

FEDMAE: FEDERATED SELF-SUPERVISED LEARNING WITH ONE-BLOCK MASKED AUTO-ENCODER

Anonymous authors

Paper under double-blind review

ABSTRACT

Latest federated learning (FL) methods started to focus on how to use unlabeled data in clients for training due to users' privacy concerns, high labeling costs, or lack of expertise. However, current Federated Semi-Supervised/Self-Supervised Learning (FSSL) approaches fail to learn large-scale images because of the limited computing resources of local clients. In this paper, we introduce a new framework FedMAE, which stands for Federated Masked AutoEncoder, to address the problem of how to utilize unlabeled large-scale images for FL. Specifically, FedMAE can pre-train one-block Masked AutoEncoder (MAE) using large images in lightweight client devices, and then cascades multiple pre-trained one-block MAEs in the server to build a multi-block ViT backbone for downstream tasks. Theoretical analysis and experimental results on image reconstruction and classification show that our FedMAE achieves superior performance compared to the state-of-the-art FSSL methods.

1 INTRODUCTION

Self-supervised learning (SSL) has attracted extensive research attention for learning representations without costly data labeling. The common practice in computer vision is to design proxy tasks to allow visual representation learning from unlabeled images (Doersch et al., 2015; Noroozi & Favaro, 2016; Zhang et al., 2016; Gidaris et al., 2018). The current state-of-the-art SSL methods integrate contrastive learning with Siamese networks to increase the similarity between two augmented versions of images (Wu et al., 2018; Chen et al., 2020b; He et al., 2020; Grill et al., 2020; Chen & He, 2021), or use autoencoder to reconstruct the input images (Dosovitskiy et al., 2020; Bao et al., 2021; He et al., 2022), while all these solutions assume training images are centrally available in cloud servers. In the real world, however, decentralized image data are exploding in growth, and data collected by different parties may not be centralized due to data privacy regulations.

Existing approaches cannot well use decentralized unlabeled data to learn a generic representation while preserving data privacy. Federated Learning (FL) (Chen et al., 2019; McMahan et al., 2017) is an emerging distributed training technique in which several clients collectively train a global model through coordinated communication. However, traditional federated learning (Kairouz et al., 2021) relies on data labels, while this assumption may be excessive as labeling data is expensive, time-consuming, and indispensable to the participation of domain experts. Decentralized data is normally unlabeled, non-independent and identically distributed (non-IID), which are critical challenges for learning from multiple clients (Li et al., 2020; Zhao et al., 2018). Motivated by these practical scenarios, many SSL methods are used in federated learning, such as FedU (Zhuang et al., 2021a), FedEMA (Zhuang et al., 2021b) and Orchestra (Lubana et al., 2022). Due to the limited computation resource of client devices, these approaches use small models, e.g., ResNet18, to train small images. The state-of-the-art methods have not revealed how lightweight clients learn unlabeled large-scale images effectively.

Recent advancement of Transformers has great achievements in computer vision Dosovitskiy et al. (2020); He et al. (2022). However, the biggest problem with applying Transformers in federated learning is a large weight, which is difficult to deploy on small computing devices for training large images. In this paper, we propose a simple, effective, and scalable framework, FedMAE, which stands for Federated Masked AutoEncoder. Our FedMAE only deploys the MAE with a one-block encoder and a one-block decoder as our pre-training model on the local clients and masks the random




Figure 1: Comparison of reconstructions of ImageNet validation images under 75% masking ratio. From left to right of each quaternion, we show the masked image, the reconstruction of the un-pretrained five-block MAE, the reconstruction of the five-block MAE pre-trained by FedMAE, and the ground-truth image.

patches in the input image and then reconstructs the missing patches in the pixel space. Although the one-block model learns limited information and performs poorly in downstream tasks, we found that cascading many pre-trained one-block MAEs together (without weights aggregation in traditional FL) can achieve stunning performance in downstream tasks. With this design, FedMAE can achieve a win-win scenario, i.e., it improves model accuracy by training large images on lightweight client devices and allowing more distributed or mobile clients to participate in training asynchronously. To get a qualitative sense of our FedMAE performance, see Figure 1, which compares the reconstruction effect of the un-pretrained five-block MAE and the five-block MAE pre-trained by FedMAE.

In summary, the main contributions of this work are summarized as follows:

- We propose FedMAE, a framework that realizes asynchronous training of large-scale unlabeled images in FL. To the best of our knowledge, we are the first to apply Transformer to FL, which shows that cascading multiple pre-trained one-block MAEs can form a high-performance pre-trained ViT model for downstream tasks.
- We model an approximate mapping for FedMAE based on the equivalence between data corruption and feature interference and give an approximate analytical solution to optimal reconstruction, which depends only on the data, revealing the data impact of the optimal reconstruction in federated autoencoding learning independent of the backbone.
- We pre-train FedMAE with large-scale datasets in FL and conduct extensive fine-tuning experiments on downstream tasks, such as image reconstruction and image classification. The results show FedMAE has huge improvement (30% to 50%) over state-of-the-art baselines.

2 RELATED WORK

2.1 SELF-SUPERVISED LEARNING AND TRANSFORMERS

Self-Supervised Learning is a recent paradigm in unsupervised representation learning, especially discriminative SSL, which exploits the invariance of different views after data augmentation, making the representations of positive pairs similar while pushing negative pairs away. In order to obtain sufficient informative negative examples in contrastive learning, the methods usually rely on large memory banks such as Moco (He et al., 2020), or large batch size such as SimCLR(Chen et al., 2020b). BYOL (Grill et al., 2020) and SimSiam (Chen & He, 2021) further minimize the need for negative samples by using a variety of strategies to prevent representation collapse.

Pre-training vision Transformers has attracted great attention recently due to the data-hungry issue. ViT (Dosovitskiy et al., 2020) predicts the 3-bit mean color of the masked patches, which is the most straightforward translation of BERT from NLP to CV. BEIT (Bao et al., 2021) uses image patches as input and visual tokens are obtained by discrete VAE instead of clustering. The current most popular way of pretraining Transformer are MAE (He et al., 2022) and MaskFeat (Wei et al., 2022), which utilize the idea of self-supervised learning to predict features of the masked area.

