcal{D}_k), \quad (1)$$

where $\mathcal{L}_k(\mathbf{w}, \mathcal{D}_k)$ is the local empirical loss of client k , which is given by:

$$\mathcal{L}_k(\mathbf{w}, \mathcal{D}_k) = \frac{1}{|\mathcal{D}_k|} \sum_{i=1}^{|\mathcal{D}_k|} \ell_k(\mathbf{w}, x_k^i, y_k^i), \quad (2)$$

where $\ell_k(\cdot)$ denotes the training loss such as cross-entropy loss. Weighting coefficient p_k associated with client k is proportional to $|\mathcal{D}_k|$ and normalized by $p_k = \frac{|\mathcal{D}_k|}{|\mathcal{D}|}$ so that $\sum_{k=1}^K p_k = 1$.

Data Condensation with Distribution Matching. Suppose a client k is tasked with learning a set of local condensed data denoted as \mathcal{S}_k . The client first initializes each class of condensed data by sampling from the local original data \mathcal{D}_k or Gaussian noise. Subsequently, the client employs a feature extractor, denoted as $h_{\mathbf{w}}(\cdot)$, which comprises all the layers preceding the last fully connected layer in a backbone model parameterized by \mathbf{w} , and extracts the feature representations from each class $c \in \{0, 1, \dots, C - 1\}$ of data in \mathcal{D}_k and \mathcal{S}_k . The means of these feature repre-

sentations are computed as below,

$$\mu_{k,c}^{\text{real}} = \frac{1}{N_{k,c}} \sum_{j=1}^{N_{k,c}} h_{\mathbf{w}}(x_{k,c}^j), \mu_{k,c}^{\text{syn}} = \frac{1}{M_{k,c}} \sum_{j=1}^{M_{k,c}} h_{\mathbf{w}}(\tilde{x}_{k,c}^j). \quad (3)$$

Here $x_{k,c}^j \in \mathcal{D}_k$ and $\tilde{x}_{k,c}^j \in \mathcal{S}_k$ represent the j -th sample of class c drawn from \mathcal{D}_k and \mathcal{S}_k , respectively. $N_{k,c}$ refers to the number of samples of class c in \mathcal{D}_k , whereas $M_{k,c}$ denotes the number of samples of class c in \mathcal{S}_k . Then client $k \in \{0, 1, \dots, K-1\}$ can learn and update \mathcal{S}_k by optimizing the DM loss as follows,

$$\arg \min_{\mathcal{S}_k} \mathcal{L}_{\text{DM}}(\mathcal{S}_k, \mathcal{D}_k) = \sum_{c=0}^{C-1} \|\mu_{k,c}^{\text{real}}(\mathcal{D}_k) - \mu_{k,c}^{\text{syn}}(\mathcal{S}_k)\|^2. \quad (4)$$

4. The Proposed Method

The overall mechanism of our proposed FedAF is illustrated in Figure 2. In each learning round, clients first engage in collaborative data condensation, updating their local condensed data. They then share this condensed data and soft labels with the server. Subsequently, the server uses this information to update the global model.

Collaborative Data Condensation. Other than the extraction of feature representations needed by DM loss in (4), we allow clients to compute class-wise mean logits and related soft labels from their local original data \mathcal{D}_k . With the latest global model updated by the server as the backbone network, clients perform the following:

$$\mathbf{v}_{k,c} = \frac{1}{N_{k,c}} \sum_{j=1}^{N_{k,c}} f_{\mathbf{w}}(x_{k,c}^j), \quad (5)$$

where $\mathbf{v}_{k,c}$ refers to the mean logit of class c computed by client k , function $f_{\mathbf{w}}(\cdot)$ denotes the global model parameterized by \mathbf{w} without the last softmax layer. Similarly, clients can compute class-wise mean logits of local condensed data $\mathbf{u}_{k,c}$ by

$$\mathbf{u}_{k,c} = \frac{1}{M_{k,c}} \sum_{j=1}^{M_{k,c}} f_{\mathbf{w}}(\tilde{x}_{k,c}^j). \quad (6)$$

To ease the presentation, we adopt the following notations for every client k :

$$\begin{aligned} \mathcal{V}_k &= [\mathbf{v}_{k,0}, \mathbf{v}_{k,1}, \dots, \mathbf{v}_{k,C-1}], \\ \mathcal{U}_k &= [\mathbf{u}_{k,0}, \mathbf{u}_{k,1}, \dots, \mathbf{u}_{k,C-1}]. \end{aligned} \quad (7)$$

