

Feature-Aware One-Shot Federated Learning via Hierarchical Token Sequences

Shudong Liu¹, Hanwen Zhang¹, Xiuling Wang², Yuesheng Zhu¹, Guibo Luo^{1*}

¹Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology,
Shenzhen Graduate School, Peking University

²Hong Kong Baptist University

{shooo, hanwen}@stu.pku.edu.cn, {zhuys, luogb}@pku.edu.cn, xiulingwang@hkbu.edu.hk

Abstract

One-shot federated learning (OSFL) reduces the communication cost and privacy risks of iterative federated learning by constructing a global model with a single round of communication. However, most existing methods struggle to achieve robust performance on real-world domains such as medical imaging, or are inefficient when handling non-IID (Independent and Identically Distributed) data. To address these limitations, we introduce FALCON, a framework that enhances the effectiveness of OSFL over non-IID image data. The core idea of FALCON is to leverage the feature-aware hierarchical token sequences generation and knowledge distillation into OSFL. First, each client leverages a pretrained visual encoder with hierarchical scale encoding to compress images into hierarchical token sequences, which capture multi-scale semantics. Second, a multi-scale autoregressive transformer generator is used to model the distribution of these token sequences and generate the synthetic sequences. Third, clients upload the synthetic sequences along with the local classifier trained on the real token sequences to the server. Finally, the server incorporates knowledge distillation into global training to reduce reliance on precise distribution modeling. Experiments on medical and natural image datasets validate the effectiveness of FALCON in diverse non-IID scenarios, outperforming the best OSFL baselines by 9.58% in average accuracy.

Code — <https://github.com/LMIAPC/FALCON>

1 Introduction

Federated Learning (FL) is a distributed collaborative training paradigm that enables the integration of data from multiple parties without exposing raw data (McMahan et al. 2017). However, conventional iterative FL frameworks require multiple rounds of parameter synchronization between clients and the central server, leading to significant communication overhead. Moreover, frequent interactions increase the risk of privacy leakage (Kairouz et al. 2021) and make the system more vulnerable to poisoning (Yazdinejad et al. 2024; Rao et al. 2025) or inference attacks (Nasr, Shokri, and Houmansadr 2019; Lyu, Yu, and Yang 2020). To alleviate these limitations, one-shot federated learning (OSFL) frameworks have been proposed (Guha, Talwalkar,

and Smith 2019), where each client uploads its local information only once, substantially reducing communication cost and mitigating privacy risks.

While OSFL models show promise in FL, their practical deployment is hindered by non-IID data distributions across clients, while insufficient efficiency in global modeling limits large-scale deployment.

Existing OSFL methods can be broadly classified into two categories: those that upload local model parameters and those that transmit surrogate data in place of original data to facilitate global model construction.

Among parameter-uploading methods, some studies focus on designing weight aggregation strategies (Tang et al. 2024; Jhunjhunwala, Wang, and Joshi 2024). However, these methods typically assume strong homogeneity among local models; the distributional shifts often lead to divergence in local models, resulting in suboptimal global performance. Bayesian ensemble approaches (Chen and Chao 2021; Hasan et al. 2024; Allouah et al. 2024) improve robustness to non-IID data by modeling and fusing the local posteriors. However, they typically incur high inference costs and require access to public data for posterior estimation. Knowledge distillation-based methods (Lin et al. 2020; Zhang et al. 2022; Kang et al. 2023; Dai et al. 2024) transfer knowledge from multiple local models to a global model using auxiliary data, such as public datasets or synthetic samples. Although they relax model homogeneity constraints, their effectiveness depends heavily on high-quality auxiliary data, limiting practical deployment and overall performance.

Surrogate data uploading methods aim to mitigate the performance degradation caused by non-IID client distributions. Dataset distillation approaches (Zhou et al. 2020; Song et al. 2023; Zhang, Liu, and Wang 2024) transmit a small amount of distilled data, which often fails to represent the full complexity of the original data distribution. Recently, distribution-estimation-based surrogate data methods have demonstrated advantages in both performance and efficiency. These methods primarily fall into two categories: (i) image-level generative-model-based approaches, and (ii) feature distribution statistical modeling approaches.

Image-level generative-model-based approaches (Kasturi and Hota 2023; Heinbaugh, Luz-Ricca, and Shao 2023; Zhang, Qi, and Zhao 2023; Yang et al. 2024; Chen et al. 2025; Yang et al. 2025) aim to reconstruct the client-side

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

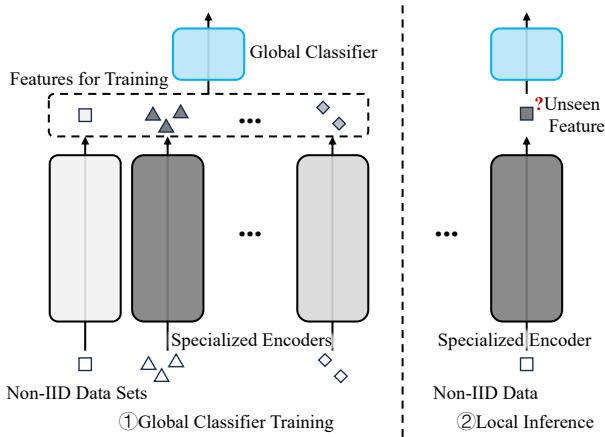


Figure 1: Illustration of the impact of locally specialized encoders on generalization. Colors indicate different locally specialized encoders; shapes indicate samples from non-IID datasets. (1) Global classifier training: Features from locally specialized encoders are aggregated to train a global classifier. (2) Local inference: When a client’s specialized encoder encounters data that is non-IID with respect to its local data, the extracted features cannot be correctly recognized by the global classifier.

data distribution on the server using GANs (Goodfellow et al. 2020), VAEs (Kingma, Welling et al. 2013), or diffusion models (Ho, Jain, and Abbeel 2020; Rombach et al. 2022), thereby improving the global model’s adaptability to non-IID data. However, GANs are prone to mode collapse and unstable training (Arjovsky, Chintala, and Bottou 2017), while VAEs often produce blurry samples. Pretrained diffusion models struggle to produce usable samples in real-world application domains such as medical imaging. Moreover, the reliance on iterative denoising steps is a key bottleneck: even with fast solvers (Song, Meng, and Ermon 2020; Liu et al. 2022), generating training samples on the server remains time-consuming, hindering practical deployment.