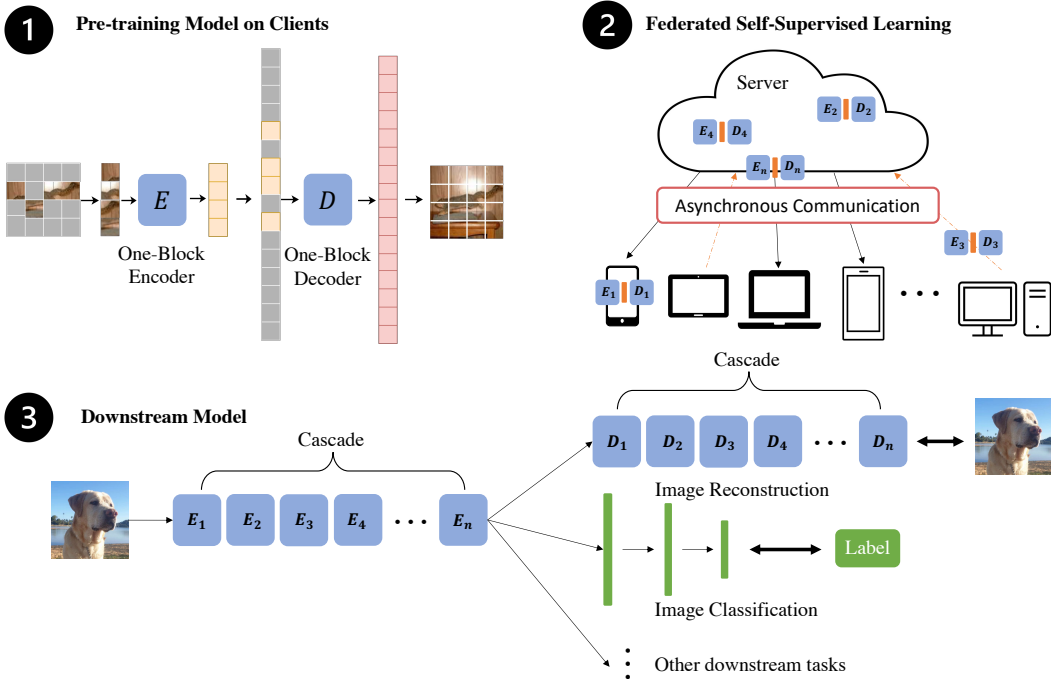


Figure 2: Overview of FedMAE Architecture. ① We use a tiny MAE with a one-block encoder and a one-block decoder as our pre-training model on clients. ② Pre-training models in FL can be trained on different clients asynchronously without average aggregation. ③ After pre-training in FL, pre-trained models are cascaded in the server to build a multi-block ViT for downstream tasks.

2.2 FEDERATED SEMI-SUPERVISED/SELF-SUPERVISED LEARNING

Semi-supervised learning and self-supervised learning are two methods in federated learning to utilize unlabelled data in client devices for training. Federated semi-supervised learning has to rely on some labeled data either in the server or in the devices (or both) to guide unlabeled data in training the global model. RSCFed (Liang et al., 2022) relies on clients with labeled data to assist in learning unlabeled data in other clients. FedMatch (Jeong et al., 2020) introduced the inter-client consistency that aims to maximize the agreement across models trained at different clients.

Federated self-supervised learning normally does not rely on any labeled data but aims at pre-training a high-performance model for downstream tasks. FedU (Zhuang et al., 2021a) is based on the self-supervised method BYOL (Grill et al., 2020), which aims for representation learning. FedEMA (Zhuang et al., 2021b) is an upgraded version of FedU, which adaptively updates online networks of clients with EMA of the global model. Orchestra Lubana et al. (2022) relies on high representational similarity for related samples and low similarity across different samples.

3 APPROACH

A masked autoencoder is a simple autoencoding method that reconstructs the original signal given its partial observation. To enable more local clients to participate in training, FedMAE’s pre-training model contains only one block of encoder and decoder, allowing many devices with insufficient computing power, e.g., mobile phones and Internet of Things devices, to participate in training on large images. Like all autoencoders, our one-block encoder maps the observed signal to a latent representation, and the one-block decoder reconstructs the original signal from the latent representation when pre-training. For the downstream model of FedMAE, we adopt a cascade design that allows multiple pre-trained one-block MAE connected in sequence to build a ViT model with multi-blocks for downstream training. Figure 2 illustrates the idea, introduced next.

3.1 PRE-TRAINED MODEL

In FedMAE, we build a tiny MAE with a one-block encoder and a one-block decoder as the pre-trained model for client devices. Our encoder is a ViT but applied only on visible, unmasked patches. Similar to a traditional ViT, our encoder embeds patches using a linear projection with additional positional embeddings and then processes the resultant set using a single Transformer block. Just as in a standard MAE, our encoder only performs on a small subset (e.g., 25%) of the whole set. Masked patches are removed, and no mask tokens are used for our encoder. The whole set is handled by a lightweight one-block decoder. The pre-trained model only has one Transformer block, which is used to predict masked patches for reconstructing original images.

3.2 DOWNSTREAM MODEL

Without a doubt, the pre-trained model has a certain ability to filter interference that masks features because of the loss of explicit information. Therefore, when the pre-trained model is transferred to the downstream ViT classifier, its performance is significantly improved compared to the unpretrained ViT, which can only obtain classification knowledge through label training and cannot mine implicit features. Moreover, the input and output of each block from the Transformer are in the same dimension, and each block in the multi-block ViT is learning how to extract feature representations. They have the same or similar learning capabilities, although there are some differences in extracting features. Due to the similarity of learning purposes for each block, FedMAE tries to cascade multiple one-block pre-trained MAEs, which will achieve higher-order implicit feature representations based on multi-block mapping to help boost and speed up the downstream task training.

3.3 TRAINING PROCESS

In Figure 2, ① shows the pre-training model of FedMAE on clients, which only contains a one-block encoder and a one-block decoder. We use unmasked patches (25% of the whole image) for encoding, which means the model is tiny that is easy to deploy in many lightweight clients and train large images. ② shows the communication between clients and the server. Due to the special one-block structure of the pre-training model and the very small learning rate of Transformer, the average operation of multiple model weights is not necessary (or effective) to boost the accuracy. Therefore, FedMAE realises pure asynchronous training, i.e., multiple pre-training models can train on different clients asynchronously without relying on each other. ③ shows cascading multiple pre-trained one-block MAEs to form a multi-block model to help train the downstream task, e.g., both the encoder and decoder are needed to perform the image reconstruction task, and only the encoder is used to produce image representations for recognition.

4 ANALYSIS

Based on the FedMAE as a system prototype, we analyze the data modeling and system mapping of federated learning with autoencoding. We analyze the relationship between the absence of explicit information in the data and feature perturbations and establish the equivalence between spatial variation in the data and feature denoising. An analytic model for federated autoencoders is given based on approximate linear mapping with a corresponding analytic solution as the linear approximation of optimal reconstruction. This analytic solution depends on the data only, revealing the data impact of the optimal reconstruction independent of the backbone in federated learning.

