

# FedUV: Uniformity and Variance for Heterogeneous Federated Learning

Ha Min Son  
 University of California, Davis

Moon-Hyun Kim  
 Hippo T&C

Tai-Myoung Chung  
 Hippo T&C

Chao Huang  
 University of California, Davis

Xin Liu  
 University of California, Davis

## Abstract

*Federated learning is a promising framework to train neural networks with widely distributed data. However, performance degrades heavily with heterogeneously distributed data. Recent work has shown this is due to the final layer of the network being most prone to local bias, some finding success freezing the final layer as an orthogonal classifier. We investigate the training dynamics of the classifier by applying SVD to the weights motivated by the observation that freezing weights results in constant singular values. We find that there are differences when training in IID and non-IID settings. Based on this finding, we introduce two regularization terms for local training to continuously emulate IID settings: (1) variance in the dimension-wise probability distribution of the classifier and (2) hyperspherical uniformity of representations of the encoder. These regularizations promote local models to act as if it were in an IID setting regardless of the local data distribution, thus offsetting proneness to bias while being flexible to the data. On extensive experiments in both label-shift and feature-shift settings, we verify that our method achieves highest performance by a large margin especially in highly non-IID cases in addition to being scalable to larger models and datasets.*

## 1. Introduction

Federated Learning (FL) [31] is a distributed learning framework that allows the training of deep neural networks with data in decentralized locations. FL is especially appealing because it achieves similar performance to centralized training in specific settings while also negating the cost of collecting data into a centralized location and allowing effective parallelization across devices [42]. However, because devices participating in FL use locally collected data for training, it is realistic to expect that the collective data

across all devices are not independent and identically distributed (non-IID). This severely degrades overall performance. Local devices optimize their copy of the model towards local optima. In non-IID settings, however, it is likely that local optima across devices are in disagreement with the true optimum of the IID setting. Thus, the direction of the gradient causes dissonance and degraded performance. This phenomena is referred to as *client drift* in the FL literature.

To combat client drift, a number of methods [1, 19, 20] introduced the idea of using the global model as a basis for regularization. More recently, some work showed it may be effective to focus on certain layers which are more prone to bias. In particular, [29] found that the final layer, referred to as the *classifier*, is the most biased in non-IID settings. They introduce a method to calibrate the classifier to offset this bias. Other work [32] used augmentation techniques [49] to regularize the classifier, while another work [34] randomly initialized and froze the weights of the classifier. Since random vectors in high dimension space are likely to be orthogonal, the output of the penultimate layer, referred to as the *representations* or *activations* of the encoder, are trained to fit these orthogonal classifiers, thus offsetting bias.

However, we note previous work have approached the non-IID problem by regularizing local models to be *less biased*, rather than directly regularizing local models to *emulate the IID setting*. In addition, many work have not been tested in a feature-shift setting, often only using the Dirichlet distribution to simulate a label-shift setting. Furthermore, while providing important insight for FL, these work generally do not take into account efficiency and scalability. State-of-the-art methods such as FedProx [20], SCAF-FOLD [15], and FedDyn [1] require distances calculations of weights between all layers of multiple models at every batch. MOON [19] and FedAlign [32] require extra forward passes at every batch. Computation and memory constraints certainly become issues as the size of models and datasets




Figure 1. Singular values of the weights of the classifier (final layer) trained on CIFAR-100. Training setup is specified in Section 4.1

grow.

**Motivation.** In most classification problems, it is standard to minimize a Cross-Entropy loss function [9] to fit a neural network’s probability distribution to the sampled probability distribution. When the sampled class distribution is imbalanced, such as in a non-IID environment, the classifier becomes biased towards some subspace regardless of the features extracted by the encoder. Motivated by the success of freezing the classifier as orthogonal vectors [34], we use the singular value decomposition to gain insight into the training dynamics of the classifier. We train a Resnet-18 [10] on CIFAR-100 [16] and take the singular value decomposition of the weight matrix ( $\mathbf{W} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ ,  $\mathbf{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_n)$ ) as shown in Fig. 1. The singular values of the classifier decrease in a non-IID setting, showing there is a difference in the training dynamic between IID and non-IID settings.

In light of this, we propose an approach to the non-IID problem in which we directly promote the emulation of an IID setting. There is a nuanced difference from previous work which regularize to decrease bias. Previous work use the global model as a source of regularization under the intuition that the global model is less biased than any single local model. However, this may not be optimal in every setting as the aggregate of biased local model may still be biased. We instead focus on regularizing local models to emulate the IID setting without the use of the global model. Namely, we present FedUV, a method that promotes IID emulation by inducing representation hyperspherical Uniformity and classifier Variance during local training. FedUV penalizes classifier probability distributions that are biased towards a subset of classes, thus promoting probability distributions in non-IID settings to match IID settings. This promotes local classifiers to be unbiased in their predictions rather than becoming biased towards local labels. Furthermore, we penalize representations that are not thoroughly spread across the hypersphere. There are two rea-

sons for this. First, this encourages the encoder to not become biased towards local features. Second, expanding the feature space allows the classifier to expand its variance in more directions. We show both regularizations are essential for improved performance. FedUV achieves state-of-the-art performance on various standard label-shift benchmark in addition to various feature-shift benchmarks while being extremely simple, efficient, and scalable.

Our contributions are summarized as follows. First, we verify that singular values of the classifier decreases as a consequence of non-IID environments in FL (Fig. 1). Second, we present two regularization terms, representation hyperspherical uniformity and classifier variance, to prevent this degradation. These regularization terms are simple yet effective. Third, we show that FedUV achieves state-of-the-art performance not only on label shift non-IID settings but also on feature shift non-IID settings. Unlike previous work, FedUV does not use the global model as a source of regularization, which requires extra forward-pass, requires extra memory, or requires matrix comparison between weights of each layer, thus being more efficient.