Feature distribution statistical modeling approaches use a unified encoder to extract semantic features from raw data, then perform statistical modeling on the features. By uploading the distribution statistics, the server can reconstruct a global feature space. Specifically, FedPFT (Beitollahi et al. 2025) models features as a Gaussian mixture, uploading estimated statistics for pseudo-sample generation. FedCGS (Guan, Zhou, and Gu 2025) aggregates global feature statistics and directly builds a Gaussian naive Bayes classifier without further training. However, features extracted by directly applying a pretrained model often perform suboptimally in real-world tasks with distribution shifts from the encoder’s pretraining domain. While local fine-tuning is a natural way to obtain a specialized encoder, it is not feasible in OSFL. Figure 1 illustrates a simplified semantic feature transmission OSFL scenario. Local specialization leads to misaligned feature spaces across clients, resulting in personalized local models rather than a true global model—thus failing to meet the generalization requirement of FL. More-

over, most existing pretrained visual encoders based on Vision Transformers (Vaswani et al. 2017; Dosovitskiy et al. 2021) accept fixed-size low-resolution inputs, which limits their ability to capture fine-grained regional details and local structures, particularly critical in high-resolution tasks such as medical imaging. Thus, with the restriction that the encoder cannot be locally fine-tuned, a key challenge is enhancing the adaptability of pretrained encoders to multi-domain data while efficiently extracting multi-scale and high-resolution information from raw data. Finally, the effectiveness of distribution-estimation-based methods is constrained by the challenge of accurately replicating client data distributions, which is rarely attainable in practice.

To overcome the above limitations, we propose FALCON, a one-shot federated learning framework designed for diverse application domains and multi-type non-IID scenarios. Specifically, FALCON employs a pretrained visual encoder on each client and introduces the proposed hierarchical scale encoding (HSE) to simultaneously extract both global semantic features and high-resolution regional features from images. This process constructs hierarchical token sequences with multi-scale semantics. Each client subsequently trains a local classifier and a multi-scale autoregressive (M-AR) transformer generator to capture the distribution of token sequences and generate the synthetic token sequences. The server receives the synthetic sequences and local classifiers from clients and applies knowledge distillation to complement generative modeling and enhance global training. Our contributions are as follows:

- We propose FALCON, a novel OSFL framework that achieves efficient global modeling on non-IID data by uploading designed compressed surrogate data and leveraging distillation-guided optimization.
- We propose HSE to extract multi-scale semantics from raw images for better domain adaptability, and design an M-AR generator to model hierarchical token sequences while preserving spatial relationships.
- We incorporate a distillation-guided global training mechanism that reduces the reliance on synthetic sequence quality and facilitates robust global modeling.
- Experiments on natural and medical image datasets under non-IID settings demonstrate the effectiveness, efficiency, and privacy-preserving capability of FALCON.

2 Method

2.1 Problem Formulation

We consider a federated learning (FL) setting consisting of a central server and N distributed clients. Each client n holds a private local dataset $\mathcal{D}_n = \{(x_i, y_i)\}_{i=1}^{|\mathcal{D}_n|}$. The server aims to train a global model $w(\cdot)$ based on information uploaded by the clients, without accessing the original data, so as to minimize the sum of empirical risks across all clients. The FL optimization objective can be formulated as:

$$\min_w \mathcal{L}(w) = \frac{1}{\sum_{j=1}^N |\mathcal{D}_j|} \sum_{n=1}^N \sum_{(x,y) \in \mathcal{D}_n} [\ell(w(x), y)], \quad (1)$$

where $\ell(\cdot, \cdot)$ denotes the task-specific loss function.

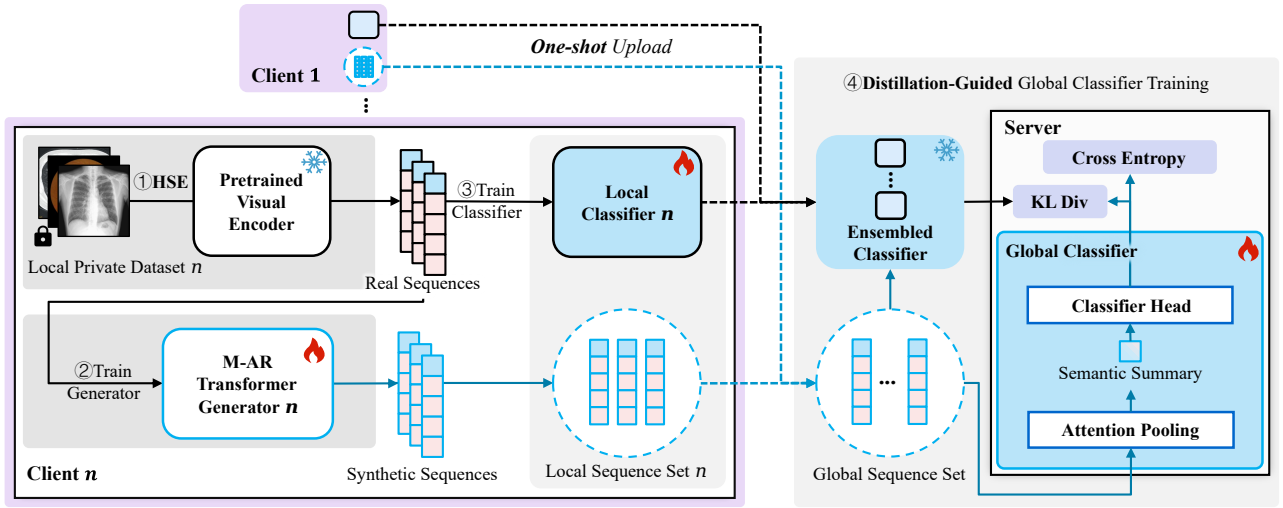


Figure 2: The framework of FALCON. Each client (1) performs hierarchical scale encoding (HSE), (2) trains a local M-AR transformer generator (3) and classifier, and uploads synthetic token sequences and the classifier to the server for (4) distillation-guided global classifier training.