4.1 PROBLEM FORMULATION

4.1.1 DATA CORRUPTION AND FEATURE INTERFERENCE

Modeling the pre-training dataset as a source dataset $D_S = \{x_{S1}, x_{S2}, \dots, x_{Sn}\}$ and downstream dataset as a target dataset $D_T = \{x_{T1}, x_{T2}, \dots, x_{Tn}\}$ with labels $Y_T = \{y_{T1}, y_{T2}, \dots, y_{Tn}\}$, we design an encoder $h(x) = z$ and a decoder $g(z) = x$. The reconstruction by the pair of $h(x), g(z)$ can be represented as $\hat{x} = g(h(x))$. We use a loss function $l(x, g(h(x)))$ to measure the error between the original and reconstructed data, while the parameters of the autoencoder are learned through D_S to minimize the reconstruction error.

Since MAE uses patch sampling, the x input loses much explicit information, degraded as \tilde{x} , which leads to a corrupted input-to-feature explicit feature knowledge representation of $g(\cdot), h(\cdot)$. At this point, the training task becomes into reconstructing the original input x from its corrupted degradation \tilde{x} by minimizing $l(x, g(h(\tilde{x})))$. While on the features, the loss of explicit information corresponds to the loss of a portion of the features, leading to the corrupted degradation of the features as \tilde{z} . Since we use uniform dimensional encoding in the features, \tilde{z} and z are still uniform dimensions, so the above corruption in the features can be modeled as an interference as $\Delta z = z - \tilde{z}$.

It can be seen that the loss of explicit information in data x is equivalent to the noise interference of feature z . We analyze this feature interference due to data loss and manage to design a mechanism to reduce this interference by hidden feature mining and train the system to learn for discovering hidden features.

4.1.2 FEDERATED LEARNING WITH VARIATION BASED AUGMENTATION

We adopt FL in the pre-training system. There are K clients and 1 server, where each client has n samples as a dataset $\mathbf{X}_k = \{x_{ki}\}_n$, where k refers to the index of the k th client. If we aggregate all the samples and format them into a package, the corresponding aggregation can be formulated as $X_k = [x_{k1}, \dots, x_{kn}]$. With MAE, the system training uses random corruption, i.e., the location of the deducted patches is different each time. Therefore, the original data on k th client then has multiple variant versions of $\tilde{\mathbf{X}}_{k,j} = \{x_{ki,j}\}_{n \times m}$ with the corresponding aggregation as $\tilde{X}_{k,j} = [\tilde{x}_{k1,j}, \dots, \tilde{x}_{kn,j}]$, where $\tilde{x}_{ki,j}$ refers to the variation result of the original input x_{ki} with the j th masked corruption.

When m times of masked corruption is applied, then $\tilde{X}_{k,j}, k \leq K, j \leq m$ can provide $K \times n \times m$ samples of variants. It is equivalent to subjecting $K \times n$ original samples to m corruption variants. It can be seen that the original data are augmented with features interfering with corruption processing to form diverse variations. This strategy can improve the sample diversity as an augmentation on the one hand and reduce the variance on the other hand so that the system can learn the impact of feature interference on the data caused by multiple explicit information loss.

4.2 FEDERATED LEARNING WITH AUTOENCODER

For a single masked autoencoder, we model the corrupted inputs as $\tilde{D}_S = \{\tilde{x}_{S1}, \tilde{x}_{S2}, \dots, \tilde{x}_{Sn}\}$, where each \tilde{x}_i is corrupted by randomly masking patches of data elements in x_i and setting the masked region to 0 with probability $p \geq 0$.

Let \hat{x} be the reconstruction of x from the corrupted \tilde{x} .

$$\hat{x} = g(h(\tilde{x})) \quad (1)$$

We set the loss function to measure the residual between x and \hat{x} under corruption as

$$l(x, \hat{x}) = l(x, g(h(\tilde{x}))) \quad (2)$$

Then the goal of the training task of the autoencoder pair is to find the optimal solution of W to reconstruct the corrupted inputs \tilde{x} that minimizes the squared reconstruction loss l .

In order to obtain an optimal autoencoder, the main concern in our design is the relationship between the stable output of the system reconstruction and the loss function, so it is of interest to study the $h(x)$ and $g(z)$ function trajectories. When an image with all zero pixels enters the system, the output of the encoder is a zero tensor, i.e., $h(0) = 0$. Consequently, the output of the decoder is also a reconstruction image with all zero pixels, i.e., $g(0) = 0$. Therefore, there exists a zero reference point for discussing the encoder and decoder operators. Then, we have the following proposition.

proposition 1. *There exists a linear equivalent mapping with W_h to the approximate encoder $h(\cdot)$.*

Proof. Expanding the nonlinear vector function $h(x)$ into a Taylor series at 0, it yields

$$h(x) = h(0) + \nabla_x h(0)(x) + \epsilon \quad (3)$$

where $\nabla_x h(0)$ denotes the gradient of operator $h(\cdot)$ at x_0 in the direction of the vector x , and ϵ a higher order infinitesimal residual. When ignoring the residual and let $\nabla_x h(0)x = W_h x$, it yields

$$h(x) \approx W_h x + h(0) \quad (4)$$

As $h(0) = 0$, $h(x)$ can be represented by the mapping W_h , and the proof is obtained. \square

Corollary 1. *There exists a linear equivalent mapping with W_g to the approximate encoder $g(\cdot)$.*

remark 1. *A practical representation of the linear mapping W_h above is essentially an autoencoder implemented using a linear neural network. In practice, this linear representation may be relatively inferior because of the omission of higher-order feature representations and residual terms. In this paper, we mainly focus on the nature of the mapping and therefore use this mapping to simplify the problem, with the main aim of discussing the relationship between the sample, the loss function, and the optimal reconstruction in federated learning.*

For federated learning based on client k , this approximate linear equivalence mapping can help to simplify the model, then approximated encoder $h(\cdot)$ can be formulated as a linear mapping W_{h_k} . Similarly, the decoder $g(\cdot)$ can be equated to another linear mapping W_{g_k} , which yields

$$\hat{x} = g_k(h_k(\tilde{x})) \approx W_{g_k} W_{h_k} \tilde{x} = W_k \tilde{x} \quad (5)$$

then, the loss function can be formulated as

$$l(x, \hat{x}) = l(x, g_k(h_k(\tilde{x}))) = \frac{1}{2n} \sum_i^n \|x_i - W_k \tilde{x}_i\|^2 \quad (6)$$

Then we have the theorem below.

