

FedSCAI: Leveraging Server and Client Alignment for Unsupervised Federated Source-Free Domain Adaptation

M. Yashwanth^{1*}, Sampath Koti^{1*}, Arunabh Singh^{1*}, Shyam Marjit^{1*}, and Anirban Chakraborty¹

¹ Indian Institute of Science, Bengaluru, India

{yashwanthm, sampathkoti, shyammarjit, anirban}@iisc.ac.in, arunabhsingh25@gmail.com

Project Page: <https://vcl-iisc.github.io/FedSCAI/>

Abstract

We address the Federated source-Free Domain Adaptation (FFreeDA) problem, with clients holding unlabeled data with significant inter-client domain gaps. The FFreeDA setup constrains the FL frameworks to employ only a pre-trained server model as the setup restricts access to the source dataset during the training rounds. Often, this source domain dataset has a distinct distribution to the clients' domains. To address the challenges posed by the FFreeDA setup, adaptation of the Source-Free Domain Adaptation (SFDA) methods to FL struggles with client-drift in real-world scenarios due to extreme data heterogeneity caused by the aforementioned domain gaps, resulting in unreliable pseudo-labels. In this paper, we introduce FedSCAI, an FL framework leveraging our proposed Server-Client Alignment (SCAI) mechanism to regularize client updates by aligning the clients' and server model's predictions. We observe an improvement in the clients' pseudo-labeling accuracy post alignment, as the SCAI mechanism helps to mitigate the client-drift. Further, we present extensive experiments on benchmark vision datasets showcasing how FedSCAI consistently outperforms state-of-the-art FL methods in the FFreeDA setup for classification tasks.

1. Introduction

Federated Learning (FL) initially proposed in [23] is a machine learning paradigm where multiple client models actively participate in the collective training of a shared global model, orchestrated by a central server. During the training phase, clients exchange their local client models with the server, thus ensuring the confidentiality of their respective local training data. The privacy-preserving nature of FL has led to its application across diverse domains, including smartphones [10, 32], Internet of Things [24, 26], and healthcare applications [2, 27, 33, 42].

In practice, clients participating in an FL framework will

*Equal contribution.




Figure 1. (a) FFreeDA Setup: Multiple clients (labeled as C_1 , C_2 , C_3 , and C_4) hold unlabeled data from same or different domains, while the central server uses a pre-trained model trained on a labeled source dataset, which differs from clients' data distributions and is unavailable during training. (b) Performance of different methods on Office-Home dataset.

be deployed in diverse environments, ingesting non-iid data leading to data heterogeneity across the client devices. Heterogeneity in data often arises due to variations in label distributions or feature distributions (covariate shift) across the clients. In their respective local training rounds, the clients optimize a task-specific objective function on their training data. Consequently, clients may drift towards the local minimum of their respective objectives. This phenomenon, 'client-drift' [16] is commonly observed when training data is non-iid across clients. As a result, client models tend to overfit their training data. Furthermore, aggregating these model parameters at the server results in a sub-optimal global model.

Prior works in FL have addressed label heterogeneity issues in [7, 16, 18, 19, 47] while covariate shift has been considered in [14, 20]. These methods assume that the client data is labeled, which is often impractical given the data privacy as well as logistical constraints associated with data annotation. Evidently, in practical real-world scenarios (as mentioned in Sec. 2.2 of the Supplementary), FL setups usually encounter substantial challenges with unlabeled client data and large domain gaps across the clients hindering the development of a robust aggregated model that generalizes well across diverse client environments. In this work, we take up the classification task in a challenging FL setting of FFreeDA introduced in [35]. In this setting, the training data samples on the clients are unlabeled

and, the server model is pre-trained on the source dataset. This dataset is unavailable during the federated training, and the FL is initialized by the pre-trained model available on the server. Along with inter-client domain shift, the source dataset used for pre-training the server model is likely to exhibit a large domain gap with each of the clients’ training data. This setting is clearly explained in Figure 1 (a).

Existing methods that require labeled source data on the server, such as DualAdapt [46], are not effective in the FFreeDA setting, where the server lacks access to labeled data. Even when access is permitted, these methods often fail to bridge significant domain gaps. More recent approaches such as LADD [35] for segmentation and FedWCA [25] for classification adopt clustering-based strategies tailored to the FFreeDA paradigm. However, their performance is dependant on the clustering quality, making them vulnerable when the cluster structure is unclear. In particular, LADD’s requirement for client-style information introduces privacy concerns that limit its applicability.

Given the unsupervised nature of FFreeDA, we turn to pseudo-labeling approaches based on source-free domain adaptation, such as SHOT [21], that allow each client to self-train using predictions from the current model. We begin with two simple baselines: Local Adaptation (LoA), where each client trains independently using pseudo-labels, and its federated counterpart FedLoA, which aggregates client models after local adaptation. As shown in Figure 1 (b), both methods perform poorly when the client domains are highly non-overlapping, due to inconsistent label assignments and client drift.

To overcome these issues, we propose FedSCAI, a new framework that replaces clustering with a Server and Client Alignment (SCAI) mechanism. Rather than grouping clients, FedSCAI aligns client predictions with those from a server model. This alignment leads to more stable and accurate pseudo-labels, thereby mitigating client drift and improving generalization across diverse domains.

Additionally, we propose a novel adaptive thresholding approach within our framework that adjusts the threshold used for determining the pseudo-labels dynamically based on the entropy of the predictions on training data samples. This further enhances the quality of pseudo labels used for training the client models across the training rounds. We discuss this in detail in Sec. 4.4.

Our key contributions are summarised as follows:

- **FedSCAI framework:** A novel framework to solve classification tasks in the FFreeDA [35] setting where only a pre-trained source model is available at the server, and clients hold unlabeled, heterogeneous data.
- **SCAI mechanism:** A server–client alignment mechanism within the FedSCAI framework that regularizes client predictions using the server model, mitigating client-drift and improving pseudo-label accuracy.

- **Adaptive Thresholding:** A dynamic approach for filtering low-confidence predictions, ensuring the reliability of pseudo-labels used for client updates.
- **Extensive Empirical Validation:** Experiments on multiple benchmarks (Office-Home, DomainNet, and Office-31) demonstrate substantial performance gains over the baselines and other state-of-the-art approaches.