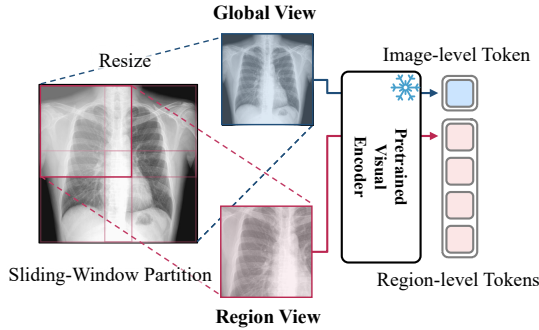


Figure 3: Illustration of hierarchical scale encoding (HSE).

2.2 FALCON Framework

FALCON enables one-shot federated learning through the transmission of compact, structured token sequences and distillation-guided training of client knowledge. As shown in Figure 2, FALCON proceeds in four steps: (1) Each client compresses its raw images into hierarchical token sequences via hierarchical scale encoding (HSE); (2) The client trains a multi-scale autoregressive (M-AR) generator on these sequences; (3) It also trains a local classifier. Both the synthetic token sequences (sampled from the generator) and the classifier are then uploaded to the server; (4) Finally, the server performs distillation-guided training using the uploaded sequences and local classifiers to optimize the global model.

2.3 Hierarchical Scale Encoding (HSE)

To capture both global discriminative information and high-resolution local details without locally fine-tuning the pre-trained visual encoder, FALCON draws inspiration from the human visual system, which first perceives the global scene and then attends to local regions. Accordingly, we design hierarchical scale encoding (HSE). As shown in Figure 3, we employ two complementary encoding paths, sending the

original image x into both a low-resolution global view and high-resolution sliding window regional views.

(1) **Global View.** The image x is resized into a global view $x^{(0)}$, which is then encoded to obtain an image-level token capturing coarse-grained semantics. We denote the process as:

$$s^{(0)} = f(x^{(0)}).$$

(2) **Sliding-window Region View.** At the original resolution, the image x is partitioned into L high-resolution patches $\{x^{(l)}\}_{l=1}^L$ via a sliding window with fixed size and stride. Each patch is independently encoded into region-level tokens capturing fine-grained semantics. We denote the process as:

$$s^{(l)} = f(x^{(l)}), \quad l = 1, \dots, L.$$

The above $1 + L$ tokens are organized into a hierarchical token sequence in a coarse-to-fine order, denoted as:

$$s = \{s^{(0)}, s^{(1)}, \dots, s^{(L)}\}.$$

The overall encoding process is denoted as:

$$s = f^{\text{HSE}}(x).$$

Each client obtains its local set of real token sequences via HSE, denoted as $\mathcal{S}_n^{\text{real}}$.

2.4 M-AR Transformer Generator

Conditional Modeling with Multi-scale Autoregression. To reflect the semantic dependencies from global to local and to reduce the complexity of modeling high-dimensional joint distributions, FALCON factorizes the class-conditional token sequence distribution using a multi-scale autoregressive (M-AR) approach:

$$p(s | y) = p(s^{(0)} | y) \cdot p(s^{(1)}, \dots, s^{(L)} | y, s^{(0)}). \quad (2)$$

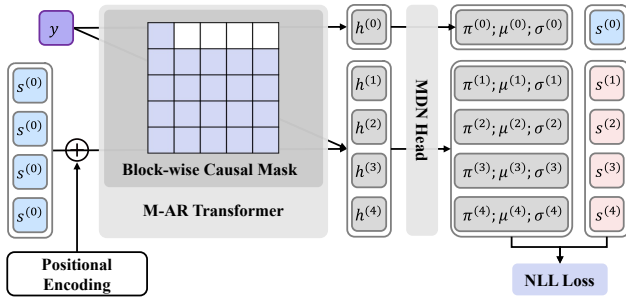


Figure 4: Architecture of the M-AR Transformer Generator.

Conditional GMM Formulation. To instantiate the above hierarchical conditional factorization, we formulate the conditional distributions of image-level and region-level tokens using Gaussian Mixture Models (GMMs). Specifically, given a class label y , the conditional distribution of the image-level token $s^{(0)}$ is assumed to follow a GMM with K components:

$$p(s^{(0)} | y) = \sum_{k=1}^K \pi_k^{(0)} \mathcal{N}(s^{(0)} | \mu_k^{(0)}, \Sigma_k^{(0)}), \quad (3)$$

where $\{\pi_k^{(0)}, \mu_k^{(0)}, \Sigma_k^{(0)}\}_{k=1}^K$ denote the mixture weights, means, and covariances of the GMM, respectively.

Subsequently, conditioned on $s^{(0)}$ and y , we model the distribution of the L region-level tokens $(s^{(1)}, \dots, s^{(L)})$ independently for each token position as:

$$p(s^{(1)}, \dots, s^{(L)} | s^{(0)}, y) = \prod_{l=1}^L \sum_{k=1}^K \pi_k^{(l)} \mathcal{N}(s^{(l)} | \mu_k^{(l)}, \Sigma_k^{(l)}), \quad (4)$$

where, for each token position l , $\{\pi_k^{(l)}, \mu_k^{(l)}, \Sigma_k^{(l)}\}_{k=1}^K$ denote the mixture weights, means, and covariances for the k -th component of the GMM.

