

FedRIR: Rethinking Information Representation in Federated Learning

Anonymous Author(s)

Abstract

Mobile and Web-of-Things (WoT) devices at the network edge generate vast amounts of data for machine learning applications, yet privacy concerns hinder centralized model training. Federated Learning (FL) allows clients (devices) to collaboratively train a shared model coordinated by a central server without transfer private data, but inherent statistical heterogeneity among clients presents challenges, often leading to a dilemma between clients' needs for personalized local models and the server's goal of building a generalized global model. Existing FL methods typically prioritize either global generalization or local personalization, resulting in a trade-off between these two objectives and limiting the full potential of diverse client data. To address this challenge, we propose a novel framework that simultaneously enhances global generalization and local personalization by **Rethinking Information Representation** in the **Federated learning process (FedRIR)**. Specifically, we introduce **Masked Client-Specific Learning (MCSL)**, which isolates and extracts fine-grained client-specific features tailored to each client's unique data characteristics, thereby enhancing personalization. Concurrently, the **Information Distillation Module (IDM)** refines the global shared features by filtering out redundant client-specific information, resulting in a purer and more robust global representation that enhances generalization. By integrating the refined global features with the isolated client-specific features, we construct enriched representations that effectively capture both global patterns and local nuances, thereby improving the performance of downstream tasks on the client. Extensive experiments across diverse datasets demonstrate that FedRIR significantly outperforms state-of-the-art FL methods, achieving up to a 3.93% improvement in accuracy while ensuring robustness and stability in heterogeneous environments.

CCS Concepts

• **Computing methodologies** → **Machine learning**; *Distributed algorithms*; • **Human-centered computing** → **Ubiquitous and mobile computing**.

Keywords

Federated Learning, Information Representation, Information Distillation, Masked Representation Learning

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, April 28–May 02, 2025, Sydney, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06

<https://doi.org/XXXXXXX.XXXXXXX>

ACM Reference Format:

Anonymous Author(s). 2024. FedRIR: Rethinking Information Representation in Federated Learning. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Mobile and Web-of-Things (WoT) devices at the network edge now generate massive amounts of data daily [9, 39], which recent advancements in deep learning have demonstrated to be crucial for achieving high performance in data-driven models [8, 10, 21]. However, data at individual clients is often scarce and biased, making it challenging to build accurate and robust models [4, 11]. Aggregating data from multiple clients into a centralized dataset might seem like a natural solution, but privacy and security concerns associated with sharing sensitive data across clients significantly hinder this approach [17, 31].

Federated Learning (FL) offers a promising solution to these challenges by enabling collaborative model training across multiple clients without the need to transfer local data [30, 38]. In FL, each client retains its private data and trains a local model, sharing only model parameters with a central server. This approach addresses privacy concerns while leveraging collective data to improve overall model performance. However, client data is typically non-IID [28, 44, 46], and the heterogeneity issue poses a significant challenge for traditional FL methods, such as FedAvg, as a single global model often fails to adequately capture the diverse data characteristics, leading to suboptimal performance across clients [20, 35].

Personalized Federated Learning (pFL) was introduced to tackle data heterogeneity across clients. It aims to build models tailored to each client by leveraging client-specific data variations to improve overall performance [18, 35]. In pFL, each client maintains a personalized component that adapts to its local data, while also contributing the parameters of a global component aimed at extracting refined global features. These global parameters are then aggregated on the server to improve generalization and address challenges associated with data scarcity. However, existing pFL methods often prioritize either global information representation, as demonstrated by FedPer [1] and FedRep [7], during local training. This singular focus tends to overlook the complementary aspect, limiting the ability to effectively balance collaborative learning with personalization. Although methods like FedProto [36] attempt to bridge this gap by employing prototypes to guide personalized feature extraction, their effectiveness heavily relies on the quality of global representations. Poor global features can misguide the personalization process, leading to suboptimal performance in both global and personalized objectives [45]. The recently proposed FedAKD [12] seeks to optimize both global and personalized models through server-side and client-side knowledge distillation, but it relies on a reference dataset accessible to both the server and clients,

which is often impractical in real-world scenarios. Therefore, it is crucial to develop methods that can simultaneously enhance global generalization and local personalization within federated learning.

To meet this need, we propose FedRIR, an innovative pFL framework that enhances both local personalization and global generalization by jointly optimizing two key components: Masked Client-Specific Learning (MCSL) and the Information Distillation Model (IDM). MCSL isolates and extracts client-specific features by applying masking techniques that emphasize the unique characteristics of each client’s data. IDM refines the shared global features by filtering out redundant or overly client-specific information, producing a purer and more robust global representation that captures common patterns across all clients. By integrating these components, FedRIR achieves an effective balance between personalization and generalization. In summary, the main contributions of this paper are as follows:

- Propose a Masked Client-Specific Learning mechanism that effectively extracts client-specific heterogeneous information, enhancing model personalization.
- Develop an Information Distillation Module that refines global features, resulting in a purer global representation and optimizing the model’s capacity for improved generalization.
- Extensive experiments demonstrate that FedRIR not only achieves superior performance and communication efficiency but also maintains robustness and stability across various heterogeneous scenarios, outperforming state-of-the-art methods.

2 Related Work

2.1 Personalized Federated Learning

Personalized Federated Learning (pFL) has been proposed to address statistical heterogeneity and improve personalization in FL by training a unique, personalized model for each client rather than relying on a single global model [35, 42].

Existing pFL approaches can be grouped into five main categories [47]: (1) Meta-learning-based pFL: Methods like Per-FedAvg [14] fine-tune the global model using local data, resulting in a more personalized model for each user. Although this method enhances personalization, it may face challenges in maintaining consistency across clients due to local data variations. (2) Regularization-based pFL: Ditto [27] uses a proximal term to incorporate global information from the global model parameters during local training. This approach helps balance personalization and global consistency, but it may not fully capture the individual nuances of each client’s data. (3) Personalized-aggregation-based pFL: Approaches like FedALA [44] use an Adaptive Local Aggregation module to dynamically merge the global model with the local model for each client. Although this technique is effective, its success depends on how well the aggregation aligns with each client’s specific needs, which can be challenging in diverse settings. (4) Model-splitting-based pFL: Methods such as FedRep, FedPer [1], and FedRoD split the model into a global feature extractor and a client-specific head. These approaches focus on enhancing either global or personalized feature representation, potentially limiting overall model performance by not fully integrating both aspects. (5) Knowledge-distillation-based

pFL: FedKD [40] utilizes knowledge distillation to train a student model guided by a teacher model, sharing only the student model to significantly reduce communication costs. Although this method is efficient, it may not adequately balance global and personalized information representation.

Despite advancements, existing pFL methods often prioritize either global or personalized information representation during local training. This focus limits their ability to achieve a balance between collaborative learning and personalization, especially in diverse and heterogeneous data environments.

2.2 Masked Representation Learning

Masked representation learning is a self-supervised technique that predicts masked components using contextual information, allowing models to learn rich and meaningful representations without labeled data. This approach has gained widespread adoption across various domains, especially in Natural Language Processing (NLP) and Computer Vision (CV) [10, 15, 19].