Upon having the class-wise mean logits obtained, clients share \mathcal{V}_k with the server. The server then updates the global class-wise mean logits \mathbf{v}_c by

$$\mathbf{v}_c \leftarrow \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{v}_{k,c}. \quad (8)$$

We then empower clients to learn and update the local condensed data \mathcal{S}_k by downloading the following from the server

$$\mathcal{V} = [\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{C-1}], \quad (9)$$

and minimizing a local training loss as below:

$$\begin{aligned} \mathcal{L}_{\text{loc}}(\mathcal{S}_k, \{\mathcal{D}_k\}_{k=0}^{K-1}) &= \mathcal{L}_{\text{DM}}(\mathcal{S}_k, \mathcal{D}_k) \\ &+ \lambda_{\text{loc}} \sum_{c=0}^{C-1} \mathcal{F}(\mathbf{u}_{k,c}(\mathcal{S}_k), \mathbf{v}_c(\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K)), \end{aligned} \quad (10)$$

where we refer to the second term on the right-hand-side of (10) as the collaborative data condensation (CDC) loss. Without loss of generality, we let $\mathcal{F}(\cdot, \cdot)$ in the above loss function be a distance metric to promote the alignment between the local knowledge $\mathbf{u}_{k,c}$ and the global knowledge \mathbf{v}_c . Inspired by recent studies [2, 7, 17], we select the Sliced Wasserstein Distance (SWD) as the choice of $\mathcal{F}(\cdot, \cdot)$. SWD serves as an effective approximation of the exact Wasserstein distance [3, 16], enabling the efficient capture of discrepancies between the knowledge distributions of locally condensed data and the original data owned by other clients. This approach allows clients not only to match the distributions of condensed and original data in the latent feature space but also ensures a matching across clients in logit space. By leveraging the CDC loss as a regularization term, each client can learn condensed data with the assistance of global insights shared by their peer clients, thereby avoiding the pitfall of biased matching towards their local data and facilitating the learning of higher-quality condensed data.

Local-Global Knowledge Matching. The data condensation process inevitably leads to a certain degree of information loss from the original data. Consequently, relying solely on the condensed data received from clients for global model training might result in limited convergence performance or even instability. To address this, we introduce a local-global knowledge matching approach, enabling the server to harness a broader spectrum of knowledge about the original data distributed across clients.

To be more specific, we let clients compute local class-wise soft labels about its original data, represented by

$$\mathcal{R}_k = [\mathbf{r}_{k,0}, \mathbf{r}_{k,1}, \dots, \mathbf{r}_{k,C-1}], \quad (11)$$

where $\mathbf{r}_{k,c} = \sigma_{\mathbf{w}}(\mathbf{v}_{k,c}, \tau)$ and $\sigma_{\mathbf{w}}(\cdot, \tau)$ denotes the last softmax layer of the global model $f_{\mathbf{w}}$ softened by a temperature τ . Clients then share \mathcal{R}_k with the server in addition to the local condensed data \mathcal{S}_k , the server subsequently averages the shared local soft labels for each class as below

$$\mathbf{r}_c \leftarrow \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{r}_{k,c}. \quad (12)$$

At the same time, the server uses the received condensed data $\mathcal{S} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_K\}$ to compute global class-wise

soft labels \mathbf{t}_c through

$$\mathbf{t}_c = \sigma_{\mathbf{w}} \left(\frac{1}{|\mathcal{S}_c|} \sum_{j=1}^{|\mathcal{S}_c|} f_{\mathbf{w}}(\tilde{x}_c^j), \tau \right), \quad (13)$$

where $\tilde{x}_c^j \in \mathcal{S}_c$ and \mathcal{S}_c denotes the condensed data received by the server and belongs to class c . Then we empower the server to match \mathbf{t}_c with \mathbf{r}_c using a Kullback–Leibler (KL) divergence-based regularization term and propose the following loss function for training the global model:

$$\mathcal{L}_{\text{glob}}(\mathbf{w}, \mathcal{S}) = \mathcal{L}_{\text{CE}}(\mathbf{w}, \mathcal{S}) + \lambda_{\text{glob}} \mathcal{L}_{\text{LGKM}}(\mathbf{w}, \mathcal{S}), \quad (14)$$

where $\mathcal{L}_{\text{CE}}(\mathbf{w}, \mathcal{S})$ denotes the cross-entropy (CE) loss and $\mathcal{L}_{\text{LGKM}}(\mathbf{w}, \mathcal{S})$ refers to the local-global knowledge matching (LGKM) regularization, which is given by

$$\mathcal{L}_{\text{LGKM}}(\mathbf{w}, \mathcal{S}) = \frac{1}{2} \left(D_{\text{KL}}(\mathcal{R} \parallel \mathcal{T}) + D_{\text{KL}}(\mathcal{T} \parallel \mathcal{R}) \right). \quad (15)$$