M-AR Transformer Architecture. As shown in Figure 4, we design a M-AR Transformer generator $g(\cdot; \theta^g)$, with each client n maintaining a local generator $g_n(\cdot; \theta_n^g)$. The core architecture consists of a Transformer equipped with a block-wise causal mask (Tian et al. 2024), which enforces a coarse-to-fine attention mechanism across the token sequence. Combining the multi-scale autoregressive factorization (see Eq.2) with conditional GMM modeling for both image-level and region-level tokens (see Eqs.3 and 4), the Transformer models the sequential dependencies among tokens at different semantic scales. At each position t in the sequence, the hidden representation $h^{(t)}$ is passed to a Mixture Density Network (MDN) head, which outputs the GMM parameters for that position.

Training. During training, each client n constructs an input sequence of length $1 + L$ for each sample: $\{y, s^{(0)}, \dots, s^{(L)}\}$. Self-attention is computed using a

block-wise causal mask, and the MDN head predicts the GMM parameters at each position. This results in a locally estimated distribution \hat{p}_n . The training objective is to minimize the negative log-likelihood (NLL) loss:

$$\mathcal{L}_{\text{NLL}}^{(n)}(\theta_n^g) = -\log \hat{p}_n(s^{(0)} | y; \theta_n^g) \cdot \hat{p}_n(s^{(1)}, \dots, s^{(L)} | y, s^{(0)}; \theta_n^g). \quad (5)$$

Construction of Synthetic Token Sequence Set. Synthetic token sequences s_n^{syn} are sampled from the estimated distribution $\hat{p}_n(\cdot; \theta_n^g)$.

2.5 Local Classifier Training

Each client trains a local classifier h_n on the real token sequence set $\mathcal{S}_n^{\text{real}}$ using a cross-entropy loss.

2.6 Distillation-Guided Global Classifier Training

The server receives all uploaded local classifiers $\{h_n\}_{n=1}^N$ and synthetic token sequence sets $\{\mathcal{S}_n^{\text{syn}}\}_{n=1}^N$, and constructs a combined training set \mathcal{S}^{syn} . It then trains the global classifier h under the guidance of knowledge distillation.

Classifier Design. The classifier includes attention pooling and a classification head. Each input token sequence s is first pooled using a learnable query vector $q \in \mathbb{R}^d$ to integrate multi-scale semantic information:

$$K = sW_k \in \mathbb{R}^{(1+L) \times d}, \quad V = sW_v \in \mathbb{R}^{(1+L) \times d}. \quad (6)$$

The attention scores are computed as:

$$\alpha = \text{softmax}\left(\frac{qK^T}{\sqrt{d}}\right). \quad (7)$$

The semantic summary is obtained as:

$$s^{\text{sem}} = \alpha V. \quad (8)$$

Finally, the semantic summary s^{sem} is passed into the classification head $h_{\text{cls}}(\cdot)$ to perform multi-class prediction:

$$\hat{y} = h_{\text{cls}}(s^{\text{sem}}). \quad (9)$$

Knowledge Distillation Loss. Since the synthetic data only indirectly optimize the FL objective, we introduce a knowledge distillation mechanism that integrates the local classifiers as teachers to correct the optimization target and compensate for limitations in distribution modeling by the generator. Specifically, for each input token sequence s , we apply temperature scaling to the softmax outputs of all local classifiers $\{h_n\}_{n=1}^N$ on s , and compute their average to obtain the ensemble teacher distribution, defined as:

$$p_T(y | s) = \frac{1}{N} \sum_{n=1}^N \text{softmax}\left(\frac{z_n(s)}{T}\right), \quad (10)$$

where $z_n(s)$ denotes the logits output of the n -th local classifier on s , and T is the distillation temperature.

The probability distribution output by the global model is:

$$p_S(y | s) = \text{softmax}\left(\frac{z_S(s)}{T}\right), \quad (11)$$

where $z_S(s)$ denotes the logits of the global classifier.

The distillation loss is defined as the Kullback–Leibler (KL) divergence between the ensemble teacher distribution and the global classifier distribution:

$$\ell_{\text{KD}} = T^2 \cdot \text{KL}(p_T(y | s) || p_S(y | s)). \quad (12)$$

Dataset	Domain	Non-IID Type	#Classes	#Clients
Tuberculosis	Med.	Feat.	2	3
PACS	Nat.	Feat.	7	4
OfficeHome	Nat.	Feat.	65	4
PMRAM	Med.	Label	4	5
Pneumonia	Med.	Label	2	5
Tiny-ImageNet	Nat.	Label	200	5

Table 1: Dataset summary. “Med.”/“Nat.” denote domains; “Feat.”/“Label” indicate types of non-IID.

Overall Loss. The final loss function is a weighted combination of the cross-entropy loss and the distillation loss:

$$\mathcal{L}_{\text{total}} = (1 - \alpha) \ell_{\text{CE}} + \alpha \ell_{\text{KD}}, \quad (13)$$

where ℓ_{CE} is the standard cross-entropy loss and α is the balancing coefficient. Guided by distillation, the global classifier is able to integrate the knowledge captured by local classifiers on real token sequences from each client.

After training, the global classifier h is distributed to all clients, forming the global model $w = f^{\text{HSE}} \cdot h$.

