

FedSA: A Unified Representation Learning via Semantic Anchors for Prototype-based Federated Learning

Yanbing Zhou¹, Xiangmou Qu², Chenlong You¹, Jiyang Zhou¹, Jingyue Tang¹, Xin Zheng¹,
Chunmao Cai^{1,3}, Yingbo Wu^{1*}

¹Chongqing University

²OPPO Research Institute

³Chongqing Changan Automobile Co., Ltd

{202124021011, wyb}@cqu.edu.cn, lokinko@oppo.com, {ycl}@stu.cqu.edu.cn

Abstract

Prototype-based federated learning has emerged as a promising approach that shares lightweight prototypes to transfer knowledge among clients with data heterogeneity in a model-agnostic manner. However, existing methods often collect prototypes directly from local models, which inevitably introduce inconsistencies into representation learning due to the biased data distributions and differing model architectures among clients. In this paper, we identify that both statistical and model heterogeneity create a vicious cycle of representation inconsistency, classifier divergence, and skewed prototype alignment, which negatively impacts the performance of clients. To break the vicious cycle, we propose a novel framework named **Federated Learning via Semantic Anchors (FedSA)** to decouple the generation of prototypes from local representation learning. We introduce a novel perspective that uses simple yet effective semantic anchors serving as prototypes to guide local models in learning consistent representations. By incorporating semantic anchors, we further propose anchor-based regularization with margin-enhanced contrastive learning and anchor-based classifier calibration to correct feature extractors and calibrate classifiers across clients, achieving intra-class compactness and inter-class separability of prototypes while ensuring consistent decision boundaries. We then update the semantic anchors with these consistent and discriminative prototypes, which iteratively encourage clients to collaboratively learn a unified data representation with robust generalization. Extensive experiments under both statistical and model heterogeneity settings show that FedSA significantly outperforms existing prototype-based FL methods on various classification tasks.

Introduction

Federated Learning (FL) has been proposed to collaboratively train a global model to address the increasingly significant privacy concerns (Li et al. 2020; Kairouz et al. 2021). Despite its advantages, traditional FL methods (McMahan et al. 2017) struggle with *statistical heterogeneity*, where heterogeneous data among clients can bias the global model towards clients with dominant data characteristics (Wang et al. 2020). To remedy this, personalized FL (Tan et al. 2022a) aims to train a personalized model for each client by

leveraging the benefits of FL rather than pursuing a global model. Many researchers have concentrated on model-based FL approaches that split the model into a body for universality and a head for personalization, requiring the same model architectures across all clients to facilitate the aggregation of the body at the server (Oh, Kim, and Yun 2021). However, *model heterogeneity* frequently arises when clients may design their own local model architectures to meet individual requirements and hardware constraints. Such model-based FL approaches not only incur significant communication costs but also risk exposing model details, which may further raise concerns about privacy and commercial proprietary information (Ye et al. 2023; Wang et al. 2023).

To address these issues, knowledge-based FL approaches have emerged as a novel FL paradigm that transfers various types of global knowledge among clients with heterogeneous data and diverse model architectures (Tan et al. 2022b). For example, knowledge distillation (KD)-based FL methods transfer logits outputs from a teacher model as global knowledge to instruct the student models (Li and Wang 2019; Lin et al. 2020; Zhang et al. 2021). However, these methods require a public dataset to align the outputs and are highly dependent on the quality of this dataset (Zhang et al. 2023a). Researchers further explore KD in a data-free manner by employing additional auxiliary models as global knowledge (Wu et al. 2022; Zhang et al. 2022), but the communication costs for transmitting the auxiliary models remain considerable. Recently, prototype learning has garnered increasing interest due to its exemplar-driven nature and intuitive interpretation (Snell, Swersky, and Zemel 2017). Prototypes, which are essentially averages of class representations, serve as abstract concepts that can effectively integrate feature representations from diverse data distributions. By transferring lightweight prototypes as global knowledge, prototype-based FL methods significantly address privacy concerns and reduce communication costs (Tan et al. 2022b,c; Zhang et al. 2024).

However, existing prototype-based FL methods directly collect prototypes from biased data distributions of clients, which inherently introduces inconsistencies into representation learning (Zhou, Zhang, and Tsang 2023). Moreover, the prototypes extracted from different model architectures have diverse scales and separation margins (Zhang et al. 2024), further exacerbating these inconsistencies. These is-

*Prof. Yingbo Wu is the corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

sues demonstrate that the traditional prototype generation is sub-optimal for heterogeneous FL since prototypes are heavily bound to the representation learning process and data distribution across clients, which can potentially lead to learning collapse (Ge et al. 2024). To investigate the impact of representation inconsistencies on the generation of prototypes, we revisit and reinterpret the biases arising from heterogeneity in feature extractors and classifiers, from the perspective of representation learning (Guo, Tang, and Lin 2023): 1) Biased feature extractors tend to produce inconsistent data representations in the semantic space; 2) Biased classifiers often diverge, learning skewed decision boundaries due to local class distributions; 3) Skewed prototype alignment occurs as local prototypes, generated from these inconsistent representations, naturally exhibit a skewed alignment in the semantic space. The naive averaging aggregation of these skewed local prototypes (Tan et al. 2022b) leads to the margin shrink problem, where the global prototypes tend to converge closely together, reducing the overall separability and effectiveness in distinguishing between different classes (Zhang et al. 2024). Moreover, the global prototypes then serve to guide the representation learning of local models, potentially creating a *vicious cycle* of representation inconsistency, classifier divergence, and skewed prototype alignment. Specifically, representation inconsistency causes classifier updates to diverge and prototype alignment to skew. Subsequently, these diverged classifiers and skewed prototypes force the feature extractors to map to more inconsistent representation space, thereby exacerbating the vicious cycle.

