

FedLPS: Heterogeneous Federated Learning for Multiple Tasks with Local Parameter Sharing

Yongzhe Jia¹, Xuyun Zhang^{2,*}, Amin Beheshti², Wanchun Dou^{1,*}

¹ State Key Laboratory for Novel Software Technology,
Department of Computer Science and Technology, Nanjing University, China

² School of Computing, Macquarie University, Australia
jiayz@smail.nju.edu.cn, {xuyun.zhang, amin.beheshti}@mq.edu.au, douwc@nju.edu.cn

Abstract

Federated Learning (FL) has emerged as a promising solution in Edge Computing (EC) environments to process the proliferation of data generated by edge devices. By collaboratively optimizing the global machine learning models on distributed edge devices, FL circumvents the need for transmitting raw data and enhances user privacy. Despite practical successes, FL still confronts significant challenges including constrained edge device resources, multiple tasks deployment, and data heterogeneity. However, existing studies focus on mitigating the FL training costs of each single task whereas neglecting the resource consumption across multiple tasks in heterogeneous FL scenarios. In this paper, we propose heterogeneous FEDerated learning with Local Parameter Sharing (FedLPS) to fill this gap. FedLPS leverages principles from transfer learning to facilitate the deployment of multiple tasks on a single device by dividing the local model into a shareable encoder and task-specific predictors. To further reduce resource consumption, a channel-wise model pruning algorithm that shrinks the footprint of local models while accounting for both data and system heterogeneity is employed in FedLPS. Additionally, a novel heterogeneous model aggregation algorithm is proposed to aggregate the heterogeneous predictors in FedLPS. We implemented the proposed FedLPS on a real FL platform and compared it with state-of-the-art (SOTA) FL frameworks. The experimental results on five popular datasets and two modern DNN models illustrate that the proposed FedLPS significantly outperforms the SOTA FL frameworks by up to 4.88% and reduces the computational resource consumption by 21.3%. Our code is available at: <https://github.com/jyzgh/FedLPS>.

Introduction

Over the past decade, there has been a remarkable surge in the generation of massive amounts of data from billions of Internet of Things (IoT) devices (Khan et al. 2021). In the context of Edge Computing (EC) environments, Federated Learning (FL) has emerged as a promising solution for processing such extensive data at the edge (McMahan et al. 2017; Khan et al. 2020). FL enables various edge devices to collaboratively optimize a global Machine Learning (ML) model with the assistance of an edge server. This approach

entails clients (i.e., edge devices) updating the ML model using their private local data, while the central server is responsible for aggregating these updated local models. In contrast to traditional centralized ML, FL facilitates data processing on distributed edge devices where the data is generated. In addition, the distributed learning nature of FL eliminates the need for transmitting raw data to the central server, thus avoiding unnecessary communication costs and enhancing user privacy (McMahan et al. 2017; Ezzeldin et al. 2023).

Despite the successful practical applications in FL, it still faces the following critical challenges: **a) Limited resource budgets of edge devices.** Popular Deep Neural Networks (DNNs) commonly possess larger than hundreds of megabytes of parameters that need to be trained (He et al. 2016), whereas the computational resources of edge devices are limited in EC environments (Kairouz et al. 2021). Training these unleashing DNN models on edge devices heavily hampered the learning efficiency of FL. **b) Multitasking on a single device.** Ideally, each edge device maintains a task-specific DNN model for each task, resulting in linear growth of training costs on edge devices. Therefore, directly deploying multiple unleashing DNN models for each specific task on edge devices is impractical (Fu et al. 2021; Ma et al. 2019; Wallingford et al. 2022). **c) Data heterogeneity and system heterogeneity.** On one hand, the numbers of data samples and the data distributions of edge devices are typically various (i.e., non-identically distributed data, non-IID data), resulting in the accuracy degradation of the global model (Gao, Yao, and Yang 2022; Luo et al. 2021). On the other hand, the system capabilities of edge devices, such as CPU, GPU, memory, battery power, etc., also can be various. The devices with weaker system capabilities (i.e., “stragglers”) will fail to finish local training, therefore hampering the learning process of FL (Gao, Yao, and Yang 2022; Jiang et al. 2022).