3 Experiments and Analyses

3.1 Experimental Settings

Datasets. Table 1 summarizes the datasets used for evaluating federated learning performance under both feature and label non-IID scenarios, covering natural and medical imaging domains. For the feature non-IID setting, we select datasets with inherent multi-source structures and treat each original subset as a separate client in the federated setup. Tuberculosis (Luo et al. 2024) consists of chest X-ray images collected from three hospitals in Montgomery (Candemir et al. 2013), China (Jaeger et al. 2014), and India. PACS (Li et al. 2017) and OfficeHome (Venkateswara et al. 2017) each contain images from four visual domains: Art Painting, Cartoon, Photo, and Sketch for PACS; and Art, Clipart, Product, and Real-World for OfficeHome. For label non-IID settings, we simulate heterogeneous label distributions across clients using PMRAM (orville 2024), Pneumonia (Kermany et al. 2018), and Tiny-ImageNet (Li et al. 2020). In these cases, we construct five clients by partitioning each dataset with a Dirichlet distribution over class labels to control label heterogeneity. To accommodate differences in dataset scale, we employ a 6:2:2 split for training, validation, and test sets in PMRAM, whereas Pneumonia adopts an 8:2 split for training and validation with the official test set for evaluation. All other datasets use a 8:1:1 split for training, validation, and testing. All images are resized to 448×448 resolution.

Baselines. To evaluate the effectiveness of the proposed method, we compare FALCON with FedAvg (McMahan et al. 2017) as well as several representative approaches from the current one-shot federated learning (OSFL) literature, including DENSE (Zhang et al. 2022), FedLMG (Yang et al. 2025), FedPFT (Beitollahi et al. 2025), and FedCGS (Guan, Zhou, and Gu 2025). Specifically, DENSE serves as a classical data-free knowledge distillation baseline. FedLMG employs a pretrained diffusion model with classifier guidance

to generate synthetic data. Both FedPFT and FedCGS are methods based on semantic feature distribution modeling.

Configurations. We use DINOv2-base (Oquab et al. 2023) as the pretrained visual encoder. For hierarchical scale encoding (HSE), both the sliding window size and stride are set to 224. The token sequence generator adopts a 4-layer architecture with a hidden dimension of 768 and 16 attention heads, using diagonal covariance and $K = 20$. It is trained for 400 epochs with a batch size of 256, using the AdamW optimizer with a learning rate of $1e-4$ and a weight decay of $1e-4$. The local classifier is trained for 60 epochs with a batch size of 256 using the Adam optimizer and a learning rate of $5e-4$. For global classifier training, the balancing coefficient α is set to 0.5 and the distillation temperature T is 4. The global classifier is trained for 60 epochs with a batch size of 256, using Adam and a learning rate of $5e-4$. For the baselines, FedAvg, DENSE, and FedLMG use a local training batch size of 64, trained for 60 epochs with the Adam optimizer and a learning rate of $5e-4$. FedPFT uploads the parameters of a diagonal-covariance Gaussian mixture model $\mathcal{G}(K)$ with $K = 20$, and all other settings follow the original paper. FedCGS also uses the DINOv2-base encoder, with all other settings consistent with the original implementation.

3.2 Main Results

Table 2 reports the classification accuracy (%) of each method under both feature and label non-IID scenarios. FALCON achieves the best or second-best results across all datasets and settings among OSFL methods. Specifically: (1) **Strong Performance under Feature non-IID.** FALCON outperforms the best OSFL baseline by 10.72% on the Tuberculosis dataset, with a performance gain of up to 39.88% compared to the weakest baseline. It also maintains consistently strong results on PACS and OfficeHome. (2) **Stable Advantage under Label non-IID.** While the performance of existing OSFL methods tends to fluctuate or degrade as the Dirichlet parameter α decreases, FALCON maintains consistently high accuracy across all α values and datasets, demonstrating strong robustness to varying degrees of label heterogeneity. (3) **Superior Results on Medical datasets.** Across six label non-IID settings on the PMRAM and Pneumonia datasets, FALCON consistently achieves the best results among OSFL methods. In these settings, FALCON outperforms the best OSFL baseline by an average of 7.83%, with improvements of 7.42%–11.22% on PMRAM and 2.05%–9.80% on Pneumonia. These results demonstrate the advantages of our hierarchical scale encoding for modeling high-resolution regional details in medical images and effectively leveraging multi-scale information from raw images. (4) **Reliable and Robust Overall Performance.** FALCON delivers stable and robust performance, achieving an average improvement of 9.58% over the strongest OSFL baseline across all settings. In contrast to other methods, which tend to suffer severe performance degradation on certain datasets or non-IID configurations, FALCON maintains high accuracy in all cases. For example, on the Tuberculosis dataset, the second-best method, FedCGS, only reaches 53.50% accuracy, whereas FALCON achieves 87.20%. Sim-

	Feature non-IID			Label non-IID									Avg.
	Tuberculosis	PACS	OfficeHome	PMRAM			Pneumonia			Tiny-ImageNet			
				0.1	0.3	0.5	0.1	0.3	0.5	0.1	0.3	0.5	
FedAvg	85.07	87.61	82.51	90.98	94.20	95.83	66.11	86.54	88.83	30.75	34.70	36.51	73.30
DENSE	<u>76.48</u>	51.63	60.74	26.93	26.53	54.15	64.12	78.75	<u>87.17</u>	21.70	26.87	30.90	50.50
FedLMG	47.32	61.42	60.19	36.99	36.54	37.98	57.87	59.97	63.09	26.83	30.22	31.41	45.82
FedPFT	72.62	96.03	83.85	<u>74.09</u>	<u>74.38</u>	<u>77.74</u>	72.65	75.27	76.74	59.40	59.70	60.78	73.61
FedCGS	53.50	97.50	90.16	64.72	69.04	62.79	<u>77.62</u>	<u>79.12</u>	76.90	81.53	<u>81.53</u>	<u>81.53</u>	<u>76.34</u>
FALCON	87.20	<u>96.42</u>	<u>89.68</u>	83.33	85.60	85.16	82.41	88.92	89.22	<u>79.93</u>	81.59	81.55	85.92

Table 2: Comparison of test accuracy (%) under feature and label non-IID settings. The first three columns show results on feature non-IID datasets; the next nine columns report results under label non-IID scenarios with Dirichlet distribution ($\alpha \in \{0.1, 0.3, 0.5\}$). The final column (“Avg.”) denotes the average across all settings. Among OSFL methods, the best and second-best results are highlighted in **bold** and underline, respectively.