In NLP, masked language models have transformed how machines understand and generate human language. The groundbreaking work by Devlin *et al.* [10] introduced BERT, which uses masked language modeling to capture bidirectional context. By randomly masking a subset of input tokens and training the model to predict them, BERT effectively learns deep contextual representations that serve as the foundation for various downstream tasks, including question answering, sentiment analysis, and machine translation. Similarly, in CV, masked representation learning has been successfully adapted to image data. He *et al.* [19] proposed the Masked Autoencoder (MAE), which improves efficiency by sparsely applying the ViT [13] encoder to the visible (unmasked) content only. MAE employs high masking ratios, typically masking 75% of the input patches, thereby reducing computational overhead and promoting the learning of more robust and generalizable features from the visible data. Beyond NLP and CV, masked representation learning has been applied in other domains such as audio processing [32], where models learn audio representations by auto-encoding masked spectrogram patches, and graph neural networks [25], where masking nodes or edges helps facilitate the learning of structural and relational information.

MCSL extends masked representation learning to federated settings by focusing on extracting client-specific representations finely tuned to each client’s unique data characteristics. By utilizing a proper masking ratio, MCSL enhances client-specific feature extraction while mitigating the risk of overfitting to local data. As a crucial component in improving personalization within pFL, MCSL leverages the strengths of masked learning to develop more efficient and privacy-preserving models in federated environments.

3 Methodology

3.1 Overview

In pFL, we consider N clients, each with its own privacy-sensitive dataset \mathcal{D}^i exhibiting statistical heterogeneity. The goal of the proposed FedRIR framework is to improve both the generalization and personalization of the overall model by refining the global and client-specific representations. This is achieved through iterative

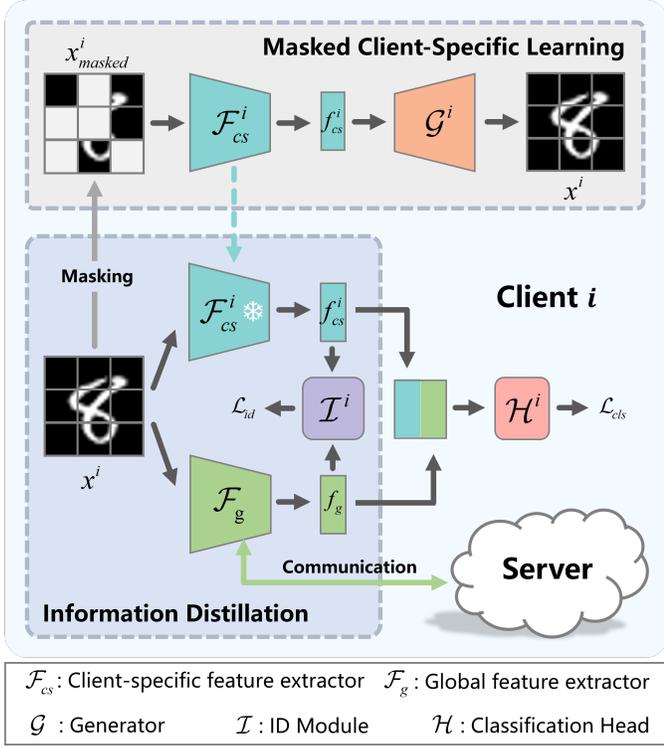


Figure 1: The workflow of the proposed FedRIR framework.

local client updates and server aggregation at each communication round. The FedRIR workflow is shown in Figure 1.

3.1.1 Local Client Update. As shown in Figure 1, the FedRIR training process consists of two key stages: Masked Client-Specific Learning (MCSL) and Information Distillation (ID).

In the MCSL stage, client-specific representations are learned. For each client i , the input data x^i undergoes a masking operation, resulting in masked data x^i_{masked} . The masked data is then processed by a client-specific feature extractor \mathcal{F}_{cs}^i to obtain client-specific features f_{cs}^i . These features are subsequently used by a generator \mathcal{G}^i to reconstruct the original input x^i .

The ID stage focuses on refining global representations by filtering out redundant information. At this stage, the client-specific feature extractor \mathcal{F}_{cs}^i is frozen to preserve the learned parameters. The original input x^i is processed by the frozen \mathcal{F}_{cs}^i to yield f_{cs}^i and by the global feature extractor \mathcal{F}_g to produce global features f_g . An Information Distillation Module (IDM) is then applied to eliminate redundant client-specific details from the global features by minimizing the mutual information between the global features and the client-specific features.

Finally, a classification head \mathcal{H}^i maps the concatenated features f_g and f_{cs}^i , forming the personalized features $f_p^i = \text{concat}(f_g, f_{cs}^i)$, to the classification label space on the client. This mapping, represented as $\hat{y}^i = \mathcal{H}^i(f_p^i)$, integrates the refined global features with the client-specific features to construct enriched representations that capture both global patterns and local nuances, enabling accurate predictions.

3.1.2 Server Aggregation. After each round of local updates, each client i sends the learned parameters of its global feature extractor \mathcal{F}_g to the server for aggregation. Following a similar approach to FedAvg, the aggregation process is defined as:

$$\bar{\mathcal{W}}_{\mathcal{F}_g} = \sum_{i=1}^N \frac{|\mathcal{D}^i|}{\sum_{i=1}^N |\mathcal{D}^i|} \mathcal{W}_{\mathcal{F}_g}^i \quad (1)$$

Here, $\bar{\mathcal{W}}_{\mathcal{F}_g}$ represents the globally aggregated parameters of all participating global feature extractors, and $|\mathcal{D}^i|$ denotes the number of data samples on the i -th client. After $\bar{\mathcal{W}}_{\mathcal{F}_g}$ is computed, it is broadcast to all clients. Each client then updates its local global feature extractor \mathcal{F}_g by replacing its parameters with the globally aggregated parameters $\bar{\mathcal{W}}_{\mathcal{F}_g}$.

3.2 Masked Client-Specific Learning

In FL, an effective client-specific feature extractor should accurately capture the unique characteristics of each client's data without simply overfitting to the training samples. To achieve this, we introduce Masked Client-Specific Learning (MCSL). MCSL encourages the client-specific feature extractor to focus on truly client-specific aspects of the data by masking parts of the input, which helps isolate and highlight the distinctive features of each client's dataset. This process ensures that the client-specific feature extractor is finely tuned to the client's unique data distribution, while also avoiding overfitting.