2. Related Work

2.1. Federated Learning (FL)

Recently, addressing non-iid issues in Federated Learning has become a rapidly evolving research area. Here, we briefly discuss a few related works. In FedAvg [23], the two main challenges explored are reducing communication costs [43] and ensuring privacy by avoiding having to share the data. FedProx [19] introduced a proximal term by penalizing the weights if they are far from the global initialized model. SCAFFOLD [16] approached this issue by considering it as an objective inconsistency, and it introduced a gradient correction term designed to function as a regularizer. Later works of [7, 16] improved upon this by introducing the dynamic regularization term. A few works include regularization methods on the client side [49], and one-shot methods where clients send the condensed data and the server trains on the condensed data [48]. In [12], an adaptive weighting scheme is considered on task-specific loss to minimize the learning from samples whose representation is negligible. Flatness-based methods based on SAM called FedSAM is introduced in [3, 31, 38]. Federated Learning with partially labelled data is considered in the works of SemiFL [6, 15]. Model personalizations in FL are considered in [39] and the distillation-based approaches are also considered in [11, 17]. The clustering approaches in FL are discussed in [8, 22] where they require retraining the models repeatedly to obtain optimal clustering in a supervised learning setup.

2.2. Federated Learning with Distributional Shifts

FL has seen few developments concerning distributional shifts. The work of [29] introduces FL with distributional shifts, framing the challenge as adversarial domain adaptation within FL settings. Personalized FL for labelled data with distributional shifts is explored in [20]. FedGen [34] adopts a distillation approach, training generators on the client side to minimize the domain gap across clients. However, existing approaches, as outlined in [14, 20, 46], share common limitations. They rely on the assumption of labelled data across clients or the availability of related domain data with the server [46]. Methods such as DualAdapt [46] depends on a labelled source dataset during federated training, and LADD [35] assumes access to client data styles in the pre-training stage. Few prior works such

as Semifl [6] use consistency regularization inspired from Fixmatch [37], where the weak and strong augmentations are used to produce better pseudo-labels.

3. Preliminary

3.1. Pseudo-Label Estimation

In the FFreeDA setting, access to the labelled source domain data D^s is not available. Instead, the federation relies on a pre-trained source model f^s trained on dataset D^s to facilitate adaptation to the target domain D^t . In such settings, the principal objective is to learn a target model f^t which could correctly predict labels y_t for the unlabelled target domain samples $x_t \in D^t$ by adaptation from the source-trained model f^s without direct access to any labelled data from the source domain D^s . The challenge lies in leveraging the knowledge encoded within the pre-trained source model f^s to adapt effectively to the characteristics of the target domain D^t . The source model f^s can be studied as a composition of two integral components: a feature extractor $g^s : X_s \rightarrow \mathbb{R}^d$, where X_s is the feature distribution of source data, and a classifier $h^s : \mathbb{R}^d \rightarrow \mathbb{R}^J$. This architecture enables the transformation of input data x from the source domain to a d -dimensional feature space, followed by classification into J classes. In other words, the source model f^s can be represented as $f^s = h^s \circ g^s$. As a result of the domain gap between the server and client models we employ a pseudo-label estimation strategy as in [21] which helps in knowledge transfer via feature prototype-based pseudo-labeling strategy, where prototypes c_j are obtained for each class akin to weighted k-means clustering, computed as:=

$$c_j = \frac{\sum_{x_t \in D^t} \delta_j(f^t(x_t)) g^t(x_t)}{\sum_{x_t \in D^t} \delta_j(f^t(x_t))}, \quad (1)$$

where $f^t = h^t \circ g^t$ denotes the target model, $\delta_j(f^t(x_t^i))$ represents the softmax probability of target instance x_t^i belonging to the j -th class. The pseudo label \hat{y}_t^i for target x_t^i is determined via the nearest prototype classifier: $\hat{y}_t^i = \arg \min_j D_f(g^t(x_t^i), c_j)$ where $D_f(a, b)$ computes the cosine distance between a and b . Therefore, given the source model $f^s = h^s \circ g^s$ and pseudo labels generated as above, SHOT freezes the hypothesis from the source $h^t = h^s$ and learns the feature encoder g^t with the full objective Eq. 2. It jointly minimizes the Information Maximization (IM) loss ($L_{\text{ent}} + L_{\text{div}}$) [13, 36] together with L_{pce} , which is the pseudo-labeling loss with the full objective as

$$L^{\text{LoA}} = L_{\text{ent}} + L_{\text{div}} + L_{\text{pce}} \quad (2)$$

$$L_{\text{ent}} = -\frac{1}{|D^t|} \sum_{x_t \in D^t} \sum_{j=1}^J \delta_j(f^t(x_t)) \log \delta_j(f^t(x_t)) \quad (3)$$

$$L_{\text{div}} = \sum_{j=1}^J \hat{p}_j \log \hat{p}_j \quad (4)$$

$$L_{\text{pce}} = -\beta \frac{1}{|D^t|} \sum_{x_t \in D^t} \sum_{j=1}^J 1_{[j=\hat{y}_t]} \log \delta_j(f^t(x_t)) \quad (5)$$

where $\hat{p}_j = \frac{1}{|D^t|} \sum_{x_t \in D^t} [\delta_j(f^t(x_t))]$ and $\beta > 0$ is a balancing hyper-parameter. For the LoA baseline, each client k optimizes the L_k^{LoA} objective independently leveraging the pre-trained server model as the source.

3.2. Federated Learning

In FL, the goal is to solve the following optimization:

$$\arg \min_{\mathbf{w} \in \mathbb{R}^n} \left(l(\mathbf{w}) \triangleq \frac{1}{K} \sum_{k=1}^K l_k(\mathbf{w}) \right) \quad (6)$$

where K is the number of clients, $l_k(\mathbf{w})$ is the client's task specific objective function and \mathbf{w} denotes model parameters. It is solved in an iterative fashion spanning over T communication rounds between the server and clients using methods such as FedAvg [23].

4. Method

4.1. Problem Setup

We assume a single server and K clients and all the clients have the same model architecture. The server has access only to the pre-trained model trained on the source data with distribution $\mathbb{P}_s(x, y)$, and the k^{th} client has access to the unlabelled dataset \mathcal{D}_k , drawn from the distribution $\mathbb{P}_m(x, y)$, where $m \in \{1, 2, \dots, M\}$, and M is the number of domains. Now, $\mathbb{P}_m(x, y) = \mathbb{P}_m(y|x) \mathbb{P}_m(x)$ for features x and labels y belonging to domain m . Through this setup we tackle variations in the marginal distribution $\mathbb{P}_m(x)$ across the clients which we refer to as *covariate shift* or domain gap i.e for any two clients k and l , having sampled data from domains k' and l' , we have $\mathbb{P}_{k'}(x) \neq \mathbb{P}_{l'}(x)$. However, the problem is poorly defined if $\mathbb{P}_s(x, y)$ and $\mathbb{P}_m(x, y)$ are completely unrelated. So, we assume the source distribution $\mathbb{P}_s(x, y)$ and the domain distribution $\mathbb{P}_m(x, y)$ follow common domain adaptation assumptions [1, 4]. We consider the scenario where the datasets on any pair of clients belong to the same or different domains.