Dataset	Full(M-AR)	w/o HSE	w/o KD	AR	NAR
Tuberculosis	87.20	83.63	84.23	86.61	86.90
PACS	96.42	96.16	95.71	96.29	96.23
OfficeHome	89.68	89.05	86.87	84.26	86.98
PMRAM	84.70	81.26	80.41	83.38	78.67
Pneumonia	86.85	82.12	84.42	72.76	85.73
Tiny-ImageNet	81.02	79.59	77.20	80.70	80.53
Avg.	85.92	82.32	82.82	81.47	82.99

Table 3: Ablation study of FALCON. “Full” denotes the complete method; “w/o HSE” disables hierarchical scale encoding; “w/o KD” disables knowledge distillation. “AR” and “NAR” represent the vanilla autoregressive and non-autoregressive generator variants, respectively.




Figure 5: Computational cost (bars, GFLOPs per sample) and average classification accuracy (dots, %) for OSFL methods involving sample generation.

ilarly, for the PMRAM dataset, over half of the competing baselines (e.g., DENSE, FedLMG) fall below 40% accuracy, while FALCON remains above 80% across all α settings.

3.3 Computational Efficiency of Sample Generation

Figure 5 compares the per-sample computational cost (in GFLOPs) and average classification accuracy (%) of OSFL methods involving sample generation, including DENSE, which generates pseudo-samples; FedLMG, which leverages a pretrained diffusion model; FedPFT, which performs direct sampling from a GMM; and our proposed method, FALCON. Notably, FedLMG incurs an extremely high cost (7,600 GFLOPs per sample) due to the expensive multi-step sampling in diffusion-based generation. In contrast, FAL-

	Tuberculosis	PACS	OfficeHome
M-AR	0.003227	0.002047	0.000905
AR	0.003868	0.002457	0.001079
NAR	0.003868	0.002151	0.000849

Table 4: Maximum mean discrepancy (MMD) between real and synthetic token sequences on feature non-IID datasets, comparing different generator variants. Lower is better.

CON reduces the per-sample cost to 0.46 GFLOPs, which is over four orders of magnitude lower, while ensuring strong expressive capacity. Although FedPFT achieves the lowest computational cost by directly sampling from a statistical feature distribution, it lacks the expressive capacity and generation fidelity of transformer-based sequence modeling. These results underscore FALCON’s computational efficiency, making it well-suited for resource-constrained or large-scale federated deployments.

3.4 Ablation Studies

We conducted ablation studies to investigate the effects of key components, including hierarchical scale encoding (HSE), the token sequence generator architecture, and knowledge distillation. Due to space constraints, for label non-IID datasets, we report the mean accuracy across different Dirichlet parameters (α) for each dataset. Table 2 reports the results of the ablation study.

Impact of Hierarchical Scale Encoding (HSE). Introducing HSE consistently improves performance across all datasets, raising the overall average by 3.60%. The improvement is particularly notable for the medical imaging benchmarks Tuberculosis, PMRAM, and Pneumonia, with improvements of 3.57%, 3.44%, and 4.73%, respectively. This demonstrates the effectiveness of HSE in capturing fine-grained information essential for complex recognition tasks.

Impact of Token Sequence Generator Architecture. The M-AR generator we adopt in FALCON achieves the highest performance among all generator variants. To further quantify the ability of each generator variant to approximate the token sequence distribution, we compute the maximum

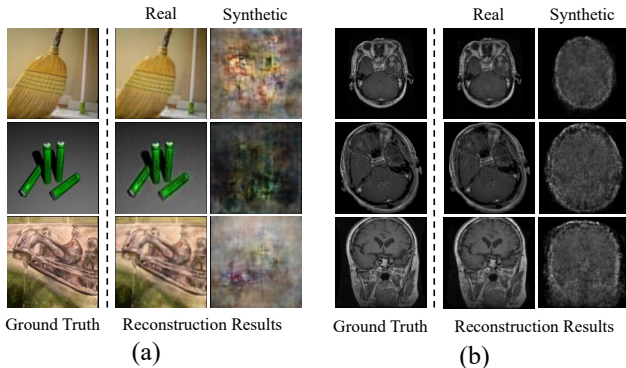


Figure 6: Visualization of reconstruction attack. (a) OfficeHome ; (b) PMRAM ($\alpha = 0.5$). For each, left: original images; middle: reconstructions from real token sequences; right: reconstructions from synthetic token sequences.

mean discrepancy (MMD) between real and synthetic token sequences on the feature non-IID datasets. We report results on Tuberculosis, PACS, and OfficeHome in Table 4. Our M-AR generator achieves the lowest MMD on most datasets, indicating a closer match to the real distribution compared to the AR and NAR generators. These results demonstrate that modeling dependencies among hierarchical tokens is critical for capturing the token sequence distribution.

Impact of Knowledge Distillation. Incorporating knowledge distillation yields a substantial improvement in average accuracy, increasing from 82.82% to 85.92%, with the largest observed improvement on a single dataset reaching 5.52%. These results highlight the critical role of knowledge distillation in enhancing global model performance.

3.5 Evaluating under Reconstruction Attacks

To evaluate potential privacy leakage, we design a reconstruction attack by training a conditional generative model to recover original images from token sequences, applying the attack to real and synthetic sequences of a selected client.