Capturing the unique aspects of local data $\mathcal{D}^i = \{(x_k^i, y_k^i)\}_{k=1}^{|\mathcal{D}^i|}$ on client i can be theoretically framed as maximizing the mutual information between the input data x^i and the client-specific representation f_{cs}^i [37]. Maximizing this mutual information encourages the feature extractor to retain as much relevant information about the input on client i as possible, thereby enhancing personalization. This can be expressed as:

$$\arg \max I(x^i; f_{cs}^i) = \arg \max \mathbb{E}_{p(x^i, f_{cs}^i)} \left[\log \frac{p(x^i | f_{cs}^i)}{p(x^i)} \right] \quad (2)$$

Since directly modeling the conditional distribution $p(x^i | f_{cs}^i)$ is intractable, we adopt an approach based on the Barber-Agakov (BA) bound [2], introducing an auxiliary distribution $q(x^i | f_{cs}^i)$ to derive a tractable lower bound for the mutual information $I(x^i; f_{cs}^i)$:

$$I_{BA} := H(x^i) + \mathbb{E}_{p(x^i, f_{cs}^i)} [\log q(x^i | f_{cs}^i)] \leq I(x^i; f_{cs}^i) \quad (3)$$

Here, $H(x^i) = \mathbb{E}_{p(x^i)} [-\log p(x^i)]$ is the entropy of the input data x^i , which depends solely on the data collection process and is independent of the feature extraction process. Consequently, we focus on maximizing the second term of I_{BA} , which represents the process of reconstructing the original input x^i from the client-specific features f_{cs}^i . To achieve this, we employ a variational autoencoder (VAE) [22] framework, constructing an encoder \mathcal{F}_{cs}^i to extract the client-specific features f_{cs}^i and a generator \mathcal{G}^i to reconstruct the input data x^i . The loss function to minimize the reconstruction error and thereby maximize $I(x^i; f_{cs}^i)$ is:

$$\mathcal{L}_{recon} = \|\mathcal{G}^i(f_{cs}^i) - x^i\|_2 \quad (4)$$

However, previous researches [19, 37] indicate that direct reconstruction from the original input x^i can lead to trivial solutions,

such as simply copying the input, which fails to capture meaningful representations. This issue is especially problematic in federated learning, where each client's data is typically heterogeneous and limited in size. To address this, we introduce random masking of portions of the input data, preventing the client-specific feature extractor \mathcal{F}_{cs}^i from overfitting. This effectively transforms the VAE into a denoising autoencoder (DAE) [37], with the modified reconstruction loss:

$$\mathcal{L}_{recon} = \left\| \mathcal{G}^i \left(\mathcal{F}_{cs}^i(x_{masked}^i) \right) - x^i \right\|_2 \quad (5)$$

3.3 Information Distillation

Once the client-specific features f_{cs}^i have been extracted, a global feature extractor \mathcal{F}_g is used to derive global features f_g from the input sample x^i . To ensure that the global feature extractor effectively captures pure and generalized features, we introduce an Information Distillation (ID) strategy. This strategy distills out redundant client-specific information from the global features, leading to two key benefits: (1) reducing the influence of non-common noise (unexpected client-specific noise) when the parameters of the global feature extractor are aggregated on the server, and (2) minimizing redundancy between the client-specific and global features, thereby improving the performance of subsequent downstream tasks when these features are concatenated. By reducing the mutual information $I(f_g, f_{cs}^i)$, we ensure that the global and client-specific features capture distinct aspects of the data, minimizing redundancy. This enhances the informativeness of both representations, leading to improved model performance. The ID process is formally defined as:

$$\arg \min I(f_{cs}^i, f_g) = \arg \min \mathbb{E}_{p(f_{cs}^i, f_g)} \left[\log \frac{p(f_g | f_{cs}^i)}{p(f_g)} \right] \quad (6)$$

However, directly solving Eq. (6) is challenging due to the intractability of the conditional distribution $p(f_g | f_{cs}^i)$. To address this, we employ a practical approach using a variational approximation $q_\theta(f_g | f_{cs}^i)$, as proposed in [3, 6], to estimate a reliable upper bound of the mutual information. Specifically, we utilize the Variational Contrastive Log-ratio Upper Bound (vCLUB) [6] to derive this bound:

$$\begin{aligned} I_{vCLUB} &:= \mathbb{E}_{p(f_{cs}^i, f_g)} \left[\log q_\theta(f_g | f_{cs}^i) \right] \\ &\quad - \mathbb{E}_{p(f_{cs}^i)} \mathbb{E}_{p(f_g)} \left[\log q_\theta(f_g | f_{cs}^i) \right] \\ &\geq I(f_{cs}^i, f_g) \end{aligned} \quad (7)$$

By minimizing this upper bound, the redundancy between f_g and f_{cs}^i is significantly reduced, resulting in more distinct and informative global and client-specific features. In practice, $q_\theta(f_g | f_{cs}^i)$ is typically implemented with neural networks, and we introduce an Information Distillation Module \mathcal{I}^i to implement and train $q_\theta(f_g | f_{cs}^i)$ with the following loss function:

$$\mathcal{L}_{id} = I_{vCLUB}(f_{cs}^i, f_g) \quad (8)$$

Then, the extracted client-specific features f_{cs}^i and the refined global features f_g are concatenated into personalized features $f_p^i = [f_{cs}^i, f_g]$, which is fed into a classification head \mathcal{H}^i to compute the personalized classification output. This classification process is

trained with the following loss functions:

$$\mathcal{L}_{cls} = \text{CrossEntropy} \left(\mathcal{H}^i(f_p^i), y^i \right) \quad (9)$$

Thus, the overall loss function at client i is:

$$\mathcal{L} = \mathcal{L}_{id} + \mathcal{L}_{cls} \quad (10)$$

In our implementation, we set the balance hyperparameters between the different loss components to 1, and the full learning procedure is detailed in Algorithm 1.

Algorithm 1 The learning procedure of FedRIR

Input: N clients $\{C^i\}_{i=1}^N$ with their local data $\{\mathcal{D}^i\}_{i=1}^N$; initial weights of $\mathcal{F}_{cs}^i, \mathcal{G}^i, \mathcal{F}_g, \mathcal{H}^i, \mathcal{I}^i$; masked ratio r ; client joining ratio ρ ; total communication rounds T .

Output: Trained model parameters of $\mathcal{F}_g, \{\mathcal{F}_{cs}^i, \mathcal{H}^i\}_{i=1}^N$.

- 1: Each client i initializes its local models $\mathcal{F}_{cs}^i, \mathcal{G}^i, \mathcal{F}_g, \mathcal{H}^i, \mathcal{I}^i$.
 - 2: **for** iteration $t = 0, \dots, T$ **do**
 - 3: Server samples a clients subset $\{C^i\}_{i=1}^{\rho \times N}$.
 - 4: Server broadcasts $\mathcal{W}_{\mathcal{F}_g}$ to all selected clients.
 - 5: **for** Client $C^i \in \{C^i\}_{i=1}^{\rho \times N}$ **in parallel do**
 - 6: **▷ Masked Client-Specific Learning**
 - 7: Mask the local data x^i with ratio r to obtain x_{masked}^i .
 - 8: Optimize \mathcal{F}_{cs}^i and \mathcal{G}^i on x_{masked}^i by minimizing \mathcal{L}_{recon} .
 - 9: **▷ Information Distillation**
 - 10: Freeze \mathcal{F}_{cs}^i .
 - 11: Compute $f_{cs}^i = \mathcal{F}_{cs}^i(x^i)$ and $f_g = \mathcal{F}_g(x^i)$.
 - 12: Update $\mathcal{F}_g, \mathcal{H}^i$ and \mathcal{I}^i by minimizing $\mathcal{L} = \mathcal{L}_{id} + \mathcal{L}_{cls}$.
 - 13: **end for**
 - 14: **▷ Server Aggregation**
 - 15: Server collects $\mathcal{W}_{\mathcal{F}_g}^i$ from all selected clients.
 - 16: Server computes $\mathcal{W}_{\mathcal{F}_g}$ using Eq. (1).
 - 17: **end for**
 - 18: **return** Trained model parameters of $\mathcal{F}_g, \{\mathcal{F}_{cs}^i, \mathcal{H}^i\}_{i=1}^N$.
-

4 Experiments

4.1 Statistically Heterogeneous Settings

We evaluate model performance under various heterogeneous data scenarios using three settings: pathological, practical, and real-world. These settings simulate different degrees of statistical heterogeneity commonly observed in federated learning environments. Figure 2 illustrates the data distribution on FMNIST [41] and Cifar10 [23] datasets in the pathological and practical settings with 20 clients.