Motivated by these insights, we propose a novel framework named **Federated Learning via Semantic Anchors** (FedSA) to break the vicious cycle. Specifically, we first introduce simple yet effective semantic anchors, serving as prototypes, to decouple the generation of prototypes from local representation learning and guide local models in learning consistent representations. Instead of collecting prototypes from biased models, we project pre-defined class anchors for all categories into the semantic space through a lightweight embedding layer, thereby obtaining semantic anchors that are independent of representation learning and well-separated. By incorporating semantic anchors, we further propose: 1) Anchor-based Regularization with Margin-enhanced Contrastive Learning (RMCL) to correct the biased feature extractors, enabling them to learn consistent prototypes that exhibit both intra-class compactness and inter-class separability; 2) Anchor-based Classifier Calibration (CC) to correct the biased classifiers, assisting them in learning consistent decision boundaries across different classes. We then update the semantic anchors with these enhanced prototypes via an Exponential Moving Average (EMA) update, which iteratively encourages clients to collaboratively learn a unified data representation in the semantic space. Our contributions are summarized as follows:

- We demonstrate that in prototype-based FL methods, heterogeneity creates a vicious cycle of representation inconsistency, classifier divergence, and skewed prototype alignment across client models.

- To break the vicious cycle, we propose a novel framework named FedSA that corrects biased feature extractors and classifiers, guiding local models in learning a unified data representation.
- We evaluate the proposed method under both statistical and model heterogeneity settings. Extensive experiments and ablation studies demonstrate the superiority of FedSA over prototype-based FL methods.

Related Work

Model-based FL Approaches

Model-based FL approaches are inspired by the training scheme that decouples the entire model into a general body and a personalized head (Kang et al. 2019; Devlin et al. 2019). For example, methods like FedRep (Collins et al. 2021), FedBABU (Oh, Kim, and Yun 2021), FedRoD (Chen and Chao 2022) and FedGC (Niu and Deng 2022) split the model into a feature extractor and a classifier. This approach allows clients to share a global feature extractor to mitigate data heterogeneity while maintaining client-specific classifiers for personalization. LG-FedAvg (Liang et al. 2020) and FedGen (Zhu, Hong, and Zhou 2021) treat the top layers as the general body for sharing while allowing the bottom layers to have different architectures. However, these methods commonly involve significant communication costs as model parameters are transmitted between clients. Recent studies initialize class anchors (Zhou, Zhang, and Tsang 2023) or use prototypes (Xu, Tong, and Huang 2023; Dai et al. 2023), serving as additional knowledge to align feature representations from different feature extractors. Nevertheless, these class anchors or prototypes are also updated or generated based on locally inconsistent representations. In this paper, we only share lightweight semantic anchors across clients in a model-agnostic manner.

Knowledge-based FL Approaches

Knowledge-based FL approaches transfer various types of global knowledge instead of model parameters, making them well-suited for model heterogeneity setting. The typical method in these approaches (Jeong et al. 2018; Li and Wang 2019; Lin et al. 2020) involves incorporating KD methods (Hinton, Vinyals, and Dean 2015) to transfer knowledge from a teacher model on the server to enhance the performance of student models on the client side. FML (Shen et al. 2020) and FedKD (Wu et al. 2022) advance KD methods in a data-free manner, where they share a small auxiliary model as global knowledge, rather than relying on a global dataset. However, the majority of these methods are not suitable for heterogeneous FL, as their effectiveness heavily relies on the quality of the public dataset or the auxiliary model. Another popular approach is to transfer abstract class representations, i.e., prototypes, between the server and clients. FedProto (Tan et al. 2022b) regularizes the local prototypes to align closer to the aggregated global prototypes. FedPCL (Tan et al. 2022c) employs prototype-wise contrastive learning to enhance the inter-class separability among prototypes. Additionally, FedTGP (Zhang et al. 2024) treats prototypes as trainable elements and enforces




Figure 1: A toy example with two clients shows the vicious cycle. Different colors represent classes, while triangles, circles and squares represent data representations, local prototypes and global prototypes, respectively. Different colored solid lines indicate decision boundaries. Black and yellow dotted arrows show the inter-class separation among prototypes, and red arrows depict the guidance of prototypes during local training. Fig. 1(a) shows that biased datasets lead feature extractors and classifiers to learn inconsistent representations and skewed decision boundaries. Fig. 1(b) shows that in FedProto, the naive averaging aggregation and local guidance reduce the separability of global prototypes and inevitably create a vicious cycle.

an adaptive margin between them to improve separability. However, these methods often overlook the biases arising from heterogeneity, which can skew local models toward learning inconsistent representations and potentially create a vicious cycle. In this paper, we introduce semantic anchors that are independent of local representation learning to guide local models in learning consistent representations.