## 2. Related Work

### 2.1. Federated Learning (FL)

The basic FL algorithm is FedAvg [31]. This algorithm progresses with the repetition of four steps. First, (step 1) a server broadcasts its global model. Second, (step 2) clients receive and train the model using their own local data. Third, (step 3) clients upload their trained local model. Fourth, (step 4) the server aggregates local models to create the next generational global model.

There are a wide range of recent work in the field of FL such as client selection [33, 39], data sharing [50], privacy [3, 14], communication efficiency [7, 46, 48], medical applications [47], and knowledge distillation [38, 52]. As we cannot hope to cover all work on FL, we mainly focus on work that address the performance degradation in non-IID settings. Work in this domain can generally be divided into two groups: 1) papers that focus on the aggregation process (step 4), and 2) efforts that focus on the local training process (step 2).

### 2.2. Aggregation Regularization

Work that focus on aggregation generally do not modify the local training scheme of FedAvg. Clients train their local model on local data using a SGD based optimizer. To address the non-IID problem, they change how the weights of the local model are aggregated. FedAvgM [11] applied a momentum term when aggregating and FedNova [44] normalized local models before averaging. Other work in Personalized FL [30, 37] used hypernetworks to generate the parameter of other client layers. PFNM [51] and

FedMA [43] introduced a Bayesian method to align potentially mixed neurons due to the permutation invariance of neural networks. PAN [22] improved efficiency by encoding position into neurons. FedDF [23] utilized generated and unlabelled data. CCVR [29] noted the largest bias exists in the classifier (the last layer) than any other layer. They used a Gaussian mixture model to create artificial data and calibrated the classifier.

### 2.3. Local Training Regularization

Work that focus on local training generally do not modify the aggregation scheme of FedAvg. To address the non-IID problem, they instead penalize local clients based on different criteria. As our work also belongs to this group, we discuss them in more detail.

FedProx [20] added a  $L_2$  regularizer between local model weights and global model weights that prevent local model weights from becoming dissimilar to the global model. FedDyn [1] added an inner product regularization term between the current round local model and previous round local model noting that regularizing solely based on the global model may cause convergence problems for local models. MOON [19] added a contrastive loss regularization term noting there should be a balance between the representational similarity between global model, the current round local model, and previous round local model.

These methods use the global model as a source of regularization. The intuition is that the global model is less biased than any single local model. However, there are two large issues concerning this approach. First, the global model can also be biased if local models are biased, since it is an averaged model. This bias may slow convergence and not improve performance. Second, efficiency and scalability becomes an issue. FedProx and FedDyn require the calculation of the  $L_2$  distance between global model and local model, and MOON requires three forward passes at every batch.

More recently, FedAlign [32] noted that it is possible to use only the local model for effective regularization. Local models employ GradAug [49] to regularize itself rather than relying on the global model. However, FedAlign still suffers in efficiency due to its reliance on an additional forward pass and estimation of second order information. FedBABU [34], on the other hand, randomly initializes and freezes the classifier to prevent bias. While this work can be categorized as Personalized Federated Learning, we include FedBABU in our discussion due to its focus on the classifier. We further explore these ideas in Section 4.2.

## 3. Proposed Method — FedUV

The motivation of FedUV comes from two main observations. First, the similarity between local models decreases further in later layers of the model [29], with the classifier

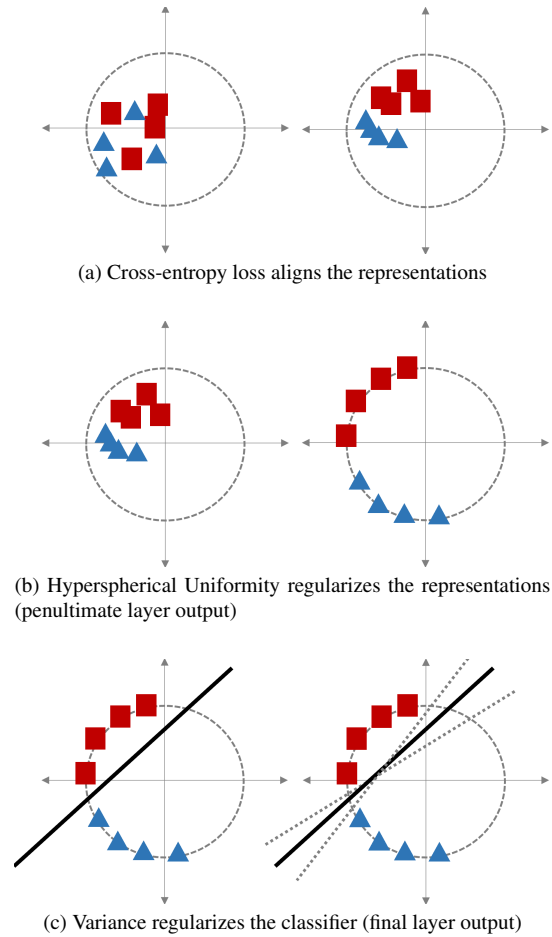


Figure 2. Uniformity regularization is applied to the penultimate layer and variance regularization is applied to the output layer.

(the final layer) being most biased towards local data. Second, as shown in Fig. 1, the singular values of the classifier of the global model degenerate in non-IID settings. Our goal then is to prevent this degeneration, promoting local models to emulate the IID setting regardless of local data distribution.