\mathcal{R} and \mathcal{T} in the above equation are defined as

$$\mathcal{R} = [\mathbf{r}_0, \mathbf{r}_1, \dots, \mathbf{r}_{C-1}], \mathcal{T} = [\mathbf{t}_0, \mathbf{t}_1, \dots, \mathbf{t}_{C-1}]. \quad (16)$$

Augmented by the local-global knowledge matching, the server gains a more comprehensive understanding of the original data distributed across clients. By ensuring the global model better retains knowledge acquired in prior rounds, this enhancement not only stabilizes the global model updates but also curbs potential overfitting, especially when dealing with less-optimal condensed data.

Model Re-sampling. We also employ a global model re-sampling approach in each step of the local data condensation. Specifically, we interpolate the global model received from the server as below

$$\mathbf{w} \leftarrow \gamma \mathbf{w} + (1 - \gamma) \tilde{\mathbf{w}}, \quad (17)$$

where $\tilde{\mathbf{w}}$ represents randomly sampled model parameters. This model re-sampling technique aims to mitigate overfitting when learning condensed data by interpolating between the parameters of the global model learned in the previous round and a benign random perturbation.

5. Experiments

Datasets. We conduct experiments to evaluate and benchmark the performance of our proposed methods on both label-skew data heterogeneity and feature-skew data heterogeneity. For label-skew scenarios, we adopt Fashion-MNIST (FMNIST) [33], CIFAR10, and CIFAR100 [18]. The FMNIST dataset consists of 70,000 grayscale images of fashion products that fall into 10 categories. The CIFAR-10 and CIFAR-100 datasets are both collections of 60,000 color images. CIFAR10 images are grouped into 10 classes with 6,000 images per class, while CIFAR100 categorizes image into 100 classes with 600 images per class. For

feature-skew scenario, we use DomainNet [27] which contains 600,000 million images in six domains: Clipart, Infograph, Painting, Quickdraw, Real, and Sketch, each domain has data of 345 classes.

Baselines. We compare FedAF against the *aggregate-then-adapt* baselines, including typical FedAvg [26] and various state-of-the-art FL algorithms designed for handling data heterogeneity: FedProx [21], FedBN [23], MOON [19], FedDyn [1], and FedGen [41]. We also consider the prior work that also employs *aggregation-free* FL, namely FedDM [34], to demonstrate the effectiveness of our proposed collaborative data condensation and local-global knowledge matching schemes.

5.1. Results for Label-skew Data Heterogeneity

Configuration and Hyperparameters. We consider $K = 10$ clients and partition the training split of each benchmark dataset into multiple data shards to simulate the local training dataset for every client. Specifically, we leverage Dirichlet distribution to partition data among clients. For every benchmark dataset, we consider three degrees of data heterogeneity, represented by $\alpha=0.02$, $\alpha=0.05$, and $\alpha=0.1$, respectively. The hyperparameter α controls the strength of heterogeneity. Notably, a smaller α implies a higher non-IID in the data distribution among clients. We choose these α values to simulate harsh scenarios of data heterogeneity that can be encountered in real world applications.

For *aggregate-then-adapt* baseline methods, we adopt 10 local epochs with local learning rate of 0.01 and local batch size of 64. For both FedDM and FedAF, we adopt a local batch size of 256, local update steps of 1000, 50 image-per-class (IPC) for local data condensation, while using 500 epochs with a batch size of 256 and learning rate of 0.001 for global model training. We initialize each class of condensed data using the average of randomly sampled local original data. For FedAF, the image learning rate is set to 1.0, 0.1, and 0.2, for CIFAR10, CIFAR100, and FMNIST, respectively. For FedDM, we adopt an image learning rate of 1.0 and clip the norm of gradients at 2.0 to ensure its stability during local data condensation. The global model re-sampling coefficient γ is set to 0.9 whereas we set $\rho=5$ in FedDM (See Appendix A for more implementation details).