Method

Problem Statement

In prototype-based FL settings, we consider a central server and m clients collaboratively training their models with heterogeneous architectures on their heterogeneous local datasets $\{\mathcal{D}_i\}_{i=1}^m$. Following FedProto (Tan et al. 2022b), we split each client i 's model w_i into a feature extractor f_i parameterized by θ_i and a classifier h_i parameterized by $\phi_i \in \mathbb{R}^{C \times D}$, where D is the dimension of the last representations and C is the number of classes. The goal is for clients to collaborate by sharing global prototypes \mathcal{P} with the central server. The overall collaborative training objective is:

$$\min_{\{\{\theta_i, \phi_i\}\}_{i=1}^m} \frac{1}{m} \sum_{i=1}^m \mathcal{L}_i(\mathcal{D}_i, \theta_i, \phi_i, \mathcal{P}) \quad (1)$$

During local training, each client i first computes its local prototype for each class c as follows:

$$P_i^c = \frac{1}{|\mathcal{D}_{i,c}|} \sum_{(x,y) \in \mathcal{D}_{i,c}} f_i(x; \theta_i) \quad (2)$$

where $\mathcal{D}_{i,c}$ denotes the subset of the local dataset \mathcal{D}_i consisting of all data samples belonging to class c . After receiving all local prototypes from clients, the server performs weighted-averaging aggregation for each class prototype to obtain the global prototypes:

$$\bar{P}^c = \frac{1}{|\mathcal{N}_c|} \sum_{i \in \mathcal{N}_c} \frac{|\mathcal{D}_{i,c}|}{N_c} P_i^c \quad (3)$$

where \mathcal{N}_c denotes the set of clients that have class c , and N_c represents the total number of data samples of class c among all clients. The server then sends the global prototypes $\mathcal{P} = \{\bar{P}^c\}_{c=1}^C$ to each client to guide local training:

$$\mathcal{L}_i = \mathcal{L}_S(h_i(f_i(x; \theta_i); \phi_i), y) + \lambda \mathcal{L}_R(P_i^c, \bar{P}^c) \quad (4)$$

where \mathcal{L}_S is a typical loss term (e.g., cross-entropy loss) in supervised learning. λ is a hyperparameter, and \mathcal{L}_R is a regularization term that minimizes the Euclidean distances between the local prototype P_i^c and the global prototype \bar{P}^c .

Motivation

Although prototype-based FL methods have achieved significant results by sharing lightweight prototypes, they have overlooked the biases arising from heterogeneity. Building on previous findings (Oh, Kim, and Yun 2021; Guo, Tang, and Lin 2023), which link the feature extractor to representation learning and the classifier to linear decision boundary learning, we seek to understand how heterogeneity affects the generation of prototypes through these components. Due to statistical heterogeneity, the dataset of each client is inherently biased. Consider different local models w_i and w_j , each trained on its respective local datasets \mathcal{D}_i and \mathcal{D}_j . As shown in Fig. 1, we use a toy example to show the existence of biases in local models.

Observation 1 (Representation inconsistency). *Given the same inputs x , the outputs $f_i(x)$ could deviate significantly from $f_j(x)$. Fig. 1(a) shows the data representations learned by different feature extractors. As local models are trained on biased datasets, we observe that the representations extracted by different local feature extractors are inconsistent in the semantic space.*

Observation 2 (Classifier divergence). *As illustrated in Fig. 1(a), we visualize the decision boundaries learned by different classifiers $h_i(\cdot)$ and $h_j(\cdot)$ for the same input. The observation confirms that inconsistent representations significantly cause local classifiers to diverge, resulting in skewed decision boundaries among clients.*

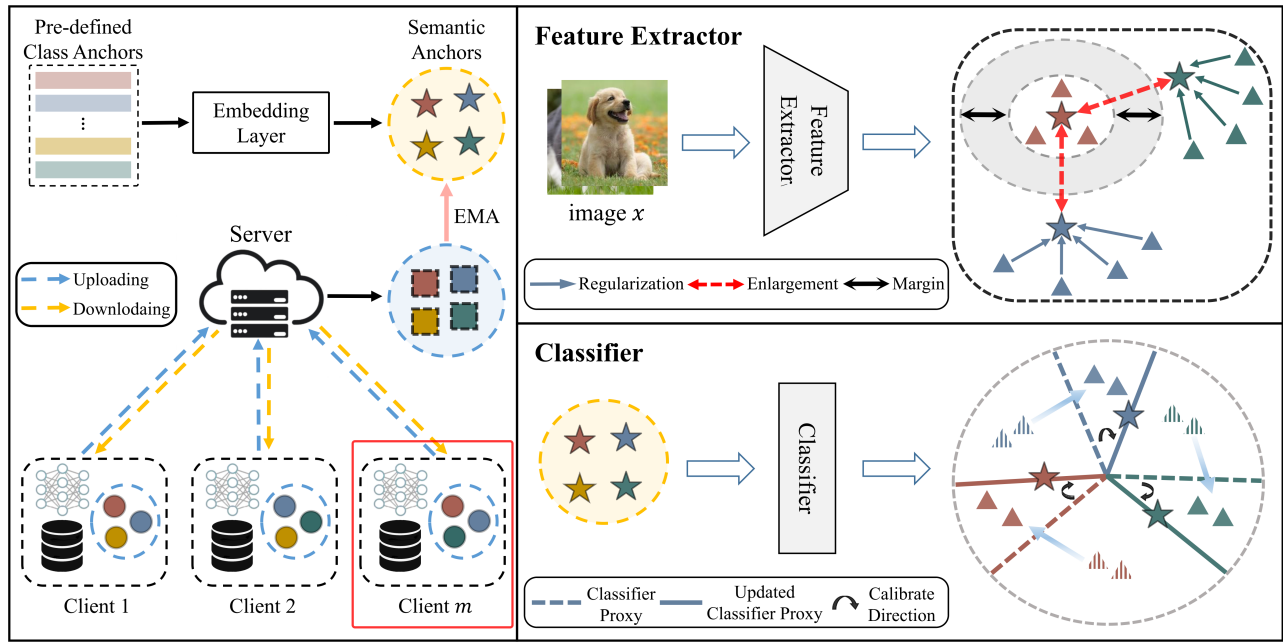


Figure 2: Framework of the proposed method FedSA. Pentagons represent semantic anchors. In the local training of client m , the semantic anchors correct the biased feature extractor and classifier via RMCL and CC, achieving intra-class compactness and inter-class separability of local prototypes, while ensuring consistent decision boundaries. On the server, we update the semantic anchors with global prototypes via an EMA update to facilitate collaboration in FL.