We achieve this by penalizing the classifier if it does not act as it would in an IID setting. We use a hinge loss to regularize the classifier if the class-wise variance of the probability distribution does not exceed a constant threshold  $c$ , which is set to be the class-wise variance of the ground truth labels in a batch with balanced classes. However, since the classifier acts as a linear hyperplane finding a separation of the representations in high dimensional space, increase in classifier variance may be limited by the representational feature imbalance in feature-shift non-IID settings. We introduce an additional regularization term to spread the representations of the encoder throughout the hypersphere. Fig. 2 is a visualization of the FedUV regularization scheme in  $R^2$  (a PyTorch-like pseudocode is included in

Appendix 1). We first promote hyperspherical uniformity in the representations of the encoder (row vectors of the penultimate layer output) then promote the class-wise variance of the probability distribution (column vectors of the last layer output). We show that both regularization terms are important to increasing the performance in label-shift and feature-shift settings.

Our overall loss function is shown in Eq. 1. Each client uses the cross-entropy loss as well as two regularization terms. The  $\mu$  and  $\lambda$  terms control the strength of each regularization (discussed in Section 4.1), and the  $f_{theta}$  and  $g_{theta}$  are the final layer and penultimate layer, respectively. We now discuss each regularization term in more detail.

$$\mathcal{L} = \mathcal{L}_{CE}(f_{\theta}(X), Y) + \mu \mathcal{L}_U(g_{\theta}(X)) + \lambda \mathcal{L}_V(f_{\theta}(X)) \quad (1)$$

### 3.1. Classifier Variance

Our goal is to encourage the classifier in any setting to emulate a classifier in an IID setting. We specifically focus on the classifier as it is the layer of the model most prone to bias in non-IID settings. We do this by inducing the class-wise variance of the classifier probability distribution to emulate an IID scenario.

Given a mini-batch  $(X, Y)$  and a model  $f_{\theta}$  that outputs a prediction  $f_{\theta}(X) = \hat{Y}$ , we create a probability distribution on our mini-batch  $\hat{P} := \text{Softmax}(\hat{Y})$ .  $\hat{P}$  is the probability distribution matrix of the entire mini-batch, where the row vectors are probabilities of a single sample and the column vectors are probabilities of a single class across the entire mini-batch. We define the variance regularization term  $\mathcal{L}_V$  as the hinge loss of the square root of the variance (standard deviation) along the column vectors of the probability matrix as seen in Eq. 2.

$$\mathcal{L}_V(\hat{Y}) = \frac{1}{D} \sum_{j=1}^D \max(0, c - \sqrt{\text{Var}(\hat{P}_j)}) \quad (2)$$

Here,  $D$  is the number of classes,  $c$  is a constant that represents the variance in an IID setting, and  $\hat{P}_j$  represents the column vectors of the probability matrix. We set  $c$  to be the average class-wise standard deviation of  $\mathbf{A}$  as shown in Eq. 3. Here,  $\mathbf{A}$  is a  $D \times D$  identity matrix ( $\text{diag}(1_1, \dots, 1_D)$ ) where  $D$  is the number of classes. This represents the ideal probability distribution of a balanced mini-batch where all classes are included.

$$c := \frac{1}{D} \sum_j^D \sqrt{\text{Var}(\mathbf{A}_j)} \quad (3)$$

In essence, we ask that the predicted probability distribution in any setting match the probability distribution of an ideal IID setting where each mini-batch contains a balanced ratio

of the included classes. Furthermore, we purposefully use a hinge loss to prevent the variance in a few classes from overpowering the remaining classes.

### 3.2. Hyperspherical Uniformity

Promoting class-wise variance in the classifier may be difficult in scenarios where the representations are clustered in a small area. For instance, in a feature-shift non-IID setting a model may be biased towards a subset of the global data manifold. This may also occur in scenarios where a client holds data from only two classes. There is no incentive for the model to spread the representations to allow larger separation with a linear hyperplane. We thus induce hyperspherical uniformity on the representations of the encoder (the output of the penultimate layer). We use the RBF kernel for its connection with the unit hypersphere [5], and its success in various deep learning applications [6, 25]. Namely, we calculate the average pairwise Gaussian energy between the row vectors of the representations as seen in Eq. 4.

$$\mathcal{L}_U(g_{\theta}(X)) = \mathbb{E}_{x, y \in g_{\theta}(X)} \left[ \exp\left(-\frac{\|x - y\|_2^2}{2\sigma}\right) \right] \quad (4)$$

$g_{\theta}(X)$  is the representation matrix of the encoder where  $x$  and  $y$  represent the row vectors in the representations.  $\sigma$  is a hyperparameter that controls the sensitivity to small distances. This is a hyperparameter that can be tuned to each setting to improve performance. We set  $\sigma$  as the median value of the squared pairwise distances to ensure hyperspherical uniformity is induced in any setting.

## 4. Experiments and Results

### 4.1. Experimental Setup

We compare FedUV with the state-of-the-art FL algorithms: FedAvg [31], FedProx [20], and MOON [19]. We also include a method *Freeze* which freezes the final layer. This is similar to FedBABU [34] without personalization. We mainly focus on these methods as they are local learning regularization. Additional baselines are presented in Appendix 5. We use the PyTorch [35] library and follow the official implementations for all available work.