Theorem 1. *The approximate optimal solution of W_k on client k can be obtained by W_k^* , in which*

$$W_k^* = \bar{X}_k \tilde{X}_k^T (\tilde{X}_k \tilde{X}_k^T)^{-1} = W_{g_k}^* W_{h_k}^* \quad (7)$$

where

$$\begin{aligned} \tilde{X}_k &= [\tilde{X}_{k,1}, \dots, \tilde{X}_{k,j}, \dots, \tilde{X}_{k,m}]_{n \times m} \\ \bar{X}_k &= [X_k, \dots, X_k]_{n \times m} \end{aligned} \quad (8)$$

Proof. With m augmented corruptions, the aggregated input of client k can be formulated as

$$\tilde{X}_k = [\tilde{X}_{k,1}, \dots, \tilde{X}_{k,j}, \dots, \tilde{X}_{k,m}]_{n \times m} \quad (9)$$

which is corresponding to the aggregated original data as $\bar{X}_k = [X_k, \dots, X_k]_{n \times m}$. Let $\hat{X} = W_k \tilde{X}$ be the aggregated reconstruction of X . Then the transformed loss function of the training with the aggregation can be formulated as

$$l(\bar{X}_k, \hat{X}_k) = \frac{1}{2} \|\bar{X}_k - W_k \tilde{X}_k\|^2 = \frac{1}{2} \text{tr} \left[(\bar{X}_k - W_k \tilde{X}_k)^T (\bar{X}_k - W_k \tilde{X}_k) \right] \quad (10)$$

This loss function is a convex function that can reach a minimum value when its derivative is 0. Therefore, with $\nabla_{W_k} l(\bar{X}, \hat{X}) = 0$, it yields

$$\begin{aligned} 2\nabla_{W_k} l(\bar{X}_k, \hat{X}_k) &= \nabla_{W_k} \text{tr} \left[(\bar{X}_k - W_k \tilde{X}_k)^T (\bar{X}_k - W_k \tilde{X}_k) \right] \\ &= \nabla_{W_k} \text{tr} (\bar{X}_k^T \bar{X}_k - \tilde{X}_k^T W_k^T \bar{X}_k - \bar{X}_k^T W_k \tilde{X}_k + \tilde{X}_k^T W_k^T W_k \tilde{X}_k) \\ &= \nabla_{W_k} \text{tr} (W_k \tilde{X}_k \tilde{X}_k^T W_k^T) - 2\nabla_{W_k} \text{tr} (W_k \bar{X}_k \tilde{X}_k^T) \\ &= 2W_k \tilde{X}_k \tilde{X}_k^T - 2\bar{X}_k \tilde{X}_k^T = 0 \end{aligned} \quad (11)$$

Solving the equation 11 yields $W_k^* = \bar{X}_k \tilde{X}_k^T (\tilde{X}_k \tilde{X}_k^T)^{-1}$ as the completion of the proof. \square

Theorem 2. *The overall approximate optimal solution of W on K clients can be obtained by W_A^* , in which*

$$W_A^* = \bar{X} \tilde{X}^T (\tilde{X} \tilde{X}^T)^{-1} \quad (12)$$

where

$$\begin{aligned} \tilde{X} &= [\tilde{X}_{1,1}, \dots, \tilde{X}_{k,j}, \dots, \tilde{X}_{K,m}]_{K \times n \times m} \\ \bar{X} &= [X, \dots, X]_{K \times n \times m} \end{aligned} \quad (13)$$

Proof. The combination of data after aggregating all variation samples can be obtained by $\tilde{X} = [\tilde{X}_{1,1}, \dots, \tilde{X}_{k,j}, \dots, \tilde{X}_{K,m}]_{K \times n \times m}$, corresponding to the same scale of the original sample data aggregation as a whole can be expressed as $\bar{X} = [X, \dots, X]_{K \times n \times m}$. With the modeling and analytic deduction similar to equation 10 and equation 11, the optimal solution can be obtained. \square

proposition 2. *When random sampling is used to corrupt the data X , the performance of the established optimal solution is proportional to the number of clients.*

remark 2. *As chunking and random sampling are applied to the raw data, each client enters each training step with a single random sample of the raw data it carries, resulting in a different variant of the data. If the sampling is different each time, then the original sample can thus be augmented to produce genes to cover different variant versions of the data. The corresponding \tilde{X} is used to solve the optimal reconstruction mapping and the corresponding \tilde{X} are then enriched with elements, resulting in mapping solutions that better cover the different inputs.*

proposition 3. *The number of variant m has its analytic upper bound as*

$$\sup m = n \times \binom{B}{b} = n \frac{B!}{b!(B-b)!} \quad (14)$$

where B refers to the number of all patches in $\{x_i\}_n$ after chunking segmentation, b the number of selected patches for encoding, and $\binom{B}{b}$ is a combination with a selection of patches from B distinct patches.

remark 3. *Although the number of versions of the variation m is limited, disordering all the versions at each step of the training can still enhance the performance of the system learning as dissimilar patches can be rearranged in different orders to form different versions of permutation to augment the data for training in each epoch.*

5 EXPERIMENTS

In this section, we demonstrate our experimental results. Section 5.1 introduces our experimental setup. Section 5.2 compares accuracies, model parameters and GFLOPs of different FSSL methods with our FedMAE. In Section 5.3, we show the scalability of the amount of data on the server for FedMAE. Moreover, we demonstrate the impact of the statistical heterogeneity and the number of total clients for both FedMAE and other FSSL methods in Appendix A.1 and A.2. Additionally, we present various ablation studies in Appendix A.3 to A.6 on many parameters (including the number of clients, the number of local epochs, the participation ratio, the number of training rounds, and the depth of the pre-trained backbone) and illustrate how they influence the training results of FedMAE. We compare the performance on mask ratio and downstream model depth in many evaluation metrics in Appendix A.7. We also provide the details about our experiment settings in A.8, and upload the experiment code as supplementary materials.

5.1 EXPERIMENTAL SETUP

Datasets We conduct our experiments using pre-training datasets and downstream datasets. For pre-training datasets, we use ImageNet (Deng et al., 2009) with 1000 classes and Mini-ImageNet with 100 classes, and both of their input size is 224x224. The total number of mini-ImageNet is 60000, which are chosen from ImageNet by the method of (Vinyals et al., 2016), while the ImageNet has 1,281,167 training images and 50000 test images. For downstream datasets, we use four public datasets, including CIFAR10, CIFAR100 (Krizhevsky et al., 2009), ROAD SIGN (Kulakin et al., 2022) and Mini-INAT2021. There are 60,000 32x32 color images in CIFAR10 and CIFAR100, where the training set has 50,000 images, and the test set has 10,000 images. CIFAR10 is used for the task of 10-class image classification, while CIFAR100 is utilized for the 100-class image classification task. On the other hand, ROAD SIGN and Mini-INAT2021 are two datasets containing large-scale images. ROAD SIGN contains 56000 images for training and 24000 images for testing. The size of each image is 224x224. There are 8 classes in ROAD SIGN, corresponding to the categories of signs in Russian traffic. Mini-INAT2021 is the more accessible version of INaturalist Dataset 2021. Compared to the original dataset which contains nearly 2.7M training images for 10000 natural species, the number of images in Mini-INAT2021 has been cut down to 50 training samples and 10 validation samples per specie for a total of 500K training images and 100K validation images. The maximum dimension of these images is 800 pixels.