Observation 3 (Skewed prototype alignment). *As shown in Fig. 1(b), although local models can learn well-separated prototypes, representation inconsistency naturally causes these prototypes to align skewed in the semantic space. The variability in margins between different local prototypes, due to model heterogeneity, further exacerbates this skewed alignment. Following the naive averaging aggregation, the margins among global prototypes tend to diminish, thus reducing the overall separability and effectiveness of the global prototypes. Moreover, the global prototypes serve to guide the training of local models. As indicated by Eq.4, the regularization term iteratively adjusts the well-separated local prototypes closer to the less effective global prototypes, impairing the training of client models and forcing them to produce more inconsistent representations. Therefore, statistical and model heterogeneity potentially creates a vicious cycle of representation inconsistency, classifier divergence, and skewed prototype alignment.*

To break this vicious cycle, we introduce simple yet effective semantic anchors to decouple the generation of prototypes from local representation learning and guide local models in learning consistent representations. Our method is motivated by the fact that prototypes are not only bound to the data distribution but also guided by the objective function in the training paradigm of empirical risk minimization. By explicitly guiding class representations toward corresponding semantic anchors that are independent of representation learning and well-separated, we can achieve more consistent and discriminative local prototypes.

Semantic Anchors

With a total of C classes in the whole dataset, we randomly initialize pre-defined class anchors $A = \{A^c\}_{c=1}^C$, where each $A^c \in \mathbb{R}^D$. Then, we project them into the semantic space through a trainable embedding layer $h_{\psi}(A)$ to improve their separability, obtaining semantic anchors $\bar{A} = \{\bar{A}^c\}_{c=1}^C \in \mathbb{R}^{C \times D}$. In each communication round t , the server sends these semantic anchors \bar{A}^t to each client to guide their local training. The goal is for clients to collaboratively learn a unified data representation that exhibits intra-class compactness and inter-class separability via the FL paradigm. To achieve this, we further propose Anchor-based Regularization with Margin-enhanced Contrastive Learning (RMCL) and Anchor-based Classifier Calibration (CC). An overview of the proposed framework is shown in Fig. 2.

Anchor-based regularization. To achieve intra-class compactness, we propose anchor-based regularization by directly minimizing the distance between class representations and corresponding semantic anchors. In Eq.4, we redefine the regularization loss as follows:

$$\mathcal{L}_R = \sum_{c \in \mathcal{D}_i} d(P_i^c, \bar{A}^c) \quad (5)$$

where d measures the Euclidean distance between the local prototype P_i^c and the semantic anchor \bar{A}^c . Although guiding class representations toward corresponding semantic anchors can achieve an intra-class compact embedding space, due to statistical and model heterogeneity, clients often extract highly diverse data representations that vary in separability and prototype margins across different classes.

Anchor-based margin-enhanced contrastive learning. Inspired by previous works (Deng et al. 2019; Chen et al. 2021; Zhang et al. 2024), we propose anchor-based margin-enhanced contrastive learning that enforces a client-specific margin d_i^* between classes to facilitate the learning of a large-margin data representation and ensure consistent prototype margins across clients. Specifically, we calculate the average margin among semantic anchors, referred to as the global margin d_{global}^t , and the average margin among local prototypes for each client, referred to as the local margin d_i^t . For client i , the local margin is computed:

$$d_i^t = \frac{1}{(N-1)^2} \sum_a \sum_{b \neq a} d(P_i^a, P_i^b) \quad (6)$$

where N is the number of classes maintained by client i . The client-specific margin d_i^* for each client is determined by taking the larger value between the local and global margins, i.e., $d_i^* = \max\{d_{global}^t, d_i^t\}$. Our goal is to decrease the distance between P_i^c and \bar{A}^c , and to increase the distance between P_i^c and $\bar{A}^{c'}$ with the client-specific margin d_i^* . Accordingly, we define the anchor-based margin-enhanced contrastive loss as follows:

$$\mathcal{L}_{MCL} = -\log \frac{e^{-(d(P_i^c, \bar{A}^c) + d_i^*)}}{e^{-(d(P_i^c, \bar{A}^c) + d_i^*)} + \sum_{c'} e^{-d(P_i^c, \bar{A}^{c'})}} \quad (7)$$

where $c' \in [C]$, $c' \neq c$. The client-specific margin adaptively guides each client toward better data representations and consistent prototype margins. Specifically, when the local model learns well-separated and clustered class representations, the client adopts its local margin d_i^t as the client-specific margin. Otherwise, the client uses the global margin d_{global}^t to encourage the local model to learn class representations with a larger margin separation. Through repeated collaborative training in FL, clients reach a consensus on the margins among local prototypes. By aggregating these uniformly adjusted local prototypes at the server, we can obtain global prototypes with enhanced separability.