**Label-shift Datasets.** **STL-10** [4] is a dataset with 10 classes with balanced 5,000 training and 8,000 validation samples. **CIFAR-100** [16] is a dataset with 100 classes with 60,000 training and 10,000 validation samples. **Tiny ImageNet** [18] is a dataset downsampled from the ImageNet [36] dataset with 200 classes with balanced 100,000 training and 10,000 validation samples. Basic data augmentation, random cropping and horizontal flipping, is used consistently throughout all datasets and methods. As with many previous work [19, 32, 34], we use the Dirichlet distribution to simulate label-shift non-IID settings. The  $\alpha$  term controls

Method	STL-10		CIFAR-100		Tiny ImageNet		PACS	HAM10000	Office-Home
	$\alpha=0.01$	$\alpha=1.0$	$\alpha=0.01$	$\alpha=1.0$	$\alpha=0.01$	$\alpha=1.0$			
FedAvg	27.6±1.6	68.5±0.7	51.6±1.7	58.3±1.2	37.5±0.9	40.2±1.2	61.9±0.6	73.7±0.3	42.2±1.1
FedProx	26.5±1.5	67.9±0.8	50.2±1.1	59.0±1.0	37.1±1.1	39.9±1.0	59.1±0.8	73.5±0.3	41.9±0.6
MOON	26.0±1.3	70.0±0.6	48.2±1.2	<b>60.5±0.7</b>	37.6±1.8	42.4±0.9	62.8±1.6	72.9±0.4	41.9±0.4
Freeze	23.7±1.5	<b>72.1±0.6</b>	51.1±1.5	59.3±1.8	<u>38.9±1.9</u>	42.6±1.6	<u>61.6±0.5</u>	73.2±0.3	41.3±1.8
FedUV	<b>30.4±1.4</b>	68.5±0.6	<b>55.7±1.0</b>	59.1±0.9	<b>40.3±0.8</b>	<b>43.2±1.5</b>	<b>65.9±0.9</b>	<b>73.9±0.5</b>	<b>45.4±1.0</b>

Table 1. Test accuracies with  $\alpha \in \{0.01, 1.0\}$  on STL-10, CIFAR-100, and Tiny ImageNet and PACS, HAM10000, and Office-Home. We **bold** the highest performing method and underline the second highest performing method.

the extent of non-IIDness with  $\alpha = 0$  being most non-IID and  $\alpha = \infty$  being IID. We use  $\alpha \in \{0.01, 1.0\}$ . We split the training datasets into a 90-10 training-validation dataset and use the original validation set as a test set.

**Feature-shift Dataests.** A more recent work [21] has shown it is important also to focus specifically on feature-shift environments. To simulate feature-shift, we use 3 datasets which includes data from various domains. **PACS** [17] is a dataset with 7 classes in 9,991 images from four domains: Photo (1,670 images), Art painting (2,048 images), Cartoon (2,344 images), and Sketch (3,929 images). We randomly sampled an equal 280 images per domain (40 images per class in each domain) for our testing dataset (1,120 images in total). **HAM10000** [40] is a dataset with 7 classes in 10,015 images from four domains. The domains are represented by different institutions which collected dermoscopic (skin disease) images. Each domain has 3363, 2259, 439, 3954 images respectively. Due to the imbalance in data, we took roughly 20% from each domain for 1955 images for our testing set. **Office-Home** [41] is a dataset with 65 classes in 15,588 images from 4 domains: Art (2,427 images), Clipart (4,365 images), Product (4,439 images), and Photo (4,357 images). We randomly sampled an equal 455 images per domain (7 images per class in each domain) for our testing dataset (1,820 images in total). We resize images from these datasets into 96x96x3 images and apply basic data augmentation (random cropping and horizontal flipping).

A detailed overview of the distributions of the datasets are provided in Appendix 2.

**Models.** We use 3 different models across the different datasets. Note that we also add a non-linear projector following [19] due to reduced performance without it. This is simply two fully-connected layers followed by a batch norm and ReLU activation. For STL-10 and PACS, we used a small CNN model (Appendix 3) with a projector with 256 neurons. For CIFAR-100 and PACS, we used a ResNet-18 [10] model, removing the initial max-pooling layer due to the small image size and adding the projector with 512 neurons. For Tiny ImageNet and Office-Home, we used a ResNet-50 model, removing the initial max-pooling layer and adding the projector with 2048 neurons. On all datasets, we report the results of a single aggregated global model.

**Hyperparameters.** Our default testing environment for our ablation is done on STL-10 ( $\alpha = 0.01$ ) with a small CNN (Appendix 3) and Office-Home with a ResNet-50, local epoch  $E = 10$ , number of clients  $\kappa = 10$  for STL-10 and  $\kappa = 4$  for Office-Home (for each of the domains), participation rate  $\rho = 1.0$ . The total aggregation rounds is chosen based on performance of the training-validation split on FedAvg. We set total aggregation rounds,  $R$ , as 100, 60, 40 for STL-10, CIFAR-100, Tiny ImageNet, and 60, 60, 40 for PACS, HAM10000, and Office-Home, respectively. All reported accuracy are averaged over 3 runs. Changes from these default settings are mentioned clearly. We use the Cross-entropy loss and SGD with learning rate of 0.01, momentum 0.9, and weight decay of 1e-5.

FedProx, MOON, and FedUV are all regularization methods, thus adds an extra parameter to balance the cross-entropy loss and the strength of regularization. We tune this parameter on the CIFAR-10 [16] dataset, since tuning this parameter in every setting of each dataset is not viable option for FL as resources are limited. For FedProx, we use  $\mu = 0.01$ , for MOON,  $\mu = 1.0$ . For FedUV, we use fixed hyperparameters  $\lambda = \frac{D}{4}$ , where  $D$  is the number of classes, and  $\mu = 0.5$  for all datasets. We report the hyperparameter tuning results in the appendix (Appendix 4). All experiments were conducted on a single RTX A5000 and one AMD EPYC 7763 processor.