Setup Our experiments consist of two phases: federated pre-training and downstream fine-tuning. In the pre-training phase, we use the Mini-ImageNet dataset if our results need to be compared with the results of other baselines. Otherwise, we use the ImageNet dataset to evaluate the effect

of various parameters on the performance of our method. The pre-training dataset is divided into K clients. When the data held by each client is assumed to be IID distributed among each other, the division of the dataset requires that each client holds images of all categories. When the data division is assumed to follow non-IID distribution, we divide the dataset by taking a sample of the Dirichlet distribution’s class priors (Hsu et al., 2019). More heterogeneous division can be made by specifying smaller Dirichlet parameter α during sampling. Unless otherwise stated, the federated learning conducted in the pre-training phase adopts the following settings. The number of training rounds is set to 200, the number of local epochs is set to 10, the total number of local clients is set to 100, and the participation ratio is set to 0.05. In the downstream fine-tuning phase, we replace the decoder part of the pre-trained backbone with the head layer specified by the downstream task and retrain the weights and biases of all layers. For all downstream tasks, the model is trained for 100 epochs at a batch size of 64. In addition, for the supervised evaluation, we fine-tuned the new model using 100% of the labeled data. For the semi-supervised evaluation, we use 10% of the labeled data for fine-tuning.

Baselines To fairly evaluate the proposed FedMAE framework, we compare it with the following state-of-the-art FSSL benchmarks. **1) Fed-SimSiam**: a naive combination of SimSiam and FL. **2) Fed-SimCLR**: a naive combination of SimCLR and FL. **3) FedU (Zhuang et al., 2021a)**: Using the divergence-aware predictor updating with self-supervised BYOL (Grill et al., 2020). **4) FedEMA (Zhuang et al., 2021b)**: Using EMA of global model to adaptively update online client networks. **5) Orchestra (Lubana et al., 2022)**: a novel clustering-based FSSL technique. For downstream tasks’ evaluation, we fine-tune the entire model using downstream labeled data (10% or 100% of the total dataset).

Table 1: Accuracy (%) in IID ($\alpha = 0$) and Non-IID ($\alpha = 1e - 1$) of FL settings on CIFAR and ROAD SIGN datasets. We use Mini-ImageNet for pre-training, and evaluate models using the popular semi-supervised fine-tuning with 10% labeled data and fully-supervised fine-tuning. For pre-training period, the local clients are 100, the training rounds of pre-training are 200, the local training epochs are 10 in each training round. For the downstream task, the total downstream training epochs are 100.

Method	CIFAR10				CIFAR100				ROAD SIGN			
	IID		Non-IID		IID		Non-IID		IID		Non-IID	
	10%	100%	10%	100%	10%	100%	10%	100%	10%	100%	10%	100%
Fed-SimSiam(%)	53.00	56.55	50.36	50.68	24.80	35.20	21.54	28.76	36.42	63.43	30.97	61.38
Fed-SimCLR(%)	60.47	64.05	58.46	61.28	29.57	41.09	27.60	38.99	46.48	79.15	45.54	76.03
FedU(%)	31.69	52.85	31.46	50.97	6.89	26.15	7.50	24.51	16.60	66.94	11.79	66.56
FedEMA(%)	32.14	55.01	33.10	53.84	8.39	28.18	7.34	28.36	11.36	68.49	16.39	69.20
Orchestra(%)	61.05	63.47	60.21	63.76	29.76	38.82	29.05	39.89	51.83	70.56	47.71	71.71
FedMAE(%)	67.35	90.92	67.24	90.57	34.56	73.33	33.97	74.11	98.24	100.00	98.19	100.00

Table 2: Comparison of FedMAE’s Model Parameters and GFLOPs with FSSL methods. We compute the model parameters and GFLOPs using images with the size of 224x224 as input.

	Fed-SimSiam	Fed-SimCLR	FedU	FedEMA	Orchestra	FedMAE (Pre-train)	FedMAE (Downstream)
Model Params	12.03M	11.70M	38.47M	38.47M	11.84M	11.62M	39.97M
GFLOPs	3.65	3.65	7.40	7.40	7.31	1.23	7.39

5.2 COMPARISON WITH FSSL METHODS

We first assess the accuracies, model parameters and GFLOPs of different FSSL methods using fine-tune protocol and transfer learning evaluation. In this experiment, we examine FedMAE’s performance in IID and Non-IID scenarios since cross-device FL is our key purpose for developing FedMAE. The pre-training dataset used in FL is Mini-ImageNet, and the downstream tasks’ datasets are CIFAR10, CIFAR100 and ROAD SIGN. Table 1 shows our FedMAE outperforms its competitors in accuracy results for three public datasets. Since both CIFAR10 and CIFAR100 are resized to 224x224 size to fit our models, we choose simple but large images of RODA SIGN, which are real 224x224. FedMAE clearly has a huge advantage of 30% to 50% improvement over the existing FSSL methods. Table 2 compares the differences between FedMAE and other FSSL methods

in terms of model parameters and GFLOPs. The pre-training model of FedMAE is not only the smallest in parameters but also the smallest in GFLOPs, which means faster transmission between local clients and the server, and fewer computing power requirements for local clients. Although the downstream model of FedMAE has a large number of parameters due to the cascade design, the computation resources required for its training and subsequent inference are acceptable among all models at present, which can be deployed for inference in local clients.

Table 3: Accuracy (%) with the amount of data on the server. The pre-training dataset is Mini-ImageNet.

	CIFAR10				CIFAR100				Mini-INAT			
Amount of data on Server	0	1000	5000	10000	0	1000	5000	10000	0	1000	5000	10000
FedMAE (%)	90.92	91.70	92.40	92.96	73.93	74.23	75.16	76.17	46.01	46.33	48.83	49.01

5.3 SCALABILITY WITH THE AMOUNT OF DATA ON THE SERVER

Since what one-block MAE can learn is very limited, we share a small amount of data on the server to help build a better downstream model. We tested the accuracies of downstream datasets when sharing 0, 1000, 5000 and 10000 unlabeled data on the server. In this scenario, we first use clients' data for federated training, after 200 training rounds, we cascade 5 one-block pre-trained MAE to build a five-block MAE. Afterward, sharing data on the server will be fed into the five-block MAE to train a better downstream model for downstream tasks. The training epochs here for sharing data are 100. Results are shown in Table 3. Although the amount of data on the server is very limited, FedMAE provides a good initialized model that can further build a better downstream model.