Anchor-based classifier calibration. Beyond correcting biased feature extractors to learn consistent representations, we also propose anchor-based classifier calibration to correct biased classifiers toward achieving consistent decision boundaries. Specifically, each client i uses the semantic anchors as inputs for its classifier, calibrating the classifiers according to the following objective:

$$\mathcal{L}_{CC} = -\frac{1}{C} \sum_{c \in C} \log \frac{e^{(\phi_{i,c}^\top \bar{A}^c)}}{\sum_{j=1}^C e^{(\phi_{i,j}^\top \bar{A}^c)}} \quad (8)$$

where \mathcal{L}_{CC} is the classifier calibration loss, which reduces the distance between the c -th class proxy and the corresponding semantic anchor. By applying this loss across different clients, we not only reduce the discrepancies in decision boundaries between different model architectures but also promote the learning of consistent representations.

Overall. By integrating all components, the semantic anchors effectively guide each client to learn a unified data representation by minimizing the overall loss:

$$\mathcal{L}_i = \mathcal{L}_S + \lambda_1 \mathcal{L}_R + \lambda_2 \mathcal{L}_{MCL} + \lambda_3 \mathcal{L}_{CC} \quad (9)$$

FedSA Framework

Following the prototype-based communication protocol, each client collects local prototypes and sends them to the server. The server then performs weighted-averaging aggregation to obtain global prototypes. To facilitate collaboration in FL, we update the semantic anchors with these global prototypes using the EMA update as follows:

$$\bar{A}^{t+1} = \alpha \bar{A}^t + (1 - \alpha) \bar{P}^t \quad (10)$$

where α is a decay factor that controls the update rate. Similar to prototype-based FL methods, FedSA transmits only compact 1D-class semantic anchors, which naturally bring benefits for both privacy preservation and communication efficiency. Moreover, by integrating the anchor-based components and the EMA update, FedSA builds a positive feedback loop that iteratively encourages clients to collaboratively learn a unified data representation that exhibits intra-class compactness and inter-class separability.

Experiments

Setup

Datasets. We consider image classification tasks and evaluate our method on three popular datasets, including Cifar10, Cifar100 (Deng et al. 2009), and Tiny-Imagenet (Chrabaszcz, Loshchilov, and Hutter 2017).

Baseline methods. To evaluate our proposed method FedSA, we consider both statistical and model heterogeneous settings and compare it with popular methods including LG-FedAvg (Liang et al. 2020), FedGen (Zhu, Hong, and Zhou 2021), FML (Shen et al. 2020), FedKD (Wu et al. 2022), FedDistill (Jeong et al. 2018), FedProto (Tan et al. 2022b), and FedTGP (Zhang et al. 2024).

Statistical heterogeneity. Like previous studies (Li, He, and Song 2021; Li et al. 2022), we simulate statistical heterogeneity among clients by using Dirichlet distribution. Specifically, we first sample $q_{c,i} \sim \text{Dir}(\beta)$ for each class c and client i . We then allocate the proportion $q_{c,i}$ of data points from class c in the dataset to client i , where $\text{Dir}(\beta)$ is the Dirichlet distribution with a concentration parameter β set to 0.1 by default. All clients use a 4-layer CNN (McMahan et al. 2017) for a homogeneous model setting.

Model heterogeneity. Following FedTGP (Zhang et al. 2024), we simulate model heterogeneity among clients by assigning different model architectures. Specifically, we denote this setting as "HtFE $_X$ ", where FE $_X$ represents the X distinct feature extractors used. Each client i is assigned the $(i \bmod X)$ th model architecture. For instance, the "HtFE $_8$ " setting includes eight architectures: 4-layer CNN, GoogleNet (Szegedy et al. 2015), MobileNet_v2 (Sandler et al. 2018), ResNet18, ResNet34, ResNet50, ResNet101, and ResNet152 (He et al. 2016), used in our main experiments. To generate feature representations with an identical feature dimension K , an average pooling layer is added after each feature extractor, with K set to 512 by default.

Implementation details. We consider two widely used FL settings, the cross-silo setting and the cross-device setting (Kairouz et al. 2021). In the cross-silo setting, we set the total number of clients to 20 with a client participation ratio ρ

Settings	Cross-silo Setting			Cross-device Setting		
Datasets	Cifar10	Cifar100	Tiny-ImageNet	Cifar10	Cifar100	Tiny-ImageNet
LG-FedAvg	89.32±0.04	48.97±0.07	35.08±0.06	87.27±0.05	43.41±0.10	28.54±0.05
FedGen	89.13±0.05	48.25±0.06	32.68±0.08	86.99±0.07	42.70±0.08	27.76±0.08
FML	89.42±0.03	48.73±0.22	34.75±0.13	87.42±0.07	42.13±0.14	27.63±0.04
FedKD	90.21±0.06	53.14±0.10	37.14±0.16	87.94±0.13	46.03±0.15	30.58±0.19
FedDistill	89.61±0.03	49.27±0.18	35.45±0.07	87.57±0.05	42.15±0.16	28.57±0.12
FedProto	90.76±0.05	53.26±0.11	36.23±0.12	87.49±0.11	44.46±0.32	28.73±0.21
FedTGP	90.23±0.08	53.71±0.13	36.43±0.08	87.98±0.16	47.41±0.11	29.96±0.17
FedSA	90.88±0.04	54.39±0.12	37.30±0.13	88.78±0.08	48.42±0.17	30.82±0.14

Table 1: The test accuracy (%) on three datasets in both cross-silo and cross-device settings under statistical heterogeneity.

of 1. In the cross-device setting, the total number of clients is set to 100 with a participation ratio ρ of 0.1. Our main experiments are conducted in the cross-device setting, where the participating clients are generally resource-constrained edge devices, more closely reflecting real-world FL scenarios.