## 4.2. Results

**Data Heterogeneity.** We study how data heterogeneity changes performance of the global model. We test the performance for  $\alpha \in \{0.01, 1\}$  for label-shift datasets, and use the defined domains for feature-shift datasets. Results are shown in Table 1.

Across most settings, FedUV outperforms other methods. Focusing first on the label-shift non-IID settings, we see that in the extreme  $\alpha = 0.01$  setting, FedUV performs best. FedProx and MOON perform worse than FedAvg as they base their regularization on the global model which may also be biased since it is averaging extremely biased local models. Freeze also does not perform well, as the frozen classifier causes the encoder to have difficulty learning in extreme settings. By emulating an IID setting, FedUV outperforms these methods.

Method	STL-10 ( $\alpha = 0.01$ )			Office-Home		
	p=0.1	p=0.5	p=1.0	p=0.25	p=0.5	p=1.0
FedAvg	17.4±3.4	21.5±1.8	27.6±1.6	39.1±2.5	41.9±1.4	42.2±1.1
FedProx	19.3±2.9	22.0±1.3	26.5±1.5	39.6±2.1	40.8±1.7	41.9±0.6
MOON	20.5±3.7	23.6±1.7	26.0±1.3	40.3±2.7	40.5±1.6	41.9±0.4
Freeze	21.5±3.6	23.7±1.6	23.7±1.5	37.6±3.9	39.5±1.1	41.3±1.8
FedUV	<b>24.9±3.1</b>	<b>30.0±1.1</b>	<b>30.4±1.4</b>	<b>41.7±2.4</b>	<b>44.8±1.3</b>	<b>45.4±1.0</b>

Table 2. Accuracy across client participation rate

In the less extreme setting of  $\alpha = 1.0$  other regularization methods generally perform well. On STL-10, Freeze performs best by a large margin. For less extreme label shift and small number of classes, a frozen classifier is beneficial as this prevents the classifier from being biased, which has been established as a desirable property. MOON performs best on CIFAR-100, and also improves performance compared to FedAvg on the Tiny ImageNet dataset. However, we find that Freeze and FedUV achieve the better performance on Tiny ImageNet most likely because each client holds more data as the dataset is much larger. When data is plentiful, it is likely that focusing specifically on the classifier is most beneficial.

In the feature-shift setting, we find that baseline regularization methods generally does not help performance. This could be because these work have focused only on label-shift settings. Freeze simply freezes the final layer, while MOON and FedProx use the global model as a source of regularization to mitigate the bias of individual clients. However, because FedUV simulates an IID setting rather than mitigating the effects of label-shift by relying on the global model as a source of regularization, performance is improved not only in label-shift settings but also feature-shift settings. As will be shown in the ablation study, the Uniformity regularizer is indeed an important factor to prevent this performance degradation.

**Client participation.** In real-world FL applications, it is realistic to expect that only a subset of all clients participate in an aggregation round. The server must triggers the aggregation process even if some clients do not participate. We study the effect of ratio of client participation,  $\rho$ , where  $\rho = 0.1$  means 10% of clients participate in the aggregation round. Across these settings, we use 100 aggregation rounds for STL-10 and 60 for Office-Home. The results are shown in Table 2.

The advantages of FedUV are clear even when the ratio of participating clients is low. Indeed, when  $\rho = 0.1$  for STL-10, FedUV performs the best by a large margin. This increase in performance occurs owing to the focus of FedUV to emulate the IID. The second best method is Freeze. Intuitively, since the classifier is the layer that is most prone to bias, this is magnified in scenarios where client participation is low, as bias of the averaged global model is further increased. By fixing the classifier as in Freeze, we limit this bias and increase performance. However, as seen with the

Method	STL-10 ( $\alpha = 0.01$ )			Office-Home		
	$E = 10$	$E = 20$	$E = 40$	$E = 10$	$E = 20$	$E = 40$
FedAvg	27.6±1.6	29.8±1.1	<b>29.1±0.3</b>	42.2±1.1	43.5±1.0	35.1±0.7
FedProx	26.5±1.5	28.6±0.8	27.4±0.6	41.9±0.6	43.3±0.5	35.7±0.5
MOON	26.0±1.3	28.9±0.7	27.4±0.4	41.9±0.4	42.0±0.4	<b>41.7±0.5</b>
Freeze	23.7±1.5	29.9±0.9	27.2±0.7	41.3±1.8	41.3±0.8	35.4±0.9
FedUV	<b>30.4±1.4</b>	<b>31.5±1.2</b>	25.9±0.9	<b>45.4±1.0</b>	<b>46.1±0.7</b>	35.9±0.8

Table 3. Accuracy across different number of local epochs

Method	STL	CIFAR	Tiny	PACS	HAM	Office
FedAvg	27.6±1.6	51.6±1.7	37.5±0.9	61.9±0.6	73.3±0.3	42.2±1.1
Freeze	23.7±1.5	54.1±1.5	38.9±1.9	61.6±0.5	73.2±0.3	41.3±1.8
Only $L_V$	29.5±1.4	54.0±1.1	38.8±1.5	60.1±0.3	72.5±0.4	41.8±1.1
Only $L_U$	24.1±1.3	52.8±1.2	37.9±1.7	63.7±0.3	73.4±0.2	43.5±1.3
FedUV	30.4±1.4	55.7±1.0	40.3±0.9	65.9±0.9	73.9±0.5	45.4±1.0

Table 4. FedUV ablation study

performance increase of FedUV, it is far more beneficial to emulate an IID setting. It is also interesting to see that when  $\rho = 0.25$  for Office-Home (1 client participating), FedProx and MOON improves performance. When the server aggregates weights from only 1 domain, it may be beneficial to use the global model as regularization. This improvement does not continue for  $\rho = 0.5$  (2 clients participating), suggesting that the average of two domains is more beneficial than the use of the global model for regularization.