Table 4: Comparison on different model depth between un-pretrained model and FedMAE with pre-trained model cascading. The pre-training dataset is Mini-ImageNet.

Model Depth		1	2	3	4	5	6	7	8	9	10
CIFAR10	Un-pretrained	66.49	71.86	75.97	78.81	80.19	81.73	82.29	82.87	82.83	83.09
	Pre-trained	72.21	80.48	85.84	88.51	90.02	91.08	91.30	91.55	91.61	91.74
CIFAR100	Un-pretrained	40.41	49.76	55.49	59.18	60.76	62.20	63.17	63.99	63.89	63.91
	Pre-trained	47.03	60.26	66.93	71.43	73.55	73.46	73.34	73.95	73.96	74.00

5.4 ABLATION STUDY

We further conducted additional experiments on the depth of the downstream model. Since the pre-trained model is an MAE with only one block, and its learning ability to extract features is very limited, so we want to know whether the pre-trained one-block encoder has a positive effect on the deeper blocks of a multi-block encoder. Therefore, we cascade different numbers of pre-trained models to compare with models of the same depth but without pre-training. The results are shown in Table 4, which indicate that the pre-trained model still has a great influence on the overall downstream model as the depth increases. Although the increase in model depth would have improved the learning ability of the model, the result is obvious that the one-block pre-trained model improves the learning ability of each block of the downstream model with multi-blocks.

6 CONCLUSION

We introduce a novel framework FedMAE, addressing the training of large-scale distributed unlabeled datasets in federated learning. Our pre-training model only contains a one-block encoder and a one-block decoder and uses a masking strategy for training data, so it is easy to deploy in various lightweight clients. Pre-training models in FL can be trained on different clients asynchronously without average aggregation. Pre-trained models are cascaded in the server to build a multi-block downstream model. Theoretical analysis and experimental results on image reconstruction and classification show that our FedMAE achieves a considerable boost compared to previous FSSL methods.

REFERENCES

- Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. In *International Conference on Learning Representations*, 2021.
- Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *International conference on machine learning*, pp. 1691–1703. PMLR, 2020a.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020b.
- Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15750–15758, 2021.
- Yang Chen, Xiaoyan Sun, and Yaochu Jin. Communication-efficient federated deep learning with layerwise asynchronous model update and temporally weighted aggregation. *IEEE transactions on neural networks and learning systems*, 31(10):4229–4238, 2019.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*, 2020.
- Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pp. 702–703, 2020.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pp. 1422–1430, 2015.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018.
- Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16000–16009, 2022.
- 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.

- Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pp. 646–661. Springer, 2016.
- Wonyong Jeong, Jaehong Yoon, Eunho Yang, and Sung Ju Hwang. Federated semi-supervised learning with inter-client consistency & disjoint learning. In *International Conference on Learning Representations*, 2020.
- Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- S. Kulakin, M. Korolev, R. Khabirov, N. Pal, and E. Bodul. Road signs categories classification dataset, 2022. URL <https://www.kaggle.com/datasets/sergeykulakin/russian-road-signs-categories-dataset>.
- 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.
- Xiaoxiao Liang, Yiqun Lin, Huazhu Fu, Lei Zhu, and Xiaomeng Li. Rscfed: Random sampling consensus federated semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10154–10163, 2022.
- Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Ekdeep Singh Lubana, Chi Ian Tang, Fahim Kawsar, Robert P. Dick, and Akhil Mathur. Orchestra: Unsupervised federated learning via globally consistent clustering. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 14461–14484. PMLR, 2022. URL <https://proceedings.mlr.press/v162/lubana22a.html>.
- 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*, pp. 1273–1282. PMLR, 2017.
- Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pp. 69–84. Springer, 2016.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016.
- Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14668–14678, 2022.

- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3733–3742, 2018.
- Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 6023–6032, 2019.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pp. 649–666. Springer, 2016.
- Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.
- Weiming Zhuang, Xin Gan, Yonggang Wen, Shuai Zhang, and Shuai Yi. Collaborative unsupervised visual representation learning from decentralized data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4912–4921, 2021a.
- Weiming Zhuang, Yonggang Wen, and Shuai Zhang. Divergence-aware federated self-supervised learning. In *International Conference on Learning Representations*, 2021b.

A APPENDIX

A.1 STATISTICAL HETEROGENEITY.

We use 100 clients and analyze three levels of heterogeneity by setting α to 0 (none/IID), 1e-1 (moderate), and 1e-3 (high). Results are shown in Figure 3. The data distribution in federated learning is dispersed and most current FSSL methods are based on contrastive learning, so SSL approaches tend to be sensitive to heterogeneity and therefore adopt certain strategies for this challenge. In contrast, FedMAE is robust and its performance generally does not degrade with increasing heterogeneity. In fact, the pre-trained model of FedMAE is an image reconstruction method that uses autoencoder, which is not sensitive to image classification, FedMAE learns how to reconstruct the image itself rather than the class features of images in the period of pre-training.

A.2 NUMBER OF CLIENTS.

Changing the number of clients also affects the performance of the pre-trained model in FL. We increased the number of local clients from 100 to 400 with IID data distribution, still selecting 5 clients with 10 epochs in each training round, and the total number of training rounds are 200. Results are provided in Figure 4. We find that, unlike other methods which significantly degrade performance as the number of clients increases, FedMAE achieves similar performance or a gentle decline, which indicates FedMAE can still learn a lot of useful knowledge on relatively dispersed and small amounts of data.

Table 5: Robustness to local epochs on CIFAR10, CIFAR100 and ROAD SIGN. The pre-training dataset is ImageNet. We find that FedMAE achieves the best accuracy when local epochs are 10.

Number of Local Epochs in each training round	1	5	10	15	20
CIFAR10 (%)	90.57	93.25	93.68	92.73	92.71
CIFAR100 (%)	73.41	76.75	77.07	75.74	75.78

A.3 LOCAL EPOCHS

The number of local epochs represents the number of times all data from each client is trained in each training round. According to Table 5, we find that increasing this federated learning attribute