4.2. Pseudo-Label Estimation in Unsupervised Federated Settings

FedSCAL aims to generate high-quality, robust pseudo labels with the help of the pre-trained server model assisting clients to train models on their unlabelled data via federation. To begin with, the estimation of pseudo-labels in a federated setting operates by deploying the LoA method to client k which then yields the local client loss L_k^{LoA} . In each communication round r , the server broadcasts the

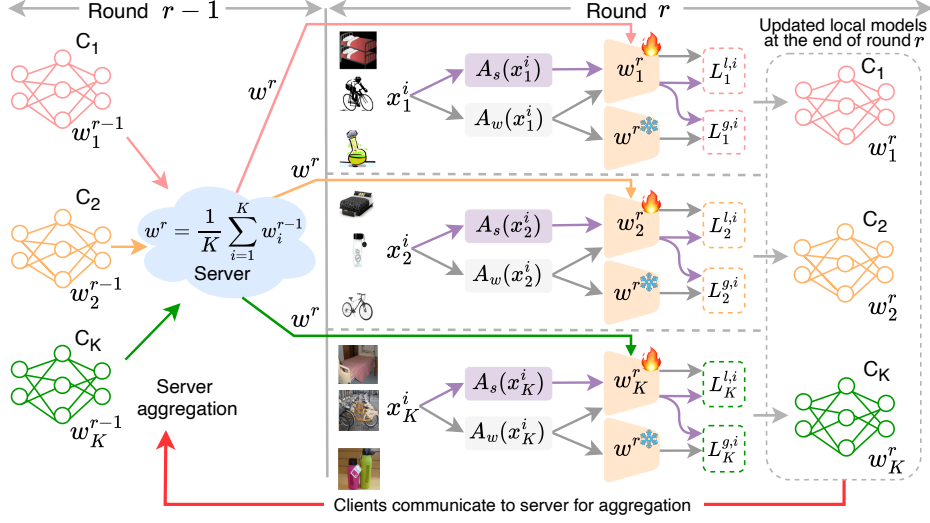


Figure 2. Overview of the proposed **FedSCAI framework with Server and Client Alignment (SCAI)** implementation. The server communicates the global model \mathbf{w}^r to the clients, the clients then train their local models (\mathbf{w}_k^r) using a weak augmentation A_w giving us the augmented image $A_w(x_k^i)$, and a strong augmentation A_s from which we get the augmented image $A_s(x_k^i)$. The clients compute the client alignment loss: $L_k^{l,i}$ and the server alignment loss: $L_k^{g,i}$ to perform the local training using the Eq. 12, and then communicate the updated local models \mathbf{w}_k^r to the server in the subsequent round for aggregation.

global model \mathbf{w}^r to participating clients who independently optimize their empirical local training loss $l_k(\mathbf{w}) := \frac{1}{|\mathcal{D}_k|} \sum_{i=1}^{|\mathcal{D}_k|} L_k^{LoA}(\mathbf{w}; x_k^i)$ where $x_k^i \in \mathcal{D}_k$. Each client k optimizes the loss for a few local epochs E and sends the model updates \mathbf{w}_k^r to the server. In federated averaging, the server simply aggregates these model updates as $\mathbf{w}^{r+1} = \frac{1}{K} \sum_{k=1}^K \mathbf{w}_k^r$. This iterates until convergence and gives us the FedLoA baseline. However, the key bottleneck with FedLoA is the degradation in pseudo-label quality due to large domain gaps across the clients which results in client-drift. To address this we propose a novel Server and Client Alignment (SCAI) mechanism within the FedSCAI framework.

4.3. Improving Pseudo-label Accuracy via SCAI

We propose an alignment of the local client prediction with the prediction of the global model via our Server and Client Alignment (SCAI) mechanism. As aforementioned, clients with domain gap suffer from client drift, leading to overfitting on their local objective losses. To avoid this we focus on utilizing predictions consistently aligned with the global model. The alignment loss for client k in a communication round is explained. Suppose $\delta(f(x_k^i; w^r))$ denotes the prediction of the server model on a data point x_k^i and similarly $\delta(f_k(x_k^i; w_k^r))^*$ denotes the client model prediction on x_k^i . Here, $f_k(\cdot)$ represents the output of the classifier and $\delta(\cdot)$ is the softmax operator. We obtain the two predictions on single weakly augmented sample $A_w(x_k^i)$ of the orig-

inal sample x_k^i as $\delta(f_k(A_w(x_k^i)))$ and $\delta(f(A_w(x_k^i)))$ respectively. Similarly, we obtain the prediction of strongly augmented sample $A_s(x_k^i)$ as $\delta(f_k(A_s(x_k^i)))$. We match the one-hot predictions of the weak augmentations from the client model f_k (i.e., $\delta(f_k(A_w(x_k^i)))$) and the server model f (i.e., $\delta(f(A_w(x_k^i)))$) using the alignment loss as described below. We denote the client alignment loss of the client model's weak and strong augmentations as $L_k^{l,i}$, where

$$L_k^{l,i} = \mathbb{I}_{\{\delta_{\max}(f_k(A_w(x_k^i))) > \tau\}} \mathcal{D}_{KL}(p(x_k^i) || q(x_k^i)) \quad (7)$$

and τ is the threshold, $q(x_k^i) := \delta(f_k(A_s(x_k^i)))$, $p(x_k^i) := \delta_{\text{oh}}(f_k(A_w(x_k^i)))$, and \mathcal{D}_{KL} is the KL divergence. The threshold τ is employed to match the most confident predictions, this is discussed in the Sec. 4.4. $\delta_{\text{oh}}(f_k(A_w(x_k^i)))$ represents the one-hot prediction of the softmax vector $\delta(f_k(A_w(x_k^i)))$. δ_{\max} denotes the maximum of the softmax outputs and \mathbb{I} is the indicator function. Averaging across the batches yields the total client alignment loss as below.*

$$L_k^l = \frac{1}{\sum_{i=1}^B \mathbb{I}_{\{\delta_{\max}(f_k(A_w(x_k^i))) > \tau\}}} \sum_{i=1}^B L_k^{l,i} \quad (8)$$

where B denotes the size of the mini-batch. Similarly, the server alignment loss denoted as $L_k^{g,i}$ is

$$L_k^{g,i} = \mathbb{I}_{\{\delta_{\max}(f(A_w(x_k^i))) > \tau\}} \mathcal{D}_{KL}(r(x_k^i) || q(x_k^i)) \quad (9)$$

*For simplicity, the parameterizations w_r and w_r^k are omitted in subsequent discussions.