Model accuracy. We first evaluate the highest accuracy of the global model achieved by each algorithm within 20 communication rounds. Table 1 and Figure 3 show that, FedAF significantly outperforms all *aggregate-then-adapt* baselines in various settings, showing notable improvements in both mean accuracy and variance. A detailed analysis in Table 1 indicates that FedAF enhances performance compared to FedAvg by up to 25.44%, 17.91%, and 31.03% on CIFAR10, CIFAR100, and FMNIST, respectively. Even against FedDyn, the leading *aggregate-then-*

Methods	$\alpha = 0.02$			$\alpha = 0.05$			$\alpha = 0.1$		
	FMNIST	CIFAR10	CIFAR100	FMNIST	CIFAR10	CIFAR100	FMNIST	CIFAR10	CIFAR100
FedAvg	56.50±5.55	39.71±1.15	30.80±2.20	69.14±5.84	46.51±3.07	33.37±0.75	82.19±5.67	56.15±4.62	39.97±1.53
FedProx	60.38±5.00	36.46±5.39	30.82±0.80	69.33±4.12	45.83±2.23	36.61±1.44	81.56±4.52	58.54±1.87	40.45±1.53
FedBN	58.26±4.28	36.53±2.52	29.73±1.73	72.91±4.69	45.13±2.18	33.73±2.15	77.33±3.07	57.67±3.21	39.84±0.20
MOON	51.33±7.00	33.32±1.13	33.41±0.70	71.41±4.08	47.41±4.59	37.90±0.80	81.61±2.68	57.62±4.99	40.24±0.68
FedDyn	69.79±5.04	45.73±3.98	35.01±2.07	75.19±5.49	57.68±1.84	39.10±0.34	84.73±2.74	59.97±2.20	41.81±1.46
FedGen	61.44±2.07	36.61±1.06	29.20±2.09	75.48±1.83	42.72±2.11	33.56±3.91	82.29±2.53	58.17±2.84	40.23±1.06
FedDM	85.36±0.96	60.28±0.82	44.15±0.30	86.08±0.68	62.97±0.96	46.27±0.98	86.65±0.31	64.88±0.35	47.05±0.13
FedAF	87.53±0.32	65.15±0.86	48.71±0.33	87.29±0.23	67.50±0.76	49.49±0.33	87.91±0.41	69.11±0.86	50.61±0.26

Table 1. Comparison of global model accuracy achieved with various FL algorithms on FMNIST, CIFAR10, and CIFAR100 datasets. FedAF consistently outperforms all baseline methods across three degrees of data heterogeneity.