Several pioneering work have made efforts to mitigate these challenges through various solutions (Caldas et al. 2018; Li et al. 2020; Wang et al. 2020b; Jiang et al. 2022; Li et al. 2021). Li et al. (Li et al. 2020) propose FedProx to address the data heterogeneity problem by incentivizing participants to preserve similarity with the global model, while addressing the system heterogeneity problem by accommodating low-end devices (i.e., devices with fewer system capabilities) to carry out a reduced number of local up-

*Corresponding authors.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

dates. However, FedProx does not involve model compression techniques, making it difficult to save storage resources or communication overhead. Recent work (Jiang et al. 2022; Li et al. 2021) propose to leverage model pruning techniques and personalized model aggregation to reduce the footprint of the ML model meanwhile mitigating the model accuracy loss caused by the heterogeneity problems. However, these approaches focus only on each single ML model used for a specific task and thus contribute limitedly to reducing resource consumption across multiple tasks.

In this paper, we propose heterogeneous FEDerated learning with Local Parameter Sharing (FedLPS), a novel FL framework for reducing resource consumption of multiple tasks in the heterogeneous FL environment. Specifically, we leverage the spirit of transfer learning (Yosinski et al. 2014; Zhuang et al. 2020) to allow the multiple tasks on a single device to share partial parameters of the ML models during the FL process. In contrast to existing FL frameworks, the proposed FedLPS reduces the resource consumption of devices not only in the context of a single task but also in the context of multiple tasks. We adopt channel-wise model pruning techniques to reduce the footprint of the local models and satisfy the strict resource budget of edge devices. Different from pioneering work that use uniform model pruning techniques (e.g., FedDrop (Caldas et al. 2018)), FedLPS generates tailored models for each participant device to mitigate data heterogeneity problem and system heterogeneity problem. Moreover, considering that aggregating the tailored models with the popular FL aggregation algorithms (e.g., FedAvg (McMahan et al. 2017)) leads to the degradation of model performance, we further design a heterogeneous aggregation algorithm for FedLPS to generate the aggregated global model. We summarize our contributions as follows:

- We propose a novel FL framework FedLPS to reduce resource consumption of edge devices that deployed with multiple tasks. By dividing the local models on edge devices into shared encoders and task-specific predictors, FedLPS reduces the training cost across multiple tasks on edge devices.
- We design a channel-wise model pruning algorithm for FedLPS to reduce the footprints of the predictors. By applying various pruning ratios, FedLPS adaptively shrinks model footprints of heterogeneous devices.
- We present a heterogeneous model aggregation algorithm for FedLPS to aggregate heterogeneous task-specific predictors. By utilizing the knowledge within the pre-trained backbone model, FedLPS efficiently aggregates the heterogeneous predictors.
- We implement the proposed FedLPS in a real-world FL platform FedML and extensively evaluate FedLPS with state-of-the-art FL frameworks. The experimental results demonstrated that FedLPS is effective in reducing resource consumption of edge devices while realizing heightened model accuracy.

Related Work

Heterogeneous Federated Learning

In the context of edge computing, federated learning is proposed to train ML models with distributed local data among edge devices (McMahan et al. 2017; Khan et al. 2020; Gao, Yao, and Yang 2022; Zhang et al. 2023). In federated learning, the raw data on edge devices will be kept locally to provide better user privacy and avoid unnecessary communication costs, and only intermediate results (e.g., parameters of models) are transmitted between the server and the devices (Khan et al. 2021). Federated learning in the edge computing environments is commonly heterogeneous in several aspects such as statistics, systems, data spaces, and models (Gao, Yao, and Yang 2022). In this paper, we focus on two of the main heterogeneous aspects: statistical heterogeneity and system heterogeneity. Statistical heterogeneity leads to non-IID distribution of data among edge devices, while system heterogeneity leads to variations in the capabilities of these devices.