*If none of the predictions cross the threshold i.e., $\sum_{i=1}^B \mathbb{I}_{\{\delta_{\max}(f_k(A_w(x_k^i))) > \tau\}} = 0$, then L_k^l is set to 0

where $r(x_k^i) := \delta_{\text{oh}}(f(A_w(x_k^i)))$. Averaging across the batches yields the total server alignment loss* as

$$L_k^g = \frac{1}{\sum_{i=1}^B \mathbb{I}_{\{\delta_{\text{max}}(f(A_w(x_k^i))) > \tau\}}} \sum_{i=1}^B L_k^{g,i} \quad (10)$$

The overall alignment loss is the sum of the server and client alignment losses and is given by

$$L_k^{SCAI} = \lambda_l * L_k^l + \lambda_g * L_k^g \quad (11)$$

Finally, our local client loss takes the form, here L_k^{LoA} refers to client k adapting the Eq. 2.

$$l_k(\mathbf{w}) = L_k^{LoA} + L_k^{SCAI} \quad (12)$$

We demonstrate the impact of the client and server alignment loss. Figure 3a shows the pseudo-label accuracy difference Δ_{pAcc} (Eq. 13) with and without SCAI loss.

$$\Delta_{pAcc} = pAcc(\text{FedSCAI}) - pAcc(\text{FedLoA}) \quad (13)$$

where $pAcc(\text{FedSCAI})$ gives the pseudo-label accuracy when we use our proposed FedSCAI with Server and Client Alignment (SCAI) mechanism and $pAcc(\text{FedLoA})$ gives the pseudo-label accuracy for the adaptation of LoA method to a federated setup. In Figure 3a, it can be seen that when we deploy our FedSCAI framework with the SCAI loss (Eq. 11), the pseudo-label accuracy keeps improving over the communication rounds.

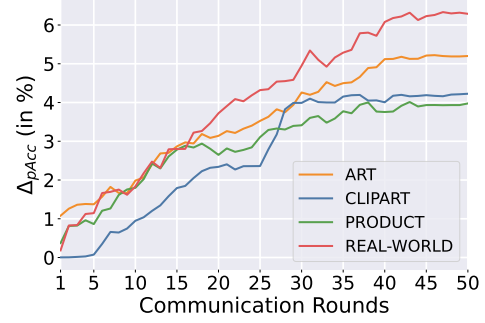
4.4. Adaptive Thresholding

In Eq. 7, we employ the threshold τ to match the most confident predictions. However, setting a fixed value for τ during the federated training may encounter certain drawbacks. Firstly, during the initial phases of training, the client models may not be very confident in their predictions, necessitating the initial relaxation of the threshold, followed by tightening as the client models' predictions grow in confidence. Another issue may arise due to client data heterogeneity; client models deployed on certain domains may have more or less confidence in their predictions than client models deployed on other domains. This motivates us to employ adaptive thresholding during the training cycles.

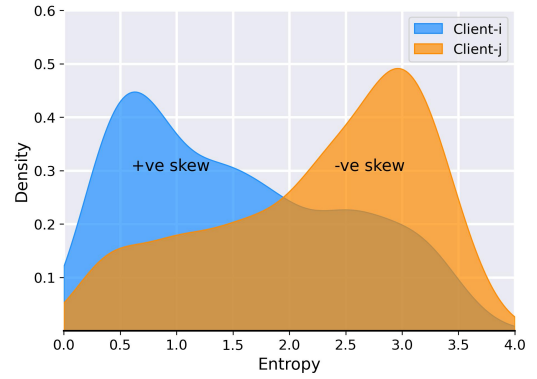
We exploit the skewness of the entropy (Eq. 14) distribution of the predictions on the client data samples, to dynamically adjust the threshold.

$$H(\delta(f_k(x_k^i))) = - \sum_j \delta_j(f_k(x_k^i)) \log(\delta_j(f_k(x_k^i))) \quad (14)$$

*If none of the samples cross the threshold i.e., $\sum_{i=1}^B \mathbb{I}_{\{\delta_{\text{max}}(f(A_w(x_k^i))) > \tau\}} = 0$, then L_k^g is set to 0.



(a) Pseudo-label prediction accuracy difference Δ_{pAcc} (Eq. 13) improves with the proposed SCAI mechanism on the Office-Home dataset, when the training is initialized with the server model pre-trained on different source domains (Art, ClipArt, Product and Real-World).



(b) Entropy density of the clients predictions on Office-Home dataset, calculated via (Eq. 14), exhibiting skewness (γ); with negative γ (e.g., $\gamma = -0.1$) indicating higher entropy and lower prediction confidence, while positive γ (e.g., $\gamma = 0.15$) reflecting lower entropy and higher prediction confidence.

Figure 3. Pseudo-label Accuracy difference and visualization of Entropy Density of Clients Predictions.

where δ_j is the softmax confidence that x_k^i belongs to class j . We denote the μ_H , M_H , and σ_H as the mean, median, and standard deviation of the entropy distribution, respectively. The skewness γ ([28] and [5]) is given by following

$$\gamma = 3 * \frac{\mu_H - M_H}{\sigma_H} \quad (15)$$

A positive value for the skewness γ (Figure. 3b) implies an entropy distribution skewed towards lower entropy values. This signals that the client model is confident in its predictions. Likewise, a negative value for γ would imply that the distribution is skewed towards higher entropy values meaning the client model is less confident in its predictions. For further enhancing the pseudo-label estimation using FedSCAI, we tighten the threshold for confident client models and relax it for the less confident ones. We start with an initial threshold of τ_{init} and then dynamically update τ during training as

$$\tau = \tau_{init} + \max(\gamma_L, \min(\gamma, \gamma_H)) \quad (16)$$

Table 1. Results on **Office-Home** (Art: **A**, ClipArt: **C**, Real-World: **R**, and Product: **P**) Dataset, while the initial server model is pre-trained on one of the 4 domains and the clients are distributed with the remaining 3 domains such that each client contains the subset of exactly one domain. We report the **Accuracy(%)** for each target domain and the average of all the domains for each pre-trained model. It can be seen that FedSCAI improves the performance significantly across all the domains. Here ‘ \rightarrow ’ denotes Source \rightarrow Target.