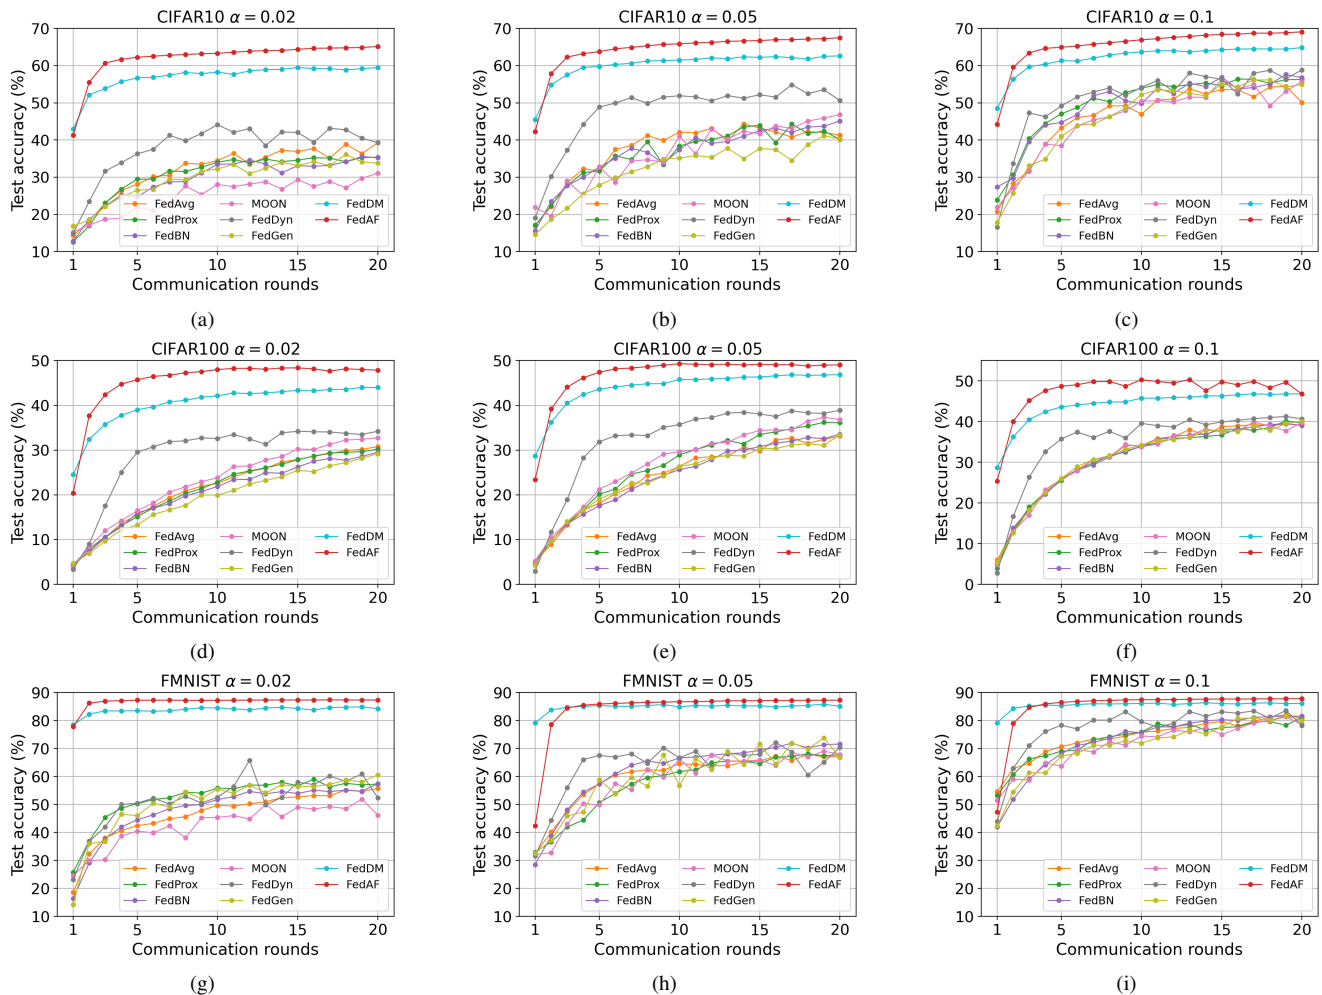


Figure 3. Comparison of convergence performance amongst baseline approaches. (a) to (c): learning curves obtained CIFAR10, (d) to (f): learning curves obtained on CIFAR100, (g) to (i): learning curves obtained on FMNIST. In addition to the the improvement in accuracy, FedAF also stands out to deliver considerably accelerated convergence speed, especially on harder dataset.

adapt baseline, FedAF maintains an edge of up to 19.43%, 13.70%, and 17.74% on the same datasets. Moreover, FedAF consistently outperforms FedDM, with accuracy advantages reaching 4.87%, 4.56%, and 2.17% on CIFAR10, CIFAR100, and FMNIST, respectively. Notably, FedAF’s performance is more pronounced under stronger data heterogeneity, such as at $\alpha = 0.02$, demonstrating the effective-

ness of our collaborative data condensation and local-global knowledge matching approaches.

Convergence Performance. We further examine and compare the convergence speed of FedAF with baselines, the learning curves resulting from three benchmark datasets are illustrated in Figure 3, which indicates that FedAF con-

Methods	DomainNet						
	C	I	P	Q	R	S	Avg
FedAvg	43.03	40.76	59.16	39.60	41.03	28.46	42.01
FedProx	44.81	43.76	60.22	38.13	41.55	29.18	42.94
FedBN	46.07	34.27	52.01	43.10	47.33	29.72	42.08
MOON	48.80	37.97	56.26	48.07	42.02	29.72	43.81
FedDyn	48.04	60.03	67.46	37.73	41.77	32.67	47.95
FedGen	42.77	37.88	54.37	37.33	42.86	25.69	40.15
FedDM	52.28	41.38	60.58	<u>62.37</u>	<u>52.45</u>	<u>46.69</u>	<u>52.62</u>
FedAF	<u>51.2</u>	<u>47.05</u>	<u>62.53</u>	64.6	52.64	50.06	54.68

Table 2. Comparison of global model accuracy across various FL algorithms on the DomainNet dataset. The domains are represented by C (Clipart), I (Infograph), P (Painting), Q (Quickdraw), R (Real), and S (Sketch), ‘‘Avg’’ denotes the average accuracy across domains. The highest and second-highest accuracies in each column are indicated by boldface and underline, respectively.




Figure 4. Comparison of convergence performance amongst baseline approaches on DomainNet dataset. FedAF also outperform the other baselines on both accuracy and convergence in feature-skew heterogeneous data distribution.