**Number of local epochs** We study the effect of number of local epochs,  $E$ , where  $E = 40$  means each clients trains for 40 epochs during a single aggregation round. Results are shown in Table 3. It would be desirable to achieve high performance with a low number of local epochs thereby reducing the burden of less powerful edge devices. We find that when  $E = 20$  performance generally improves. However, when the number of local epochs increase to 40, the accuracy of all methods decrease, suggesting overfitting occurs with a high number of local epochs. The performance of FedUV significantly degraded when local epoch is high. This suggests high local epochs cause local models to converge to local optima, harming performance. On the other hand, MOON and FedProx did not show a large drop in performance. This can be explained by the fact that these methods rely on the global model for regularization. Because of this, local models are less likely to converge to local optima as compared to FedUV.

**Ablation Study.** We study the effect of hyperspherical uniformity and classifier variance separately. We also compare Freeze which creates an orthogonal classifier by randomly initializing and freezing the weights. Results are shown in Table 4. For the label-shift datasets (STL-10, CIFAR-100, Tiny ImageNet), we find that hyperspherical uniformity regularization, *Only  $L_U$* , does not improve performance significantly, whereas classifier variance regularization, *Only  $L_V$* , performs much better than Freeze on the STL-10 dataset while being competitive in other datasets. These results




Figure 3. The training loss for FedAvg and FedUV across different settings

suggests that label-shift settings, focusing on the final layer is more important than focusing on the feature representations of the encoder, and that emulating an IID label distribution is more desirable than a fixed classifier.

Furthermore, for the feature-shift datasets (PACS, HAM10000, Office-Home), we find that variance regularization does not improve performance significantly, whereas uniformity regularization improves performance compared to both *Only*  $L_V$  and Freeze. This suggests that in feature-shift settings, when clients have vastly differing feature distributions, it is beneficial to encourage features to be more spread. Intuitively, this prevents any local encoder from heavy bias towards a subspace, which is an aspect other regularization methods do not incorporate. The performance of FedUV across different models is shown in Appendix 5.

Overall, these results suggests that, in label-shift situations, focusing specifically on the classifier through variance regularization is more important, while in feature-shift scenarios, focusing specifically on the feature space representations of the encoder through hyperspherical uniformity is more important. However, both regularization terms are important to emulate the IID setting regardless of data distribution. The combined approach in FedUV shows further improved performance across these various settings.

**Convergence.** We study the convergence of FedUV. Graphs of the training loss are shown in Fig. 3. FedUV shows quicker convergence when compared to FedAvg, often reaching a lower loss at earlier training iterations. This is intuitive considering the performance improvement FedUV provides over FedAvg. We also find that in the extreme case of STL-10  $\alpha = 0.01$ , the loss curve is much smoother for FedUV when compared to FedAvg. This shows FedUV stabilizes training.

## 5. Discussion

### 5.1. Preventing Classifier Bias

Our method of emulating the IID setting by encouraging classifier variance is not the only method to prevent classifier bias. When focusing on label-shift, many loss function weighting and specialized sampling techniques have been considered [13]. There have also been interesting applica-

tions in representation learning that use latent space representations to sample training data such that the predicted probability distribution is uniform [2]. These are interesting directions to explore for addressing the non-IID problem. The biggest constraint using these approaches in FL however, is privacy. The data and class distributions cannot be known beforehand, thus cannot be set to match each client. In addition, parameters that control sampling or loss weighting may introduce security vulnerabilities.

Another straightforward method would be to randomly initialize and freeze the classifier. Freezing the classifier would also prevent biases forming in the probability distribution of the classifier. Indeed, this is the intuition behind FedBABU [34] and the Freeze method we have used as our baseline, which initializes and freezes the classifier to orthogonal vectors. Freezing weights also prevents the degeneration of singular values in non-IID settings, a phenomenon shown in Fig. 1.

However, it may not be optimal to separate classes with orthogonal vectors in every setting. A number of studies [25, 28] have shown that when the number of neurons exceeds the dimensions of data manifold, promoting orthogonality is problematic. This problem arises when there is a large number of classes with data that lives in a lower subspace. Another alternative would be to regularize the weights directly using the spectral norm or the Frobenius norm. These methods differ from FedUV as we regularize in the representation space rather than in the weight space. The main reason is efficiency. We re-use the representations obtained in the forward-pass of the model without directly accessing the weight parameters.

### 5.2. Hyperspherical Uniformity

The goals of hyperspherical uniformity in FedUV are to mitigate representational bias and to create separation between representations such that the classifier has more directions to increase its variance. Note that this regularizer focuses on feature distributions, not class distributions. This is shown in Figure 4. Features are aligned regardless of class distribution. Empirically, we have shown that this regularization is beneficial in feature-shift settings. There are alternatives to this, such as spreading representations along a hypercube. However, hyperspherical uniformity is gain-




Figure 4. t-SNE on the features (output of the penultimate layer) of STL-10 (Class #1). Experiment setup is equivalent to main text. Clients and class were selected at random.