Method	(a) Initial Sever Model is pre-trained on Art				(b) Initial Sever Model is pre-trained on Clipart				(c) Initial Sever Model is pre-trained on Product				(d) Initial Sever Model is pre-trained on Real-World			
	A \rightarrow C	A \rightarrow P	A \rightarrow R	Avg.	C \rightarrow A	C \rightarrow P	C \rightarrow R	Avg.	P \rightarrow A	P \rightarrow C	P \rightarrow R	Avg.	R \rightarrow A	R \rightarrow C	R \rightarrow P	Avg.
LoA	44.95	79.75	84.86	69.85	73.86	80.45	84.03	79.45	69.54	46.42	85.98	67.31	74.85	48.12	85.03	69.34
FedLoA	56.69	83.56	86.38	75.55	78.97	84.33	86.24	83.18	75.92	52.82	87.08	71.94	77.55	55.01	86.68	73.08
FedProx	57.95	83.48	86.26	75.90	79.81	85.10	86.12	83.67	73.98	54.93	85.71	71.54	77.12	57.06	85.92	73.38
FedMOON	56.23	85.08	87.13	76.15	79.42	85.67	87.34	84.15	75.66	53.34	86.78	71.93	78.86	55.26	86.59	73.57
FedWCA	53	84.02	86.72	74.58	78.88	84.83	86.68	83.46	75.43	52.39	86.95	71.59	76.59	54.25	86.12	72.32
LADD	37.50	65.21	71.23	57.98	53.90	66.24	70.42	63.52	61.21	42.46	79.09	60.92	66.30	40.80	78.36	61.82
Dual Adapt	45.85	76.56	84.11	68.84	70.58	78.06	81.99	76.88	68.92	44.33	84.73	65.99	73.79	46.17	84.1	68.02
FedSCAI (Ours)	61.22	87.16	88.76	79.05	81.13	88.19	88.69	86.01	79.39	58.48	88.81	75.56	79.68	60.96	89.12	76.59

To ensure that we are not left with too few prediction class labels for a very high threshold τ while also ensuring we do not sample noisy predictions for a very small τ , we clip the threshold in the range of $[\gamma_L, \gamma_H]$. The detailed algorithm is provided in Sec. 2 of the Supplementary. The impact of the adaptive threshold is demonstrated in Table 4.

5. Experiments

5.1. Datasets and Setup

For all our experiments, we choose an ImageNet pre-trained ViT small model. We consider the cross-silo setting inspired from FedBN [20] and pFedPG [44], where a data domain is assigned to a client. We consider an even difficult setting, wherein each domain is distributed across a small number of clients and set a participation rate (< 1) allowing only a few clients to participate during each communication round. We consider Office-Home and DomainNet datasets for our experiments. **Office-Home** [40] dataset has four domains namely Art, Clipart, Product, and Real-World. While we consider one domain as a source, the other 3 domains are distributed among 15 clients (5 each for a particular domain). The participation rate is set to 0.3 for the experiment on *Office-Home* dataset. **Domain-Net** [9] dataset has six domains namely Clipart, Infograph, Painting, Quickdraw, Real, and Sketch. We follow the protocol of [20] for adapting the Domain-Net dataset to the federated setup. While we consider one domain as a source, the other 5 domains are distributed among 10 clients (2 each for a particular domain). The participation rate is set to 0.3. Local training epochs for each client are kept at 5 for both datasets with a learning rate of 0.3.

Augmentations: Our weak augmentations consist of resizing the image to a uniform size, followed by a random resized crop and horizontal flip, which are then transformed

to tensors and normalized. For strong augmentations, we randomly choose a pair of augmentations to be applied sequentially to the image from a diverse pool. This augmentation pool includes transformations as used in FixMatch [37].

5.2. Baselines and Evaluation

To establish baselines, we include LoA (Section 3.1) and FedLoA (Section 4.2), and compare against LADD [35], DualAdapt [45], FedProx [19], FedMOON [18], FedWCA [25], and our proposed FedSCAI. We report the best average accuracy (%), by averaging the accuracies of clients belonging to each domain. By testing on the client data, our evaluation captures both performance and generalization capabilities. The hyper-parameter settings and implementation details are provided in the Sec. 7 of the supplementary material.

5.3. Results and Discussion

Office-Home: In Table 1, we present the results for the Office-Home dataset, for all the combinations of server and client domains. For example, Table 1(a) shows that the server has the model pre-trained on the Art domain, and the clients are distributed with Clipart, Product, and Real-World data. It can be seen that LADD performs inferior to other baseline methods. The LoA baseline also performs inferior to baselines such as FedLoA, key issue being the training on data limited to a single client. In the case of FedLoA, the performance degrades due to client-drift. FedSCAI performs the best as it mitigates the client-drift by employing the Server and Client Alignment(SCAI) mechanism and also improves the generalization performance. FedSCAI achieves 3.50% improvement over the FedLoA baseline for the Art pre-trained server model (Table 1(a)), and 3.51% improvement for the Real-World pre-trained server model (Table 1(d)). For the Clipart and Product

Table 2. Results on **DomainNet** Dataset (Clipart: **C**, Infograph: **I**, Product: **P**, Sketch: **S**, Quickdraw: **Q**, and Real: **R**), while the initial server model is pre-trained on one of the 6 domains and clients are distributed with the remaining 5 domains such that each client contains the subset of exactly one domain. We report **accuracy(%)** for each target domain and the average of all domains for each pre-trained model. It is evident that adding our SCAI regularizer boosts the performance across all the domains. Here ‘ \rightarrow ’ denotes Source \rightarrow Target.