Configuration	$\alpha=0.02$	$\alpha=0.05$	$\alpha=0.1$
IPC=10	53.39 \pm 2.09	55.33 \pm 0.81	56.15 \pm 0.42
IPC=20	58.56 \pm 0.55	60.89 \pm 0.11	61.79 \pm 0.59
IPC=50	65.15 \pm 0.86	67.50 \pm 0.76	69.11 \pm 0.86
IPC=80	67.94 \pm 1.18	70.07 \pm 0.45	70.72 \pm 0.37
IPC=100	69.14 \pm 0.56	71.27 \pm 0.58	71.66 \pm 0.37

Table 3. Impact of IPC on the global model accuracy for learning CIFAR10 under three different degrees of heterogeneity.

sistently surpasses other baseline methods in convergence speed, particularly under significant data heterogeneity. For instance, at $\alpha = 0.02$, FedAF achieves the highest accuracy of other *aggregate-then-adapt* baselines within just two rounds. Compared to FedDM, FedAF also maintains a similar edge; on CIFAR10 with $\alpha = 0.02$, while FedDM reaches a mean accuracy of 60% in fifteen rounds, FedAF attains this target in only three rounds, marking an 80% increase in convergence speed.

5.2. Result for Feature-skew Data Heterogeneity

Configuration and Hyperparameters. For the experiments in feature-skew scenario, we follow [23] to form

a sub-dataset of DomainNet comprising only the top ten most frequent classes across all domains. We configure six clients, each holding data from a unique domain, to mimic real-world scenarios such as different hospitals using distinct imaging protocol and equipment, influencing the data heterogeneity. All algorithms are run for ten communication rounds to compare the resulting global model’s accuracy for every domain and the average accuracy across domains. For *aggregate-then-adapt* baselines, we maintain the same hyperparameters as in the label-skew scenarios. Both FedDM and FedAF use an image learning rate of 1.0, with the other settings such as batch size and number of local steps consistent with those in the label-skew scenarios (see Appendix A for more details).

Model Accuracy and Convergence Performance. From Table 2 and Figure 4, it is evident that FedAF outperforms all other baseline methods in average accuracy and related convergence performance across all domains. This comparison highlights FedAF’s consistent superiority, ranking as either the top or the second-best performer in every domain. Similar to the label-skew scenarios, *aggregate-then-adapt* FL approaches are found to be less effective compared to FedAF. Moreover, FedAF not only demonstrates higher accuracy but also exhibits faster convergence performance than FedDM. For instance, FedAF achieves the accuracy level in just two rounds that FedDM requires ten rounds to reach, indicating an 80% acceleration. These advantages underscore the benefit of integrating knowledge from other domains through our collaborative data condensation and local-global knowledge matching strategies, validating the effectiveness of these approaches.

5.3. Performance Analysis of FedAF

Impact of IPC. We conduct an ablation study on CIFAR10 to examine how variations in IPC might potentially influence the performance of FedAF under various degrees of data heterogeneity. We summarize the results in Table 3, and also illustrate and compare the resulting learning performance in Figure 5. It can be concluded that higher IPC values generally lead to higher accuracy. Besides, for each IPC value, the resulting model accuracy does not vary significantly over different values of α (with the gap between the highest and the lowest being 3.96%). From Figure 5d we find that although remarkable performance gain can be observed when IPC ranges from 10 to 50, this improvement starts to become marginal using an even higher IPC. Considering the increase in communication cost is approximately linear with the value of IPC. On the other hand, the lower the compression ratio (i.e., the ratio between the amount of condensed and original data), the higher the privacy retention [8]. Therefore, an IPC value lower than 50 should be generally considered as a trade-off between performance, communication cost, and privacy.




Figure 5. Impact of IPC on the learning performance of FedAF on CIFAR10. (a) to (c): the resulting learning curves. (d): improvement in model accuracy compared to FedAvg. Generally, a higher IPC correlates with enhanced performance, with all tested IPC values demonstrating improvements over FedAvg.

Configuration	$\alpha=0.02$	$\alpha=0.05$	$\alpha=0.1$
FedAF	65.15 \pm 0.86	67.50 \pm 0.76	69.11 \pm 0.86
w/o CDC	64.16 \pm 0.83	65.88 \pm 0.93	67.90 \pm 0.53
w/o LGKM	64.12 \pm 0.85	66.27 \pm 1.31	68.14 \pm 0.81
FedDM	60.28 \pm 0.82	62.97 \pm 0.96	64.88 \pm 0.35

Table 4. Impact of core design on the global model accuracy for learning CIFAR10 under three different degrees of heterogeneity. “w/o CDC” denotes the FedAF without collaborative data condensation, “w/o LGKM” denotes the FedAF without the local-global knowledge matching.