4.1.1 Pathological Setting. This setting simulates extreme heterogeneity by assigning each client a subset of classes, resulting in non-overlapping labels among clients. For FMNIST and Cifar100 [23], each client receives data from only 2 out of 10 and 10 out of 100 classes, respectively, following the approach of FedCP [45]. This setup emphasizes the model's ability to learn in highly fragmented environments.

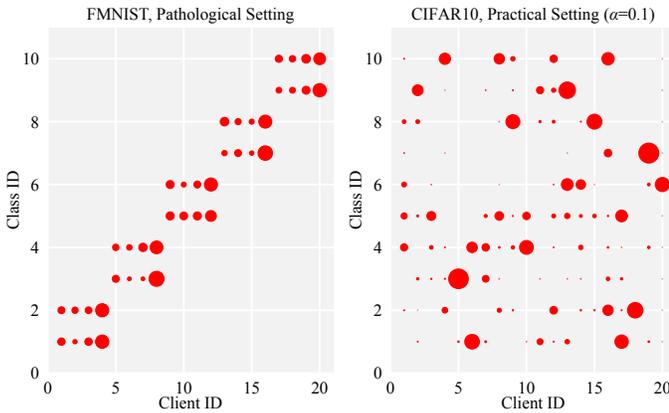


Figure 2: The data distribution on FMNIST and Cifar10 datasets in pathological and practical settings with 20 clients. The size of a circle indicates the number of local data.

4.1.2 *Practical Setting.* This setting evaluates performance in scenarios that closely resemble real-world data distribution using the Dirichlet distribution, as implemented in previous works [26, 45]. This setting is applied to the Cifar10 and MNIST [24] datasets, and we sample $q_{c,i} \sim \text{Dir}(\alpha)$ to determine the proportion of samples from class c allocated to client i . For our experiments, without explicitly stated, α is set to 0.1.

4.1.3 *Real-world Setting.* The real-world setting involves heterogeneous datasets with significant differences in data distributions across clients due to varying acquisition conditions. We utilize two datasets: (1) OfficeCaltech10 [16], which includes four distinct data sources, three from Office-31 and one from Caltech-256, each acquired using different camera devices or in different environments with various backgrounds. (2) DomainNet [34], which comprises images from six different sources. Each client in our experiments is assigned data from only one of these sources, resulting in non-IID data across clients.

4.2 Implementation Details

Following the setup in FedRoD, unless otherwise specified, our experiments involve 20 clients with a participation ratio of $\rho = 1$, and each client’s data is divided into a training set (75%) and a test set (25%). Based on our hyper-parameter analysis, we set the masking ratio $r = 0.6$. Consistent with previous works [29, 43], we use a simple CNN for the feature extractors \mathcal{F} , composed of two convolutional layers, each followed by Batch Normalization, ReLU activation, and Max Pooling layers. The classification head \mathcal{H} includes a fully connected layer, while the information distillation module \mathcal{I} consists of four fully connected layers, each followed by a ReLU activation function. The generator \mathcal{G} uses two deconvolution layers. Additionally, each task is trained for 1000 communication rounds using the Adam optimizer with a learning rate of 5×10^{-4} , a batch size of 100, and one epoch per local training round. Additionally, all experiments are conducted on a machine equipped with an Intel Xeon Gold 5220 CPU @ 2.20GHz and an NVIDIA GeForce RTX 3090 GPU. Additionally, each experiment is repeated three

times to ensure statistical significance, and we report the mean and standard deviation of the results.

4.3 Performance Comparison

We evaluate the performance of FedRIR in terms of effectiveness, scalability, stability, and hyper-parameter impact, comparing it with eleven state-of-the-art federated learning methods. These include Local Training (each client trains independently without server communication), FedAvg [30], per-FedAvg [14], FedProto [36], Ditto [27], FedRep [7], FedRoD [5], FedBABU [33], FedALA [44], FedKD [40], and FedCP [45]. The results across diverse heterogeneous scenarios are shown in Table 1.

4.4 Effectiveness

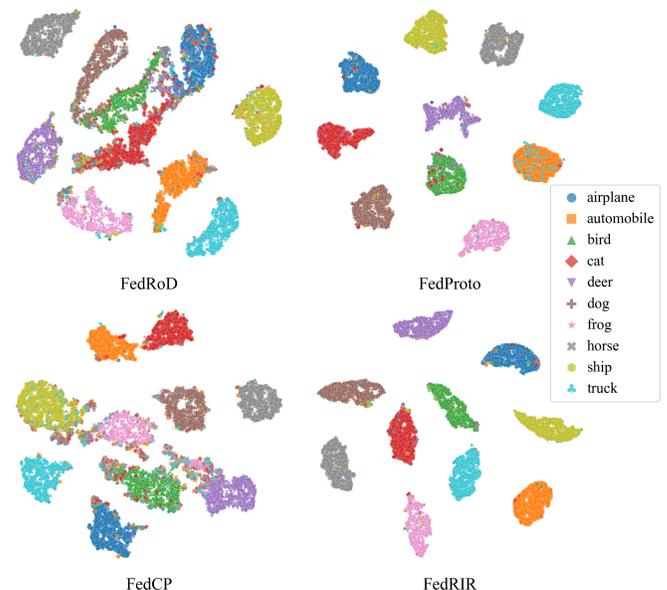


Figure 3: t-SNE visualization of global features extracted by FedRoD, FedProto, FedCP and FedRIR on Cifar10 dataset, with color represents a different class.

4.4.1 *Global Generalization Effectiveness.* As demonstrated in Table 1, FedRIR consistently outperforms Local Training, effectively capturing and utilizing shared knowledge among clients to enhance performance beyond what isolated local models can achieve. For instance, on the real-world datasets OfficeCaltech10 and DomainNet, FedRIR surpasses Local Training by 4.57% and 3.93%, respectively, highlighting its ability to leverage federated learning to improve generalization across diverse data distributions. When compared to exact global methods like FedAvg, FedRIR shows substantial advantages, particularly in non-IID environments where FedAvg struggles with conflicting updates from heterogeneous client data, resulting in poor generalization. For example, on Cifar100, FedAvg achieves only 26.50%, significantly lower than FedRIR’s performance.