Unless explicitly specified, we follow previous studies (Zhang et al. 2023b) by running one local epoch of training on each client per round, using a batch size of 10 and a learning rate $\lambda = 0.01$ for 1000 communication rounds. Each client’s local dataset is split into a training set (75%) and a test set (25%), with the test set used to evaluate performance. We run three trials for all experiments and report the mean and standard deviation. For our FedSA, we set $\lambda_1 = 0.1$ (the same as in prototype-based FL methods), $\lambda_3 = 1$, and $\alpha = 0.9999$. We use smaller $\lambda_2 = 0.01$ for statistical heterogeneity, and $\lambda_2 = 1$ for model heterogeneity. Please refer to the Appendix for more results and details.

Performance

Impact of statistical heterogeneity. Table 1 shows the accuracy of all methods in both cross-silo and cross-device settings under statistical heterogeneity. Comparing different FL methods, we observe that FedSA consistently outperforms all other methods across all tasks. Specifically, FedSA outperforms the baselines on three datasets by up to 6.29%. In more complex tasks, the increased number of classes leads to more skewed prototypes during aggregation, exacerbating the vicious cycle. FedSA improves prototype-based FL methods by up to 3.96% through the incorporation of lightweight but innovative semantic anchors serving as prototypes. This improvement demonstrates the effectiveness of semantic anchors in addressing statistical heterogeneity, enabling local clients to learn consistent representations.

Impact of model heterogeneity. Table 2 shows the accuracy on Cifar100 in the cross-device setting under both statistical and model heterogeneity. As the number of feature extractors X increases, we observe a significant degradation in the performance of prototype-based FL methods compared to other methods. This confirms that prototype-based methods heavily rely on the quality of data representation, and the naive aggregation of prototypes with varying margins, such as FedProto, leads to catastrophic performance

Settings	HtFE ₂	HtFE ₃	HtFE ₅	HtFE ₈
LG-FedAvg	41.22±0.13	40.59±0.42	39.66±0.23	35.24±0.13
FedGen	40.05±0.18	38.47±0.29	39.48±0.28	34.63±0.44
FML	40.62±0.06	40.19±0.09	36.39±0.13	34.47±0.09
FedKD	42.19±0.09	40.53±0.07	37.47±0.21	35.08±0.12
FedDistill	41.58±0.15	40.92±0.36	39.98±0.13	35.26±0.48
FedProto	35.63±0.22	29.89±0.48	27.75±0.64	16.31±0.32
FedTGP	42.14±0.29	37.97±0.57	38.66±0.42	32.04±0.74
FedSA	43.54±0.21	41.75±0.16	40.96±0.12	35.68±0.17

Table 2: The test accuracy (%) on Cifar100 in the cross-device setting under both statistical and model heterogeneity.

	$E = 1$	$E = 5$	$E = 10$
LG-FedAvg	35.24±0.13	34.98±0.15	34.78±0.09
FedGen	34.63±0.44	33.89±0.39	32.95±0.06
FML	34.47±0.09	35.22±0.05	35.02±0.31
FedKD	35.08±0.12	37.45±0.57	37.16±0.11
FedDistill	35.24±0.48	36.15±0.21	36.17±0.17
FedProto	16.31±0.32	23.96±0.55	22.88±0.45
FedTGP	32.04±0.74	34.86±0.82	34.51±0.25
FedSA	35.68±0.17	37.57±0.27	37.43±0.13

Table 3: The test accuracy (%) with different number of local epochs on Cifar100 in the HtFE₈ setting.

degradation. Although FedTGP enhances the separability of prototypes through margin-enhanced contrastive learning, it fails to achieve better performance due to its neglect of biases in local clients. Using our RMCL and CC, FedSA can improve FedProto by up to 19.37% and FedTGP by up to 3.78%, demonstrating that FedSA is more robust and less affected by model heterogeneity.

Number of Local Epochs

Clients can perform more local epochs to alleviate communication costs, albeit at the expense of exacerbating the biases across clients (Wang et al. 2020). The results of different local epochs are shown in Table 3. We set the num-

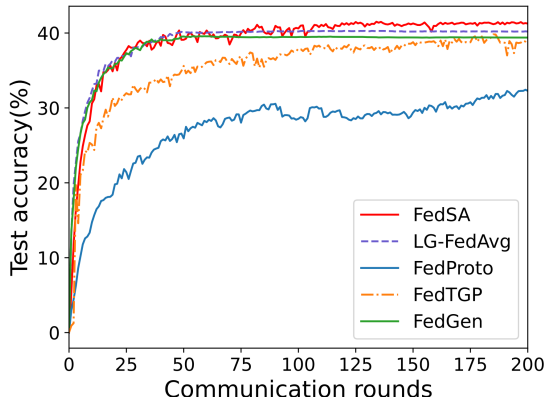


Figure 3: The test accuracy (%) on Cifar100 in the cross-silo setting using the HtFE₈ setting under both statistical and model heterogeneity.

ber of communication rounds to 500 for $E = 5$ and 200 for $E = 10$. We note that as the number of local epochs increases, the performance of model-based methods declines, whereas knowledge-based methods show improvement. This improvement is attributed to the increased number of local epochs enhancing the model’s capability to extract more fine-grained knowledge. There remains a significant performance gap between prototype-based FL methods and other knowledge-based methods. In contrast, our method FedSA maintains better performance, demonstrating its effectiveness with different numbers of local epochs.