Method	(a) Initial Server Model is pre-trained on Clipart						(b) Initial Server Model is pre-trained on Infograph						(c) Initial Server Model is pre-trained on Painting					
	C \rightarrow I	C \rightarrow P	C \rightarrow Q	C \rightarrow R	C \rightarrow S	Avg.	I \rightarrow C	I \rightarrow P	I \rightarrow Q	I \rightarrow R	I \rightarrow S	Avg.	P \rightarrow C	P \rightarrow I	P \rightarrow Q	P \rightarrow R	P \rightarrow S	Avg.
LoA	49.00	93.90	70.5	95.57	86.77	79.15	80.24	92.21	54.32	94.35	82.30	80.80	79.93	49.43	33.22	94.88	83.30	68.15
FedLoA	52.56	93.69	73.49	96.13	87.03	80.58	80.66	91.62	70.79	95.95	83.86	84.58	81.93	54.42	68.91	95.80	81.06	76.42
FedProx	56.78	92.75	70.35	95.89	87.95	80.74	80.63	91.58	61.71	95.86	83.59	82.67	79.15	54.50	52.44	95.17	82.54	72.81
FedMOON	56.31	93.06	64.75	96.21	88.15	79.7	81.77	92.46	56.25	96.34	83.00	81.97	78.04	52.92	36.24	96.09	82.58	69.17
FedWCA	57.13	93.77	73.37	96.19	88.5	81.79	80.48	92.20	67.44	96.21	82.48	83.76	80.36	54.49	64.03	95.94	80.80	75.12
LADD	53.66	91.28	34.69	95.83	82.72	71.64	66.11	80.52	26.77	85.21	65.59	64.84	60.94	36.67	10.14	94.33	72.86	54.99
DualAdapt	54.46	90	35.36	95.59	85.33	72.15	67.35	81.09	19.63	86.52	69.37	64.79	73.05	47.17	10	95.01	80.29	61.1
FedSCAI (Ours)	60.1	94.78	78.8	96.53	92.72	84.58	87.09	94.78	76.01	96.42	88.94	88.65	85.96	59.23	61.43	96.60	88.78	78.4
Method	(d) Initial Server Model is pre-trained on Quickdraw						(e) Initial Server Model is pre-trained on Real						(f) Initial Server Model is pre-trained on Sketch					
	Q \rightarrow C	Q \rightarrow I	Q \rightarrow P	Q \rightarrow R	Q \rightarrow S	Avg.	R \rightarrow C	R \rightarrow P	R \rightarrow Q	R \rightarrow I	R \rightarrow S	Avg.	S \rightarrow C	S \rightarrow P	S \rightarrow Q	S \rightarrow R	S \rightarrow I	Avg.
LoA	69.68	14.91	39.15	56.38	75.25	51.07	70.46	85.87	59.43	46.51	73.90	67.23	85.20	92.70	77.60	95.47	48.58	79.91
FedLoA	45.21	29.56	50.02	64.08	52.35	48.24	72.58	86.09	70.43	45.57	76.09	70.15	86.33	93.38	72.91	96.29	56.02	80.99
FedProx	59.74	35.76	51.33	57.34	60.86	53.01	71.79	86.75	51.97	48.55	73.49	66.51	87.2	91.62	67.72	95.57	53.94	79.21
FedMOON	60.34	30.44	52.90	66.35	56.41	53.29	76.73	90.24	58.71	53.35	79.6	71.73	86.48	93.64	68.83	96.18	54.75	79.98
FedWCA	41.49	32.70	46.22	54.74	43.29	43.68	72.51	87.63	64.51	48.63	74.83	69.62	86.90	92.80	73.04	95.99	51.33	80.02
LADD	52.72	25.94	37.19	50.80	52.95	43.92	51.79	74.70	10.90	32.01	61.64	46.21	79.54	89.83	42.67	95.26	53.49	72.15
DualAdapt	59.78	22.02	40.68	51.85	49.58	44.78	69.72	86.14	13.43	46.54	73.36	57.84	86.29	91.67	39.53	95.76	56.08	73.86
FedSCAI (Ours)	65.37	40.15	65.74	66.03	65.75	60.61	81.74	92.33	76.84	53.45	86.03	78.08	89.96	95.02	77.61	96.53	58.89	83.60

pre-trained server models, FedSCAI yields average performance gains of 2.83% and 3.62%, respectively, compared to FedLoA (Table 1(b,c)). Table 1 further compares FedSCAI with FedProx [19] and Model Contrastive (MOON) [18]. FedSCAI consistently outperforms both FedProx and FedMOON across all target domains. Specifically, with the Art pre-trained server model, FedSCAI delivers an average improvement of 3.15% over FedProx and 2.90% over FedMOON. Comparable improvements are observed over other methods, including FedWCA, LADD, and Dual Adapt.

DomainNet: The results on the DomainNet dataset are presented in Table 2. With the Clipart pre-trained server model (Table 2(a)), the proposed method achieves an average improvement of 4% over the FedLoA baseline, with consistent gains across all target domains and a maximum increase of 7.54% on the Infograph domain. For the Infograph pre-trained server model (Table 2(b)), our approach yields an average improvement of 4.07% over FedLoA, reaching up to 6.43% on the Clipart domain. Using the Painting pre-trained server model (Table 2(c)), the method delivers an average improvement of 1.98%, peaking at 7.72% on the Sketch domain. When initialized with the Quickdraw pre-trained server model (Table 2(d)), the proposed method achieves a substantial average improvement of 9.54% over the LoA baseline, showing its effectiveness.

Comparable performance gains are observed with server models pre-trained on other domains such as Real and Sketch, demonstrating the robustness and adaptability of the approach across diverse visual domains. Furthermore, FedSCAI consistently outperforms competing federated learning algorithms, including FedProx and FedMOON, in all evaluated settings. For example, with the Real domain pre-trained server model, FedSCAI achieves an average improvement of 11.57% over FedProx and 6.35% over FedMOON. These consistent improvements across initializations, including Painting, Quickdraw, and Clipart, show the generalization ability of the proposed method for federated learning with heterogeneous data under the FFreeDA setup. Office-31 results are in Sec. 6 of the supplementary.

5.4. Analysis of SCAI Mechanism

Server and Client Alignment Losses: We conduct a detailed ablation study on our proposed SCAI loss. We consider the Office-Home dataset for this experiment. As illustrated in Table 3, each component of our method contributes significantly to the performance gain. Introducing the server alignment loss boosts the FedSCAI performance to around $\approx 1.90\%$ over FedSCAI with only the client alignment loss.

Varying the threshold τ_{init} : For all experiments, we set $\tau_{init} = 0.8$, $\gamma_L = -0.1$ and $\gamma_H = 0.15$ as a design choice. This configuration enables fine-grained threshold

Table 3. Ablation study of SCAI Losses on **Office-Home** dataset.

L_k^i	L_k^o	A→{C, P, R}	C→{A, P, R}	P→{C, A, R}	R→{C, P, A}	# Avg.
✗	✗	75.55	83.18	71.94	73.08	75.93
✓	✗	77.13	84.32	74.12	74.25	77.46
✗	✓	77.33	84.63	74.28	74.41	77.66
✓	✓	79.05	86.01	75.56	76.59	79.30

'→' denotes Source → Target; Art (A), ClipArt (C), Product (P), and Real-World (R).

Table 4. SCAI with Adaptive vs Fixed Thresholding.

Thresholding	C→ {R,I,P,Q,S}	I→ {R,C,P,Q,S}	P→ {R,I,C,Q,S}	Q→ {R,I,P,C,S}	R→ {C,I,P,Q,S}	S→ {R,I,P,C,Q}
Fixed	82.77	87.77	78.19	58.63	75.81	83.4
Adaptive	84.58	88.65	78.40	60.61	78.08	83.60

'→' denotes Source → Target; Clipart: C, Infograph: I, Product: P, Sketch: S, Quickdraw: Q, and Real: R

adjustments, resulting in lower and upper bounds of 0.7 and 0.95, respectively. If larger clipping values were used (e.g., $\gamma_L = \gamma_H = 0.6$), the threshold would span from 0.2 to 1.4. At the upper bound of 1.4, the SCAI loss remains inactive, whereas at the lower bound of 0.2, it may be triggered by noisy predictions, potentially degrading performance. The rationale for selecting $\tau_{init} = 0.8$ was to mitigate the influence of noisy samples at the onset of training. To illustrate the impact of varying τ_{init} while keeping γ_L and γ_H fixed, we present an experiment in Fig. 4 using the Office-Home dataset. In this figure, we plot the average accuracy of all the clients when the training is initialized with the server model pre-trained on different source domains such as Art, Clipart, Product, and Real-World.