Figure 3 provides further insight into the superior global feature separation achieved by FedRIR. When compared to pFL methods

Table 1: The average test accuracy (%) across different federated learning methods in various heterogeneous scenarios, along with the number of communication parameters used by each method.

Scenarios	Pathological		Practical ($\alpha = 0.1$)		Real-world		# of Comm. Params.
Datasets	Cifar100	FMNIST	MNIST	Cifar10	OfficeCaltech10	DomainNet	-
Local Training	69.41±0.22	86.36±0.04	99.16±0.02	89.27±0.25	68.66±0.34	57.99±0.03	-
FedAvg	26.50±0.55	86.01±0.15	97.49±0.32	53.41±0.34	68.14±0.45	54.92±0.33	5.597M
Per-FedAvg	56.61±0.19	86.23±0.16	98.30±0.24	85.93±0.12	66.61±0.56	50.46±0.64	5.597M
FedProto	68.08±0.27	84.08±0.04	99.18±0.02	86.94±0.16	66.81±0.56	55.94±0.18	5.120K
Ditto	69.44±0.21	86.14±0.05	99.15±0.03	89.61±0.11	68.98±0.13	55.02±0.14	5.597M
FedRep	67.97±0.32	85.86±0.07	99.34±0.03	<u>90.14±0.03</u>	69.50±0.59	58.38±0.03	5.592M
FedRoD	64.04±0.65	86.43±0.20	<u>99.49±0.02</u>	89.64±0.03	<u>69.66±0.63</u>	57.49±0.39	5.597M
FedBABU	60.34±0.14	84.26±0.25	98.92±0.00	88.05±0.14	68.24±0.61	56.16±0.38	5.592M
FedALA	63.31±0.48	<u>86.65±0.01</u>	99.32±0.01	89.42±0.11	67.77±0.65	55.92±0.39	5.597M
FedKD	<u>69.61±0.17</u>	86.12±0.01	98.81±0.03	89.77±0.11	68.19±1.05	57.20±0.03	<u>5.590M</u>
FedCP	65.15±0.31	85.85±0.09	99.44±0.05	89.86±0.04	67.82±0.41	<u>59.02±0.02</u>	6.124M
FedRIR	72.07±0.12	90.58±0.07	99.69±0.01	91.64±0.03	73.23±0.59	61.92±0.12	5.592M
Δ SOTA	↑ 2.46	↑ 3.93	↑ 0.20	↑ 1.50	↑ 3.57	↑ 2.90	

that emphasize global feature enhancement, such as FedRoD and FedProto, as well as methods that explore both global and personalized representations, such as FedCP, FedRIR demonstrates clearer and more distinct clustering across classes. The minimal overlap between different categories in FedRIR’s global features reflects a well-separated and interpretable global representation. In contrast, FedRoD and FedCP show significant class overlap, indicating suboptimal global feature extraction. Although FedProto achieves clearer boundaries between some classes, misclassifications are evident, particularly where certain truck samples (light blue club suit symbols ♣) are incorrectly clustered with automobiles (orange squares □). FedRIR’s use of information distillation strategy plays a critical role in addressing these issues by filtering out redundant client-specific information to refine the shared global model, ultimately resulting in superior generalization across clients.

4.4.2 Local Personalization Effectiveness. FedRIR stands out for its strong personalization capabilities, surpassing methods like Per-FedAvg, which struggles to balance local personalization and global accuracy. This issue is particularly evident in non-IID settings, where Per-FedAvg achieves only 56.61% accuracy on Cifar100. In contrast, FedRIR combines individual client characteristics with refined global knowledge, enabling superior model adaptation and personalization across diverse data. Methods like Ditto fall short in personalization compared to FedRIR, as they struggle to capture the nuances of client-specific data. This limitation is evident in Ditto’s lower accuracy on real-world datasets like DomainNet (55.02%), whereas FedRIR achieves 61.92%.

Similarly, approaches such as FedRep and FedPer, despite focusing on personalization, fail to fully leverage the global aggregated information from the federated learning paradigm to enhance personalization. This often results in neglecting essential global insights or facing conflicts between global and personalized objectives, leading to lower performance (e.g., FedRep: 58.38% on DomainNet). Furthermore, FedALA’s adaptive local aggregation aims to merge

global and local models, but its performance (67.77% on OfficeCaltech10) is limited by the alignment between the aggregation and the specific needs of each client. FedKD uses knowledge distillation to reduce communication costs but compromises accuracy on diverse datasets, achieving only 57.20% on DomainNet compared to FedRIR’s 61.92%. Similarly, FedCP separates global and personalized information without a robust refinement mechanism, leading to lower accuracy on Cifar100 (65.15%) compared to FedRIR’s 72.07%. These comparisons highlight FedRIR’s ability to optimize global and personalized representations concurrently, consistently delivering superior personalization across diverse datasets.

4.4.3 Communication Overhead Effectiveness. In addition to its superior performance, FedRIR maintains efficient communication overhead, which is critical in federated learning environments. Despite incorporating more advanced personalization mechanisms, FedRIR’s communication costs are comparable to simpler methods like FedAvg, FedRep, and Ditto. By focusing on refining only the most essential features through MCSL and ID, FedRIR minimizes the amount of redundant or unnecessary data that needs to be transmitted between clients and the server. This makes it practical and scalable for real-world deployments, where communication efficiency is a crucial factor. For example, methods like FedKD, which uses knowledge distillation to reduce communication costs, may sacrifice model accuracy, especially in heterogeneous settings, whereas FedRIR achieves both high accuracy and comparable communication overhead, ensuring that its strong performance can be maintained in large-scale federated learning systems without imposing excessive communication burdens.

Overall, FedRIR’s ability to simultaneously enhance generalization and personalization, while maintaining efficient communication, positions it as a robust and scalable solution for federated learning in diverse and challenging environments.

Table 2: The averaged test accuracy (%) of Cifar100 classification for client amounts scalability.

	$N = 10$	$N = 20$	$N = 50$	$N = 100$
FedAvg	30.26±0.02	30.10±0.12	21.17±0.05	14.14±0.06
Per-FedAvg	45.54±0.03	43.96±0.04	37.43±0.08	33.58±0.09
FedProto	48.09±0.02	46.15±0.06	44.80±0.06	41.68±0.08
Ditto	46.33±0.04	45.24±0.05	43.28±0.10	41.09±0.10
FedRep	49.60±0.03	47.87±0.09	45.06±0.08	40.88±0.09
FedRoD	49.18±0.03	50.05±0.26	46.81±0.08	42.31±0.09
FedBABU	49.42±0.02	49.63±0.29	43.36±0.07	37.20±0.06
FedALA	52.85±0.03	48.63±0.04	45.09±0.08	40.66±0.09
FedKD	51.68±0.04	50.31±0.05	47.30±0.08	44.29±0.10
FedCP	50.21±0.01	48.38±0.47	45.89±0.04	41.85±0.07
FedRIR	56.86±0.04	54.75±0.05	53.63±0.03	50.02±0.04
Δ SOTA	\uparrow 4.01	\uparrow 4.44	\uparrow 6.33	\uparrow 5.73

4.5 Scalability

To assess the scalability of FedRIR, we adopted the methodology from FedCP, partitioning the Cifar100 dataset into 10, 20, 50, and 100 sub-datasets to simulate 10, 20, 50, and 100 clients, respectively. As shown in Table 2, FedRIR consistently outperforms state-of-the-art methods across various client numbers, maintaining strong performance even as the number of clients increases and data becomes increasingly fragmented. FedRIR’s accuracy remains robust, decreasing moderately from 56.86% with 10 clients to 50.02% with 100 clients, a decline of only 6.84 percentage points. In comparison, other methods suffer more pronounced performance drops.