Communication Efficiency

Fig. 3 shows the accuracy in the cross-silo setting using the HtFE₈ settings under both statistical and model heterogeneity. We observe that the model-based method LG-FedAvg and the KD-based method FedGen show faster convergence compared to prototype-based methods. The convergence curve of prototype-based methods FedProto and FedTGP initially exhibits significant volatility, due to the inadequate feature extraction capabilities of heterogeneous models in the early stages of training. Moreover, prototypes are more susceptible to the impacts of statistical and model heterogeneity, resulting in lower accuracy for prototype-based methods. By guiding class representations toward well-separated semantic anchors, FedSA can effectively increase accuracy without slowing down convergence, while also ensuring privacy preservation.

Ablation and Hyperparameter Studies

In Table 4, we evaluate the efficacy of each component in our method on Cifar100 in the HtFE₈ setting. In FedSA, EMA and \mathcal{L}_R serve as default components. Without embedding projection, data representations are directly regularized by randomly initialized, pre-defined class anchors. The accuracy surpasses that of FedProto by 13.23%, validating our motivation to decouple the generation of prototypes from local representation learning. To enhance the separability of semantic anchors, we project pre-defined class anchors through an embedding layer and guide feature extractors to

EMA	\mathcal{L}_R	ER	\mathcal{L}_{MCL}	\mathcal{L}_{CC}	Acc (%)
✓	✓				29.54
✓	✓	✓			30.80 (+1.26)
✓	✓	✓		✓	30.61 (+1.07)
✓	✓	✓	✓		32.82 (+3.28)
✓	✓	✓	✓	✓	35.68 (+6.14)

Table 4: Ablation studies on the key components of our proposed FedSA on Cifar100 in the HtFE₈ setting. \mathcal{L}_R : Anchor-based Regularization. ER: Embedding Projection. \mathcal{L}_{MCL} : Anchor-based Margin-enhanced Contrastive Learning. \mathcal{L}_{CC} : Anchor-based Classifier Calibration.

	\mathcal{L}_{MCL}				\mathcal{L}_{CC}			
	0.1	1	5	10	0.1	1	5	10
Acc (%)	33.42	35.68	35.54	35.12	33.26	35.68	34.66	34.29

Table 5: The test accuracy (%) on Cifar100 in the HtFE₈ setting with different λ_2 or λ_3 . We set $\lambda_2 = 1$ and $\lambda_3 = 1$ by default.

learn consistent and large margins between prototypes using \mathcal{L}_{MCL} . This approach can improve performance by up to 3.28% and outperforms FedTGP. Furthermore, we correct biased classifiers by utilizing \mathcal{L}_{CC} to ensure consistent decision boundaries. Combining all components, FedSA guides local clients to learn a unified data representation, achieving a significant improvement of 6.14% over the baseline. We also notice that FedSA without \mathcal{L}_{MCL} shows a marginal improvement in performance, highlighting the importance of learning consistent representations.

Table 5 shows the accuracy of our method with varying hyperparameters λ_2 and λ_3 . FedSA performs better with λ_2 ranging from 1 to 5. However, there is a slight drop in accuracy when $\lambda_2 = 10$. This drop is attributed to a larger λ_2 continuously expanding the margins among prototypes, where overly large margins result in unstable prototype guidance for clients. For λ_3 , the highest accuracy is achieved at 1, but increasing this hyperparameter leads to excessive classifier optimization, which compromises the model’s generalizability and significantly degrades performance. Notably, even with $\lambda_2 = 0.1$ or $\lambda_3 = 0.1$, FedSA still achieves an accuracy of at least 33.26%, outperforming other prototype-based FL methods.

Conclusion

In this paper, we propose a novel framework named FedSA, which introduces simple yet effective semantic anchors to decouple the generation of prototypes. By incorporating semantic anchors, we further propose RMCL and CC to correct local models, enabling them to learn consistent representations and decision boundaries. We then update the semantic anchors with global prototypes via an EMA update to encourage clients to learn a unified data representation. Extensive experiments and ablation studies demonstrate the superiority of FedSA over prototype-based FL methods.

Acknowledgments

This work was supported by Chongqing Natural Science Foundation Innovation and Development Joint Fund (CSTB2023NSCQ-LZX0109) and Fundamental Research Funds for the Central Universities (No.2023CDJQCZX-001). The authors thank Jinghua Zhou and Ping Jiang for their invaluable support.