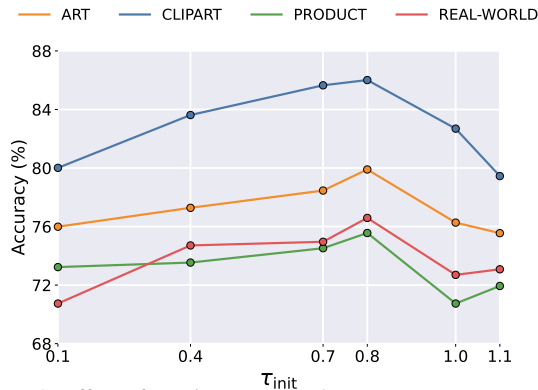


Figure 4. Effect of varying τ_{init} on the average accuracy attained by all the clients when the training is initialized with the server model pre-trained on Art, Clipart, Product, and Real-World.

Adaptive Vs Fixed Threshold: In Table 4, we report average accuracies on *Domain-Net* for the proposed adaptive threshold compared to the fixed threshold. We demonstrate an improvement of 1.81% when the initial server model is pre-trained with Clipart and an improvement of 1.98% when the initial server model is pre-trained using Quickdraw with adaptive thresholding. Similar improvements can be observed for other domains with adaptive thresholding improving performance for all the transfer settings. **On the Choice of LoA:** In our main experiments, we used LoA method as SHOT due to its simplicity and as well

as to ensure fair comparison with SOTA methods such as FedWCA [25], which are designed atop of it. Here, we showcase experiments with alternative LoA methods such as BMD [30] and UCon [41] in Tables 5 and 6, respectively. This demonstrates that FedSCAI framework works independently of the underlying LoA method. Due to space constraints, we provide only the average accuracy from each source; the detailed splits of every source-target pairs and the Domain-Net results are provided in the Sec. 4 of the supplementary material.

Table 5. Comparison of FedSCAI and FedLoA using **BMD** as the underlying LoA method.

Method	A→{C, P, R}	C→{A, P, R}	P→{C, A, R}	R→{C, P, A}	# Avg.
FedLoA	77.52	83.98	74.35	74.72	77.64
FedSCAI	80.47	86.47	77.15	76.50	80.15

'→' denotes Source → Target; Art (A), ClipArt (C), Product (P), and Real-World (R).

Table 6. Comparison of FedSCAI and FedLoA using **UCon** as the underlying LoA method.

Method	A→{C, P, R}	C→{A, P, R}	P→{C, A, R}	R→{C, P, A}	# Avg.
FedLoA	72.00	78.55	67.25	71.85	72.41
FedSCAI	73.93	81.77	71.77	74.58	75.51

'→' denotes Source → Target; Art (A), ClipArt (C), Product (P), and Real-World (R).

Further details on the computation/communication is provided in the Sec. 9 of the supplementary.

6. Limitations and Conclusion

Our work introduces FedSCAI, a novel framework for federated learning (FL) designed to address the classification task under the FFreeDA setup. We observe that existing baselines such as FedLoA can suffer from client-drift issues in this setting, primarily due to domain shifts and lack of supervision. To tackle this, FedSCAI incorporates a Server and Client Alignment (SCAI) mechanism, which enables clients to align with the server’s pre-trained model while preserving local adaptation. In doing so, FedSCAI not only mitigates client drift but also presents a clustering-free alternative that addresses the limitations of imperfect clustering in prior methods. The effectiveness of our framework is demonstrated through extensive experiments on the Office-Home, DomainNet, and Office-31 datasets, showing consistent and substantial improvements over baseline and state-of-the-art approaches tailored for this setup. A key limitation of our current approach is that it assumes a predefined number of classes across clients. While this assumption simplifies the setup, it may limit applicability in environments where class distributions evolve over time. Addressing this challenge is an important direction for future work.

Acknowledgement: This work is supported by the P3DX project, seed funded by the Ministry of Electronics and Information Technology (MeitY). The authors would also like to acknowledge compute and travel supports received from Kotak-IISc AI-ML Centre (KIAC), IISc and the PMRF fellowship.