As the number of clients grows, the performance gap between FedRIR and competing methods widens, highlighting FedRIR’s superior scalability. For example, with 10 clients, FedRIR outperforms the second-best method, FedALA, by 4.01 percentage points. However, this margin expands to 5.73 percentage points with 100 clients, demonstrating FedRIR’s resilience against increasing client heterogeneity and data fragmentation.

This trend underscores the effectiveness of FedRIR’s design, particularly the combination of Masked Client-Specific Learning (MCSL) and Information Distillation (ID), which enable it to handle the complexities associated with scaling to larger numbers of clients more effectively than other approaches. The results affirm that FedRIR not only scales well but also becomes increasingly advantageous in large-scale federated learning environments, maintaining stability and accuracy where other methods falter.

4.6 Stability

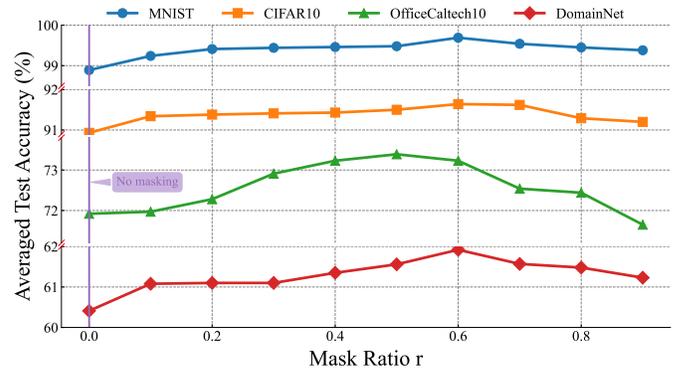
We also evaluate the stability of our FedRIR in client dropout scenarios, which are commonly assessed in many pFL methods [33, 43, 45]. This situation is prevalent in real-world settings as some clients may accidentally drop out due to battery depletion or unstable network connections, resulting in not all clients participating in every communication round. Following the strategy used in previous work [45], we vary the client joining ratio ρ during each communication round on the FMNIST dataset in the training phase. Specifically, in our experiment, the client joining ratio ρ is randomly

Table 3: The averaged test accuracy (%) of FMNIST classification with different participation ratios.

Algorithm	$\rho = 1.0$	$\rho \in [0.5, 1]$	$\rho \in [0.1, 1]$
Per-FedAvg	93.19±0.12	92.53±0.54	92.46±0.47
FedProto	92.19±0.04	91.77±0.08	90.36±0.57
Ditto	95.98±0.01	95.36±0.43	95.02±1.01
FedRep	95.47±0.03	95.04±0.04	94.65±0.03
FedROD	95.57±0.02	95.22±0.32	94.48±0.93
FedBABU	95.21±0.13	95.21±0.13	95.28±0.15
FedALA	95.56±0.00	95.39±0.12	95.15±0.11
FedKD	95.17±0.03	94.92±0.05	94.21±0.08
FedCP	95.19±0.09	95.03±0.12	94.97±0.08
FedRIR	97.51±0.00	97.48±0.04	97.43±0.02
Δ SOTA	\uparrow 1.53	\uparrow 2.09	\uparrow 2.15

sampled within a given range, such as [0.5, 1.0] and [0.1, 1.0], for each training epoch. As shown in Table 3, although all methods exhibit performance drops as ρ becomes more variable, FedRIR consistently maintains superior performance among the pFL methods, with the smallest performance decline. Additionally, the standard deviations of the results obtained by per-FedAvg, FedProto, Ditto, and FedRoD increase under unstable conditions, for example, when $\rho \in [0.1, 1.0]$.

4.7 Effect of Hyper-parameters

**Figure 4: The impact of different mask ratios on the performance of the FedRIR method.**

4.7.1 Effect of Mask Ratio r . We conducted a series of experiments focusing on the Masked Client-Specific Learning (MCSL) component by applying various mask ratios. The results demonstrate three key observations: first, as the mask ratio increases, the test accuracy initially improves for all datasets, but after reaching an optimal point, the accuracy either stabilizes or slightly declines for some datasets. Second, using masking consistently outperforms no masking across all datasets, validating the effectiveness of our proposed Masked Client-Specific Learning approach. Third, despite the optimal mask ratio varying slightly among datasets, a mask ratio of 0.6

provides the best trade-off, maximizing test accuracy and ensuring robust performance across different datasets. This balance confirms that a mask ratio of 0.6 is optimal for enhancing model performance and stability in diverse scenarios. Based on these findings, we have chosen 0.6 as the mask ratio for all subsequent experiments.

4.7.2 Effect of Heterogeneity Degree α . We conduct experiments on the FMNIST dataset under varying degrees of heterogeneity by adjusting the parameter α in the practical setting. The smaller the α is, the more likely the clients hold samples from one class. The averaged test accuracy for different values of α is presented in Table 4. The results show that as α increases, indicating less heterogeneity, the performance of FedAvg improves steadily, suggesting that FedAvg benefits from more homogeneous data distributions across clients. Conversely, all other pFL methods exhibit a decline in performance as α increases. This trend highlights their struggle with more homogeneous data distributions, where the need for strong personalization is reduced. Notably, FedRIR shows significant improvement over other methods, demonstrating superior robustness across all degrees of heterogeneity. As α increases, the performance gap between FedRIR and the second-best method widens, confirming the efficacy of FedRIR in handling varying degrees of data heterogeneity.

Table 4: The averaged test accuracy (%) of the FMNIST classification for heterogeneity.