Method	STL (10-Layer)	CIFAR (ResNet-18)	Tiny (ResNet-50)	PACS (10-Layer)	HAM (ResNet-18)	Office (ResNet-50)
FedAvg	26	57	728	47	131	359
FedProx	27	67	823	48	135	371
MOON	37	88	1186	70	145	493
Freeze	26	56	728	47	129	351
FedUV	27	58	755	48	132	362

Table 5. Average time (in seconds) per aggregation round.

ing popularity in the field of deep learning due to its desirable properties [6, 24–27]. While more traditional regularization approaches such as the  $L_2$  regularization have seen success in various applications in the wider field of machine learning and statistics [8, 12], neural networks are special due to their high over-parameterization. Using  $L_2$  regularization focuses only on weights rather than the interactions between over-parameterized neurons. Hyperspherical uniformity regularizes based on these interactions and also has clear geometric properties [28]. We point readers to [28] for a recent in-depth discussion on its development in deep learning.

Furthermore, we choose the RBF kernel for its efficiency. In particular, the kernel method allows us to first calculate the pairwise  $L_2$  norms of the row vectors before using the kernel function. It also allows for easy differentiation on automatic differentiation software, which facilitate the training of neural networks.

### 5.3. Convergence of FedUV

We have shown in Fig. 3 that in our extensive testing environments, FedUV converges stably. In the general case, we know that the Uniformity regularizer converges weakly to the uniform distribution, since we are minimizing the average pairwise gaussian potential between sample points. As the number of samples approaches infinity, the loss will converge weakly to the uniform distribution according to Wang and Isola [45]. For the Variance regularizer, however, the hinge loss of the regularizer is not smooth. Because of this, it is non-trivial to prove convergence. This may become much simpler when replaced with a smooth surrogate loss. We show in Appendix 5 the performance of settings in

which squared hinge and no hinge loss is used.

### 5.4. Efficiency and Scalability

One of the most important aspects in FL is efficiency. Edge devices cannot be expected to have powerful hardware that allows heavy computation in real-world scenarios. We unfortunately find that many regularization methods are not very scalable. While this is not a problem for small datasets and models, as dataset size and model size both increase at a rapid rate, these models become more computationally expensive.

We compare the wall-clock time of different methods as shown in Table 5. Note that the time extended is per aggregation round with 10 local epochs. FedProx requires  $L_2$  norm calculations between each layer of local model and global model at each batch, while MOON requires 3 forward passes at each batch as well as the storage of 3 full models. Though this may not be a problem with small models, FedProx and MOON are not very scalable methods as the size of models increases. FedUV is a simple yet efficient regularization technique as it only requires few matrix operations with no weight parameter access, extra memory, or extra forward passes. The most expensive operation is the pairwise distance calculation which is negligible even with less powerful hardware.

## 6. Conclusion

In this work, we introduced a new approach to address the non-IID problem in FL by encouraging local models to emulate the IID setting. Rather than regularizing to reduce bias, we emulate the IID setting by promoting variance in the predicted dimension-wise probability distribution. We also promote hyperspherical uniformity on the representations of the encoder to allow the classifier to increase variance in more directions. Overall, our method improves performance by a large margin throughout our extensive experiments while being the most efficient and scalable among regularization methods.

**Acknowledgements** The work was partially supported by USDA-020-67021-32855 and NSF OIA-2134901.



## References

- [1] Durmus Alp Emre Acar, Yue Zhao, Ramon Matas Navarro, Matthew Mattina, Paul N Whatmough, and Venkatesh Saligrama. Federated learning based on dynamic regularization. *arXiv preprint arXiv:2111.04263*, 2021. 1, 3
- [2] Alexander Amini, Ava P Soleimany, Wilko Schwarting, Sangeeta N Bhatia, and Daniela Rus. Uncovering and mitigating algorithmic bias through learned latent structure. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 289–295, 2019. 7
- [3] Anda Cheng, Peisong Wang, Xi Sheryl Zhang, and Jian Cheng. Differentially private federated learning with local regularization and sparsification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10122–10131, 2022. 2
- [4] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011. 4
- [5] Henry Cohn and Abhinav Kumar. Universally optimal distribution of points on spheres. *Journal of the American Mathematical Society*, 20(1):99–148, 2007. 4
- [6] Tim R Davidson, Luca Falorsi, Nicola De Cao, Thomas Kipf, and Jakub M Tomczak. Hyperspherical variational auto-encoders. *arXiv preprint arXiv:1804.00891*, 2018. 4, 8
- [7] Enmao Diao, Jie Ding, and Vahid Tarokh. Heteroff: Computation and communication efficient federated learning for heterogeneous clients. *arXiv preprint arXiv:2010.01264*, 2020. 2
- [8] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. 8
- [9] Irving John Good. Rational decisions. In *Breakthroughs in statistics*, pages 365–377. Springer, 1992. 2
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 5
- [11] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019. 2
- [12] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 8
- [13] Justin M Johnson and Taghi M Khoshgoftaar. Survey on deep learning with class imbalance. *Journal of Big Data*, 6(1):1–54, 2019. 7
- [14] Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. Byzantine-robust learning on heterogeneous datasets via bucketing. *arXiv preprint arXiv:2006.09365*, 2020. 2
- [15] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pages 5132–5143. PMLR, 2020. 1
- [16] Alex Krizhevsky. Learning multiple layers of features from tiny images, 2009. 2, 4, 5
- [17] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017. 5
- [18] Fei-Fei Li, Andrej Karpathy, and Justin Johnson. Tiny imagenet, 2014. 4
- [19] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10713–10722, 2021. 1, 3, 4, 5
- [20] 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, 3, 4
- [21] Xiaoxiao Li, Meirui Jiang, Xiaofei Zhang, Michael Kamp, and Qi Dou. Fedbn: Federated learning on non-iid features via local batch normalization. *arXiv preprint arXiv:2102.07623*, 2021. 5
- [22] Xin-Chun Li, Yi-Chu Xu, Shaoming Song, Bingshuai Li, Yinchuan Li, Yunfeng Shao, and De-Chuan Zhan. Federated learning with position-aware neurons. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10082–10091, 2022. 3
- [23] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. *Advances in Neural Information Processing Systems*, 33:2351–2363, 2020. 3
- [24] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheraface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017. 8
- [25] Weiyang Liu, Rongmei Lin, Zhen Liu, Lixin Liu, Zhiding Yu, Bo Dai, and Le Song. Learning towards minimum hyperspherical energy. *Advances in neural information processing systems*, 31, 2018. 4, 7
- [26] Weiyang Liu, Zhen Liu, Zhiding Yu, Bo Dai, Rongmei Lin, Yisen Wang, James M Rehg, and Le Song. Decoupled networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2771–2779, 2018.
- [27] Weiyang Liu, Zhen Liu, James M Rehg, and Le Song. Neural similarity learning. *Advances in Neural Information Processing Systems*, 32, 2019. 8
- [28] Weiyang Liu, Rongmei Lin, Zhen Liu, Li Xiong, Bernhard Schölkopf, and Adrian Weller. Learning with hyperspherical uniformity. In *International Conference On Artificial Intelligence and Statistics*, pages 1180–1188. PMLR, 2021. 7, 8
- [29] Mi Luo, Fei Chen, Dapeng Hu, Yifan Zhang, Jian Liang, and Jiashi Feng. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. *Advances in Neural Information Processing Systems*, 34:5972–5984, 2021. 1, 3