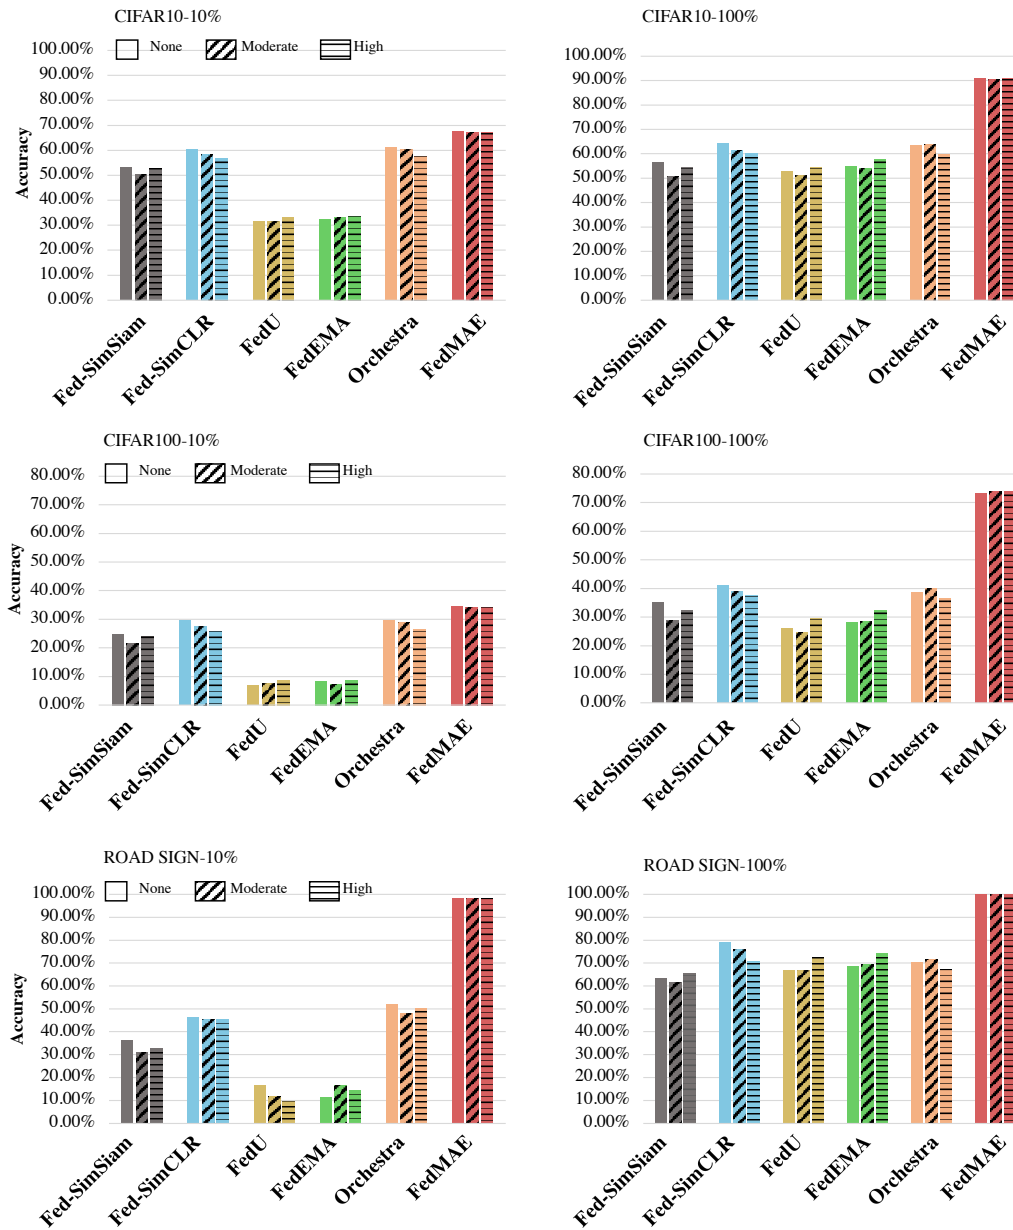


Figure 3: Sensitivity of Statistical Heterogeneity on CIFAR10, CIFAR100 and ROAD SIGN. The pre-training dataset is Mini-ImageNet. FedMAE does not lose performance with increased heterogeneity.

can lead to an improvement in fine-tuning accuracy and the best result is achieved by setting the number of local epochs to 10. However, local training may be limited to a small number of epochs because of the resource limitations on clients. If we need to take this into account, it might be more appropriate to set the number of local epochs to 5, as the improvement from increasing the number of local epochs from 5 to 10 is very small.

A.4 PARTICIPATION RATIO

The participation ratio indicates the number of clients connecting to the server in each training round. In this experiment, our method follows the aggregation strategy of other FSSL methods to average




Figure 4: Scalability with the number of clients on CIFAR10, CIFAR100 and ROAD SIGN. The pre-training dataset is Mini-ImageNet. FedMAE outperforms other FSSL methods and is generally robust to a number of clients in the federation.

Table 6: Sensitivity to participation ratio on CIFAR10 and CIFAR100. The pre-training dataset is ImageNet. While the accuracies of clients decrease at smaller participation ratios normally, we find FedMAE is almost no degradation.

Number of Clients in each training round	5	10	15	20	25
CIFAR10 (%)	91.77	91.57	91.98	91.56	91.66
CIFAR100 (%)	75.68	75.32	75.77	75.75	75.02

the client models trained in each round, while Table 6 shows that simply increasing this federated learning attribute does not produce a boost in fine-tuning accuracy as in other FSSL methods which uses CNN architecture. The reason why this attribute has no effect to the results is that the pre-train model is very short which only has one block autoencoder, and the learning rate is very small. If we increase the depth of the pre-trained backbone and the participation ratio, the fine-tuning accuracy may also be improved.

A.5 NUMBER OF TRAINING ROUNDS

According to Table 7, the fine-tuning accuracy for the downstream tasks increases as the number of training rounds increases in pre-training. This increase trend is more pronounced in the first 1200 rounds, while eases off in the later rounds. Therefore, for our method FedMAE, we consider a good Nash equilibrium between training cost and fine-tuning accuracy is 1200 training rounds.

Table 7: Accuracy with the number of training rounds for pre-training on CIFAR10 and CIFAR100. The pre-training dataset is ImageNet. The accuracy will improve with the number of training rounds and finally reach a stable value.

Number of Training Rounds	200	400	600	800	1000
CIFAR10 (%)	91.44	92.43	92.49	92.17	92.76
CIFAR100 (%)	75.11	75.70	76.07	76.19	76.01
Number of Training Rounds	1200	1400	1600	1800	2000
CIFAR10 (%)	92.98	92.83	92.99	93.18	93.02
CIFAR100 (%)	76.42	76.26	76.20	76.46	76.48

Table 8: The influence of the pre-trained model depth (i.e. the number of pre-trained Transformer blocks in downstream model) on fine-tuning accuracy. The pre-training dataset is Mini-ImageNet.