	$\alpha = 0.1$	$\alpha = 0.3$	$\alpha = 0.5$	$\alpha = 1.0$
FedAvg	71.99±0.24	77.41±0.01	79.16±0.12	79.56±0.06
Per-FedAvg	93.19±0.12	86.46±0.05	84.07±0.03	82.94±0.03
FedProto	92.19±0.04	85.24±0.07	80.50±0.06	79.20±0.00
Ditto	<u>95.98±0.01</u>	<u>90.07±0.02</u>	86.90±0.04	<u>85.20±0.02</u>
FedRep	95.47±0.03	89.40±0.05	86.89±0.03	84.71±0.05
FedRoD	95.57±0.02	89.49±0.15	<u>87.39±0.13</u>	85.14±0.19
FedBABU	95.21±0.13	87.95±0.57	85.72±0.01	83.55±0.31
FedALA	95.56±0.00	89.52±0.11	86.10±0.10	84.90±0.08
FedKD	95.17±0.03	89.29±0.03	85.92±0.06	84.60±0.06
FedCP	95.19±0.09	88.75±0.06	85.47±0.09	84.13±0.08
FedRIR	97.51±0.00	92.53±0.02	90.72±0.03	89.66±0.03
Δ SOTA	$\uparrow 1.53$	$\uparrow 2.46$	$\uparrow 3.33$	$\uparrow 4.46$

4.8 Ablation Study

The ablation study in Table 5 highlights the importance of the key components, Masked Client-Specific Learning (MCSL) and Information Distillation (ID) in FedRIR, along with the impact of setting the masking ratio r to zero. When the masking mechanism is disabled (FedRIR $r = 0$), performance declines across all datasets, particularly on Cifar100, where accuracy drops from 72.07% to 69.28%, emphasizing that masking helps prevent overfitting and improves feature extraction. Removing MCSL (FedRIR w/o MCSL) results in an even larger performance decrease, particularly on Cifar100 (to 68.88%) and OfficeCaltech10 (to 71.44%), demonstrating that client-specific feature learning is essential for handling diverse data distributions. Interestingly, on OfficeCaltech10, which represents

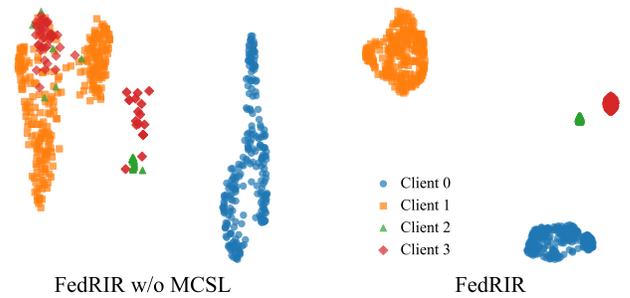


Figure 5: t-SNE visualization of client-specific features extracted by FedRIR w/o MCSL and FedRIR on OfficeCaltech10 dataset, with color represents a different client.

a real-world dataset with significantly heterogeneous client data, the absence of ID (FedRIR w/o ID) yields slightly better results (72.38%) compared to removing MCSL. This suggests that in real-world datasets like OfficeCaltech10, where client data distributions vary widely, MCSL plays a more crucial role in achieving personalization, as capturing the unique characteristics of each client’s data is more important than refining global features. Thus, while ID is essential for aligning global knowledge, in datasets with high client-level variability, such as real-world scenarios, MCSL is the key driver of performance by tailoring the model to each client’s specific data, which is evident in Figure 5. Without MCSL (left), client-specific features are poorly separated, leading to significant overlap between client data. In contrast, FedRIR (right) shows clear clustering of client features, highlighting the importance of MCSL in handling diverse data distributions effectively.

Table 5: Ablation Study

Settings	Pathological	Practical	Real-world
Datasets	Cifar100	Cifar10	OfficeCaltech10
FedRIR	72.07±0.12	91.64±0.03	73.23±0.59
FedRIR ($r = 0$)	69.28±0.04	90.93±0.02	71.92±0.56
FedRIR w/o MCSL	68.88±0.03	90.76±0.10	71.44±0.48
FedRIR w/o ID	68.37±0.06	90.68±0.14	72.38±1.46

5 Conclusion

In this paper, we introduced FedRIR, a novel framework for federated learning that balance personalization and generalization through Masked Client-Specific Learning (MCSL) and Information Distillation (ID). Extensive experiments across pathological, practical, and real-world scenarios demonstrate that FedRIR consistently outperforms state-of-the-art methods in feature representation and classification accuracy. The key innovation lies in client-specific representation learning and refining global shared information, enhancing personalization and generalization. FedRIR also shows remarkable scalability and stability with varying client participation and comparable communication overhead, making it a robust and efficient solution for federated learning.