- [30] Xiaosong Ma, Jie Zhang, Song Guo, and Wenchao Xu. Layer-wised model aggregation for personalized federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10092–10101, 2022. [2](#)
- [31] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. [1](#), [2](#), [4](#)
- [32] Matias Mendieta, Taojiannan Yang, Pu Wang, Minwoo Lee, Zhengming Ding, and Chen Chen. Local learning matters: Rethinking data heterogeneity in federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8397–8406, 2022. [1](#), [3](#), [4](#)
- [33] Lokesh Nagalapatti, Ruhi Sharma Mittal, and Ramasuri Narayanam. Is your data relevant?: Dynamic selection of relevant data for federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022. [2](#)
- [34] Jaehoon Oh, Sangmook Kim, and Se-Young Yun. Fedbabu: Towards enhanced representation for federated image classification. *arXiv preprint arXiv:2106.06042*, 2021. [1](#), [2](#), [3](#), [4](#), [7](#)
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. [4](#)
- [36] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. [4](#)
- [37] Aviv Shamsian, Aviv Navon, Ethan Fetaya, and Gal Chechik. Personalized federated learning using hypernetworks. In *International Conference on Machine Learning*, pages 9489–9502. PMLR, 2021. [2](#)
- [38] Yiqing Shen, Yuyin Zhou, and Lequan Yu. Cd2-pfed: Cyclic distillation-guided channel decoupling for model personalization in federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10041–10050, 2022. [2](#)
- [39] Minxue Tang, Xuefei Ning, Yitu Wang, Jingwei Sun, Yu Wang, Hai Li, and Yiran Chen. Fedcor: Correlation-based active client selection strategy for heterogeneous federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10102–10111, 2022. [2](#)
- [40] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1):1–9, 2018. [5](#)
- [41] 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. [5](#)
- [42] Joost Verbraeken, Matthijs Wolting, Jonathan Katzy, Jeroen Kloppenburg, Tim Verbelen, and Jan S. Rellermeyer. A survey on distributed machine learning, 2019. [1](#)
- [43] Hongyi Wang, Mikhail Yurochkin, Yuekai Sun, Dimitris Papaliopoulos, and Yasaman Khazaeni. Federated learning with matched averaging. *arXiv preprint arXiv:2002.06440*, 2020. [3](#)
- [44] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33:7611–7623, 2020. [2](#)
- [45] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR, 2020. [8](#)
- [46] Donglei Wu, Xiangyu Zou, Shuyu Zhang, Haoyu Jin, Wen Xia, and Binxing Fang. Smartidx: Reducing communication cost in federated learning by exploiting the cnns structures. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4254–4262, 2022. [2](#)
- [47] An Xu, Wenqi Li, Pengfei Guo, Dong Yang, Holger R Roth, Ali Hatamizadeh, Can Zhao, Daguang Xu, Heng Huang, and Ziyue Xu. Closing the generalization gap of cross-silo federated medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20866–20875, 2022. [2](#)
- [48] Haibo Yang, Minghong Fang, and Jia Liu. Achieving linear speedup with partial worker participation in non-iid federated learning. *arXiv preprint arXiv:2101.11203*, 2021. [2](#)
- [49] Taojiannan Yang, Sijie Zhu, and Chen Chen. Gradaug: A new regularization method for deep neural networks. *Advances in Neural Information Processing Systems*, 33:14207–14218, 2020. [1](#), [3](#)
- [50] Tehrim Yoon, Sumin Shin, Sung Ju Hwang, and Eunho Yang. Fedmix: Approximation of mixup under mean augmented federated learning. *arXiv preprint arXiv:2107.00233*, 2021. [2](#)
- [51] Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Nghia Hoang, and Yasaman Khazaeni. Bayesian nonparametric federated learning of neural networks. In *International Conference on Machine Learning*, pages 7252–7261. PMLR, 2019. [2](#)
- [52] Lin Zhang, Li Shen, Liang Ding, Dacheng Tao, and Ling-Yu Duan. Fine-tuning global model via data-free knowledge distillation for non-iid federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10174–10183, 2022. [2](#)