Model Depth	Downstream Model Blocks					Accuracy (%)	
	Block 1	Block 2	Block 3	Block 4	Block 5	CIFAR10	CIFAR100
1	Un-pretrained	-	-	-	-	66.49	40.41
	Pre-trained	-	-	-	-	72.21	47.03
2	Un-pretrained	Un-pretrained	-	-	-	71.86	49.76
	Pre-trained	Un-pretrained	-	-	-	77.28	54.90
	Pre-trained	Pre-trained	-	-	-	80.48	60.26
3	Un-pretrained	Un-pretrained	Un-pretrained	-	-	75.97	55.49
	Pre-trained	Un-pretrained	Un-pretrained	-	-	79.97	61.04
	Pre-trained	Pre-trained	Un-pretrained	-	-	83.45	65.26
	Pre-trained	Pre-trained	Pre-trained	-	-	85.84	66.93
4	Un-pretrained	Un-pretrained	Un-pretrained	Un-pretrained	-	78.81	59.18
	Pre-trained	Un-pretrained	Un-pretrained	Un-pretrained	-	82.79	64.08
	Pre-trained	Pre-trained	Un-pretrained	Un-pretrained	-	85.09	67.38
	Pre-trained	Pre-trained	Pre-trained	Un-pretrained	-	87.67	70.19
	Pre-trained	Pre-trained	Pre-trained	Pre-trained	-	88.51	71.43
5	Un-pretrained	Un-pretrained	Un-pretrained	Un-pretrained	Un-pretrained	80.19	60.76
	Pre-trained	Un-pretrained	Un-pretrained	Un-pretrained	Un-pretrained	84.61	66.04
	Pre-trained	Pre-trained	Un-pretrained	Un-pretrained	Un-pretrained	85.75	68.70
	Pre-trained	Pre-trained	Pre-trained	Un-pretrained	Un-pretrained	87.29	71.31
	Pre-trained	Pre-trained	Pre-trained	Pre-trained	Un-pretrained	89.25	73.00
	Pre-trained	Pre-trained	Pre-trained	Pre-trained	Pre-trained	90.02	73.55

A.6 DEPTH OF PRE-TRAINED BACKBONE

The results of the previous ablation study show that the fine-tuning accuracy is influenced by the depth of the downstream model, and our method FedMAE still has an effect on the downstream model as the depth increases. Based on this finding, we continue to test whether changing the number of pre-trained blocks in the downstream model could also affect the fine-tuning accuracy for downstream tasks. The model depth in this session is from one to five. The experimental results in Table 8 demonstrate that every pre-trained Transformer block in the cascaded encoder contributes to the improvement of fine-tuning accuracy. As the depth of the model increases, the more pre-trained blocks are cascaded in downstream models, the higher fine-tuning accuracy the downstream model has.

Table 9: The impact of special designs in FedMAE on the computational speed and the fine-tuning (FT) accuracy. The pre-training dataset is Mini-ImageNet, and the downstream dataset is Mini-INat2021.

Participation Ratio	Mask Ratio	Pre-trained Model Depth	Downstream Model Depth	Pre-trained Model Params	Pre-train GFLOPs	FT Acc	Pre-train Time	Speed Up
1/100	0.75	5 blocks	5 blocks	39.97M	2.67	53.45%	8h	1x
1/100	0.75	1 block	5 blocks	11.62M	1.23	48.48%	7.6h	1.05x
5/100	-	1 block	5 blocks	11.62M	2.37	48.02%	43.2h	1x
5/100	0.75	1 block	5 blocks	11.62M	1.23	48.32%	38.5h	1.12x

A.7 SPEEDUP ANALYSIS

The experimental results listed previously show that our method FedMAE not only outperforms other FSSL methods in terms of fine-tuning accuracy but also in terms of computational speed. We also conduct additional experiments to analyze which of FedMAE’s designs achieves speedup and how much. The analysis results are demonstrated in Table 9. On the one hand, the results show that pre-training one-block MAE instead of the five-block MAE on each client is an effective design to achieve speedup. Although the fine-tuning accuracy becomes worse, the total pre-train time is decreased, and the computation resources are significantly reduced. On the other hand, the masking strategy applied in MAE also leads to a boost in computational speed. The speedup caused by masking is more noticeable and does not affect the fine-tuning accuracy.

Table 10: Environment Setting.

Config	Details
Server GPU Count	8
Server GPU Type	RTX A5000 (24GB)
Server CPU Type	AMD EPYC 7513 32-core
CUDA	11.3
Framework	PyTorch

Table 11: Federated Pre-training Settings. The first sub-table shows the default setting for federated learning, and the second sub-table shows the default setting for local training on clients.

FL Config	Details
communication rounds	200
local epoch	10
total clients	100
participant ratio	0.05

Training Config	Details
batch size	256 / Mini, 512 / ImageNet
base learning rate	1.5e-4
learning rate schedule	cosine decay (Loshchilov & Hutter, 2016)
optimizer	AdamW (Loshchilov & Hutter, 2017)
optimizer momentum	$\beta_1 = 0.9, \beta_2 = 0.999$ (Chen et al., 2020a)
weight decay	0.05
augmentation	RandomResizedCrop

A.8 EXPERIMENT DETAILS

Model Architecture The architecture of the model trained on the clients follows the MAE architecture in the original paper (He et al., 2022). The model consists of an encoder and a decoder, and uses a linear projection layer in the middle to connect the two parts. However, unlike the original MAE, the encoder and decoder depths in our client model are the same and minimal, i.e. there is only one Transformer block for each part (Vaswani et al., 2017). On the other hand, the final pre-trained model used for various downstream tasks follows the typical ViT architecture (Dosovitskiy et al., 2020). We cascade multiple pre-trained model for downstream training but without masking training images.

Environment Configuration All the experiments were conducted in the same experimental environment. The specific environment configurations are given in Table 10.

Federated Pre-Training Table 11 shows the default settings of the federated pre-training phase. For the FL side, we adopt this set of settings that are used generically in many FSSL papers. For the training side, we do not change the default MAE pre-training setting specified in the original paper, as we find that this set of settings does a good enough job of training our client model as well. Besides, we use different batch sizes when the dataset varies.

Table 12: Fine-tuning Setting.

Config	Details
batch size	64
base learning rate	1e-3
training epochs	100
learning rate schedule	cosine decay
warmup epochs (Goyal et al., 2017)	5
optimizer	AdamW
optimizer momentum	$\beta_1 = 0.9, \beta_2 = 0.999$
weight decay	0.05
layer-wise lr decay	0.75 (Clark et al., 2020)
augmentation	RandAug (9, 0.5) (Cubuk et al., 2020)
mixup (Zhang et al., 2017)	0.8
cutmix (Yun et al., 2019)	1.0
drop path (Huang et al., 2016)	0.1
label smoothing (Szegedy et al., 2016)	0.1

Fine-Tuning The default settings of the fine-tuning phase are in Table 12. Similar to the federated pre-training phase, we also follow the default settings in the original paper to fine-tune our model. However, since the size of benchmark datasets for the downstream task is not large, we use a small batch size for the best fine-tuning result.