References

- [1] ARIVAZHAGAN, M. G., AGGARWAL, V., SINGH, A. K., AND CHOUDHARY, S. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818* (2019).
- [2] BARBER, D., AND AGAKOV, F. The im algorithm: a variational approach to information maximization. *Advances in Neural Information Processing Systems 16*, 320 (2004), 201.
- [3] BELGHAZI, M. I., BARATIN, A., RAJESHWAR, S., OZAI, S., BENGIO, Y., COURVILLE, A., AND HJELM, D. Mutual information neural estimation. In *International Conference on Machine Learning* (2018), PMLR, pp. 531–540.
- [4] CHANG, Q., QU, H., ZHANG, Y., SABUNCU, M., CHEN, C., ZHANG, T., AND METAXAS, D. N. Synthetic learning: Learn from distributed asynchronous discriminator gan without sharing medical image data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2020), pp. 13856–13866.
- [5] CHEN, H.-Y., AND CHAO, W.-L. On bridging generic and personalized federated learning for image classification. *arXiv preprint arXiv:2107.00778* (2021).
- [6] CHENG, P., HAO, W., DAI, S., LIU, J., GAN, Z., AND CARIN, L. Club: A contrastive log-ratio upper bound of mutual information. In *International Conference on Machine Learning* (2020), PMLR, pp. 1779–1788.
- [7] COLLINS, L., HASSANI, H., MOKHTARI, A., AND SHAKKOTTAI, S. Exploiting shared representations for personalized federated learning. In *Proceedings of the 38th International Conference on Machine Learning* (18–24 Jul 2021), M. Meila and T. Zhang, Eds., vol. 139 of *Proceedings of Machine Learning Research*, PMLR, pp. 2089–2099.
- [8] DENG, J., DONG, W., SOCHER, R., LI, L.-J., LI, K., AND FEI-FEI, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition* (2009), IEEE, pp. 248–255.
- [9] DENG, R., DU, X., LU, Z., DUAN, Q., HUANG, S.-C., AND WU, J. Hsfl: Efficient and privacy-preserving offloading for split and federated learning in iot services. In *2023 IEEE International Conference on Web Services (ICWS)* (2023), IEEE, pp. 658–668.
- [10] DEVLIN, J., CHANG, M.-W., LEE, K., AND TOUTANOVA, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [11] DHRUVA, S. S., ROSS, J. S., AKAR, J. G., CALDWELL, B., CHILDERS, K., CHOW, W., CIACCIO, L., COPLAN, P., DONG, J., DYKHOFF, H. J., ET AL. Aggregating multiple real-world data sources using a patient-centered health-data-sharing platform. *NPJ digital medicine 3*, 1 (2020), 60.
- [12] DING, K., FENG, X., AND YU, H. To be global or personalized: Generalized federated learning with cooperative adaptation for data heterogeneity. *Knowledge-Based Systems 301* (2024), 112317.
- [13] DOSOVITSKIY, A., BEYER, L., KOLESNIKOV, A., WEISSENBORN, D., ZHAI, X., UNTERTHINER, T., DEGHANI, M., MINDERER, M., HEIGOLD, G., GELLY, S., ET AL. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [14] FALLAH, A., MOKHTARI, A., AND OZDAGLAR, A. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. In *Advances in Neural Information Processing Systems* (2020), H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33, Curran Associates, Inc., pp. 3557–3568.
- [15] FANG, Y., WANG, W., XIE, B., SUN, Q., WU, L., WANG, X., HUANG, T., WANG, X., AND CAO, Y. Eva: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 19358–19369.
- [16] GHIFARY, M., BALDUZZI, D., KLEIJN, W. B., AND ZHANG, M. Scatter component analysis: A unified framework for domain adaptation and domain generalization. *IEEE transactions on pattern analysis and machine intelligence 39*, 7 (2016), 1414–1430.
- [17] GUAN, H., YAP, P.-T., BOZOKI, A., AND LIU, M. Federated learning for medical image analysis: A survey. *Pattern Recognition* (2024), 110424.
- [18] HANZELY, F., HANZELY, S., HORVÁTH, S., AND RICHTÁRIK, P. Lower bounds and optimal algorithms for personalized federated learning. *Advances in Neural Information Processing Systems 33* (2020), 2304–2315.
- [19] HE, K., CHEN, X., XIE, S., LI, Y., DOLLÁR, P., AND GIRSHICK, R. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 16000–16009.
- [20] HUANG, Y., CHU, L., ZHOU, Z., WANG, L., LIU, J., PEI, J., AND ZHANG, Y. Personalized cross-silo federated learning on non-iid data. In *Proceedings of the AAAI Conference on Artificial Intelligence* (2021), vol. 35, pp. 7865–7873.
- [21] KAPLAN, J., MCCANDLISH, S., HENIGHAN, T., BROWN, T. B., CHESSE, B., CHILD, R., GRAY, S., RADFORD, A., WU, J., AND AMODEI, D. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361* (2020).
- [22] KINGMA, D. P., AND WELLMING, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [23] KRIZHEVSKY, A., HINTON, G., ET AL. Learning multiple layers of features from tiny images. *Technical Report* (2009).
- [24] LECUN, Y., BOTTOU, L., BENGIO, Y., AND HAFFNER, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE 86*, 11 (1998), 2278–2324.
- [25] LI, J., WU, R., SUN, W., CHEN, L., TIAN, S., ZHU, L., MENG, C., ZHENG, Z., AND WANG, W. What’s behind the mask: Understanding masked graph modeling for graph autoencoders. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2023), pp. 1268–1279.
- [26] LI, Q., HE, B., AND SONG, D. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 10713–10722.
- [27] LI, T., HU, S., BEIRAMI, A., AND SMITH, V. Ditto: Fair and robust federated learning through personalization. In *Proceedings of the 38th International Conference on Machine Learning* (18–24 Jul 2021), M. Meila and T. Zhang, Eds., vol. 139 of *Proceedings of Machine Learning Research*, PMLR, pp. 6357–6368.
- [28] LI, T., SAHU, A. K., TALWALKAR, A., AND SMITH, V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine 37*, 3 (2020), 50–60.
- [29] LUO, M., CHEN, F., HU, D., ZHANG, Y., LIANG, J., AND FENG, J. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. *Advances in Neural Information Processing Systems 34* (2021), 5972–5984.
- [30] MCMAHAN, B., MOORE, E., RAMAGE, D., HAMPSON, S., AND Y. ARCAS, B. A. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics* (2017), PMLR, pp. 1273–1282.
- [31] NGUYEN, D. C., DING, M., PATHIRANA, P. N., SENEVIRATNE, A., LI, J., AND POOR, H. V. Federated learning for internet of things: A comprehensive survey. *IEEE Communications Surveys & Tutorials 23*, 3 (2021), 1622–1658.
- [32] NIZUMI, D., TAKEUCHI, D., OHISHI, Y., HARADA, N., AND KASHINO, K. Masked spectrogram modeling using masked autoencoders for learning general-purpose audio representation. In *HEAR: Holistic Evaluation of Audio Representations* (2022), PMLR, pp. 1–24.
- [33] OH, J., KIM, S., AND YUN, S.-Y. FedBABU: Toward enhanced representation for federated image classification. In *International Conference on Learning Representations* (2022).
- [34] PENG, X., BAI, Q., XIA, X., HUANG, Z., SAENKO, K., AND WANG, B. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision* (2019), pp. 1406–1415.
- [35] T DINH, C., TRAN, N., AND NGUYEN, J. Personalized federated learning with moreau envelopes. *Advances in Neural Information Processing Systems 33* (2020), 21394–21405.
- [36] TAN, Y., LONG, G., LIU, L., ZHOU, T., LU, Q., JIANG, J., AND ZHANG, C. Fedproto: Federated prototype learning across heterogeneous clients. In *Proceedings of the AAAI Conference on Artificial Intelligence* (2022), vol. 36, pp. 8432–8440.
- [37] VINCENT, P., LAROCHELLE, H., LAJOIE, I., BENGIO, Y., MANZAGOL, P.-A., AND BOTTOU, L. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research 11*, 12 (2010).
- [38] WANG, H., JIA, Y., ZHANG, M., HU, Q., REN, H., SUN, P., WEN, Y., AND ZHANG, T. Feddse: Distribution-aware sub-model extraction for federated learning over resource-constrained devices. In *Proceedings of the ACM on Web Conference 2024* (2024), pp. 2902–2913.
- [39] WANG, K., HE, Q., CHEN, F., CHEN, C., HUANG, F., JIN, H., AND YANG, Y. Flexifed: Personalized federated learning for edge clients with heterogeneous model architectures. In *Proceedings of the ACM Web Conference 2023* (2023), pp. 2979–2990.
- [40] WU, C., WU, F., LYU, L., HUANG, Y., AND XIE, X. Communication-efficient federated learning via knowledge distillation. *Nature Communications 13*, 1 (2022), 2032.
- [41] XIAO, H., RASUL, K., AND VOLLGRAF, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017).
- [42] YI, L., WANG, G., LIU, X., SHI, Z., AND YU, H. Fedgh: Heterogeneous federated learning with generalized global header. In *Proceedings of the 31st ACM International Conference on Multimedia* (2023), pp. 8686–8696.
- [43] ZHANG, J., HUA, Y., WANG, H., SONG, T., XUE, Z., MA, R., CAO, J., AND GUAN, H. Gpfl: Simultaneously learning global and personalized feature information for personalized federated learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 5041–5051.
- [44] ZHANG, J., HUA, Y., WANG, H., SONG, T., XUE, Z., MA, R., AND GUAN, H. Fedala: Adaptive local aggregation for personalized federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence* (2023), vol. 37, pp. 11237–11244.
- [45] ZHANG, J., HUA, Y., WANG, H., SONG, T., XUE, Z., MA, R., AND GUAN, H. Fedcp: Separating feature information for personalized federated learning via conditional policy. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2023), pp. 3249–3261.
- [46] ZHANG, J., LI, S., HUANG, H., WANG, Z., FU, X., HONG, D., GUPTA, R. K., AND SHANG, J. How few davids improve one goliath: Federated learning in resource-skewed edge computing environments. In *Proceedings of the ACM on Web Conference 2024* (2024), pp. 2976–2985.
- [47] ZHANG, J., LIU, Y., HUA, Y., WANG, H., SONG, T., XUE, Z., MA, R., AND CAO, J. Pflib: Personalized federated learning algorithm library. *arXiv preprint arXiv:2312.04992* (2023).