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Vaughan. A theory of learning from different domains. *Machine Learning*, 79:151–175, 2010. 3
- [2] Theodora S Brisimi, Ruidi Chen, Theofanie Mela, Alex Olshevsky, Ioannis Ch Paschalidis, and Wei Shi. Federated learning of predictive models from federated electronic health records. *International journal of medical informatics*, 112:59–67, 2018. 1
- [3] Debora Caldarola, Barbara Caputo, and Marco Ciccone. Improving generalization in federated learning by seeking flat minima. In *European Conference on Computer Vision*, pages 654–672. Springer, 2022. 2
- [4] Shai Ben David, Tyler Lu, Teresa Luu, and David Pal. Impossibility theorems for domain adaptation. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 129–136, Chia Laguna Resort, Sardinia, Italy, 2010. PMLR. 3
- [5] José Moral de la Rubia. Standardized distance from the mean to the median as a measure of skewness. *Open J. Stat.*, 13 (03):359–378, 2023. 5
- [6] Enmao Diao, Jie Ding, and Vahid Tarokh. Semifl: Semi-supervised federated learning for unlabeled clients with alternate training. In *Advances in Neural Information Processing Systems*, 2022. 2, 3
- [7] 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
- [8] Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. An efficient framework for clustered federated learning. *Advances in Neural Information Processing Systems*, 33:19586–19597, 2020. 2
- [9] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2066–2073, 2012. 6
- [10] Andrew Hard, Chloé M Kiddon, Daniel Ramage, Françoise Beaufays, Hubert Eichner, Kanishka Rao, Rajiv Mathews, and Sean Augenstein. Federated learning for mobile keyboard prediction, 2018. 1
- [11] Yuting He, Yiqiang Chen, Xiaodong Yang, Yingwei Zhang, and Bixiao Zeng. Class-wise adaptive self distillation for heterogeneous federated learning. 2022. 2
- [12] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Federated visual classification with real-world data distribution. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pages 76–92. Springer, 2020. 2
- [13] Weihua Hu, Takeru Miyato, Seiya Tokui, Eiichi Matsumoto, and Masashi Sugiyama. Learning discrete representations via information maximizing self-augmented training. In *International conference on machine learning*, pages 1558–1567. PMLR, 2017. 3
- [14] Wenke Huang, Mang Ye, Zekun Shi, He Li, and Bo Du. Rethinking federated learning with domain shift: A prototype view. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16312–16322, 2023. 1, 2
- [15] Wonyong Jeong, Jaehong Yoon, Eunho Yang, and Sung Ju Hwang. Federated semi-supervised learning with inter-client consistency & disjoint learning. In *International Conference on Learning Representations (ICLR) 2021*. International Conference on Learning Representations (ICLR), 2021. 2
- [16] 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
- [17] Gihun Lee, Minchan Jeong, Yongjin Shin, Sangmin Bae, and Se-Young Yun. Preservation of the global knowledge by not-true distillation in federated learning. In *Advances in Neural Information Processing Systems*, 2022. 2
- [18] 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, 6, 7
- [19] 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. 1, 2, 6, 7
- [20] 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. 1, 2, 6
- [21] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning (ICML)*, pages 6028–6039, 2020. 2, 3
- [22] Guodong Long, Ming Xie, Tao Shen, Tianyi Zhou, Xi-anzhi Wang, and Jing Jiang. Multi-center federated learning: clients clustering for better personalization. *World Wide Web*, 26(1):481–500, 2023. 2
- [23] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. 1, 2, 3
- [24] Jed Mills, Jia Hu, and Geyong Min. Communication-efficient federated learning for wireless edge intelligence in iot. *IEEE Internet of Things Journal*, 7(7):5986–5994, 2019. 1
- [25] Junki Mori, Kosuke Kihara, Taiki Miyagawa, Akinori F Ebihara, Isamu Teranishi, and Hisashi Kashima. Federated source-free domain adaptation for classification: Weighted cluster aggregation for unlabeled data. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 6879–6889. IEEE, 2025. 2, 6, 8
- [26] Dinh C Nguyen, Ming Ding, Pubudu N Pathirana, Aruna Seneviratne, Jun Li, and H Vincent Poor. Federated learning

- for internet of things: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 23(3):1622–1658, 2021. 1
- [27] Dinh C Nguyen, Quoc-Viet Pham, Pubudu N Pathirana, Ming Ding, Aruna Seneviratne, Zihuai Lin, Octavia Dobre, and Won-Joo Hwang. Federated learning for smart healthcare: A survey. *ACM Computing Surveys (CSUR)*, 55(3): 1–37, 2022. 1
- [28] K Pearson. III. contributions to the mathematical theory of evolution. *Philos. Trans. R. Soc. Lond. A*, 185(0):71–110, 1894. 5
- [29] Xingchao Peng, Zijun Huang, Yizhe Zhu, and Kate Saenko. Federated adversarial domain adaptation. In *International Conference on Learning Representations*, 2019. 2
- [30] Sanqing Qu, Guang Chen, Jing Zhang, Zhijun Li, Wei He, and Dacheng Tao. Bmd: A general class-balanced multicentric dynamic prototype strategy for source-free domain adaptation. In *European Conference on Computer Vision*, pages 165–182. Springer, 2022. 8
- [31] Zhe Qu, Xingyu Li, Rui Duan, Yao Liu, Bo Tang, and Zhuo Lu. Generalized federated learning via sharpness aware minimization. In *International Conference on Machine Learning*, pages 18250–18280. PMLR, 2022. 2
- [32] Swaroop Ramaswamy, Rajiv Mathews, Kanishka Rao, and Françoise Beaufays. Federated learning for emoji prediction in a mobile keyboard. *arXiv preprint arXiv:1906.04329*, 2019. 1
- [33] Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N Galtier, Bennett A Landman, Klaus Maier-Hein, et al. The future of digital health with federated learning. *NPJ digital medicine*, 3(1):1–7, 2020. 1
- [34] Hywoon Seo, Jihong Park, Seungeun Oh, Mehdi Bennis, and Seong-Lyun Kim. 16 federated knowledge distillation. *Machine Learning and Wireless Communications*, page 457, 2022. 2
- [35] Donald Shenaj, Eros Fanì, Marco Toldo, Debora Caldarola, Antonio Tavera, Umberto Michieli, Marco Ciccone, Pietro Zanuttigh, and Barbara Caputo. Learning across domains and devices: Style-driven source-free domain adaptation in clustered federated learning. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 444–454, 2023. 1, 2, 6
- [36] Yuan Shi and Fei Sha. Information-theoretical learning of discriminative clusters for unsupervised domain adaptation. In *Proceedings of the 29th International Conference on Machine Learning*, pages 1275–1282, 2012. 3
- [37] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 3, 6
- [38] Yan Sun, Li Shen, Tiansheng Huang, Liang Ding, and Dacheng Tao. Fedspeed: Larger local interval, less communication round, and higher generalization accuracy. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [39] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 2
- [40] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5018–5027, 2017. 6
- [41] Gezheng Xu, Hui Guo, Li Yi, Charles Ling, Boyu Wang, and Grace Yi. Revisiting source-free domain adaptation: a new perspective via uncertainty control. In *The Thirteenth International Conference on Learning Representations*, 2025. 8
- [42] Jie Xu, Benjamin S Glicksberg, Chang Su, Peter Walker, Jiang Bian, and Fei Wang. Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5: 1–19, 2021. 1
- [43] Sarika Yadav and Rama Shankar Yadav. A review on energy efficient protocols in wireless sensor networks. *Wireless Networks*, 22(1):335–350, 2016. 2
- [44] Fu-En Yang, Chien-Yi Wang, and Yu-Chiang Frank Wang. Efficient model personalization in federated learning via client-specific prompt generation, 2023. 6
- [45] Chun-Han Yao, Boqing Gong, Hang Qi, Yin Cui, Yukun Zhu, and Ming-Hsuan Yang. Federated multi-target domain adaptation. In *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1081–1090, 2022. 6
- [46] Chun-Han Yao, Boqing Gong, Hang Qi, Yin Cui, Yukun Zhu, and Ming-Hsuan Yang. Federated multi-target domain adaptation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1424–1433, 2022. 2
- [47] M. Yashwanth, Gaurav Kumar Nayak, Harsh Rangwani, Arya Singh, R. Venkatesh Babu, and Anirban Chakraborty. Minimizing layerwise activation norm improves generalization in federated learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2287–2296, 2024. 1
- [48] Yanlin Zhou, George Pu, Xiyao Ma, Xiaolin Li, and Dapeng Wu. Distilled one-shot federated learning. *arXiv preprint arXiv:2009.07999*, 2020. 2
- [49] 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. 2