Figure 6 visualizes reconstruction results for both OfficeHome and PMRAM ($\alpha = 0.5$). For synthetic token sequences, we present the reconstruction with the highest SSIM among all generated samples for each original image, corresponding to a conservative, worst-case evaluation of potential privacy risk. When using real token sequences, the reconstructed images retain substantial correspondence to the original inputs, indicating that detailed visual information can still be recovered. In contrast, reconstructions from synthetic token sequences are largely unrecognizable, exhibiting significant loss of structural and semantic content, demonstrating strong privacy protection.

For quantitative evaluation, we report the average peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) for each dataset. For real token sequences, the average similarity between original and reconstructed images is computed directly. For synthetic token sequences, due to the lack of one-to-one correspondence, we report the highest PSNR and SSIM across all reconstructions for each image, and average these values. As shown in Table 5, constructions

Token Source	OfficeHome		PMRAM ($\alpha = 0.5$)	
	PSNR	SSIM	PSNR	SSIM
Real	27.18	0.8111	28.75	0.8176
Synthetic	13.30	0.1534	18.37	0.4368

Table 5: Average PSNR and SSIM for reconstructions from real and synthetic token sequences on OfficeHome and PMRAM ($\alpha = 0.5$).

from synthetic tokens yield much lower similarity scores on both datasets, indicating the privacy-preserving advantage of FALCON. Notably, the SSIM for synthetic tokens is higher on PMRAM than on OfficeHome, consistent with the greater homogeneity and structural repetition of medical images. This facilitates generative modeling of global anatomical features, yet does not compromise privacy, as patient-specific details remain unrecognizable.

3.6 Communication Cost Analysis

In FALCON, each client can either upload the synthetic token sequences or the generator model, depending on which option yields a lower communication cost under the given configuration. In our experiments, the generator model size is approximately 144.33 MB, and each 5×768 token sequence occupies about 15 KB with 32-bit storage. Uploading token sequences is thus more efficient when the total number of sequences is less than about 9,850. The per-client communication cost is therefore bounded by the combined size of the generator and local classifier, reaching up to 166.86 MB in the worst case. In practice, the actual cost is often substantially lower, as FALCON enables clients to flexibly select the upload strategy that best suits their data volume and communication constraints.

4 Conclusion

We propose FALCON, a one-shot federated learning framework tailored for diverse application domains and multi-type non-IID scenarios. FALCON introduces HSE to extract hierarchical token sequences without fine-tuning the encoder, and employs an M-AR transformer to model the distribution of token sequences. Combined with distillation-guided global training, the server learns a high-performing classifier. Extensive experiments show that FALCON delivers robust and competitive results across diverse tasks, consistently matching or surpassing the best existing OSFL approaches. Our method significantly enhances the practical utility of pretrained encoders for federated learning in real-world settings, improving semantic retention and demonstrating strong scalability. These advantages position FALCON as a promising paradigm for federated classification. To further reinforce these benefits, key directions for future work are to improve encoding effectiveness while ensuring all clients adopt a consistent encoding strategy, and to explore more powerful generative modeling paradigms with improved effectiveness and efficiency.

Acknowledgments

This work is supported by National Key Research and Development Program of China (2024YFE0203100), Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology (Grant No. 2024B1212010006), Shenzhen Science and Technology Program (JCYJ20230807120800001), and the project supplementary funding of National Innovation 2030 Major S&T Project of China (2020AAA0104203).

References

- Allouah, Y.; Dhasade, A.; Guerraoui, R.; Gupta, N.; Kermarrec, A.-M.; Pinot, R.; Pires, R.; and Sharma, R. 2024. Revisiting Ensembling in One-Shot Federated Learning. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 68500–68527. Curran Associates, Inc.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein GAN. arXiv:1701.07875.
- Beitollahi, M.; Bie, A.; Hemati, S.; Brunswic, L. M.; Li, X.; Chen, X.; and Zhang, G. 2025. Foundation Models Meet Federated Learning: A One-shot Feature-sharing Method with Privacy and Performance Guarantees. *Transactions on Machine Learning Research*.
- Candemir, S.; Jaeger, S.; Palaniappan, K.; Musco, J. P.; Singh, R. K.; Xue, Z.; Karargyris, A.; Antani, S.; Thoma, G.; and McDonald, C. J. 2013. Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration. *IEEE transactions on medical imaging*, 33(2): 577–590.
- Chen, H.; Li, H.; Zhang, Y.; Bi, J.; Zhang, G.; Zhang, Y.; Torr, P.; Gu, J.; Krompass, D.; and Tresp, V. 2025. Fedbip: Heterogeneous one-shot federated learning with personalized latent diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 30440–30450.
- Chen, H.-Y.; and Chao, W.-L. 2021. FedBE: Making Bayesian Model Ensemble Applicable to Federated Learning. In *ICLR*.
- Dai, R.; Zhang, Y.; Li, A.; Liu, T.; Yang, X.; and Han, B. 2024. Enhancing One-Shot Federated Learning Through Data and Ensemble Co-Boosting. In *The Twelfth International Conference on Learning Representations*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.
- Guan, Z.; Zhou, Y.; and Gu, X. 2025. Capture global feature statistics for one-shot federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 16942–16950.
- Guha, N.; Talwalkar, A.; and Smith, V. 2019. One-Shot Federated Learning. *CoRR*, abs/1902.11175.
- Hasan, M.; Zhang, G.; Guo, K.; Chen, X.; and Poupart, P. 2024. Calibrated One Round Federated Learning with Bayesian Inference in the Predictive Space. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(11): 12313–12321.
- Heinbaugh, C. E.; Luz-Ricca, E.; and Shao, H. 2023. Data-free one-shot federated learning under very high statistical heterogeneity. In *The Eleventh International Conference on Learning Representations*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Jaeger, S.; Candemir, S.; Antani, S.; Wang, Y.-X. J.; Lu, P.-X.; and Thoma, G. 2014. Two public chest X-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative imaging in medicine and surgery*, 4(6): 475.
- Jhunjhunwala, D.; Wang, S.; and Joshi, G. 2024. Fed-Fisher: Leveraging Fisher Information for One-Shot Federated Learning. In Dasgupta, S.; Mandt, S.; and Li, Y., eds., *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, 1612–1620. PMLR.
- Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A. N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. 2021. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2): 1–210.
- Kang, M.; Chikontwe, P.; Kim, S.; Jin, K. H.; Adeli, E.; Pohl, K. M.; and Park, S. H. 2023. One-Shot Federated Learning on Medical Data Using Knowledge Distillation with Image Synthesis and Client Model Adaptation. In Greenspan, H.; Madabhushi, A.; Mousavi, P.; Salcudean, S.; Duncan, J.; Syeda-Mahmood, T.; and Taylor, R., eds., *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, 521–531. Cham: Springer Nature Switzerland.
- Kasturi, A.; and Hota, C. 2023. OSGAN: One-shot distributed learning using generative adversarial networks: A. Kasturi, C. Hota. *The Journal of Supercomputing*, 79(12): 13620–13640.
- Kermany, D. S.; Goldbaum, M.; Cai, W.; Valentim, C. C.; Liang, H.; Baxter, S. L.; McKeown, A.; Yang, G.; Wu, X.; Yan, F.; et al. 2018. Identifying medical diagnoses and treatable diseases by image-based deep learning. *cell*, 172(5): 1122–1131.
- Kingma, D. P.; Welling, M.; et al. 2013. Auto-encoding variational bayes.
- Li, A.; Sun, J.; Wang, B.; Duan, L.; Li, S.; Chen, Y.; and Li, H. 2020. Lotteryfl: Personalized and communication-efficient federated learning with lottery ticket hypothesis on non-iid datasets. *arXiv preprint arXiv:2008.03371*.

- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2017. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, 5542–5550.
- Lin, T.; Kong, L.; Stich, S. U.; and Jaggi, M. 2020. Ensemble Distillation for Robust Model Fusion in Federated Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 2351–2363. Curran Associates, Inc.
- Liu, L.; Ren, Y.; Lin, Z.; and Zhao, Z. 2022. Pseudo Numerical Methods for Diffusion Models on Manifolds. In *The Tenth International Conference on Learning Representations*. OpenReview.net.
- Luo, G.; Zhang, H.; Wang, X.; Chen, M.; and Zhu, Y. 2024. MPCPA: Multi-Center Privacy Computing with Predictions Aggregation based on Denoising Diffusion Probabilistic Model. *arXiv preprint arXiv:2403.07838*.
- Lyu, L.; Yu, H.; and Yang, Q. 2020. Threats to federated learning: A survey. *arXiv preprint arXiv:2003.02133*.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Nasr, M.; Shokri, R.; and Houmansadr, A. 2019. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In *IEEE symposium on security and privacy*, 739–753.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- orville. 2024. PMRAM: Bangladeshi Brain Cancer - MRI Dataset. <https://www.kaggle.com/datasets/orvile/pmram-bangladeshi-brain-cancer-mri-dataset>. Accessed: 2025-07-31.
- Rao, B.; Zhang, J.; Wu, D.; Zhu, C.; Sun, X.; and Chen, B. 2025. Privacy Inference Attack and Defense in Centralized and Federated Learning: A Comprehensive Survey. *IEEE Transactions on Artificial Intelligence*, 6(2): 333–353.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, R.; Liu, D.; Chen, D. Z.; Festag, A.; Trinitis, C.; Schulz, M.; and Knoll, A. 2023. Federated learning via decentralized dataset distillation in resource-constrained edge environments. In *2023 International Joint Conference on Neural Networks (IJCNN)*, 1–10. IEEE.
- Tang, Z.; Zhang, Y.; Dong, P.; Cheung, Y.-m.; Zhou, A. C.; Han, B.; and Chu, X. 2024. FuseFL: One-Shot Federated Learning through the Lens of Causality with Progressive Model Fusion. In *Advances in Neural Information Processing Systems*, volume 37, 28393–28429.
- Tian, K.; Jiang, Y.; Yuan, Z.; Peng, B.; and Wang, L. 2024. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in neural information processing systems*, 37: 84839–84865.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5018–5027.
- Yang, M.; Su, S.; Li, B.; and Xue, X. 2024. Feddeo: Description-enhanced one-shot federated learning with diffusion models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 6666–6675.
- Yang, M.; Su, S.; Li, B.; and Xue, X. 2025. One-Shot Heterogeneous Federated Learning with Local Model-Guided Diffusion Models. In *Forty-second International Conference on Machine Learning*.
- Yazdinejad, A.; Dehghantanha, A.; Karimipour, H.; Srivastava, G.; and Parizi, R. M. 2024. A Robust Privacy-Preserving Federated Learning Model Against Model Poisoning Attacks. *IEEE Transactions on Information Forensics and Security*, 19: 6693–6708.
- Zhang, J.; Chen, C.; Li, B.; Lyu, L.; Wu, S.; Ding, S.; Shen, C.; and Wu, C. 2022. DENSE: Data-Free One-Shot Federated Learning. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 21414–21428. Curran Associates, Inc.
- Zhang, J.; Liu, S.; and Wang, X. 2024. One-shot federated learning via synthetic distiller-distillate communication. *Advances in Neural Information Processing Systems*, 37: 102611–102633.
- Zhang, J.; Qi, X.; and Zhao, B. 2023. Federated Generative Learning with Foundation Models. *arXiv:2306.16064*.
- Zhou, Y.; Pu, G.; Ma, X.; Li, X.; and Wu, D. 2020. Distilled one-shot federated learning. *arXiv preprint arXiv:2009.07999*.