Impact of Core Designs. The collaborative data condensation and local-global knowledge matching act as two fundamental techniques of FedAF. We conduct additional experiments on CIFAR10 to further reveal how these two techniques contribute to performance improvement. Specifically, we compare the full FedAF with two other configurations where each fundamental technique is not utilized. We also compare it with FedDM, where none of these techniques exist. As shown in Table 4, the mean accuracy resulting from using just one of the fundamental techniques alone still shows noticeable improvement over that obtained with FedDM. Moreover, the full FedAF exhibits further improvement in the mean accuracy. These results verify that by promoting the utilization of additional knowledge derived from data distributed across clients, FedAF indeed can empower clients to learn higher quality condensed data and train the global model with improved performance, which eventually contributes to the enhancement in the overall learning performance.

Impact of Model Re-sampling. As the value of model re-sampling coefficient γ in (17) controls the interpolation between the parameters of the global model received from the server and those parameters randomly sampled, we analyze the effect of this technique by running ten rounds of FedAF on CIFAR10 with $\alpha=0.1$ and comparing the resulting accuracy over different choices of γ . Note that a larger γ indicates the re-sampled model tends to retain more knowledge learned in the global model from the previous round, while $\gamma=0$ implies clients use a model with randomly initialized parameters for data condensation. As expected, Ta-

γ	0.2	0.5	0.8	0.9	1.0
Accuracy	61.30	63.52	66.15	66.97	64.92

Table 5. The influence of global model re-sampling on learning performance. The first row lists the values of γ tested and the second row reports the corresponding accuracy.

ble 5 shows that larger γ values generally lead to higher model accuracy. However, the optimal γ is found to be 0.9, rather than the maximal possible value of 1.0, verifying that a suitable random perturbation to the true global model can indeed regulate the data condensation and enhance learning performance. Notably, $\gamma=0.8$ also yields comparable accuracy, suggesting that FedAF’s performance is robust to variations in γ . Additionally, we experimented with $\gamma=0$. However, using purely random weight parameters, the model struggled to stabilize the data condensation process, leading to its exclusion from the further comparison.

6. Conclusion

This paper presents FedAF, a novel FL algorithm designed to tackle data heterogeneity with an aggregation-free framework. FedAF grants clients and the server richer insights about the original data distributed across clients through collaborative data condensation and local-global knowledge matching. This strategy effectively addresses cross-client data heterogeneity and boosts the learning performance of both condensed data and the global model. Consequently, FedAF demonstrates considerable improvement over state-of-the-art FL methods in terms of model accuracy and convergence speed.

Acknowledgement. This work is supported by the Agency for Science, Technology and Research (A*STAR) under its IAF-ICP Programme (Award No: I2301E0020). This work is also supported by the National Research Foundation, Singapore under its AI Singapore Programme (Award No: AISG2-TC-2021-003). This work is also partially supported by A*STAR Central Research Fund “A Secure and Privacy Preserving AI Platform for Digital Health”. This work is also supported by A*STAR Career Development Fund (No. C222812010).