Transfer Learning

Transfer learning is a promising machine learning methodology for transferring knowledge across different domains following different probability distributions (Yosinski et al. 2014; Long et al. 2017; Zhuang et al. 2020; Tan et al. 2023). Di et al. (Di et al. 2017) propose to transfer the knowledge of images that are taken from a certain location, aiming to alleviate the adverse impact caused by various conditions such as different weather and illumination conditions in transportation applications. Yu et al. (Yu et al. 2022) propose SPATL for addressing the resource consumption problem and data heterogeneity problem in FL. In SPATL, transfer learning is adopted to address the data heterogeneity problem by transferring the knowledge of a shared encoder to the predictors on heterogeneous clients. Tu et al. (Tu et al. 2021) propose FedDL to capture the potential relationships between users and transfer knowledge between the related users in FL, aiming to improve the performance of Human Activity Recognition (HAR) task with unbalanced and sparse user data. However, few existing work explore the transferability of ML models for multiple tasks on edge devices in FL. In contrast, we fill this gap in this paper and demonstrate the feasibility of leveraging transfer learning to reduce resource consumption of edge devices deployed with multiple tasks.

Model Pruning

Model pruning techniques are proposed to accelerate the training and inference processes of DNN by removing the redundant parameters and structures in the DNN model (Liu et al. 2017; He, Zhang, and Sun 2017; Li et al. 2021; Jiang et al. 2022; Ye et al. 2023). Liu et al. (Liu et al. 2017) propose the network slimming scheme for Convolutional Neural Networks (CNNs), aiming to identify and remove insignificant parameters in CNNs by imposing L1 regularization on the scaling factors in batch normalization layers. Caldas et al. (Caldas et al. 2018) propose FedDrop to reduce the computational burden of local training and the cor-

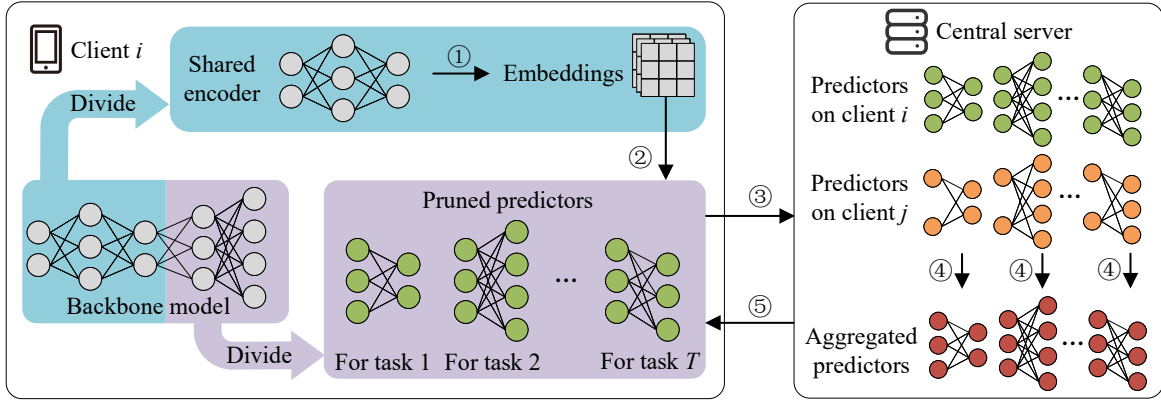


Figure 1: Overview of the proposed FedLPS framework. In FedLPS, the backbone model within each client is divided into the shared encoder and task-specific predictors. The predictors subsequently pruned to reduce resource consumption. During the training process, the encoder parameters remain frozen, while the predictor parameters are updated to handle specific tasks and transmitted between the central server and the client. To elaborate, ① local data for each task is fed into the encoder to generate embeddings. ② The task-specific predictors utilize these embeddings to update their parameters. ③ The client sends the updated predictors to the central server. ④ The central server aggregates the predictors that have been updated on different clients but belong to the same task. ⑤ The central server sends the aggregated predictors back to the clients for further training rounds.