References

- Chen, H.-Y.; and Chao, W.-L. 2022. On Bridging Generic and Personalized Federated Learning for Image Classification. In *International Conference on Learning Representations*.
- Chen, S.; Niu, G.; Gong, C.; Li, J.; Yang, J.; and Sugiyama, M. 2021. Large-margin contrastive learning with distance polarization regularizer. In *International Conference on Machine Learning*, 1673–1683. PMLR.
- Chrabaszcz, P.; Loshchilov, I.; and Hutter, F. 2017. A down-sampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*.
- Collins, L.; Hassani, H.; Mokhtari, A.; and Shakkottai, S. 2021. Exploiting shared representations for personalized federated learning. In *International conference on machine learning*, 2089–2099. PMLR.
- Dai, Y.; Chen, Z.; Li, J.; Heinecke, S.; Sun, L.; and Xu, R. 2023. Tackling data heterogeneity in federated learning with class prototypes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 7314–7322.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4690–4699.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186.
- Ge, Y.; Nie, Q.; Huang, Y.; Liu, Y.; Wang, C.; Zheng, F.; Li, W.; and Duan, L. 2024. Beyond prototypes: Semantic anchor regularization for better representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1887–1895.
- Guo, Y.; Tang, X.; and Lin, T. 2023. Fedbr: Improving federated learning on heterogeneous data via local learning bias reduction. In *International Conference on Machine Learning*, 12034–12054. PMLR.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Jeong, E.; Oh, S.; Kim, H.; Park, J.; Bennis, M.; and Kim, S.-L. 2018. Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data. *arXiv preprint arXiv:1811.11479*.
- Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A. N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. 2021. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2): 1–210.
- Kang, B.; Xie, S.; Rohrbach, M.; Yan, Z.; Gordo, A.; Feng, J.; and Kalantidis, Y. 2019. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*.
- Li, D.; and Wang, J. 2019. Fedmd: Heterogenous federated learning via model distillation. *arXiv preprint arXiv:1910.03581*.
- Li, Q.; Diao, Y.; Chen, Q.; and He, B. 2022. Federated learning on non-iid data silos: An experimental study. In *2022 IEEE 38th international conference on data engineering (ICDE)*, 965–978. IEEE.
- Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10713–10722.
- Li, T.; Sahu, A. K.; Talwalkar, A.; and Smith, V. 2020. Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine*, 37(3): 50–60.
- Liang, P. P.; Liu, T.; Ziyin, L.; Allen, N. B.; Auerbach, R. P.; Brent, D.; Salakhutdinov, R.; and Morency, L.-P. 2020. Think locally, act globally: Federated learning with local and global representations. *arXiv preprint arXiv:2001.01523*.
- Lin, T.; Kong, L.; Stich, S. U.; and Jaggi, M. 2020. Ensemble distillation for robust model fusion in federated learning. *Advances in neural information processing systems*, 33: 2351–2363.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Niu, Y.; and Deng, W. 2022. Federated learning for face recognition with gradient correction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 1999–2007.
- Oh, J.; Kim, S.; and Yun, S.-Y. 2021. Fedbabu: Towards enhanced representation for federated image classification. *arXiv preprint arXiv:2106.06042*.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4510–4520.
- Shen, T.; Zhang, J.; Jia, X.; Zhang, F.; Huang, G.; Zhou, P.; Kuang, K.; Wu, F.; and Wu, C. 2020. Federated mutual learning. *arXiv preprint arXiv:2006.16765*.

- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1–9.
- Tan, A. Z.; Yu, H.; Cui, L.; and Yang, Q. 2022a. Towards personalized federated learning. *IEEE transactions on neural networks and learning systems*, 34(12): 9587–9603.
- Tan, Y.; Long, G.; Liu, L.; Zhou, T.; Lu, Q.; Jiang, J.; and Zhang, C. 2022b. Fedproto: Federated prototype learning across heterogeneous clients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 8432–8440.
- Tan, Y.; Long, G.; Ma, J.; Liu, L.; Zhou, T.; and Jiang, J. 2022c. Federated learning from pre-trained models: A contrastive learning approach. *Advances in neural information processing systems*, 35: 19332–19344.
- Wang, J.; Liu, Q.; Liang, H.; Joshi, G.; and Poor, H. V. 2020. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33: 7611–7623.
- Wang, L.; Wang, M.; Zhang, D.; and Fu, H. 2023. Model barrier: A compact un-transferable isolation domain for model intellectual property protection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20475–20484.
- Wu, C.; Wu, F.; Lyu, L.; Huang, Y.; and Xie, X. 2022. Communication-efficient federated learning via knowledge distillation. *Nature communications*, 13(1): 2032.
- Xu, J.; Tong, X.; and Huang, S.-L. 2023. Personalized federated learning with feature alignment and classifier collaboration. *arXiv preprint arXiv:2306.11867*.
- Ye, M.; Fang, X.; Du, B.; Yuen, P. C.; and Tao, D. 2023. Heterogeneous federated learning: State-of-the-art and research challenges. *ACM Computing Surveys*, 56(3): 1–44.
- Zhang, J.; Guo, S.; Guo, J.; Zeng, D.; Zhou, J.; and Zomaya, A. Y. 2023a. Towards data-independent knowledge transfer in model-heterogeneous federated learning. *IEEE Transactions on Computers*, 72(10): 2888–2901.
- Zhang, J.; Guo, S.; Ma, X.; Wang, H.; Xu, W.; and Wu, F. 2021. Parameterized knowledge transfer for personalized federated learning. *Advances in Neural Information Processing Systems*, 34: 10092–10104.
- Zhang, J.; Liu, Y.; Hua, Y.; and Cao, J. 2024. Fedtgp: Trainable global prototypes with adaptive-margin-enhanced contrastive learning for data and model heterogeneity in federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 16768–16776.
- Zhang, J.; Liu, Y.; Hua, Y.; Wang, H.; Song, T.; Xue, Z.; Ma, R.; and Cao, J. 2023b. PFLlib: Personalized Federated Learning Algorithm Library. *arXiv preprint arXiv:2312.04992*.
- Zhang, L.; Shen, L.; Ding, L.; Tao, D.; and Duan, L.-Y. 2022. 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*, 10174–10183.
- Zhou, T.; Zhang, J.; and Tsang, D. H. 2023. FedFA: Federated learning with feature anchors to align features and classifiers for heterogeneous data. *IEEE Transactions on Mobile Computing*.
- Zhu, Z.; Hong, J.; and Zhou, J. 2021. Data-free knowledge distillation for heterogeneous federated learning. In *International conference on machine learning*, 12878–12889. PMLR.