References

- [1] Durmus Alp Emre Acar, Yue Zhao, Ramon Matas Navarro, Matthew Mattina, Paul N Whatmough, and Venkatesh Saligrama. Federated learning based on dynamic regularization. *arXiv preprint arXiv:2111.04263*, 2021. [1](#), [2](#), [5](#)
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 214–223. PMLR, 2017. [4](#)
- [3] Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and radon wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51:22–45, 2015. [4](#)
- [4] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A Efros, and Jun-Yan Zhu. Dataset distillation by matching training trajectories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4750–4759, 2022. [2](#), [3](#)
- [5] Hong-You Chen and Wei-Lun Chao. FedBE: Making bayesian model ensemble applicable to federated learning. In *International Conference on Learning Representations*, 2021. [1](#), [2](#)
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. IEEE, 2009. [1](#)
- [7] Ishan Deshpande, Ziyu Zhang, and Alexander G Schwing. Generative modeling using the sliced wasserstein distance. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3483–3491, 2018. [4](#)
- [8] Tian Dong, Bo Zhao, and Lingjuan Lyu. Privacy for free: How does dataset condensation help privacy? In *International Conference on Machine Learning*, pages 5378–5396. PMLR, 2022. [3](#), [7](#)
- [9] Chun-Mei Feng, Bangjun Li, Xinxing Xu, Yong Liu, Huazhu Fu, and Wangmeng Zuo. Learning federated visual prompt in null space for mri reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8064–8073, 2023. [2](#)
- [10] Liang Gao, Huazhu Fu, Li Li, Yingwen Chen, Ming Xu, and Cheng-Zhong Xu. Feddc: Federated learning with non-iid data via local drift decoupling and correction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10112–10121, 2022. [1](#), [2](#)
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. [3](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#)
- [13] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019. [1](#), [2](#)
- [14] Wenke Huang, Mang Ye, and Bo Du. Learn from others and be yourself in heterogeneous federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10143–10153, 2022. [2](#)
- [15] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, pages 5132–5143. PMLR, 2020. [1](#), [2](#)
- [16] Soheil Kolouri, Yang Zou, and Gustavo K Rohde. Sliced wasserstein kernels for probability distributions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5258–5267, 2016. [4](#)
- [17] Soheil Kolouri, Phillip E. Pope, Charles E. Martin, and Gustavo K. Rohde. Sliced wasserstein auto-encoders. In *International Conference on Learning Representations*, 2019. [4](#)
- [18] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Toronto, ON, Canada, 2009. [5](#)
- [19] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10713–10722, 2021. [1](#), [2](#), [5](#)
- [20] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated learning on non-iid data silos: An experimental study. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 965–978. IEEE, 2022. [1](#)
- [21] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020. [2](#), [5](#)
- [22] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of fedavg on non-iid data. In *International Conference on Learning Representations*, 2020. [1](#)
- [23] Xiaoxiao Li, Meirui JIANG, Xiaofei Zhang, Michael Kamp, and Qi Dou. FedBN: Federated learning on non-IID features via local batch normalization. In *International Conference on Learning Representations*, 2021. [2](#), [5](#), [7](#), [3](#)
- [24] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. *Advances in Neural Information Processing Systems*, 33:2351–2363, 2020. [1](#), [2](#)
- [25] Ping Liu, Xin Yu, and Joey Tianyi Zhou. Meta knowledge condensation for federated learning. In *The Eleventh International Conference on Learning Representations*, 2023. [2](#), [3](#)
- [26] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. [1](#), [2](#), [5](#)
- [27] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019. [5](#)

- [28] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9 (86):2579–2605, 2008. [1](#)
- [29] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33:7611–7623, 2020. [1](#), [2](#)
- [30] Jianyu Wang, Zachary Charles, Zheng Xu, Gauri Joshi, H Brendan McMahan, Maruan Al-Shedivat, Galen Andrew, Salman Avestimehr, Katharine Daly, Deepesh Data, et al. A field guide to federated optimization. *arXiv preprint arXiv:2107.06917*, 2021. [1](#)
- [31] Kai Wang, Bo Zhao, Xiangyu Peng, Zheng Zhu, Shuo Yang, Shuo Wang, Guan Huang, Hakan Bilen, Xinchao Wang, and Yang You. Cafe: Learning to condense dataset by aligning features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12196–12205, 2022. [2](#), [3](#)
- [32] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018. [2](#), [3](#)
- [33] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. [5](#)
- [34] Yuanhao Xiong, Ruochen Wang, Minhao Cheng, Felix Yu, and Cho-Jui Hsieh. FedDM: Iterative distribution matching for communication-efficient federated learning. In *Workshop on Federated Learning: Recent Advances and New Challenges (in Conjunction with NeurIPS 2022)*, 2022. [2](#), [3](#), [5](#)
- [35] Lin Zhang, Li Shen, Liang Ding, Dacheng Tao, and Ling-Yu Duan. Fine-tuning global model via data-free knowledge distillation for non-iid federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10174–10183, 2022. [1](#)
- [36] Bo Zhao and Hakan Bilen. Dataset condensation with differentiable siamese augmentation. In *International Conference on Machine Learning*, pages 12674–12685. PMLR, 2021. [2](#), [3](#)
- [37] Bo Zhao and Hakan Bilen. Dataset condensation with distribution matching. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6514–6523, 2023. [2](#), [3](#), [1](#)
- [38] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In *International Conference on Learning Representations*, 2021. [2](#)
- [39] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018. [1](#)
- [40] Hangyu Zhu, Jinjin Xu, Shiqing Liu, and Yaochu Jin. Federated learning on non-iid data: A survey. *Neurocomputing*, 465:371–390, 2021. [1](#)
- [41] Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous federated learning. In *International conference on machine learning*, pages 12878–12889. PMLR, 2021. [1](#), [2](#), [5](#)