responding communication costs of FL. FedDrop leverages lossy compression techniques to shrink the footprint of the ML model and generates identical compact local models for all devices. Jiang et al. (Jiang et al. 2022) propose FedMP to address system heterogeneity problem meanwhile saving communication bandwidth. FedMP adopts a multi-armed bandit-based online learning algorithm to calculate personalized pruning ratios for heterogeneous edge devices and a Residual Recovery Synchronous Parallel (R2SP) scheme to aggregate parameters. However, most of the existing work is on model pruning of native ML models on edge devices, and very little work has focused on transferable models applicable to multi-task scenarios. In contrast, we design an adaptive channel-wise model pruning algorithm for the transferable models in FedLPS to reduce unnecessary resource consumption in multi-task scenarios.

Design of FedLPS

In this section, we first present an overview of the proposed FedLPS. Then, we describe how transfer learning can be used to train models for multiple tasks. Subsequently, we employ an adaptive channel-wise model pruning approach to reduce the resource consumption caused by training task-specific predictors. Finally, we present a heterogeneous model aggregation algorithm for aggregating heterogeneous predictors updated by different clients.

Overview

In this paper, we propose FedLPS to efficiently train multiple task-specific models on individual clients in the con-

text of FL. Fig. 1 illustrates the workflow of the proposed FedLPS framework. Distinct from existing FL frameworks that optimize a specialized model for each task on the client, in FedLPS, the backbone model w^i in individual client $i \in \mathcal{C}$ is partitioned into a shared encoder w_e^i , and $|\mathcal{T}|$ task-specific predictors $w_p^{i,t}$ for each task $t \in \mathcal{T}$, where \mathcal{T} denote the task set. The backbone model can be either trained on the public dataset or trained on local data of arbitrary tasks on the client and is accessible to both the client and the central server. During the federated training process, the encoder w_e^i will be frozen while the task-specific predictor $w_p^{i,t}$ will be updated on the local data of the task t .

We outline the federated training steps as follows: ① The local data for each task is fed into the encoder w_e^i to generate embeddings. ② The task-specific predictors $w_p^{i,t}$ utilize these embeddings generated by local data of tasks t to update to update their parameters. ③ Each client i sends the updated predictors $w_p^{i,t}$ to the central server. ④ The central server aggregates the predictors that have been updated on different clients (e.g., predictors on client i and predictors on client j in Fig. 1) but belong to the same task t . ⑤ The central server sends the aggregated predictors back to the clients for further training rounds.

Local Parameter Sharing across Multiple Tasks

Existing FL frameworks (e.g., FedAvg (McMahan et al. 2017), FedDrop (Caldas et al. 2018), Hermes (Li et al. 2021)) commonly adopt an approach where a specific model is optimized for each FL task on the client. However, this practice of optimizing multiple models on the clients leads

to significant resource consumption, especially for edge devices with limited capabilities. Although some efforts (Li et al. 2020; Jiang et al. 2022) have been made to reduce the training overhead on the client, few of them focus on effectively reducing training overhead across multiple FL tasks. In this subsection, we propose a novel federated training method that leverages local parameter sharing on the clients to mitigate the resource consumption associated with training multiple models for multiple tasks, meanwhile maintaining a satisfying model accuracy.¹

Inspired by the spirit of the transfer learning (Zhuang et al. 2020; Weiss, Khoshgoftaar, and Wang 2016; Yosinski et al. 2014), we explore enabling the multiple tasks on each client to share a part of model parameters, thus effectively reducing the training overhead. In parameter-sharing-based transfer learning, the lower layers of the neural networks capture more generalized features, making them suitable for sharing across multiple tasks. Conversely, the upper layers tend to capture higher-level abstract features, making them more task-specific. Building upon this observation, we divide the local backbone model w on each client into a shareable encoder w_e and multiple task-specific predictors w_p^t .² The shared encoder w_e comprises the first n layers of the backbone model, while the task-specific predictors w_p^t consist of the remaining $N - n$ layers, where n is a tunable hyper-parameter and the N represents the total number of layers in the backbone model. In order to facilitate knowledge transfer across various tasks, the weights of shared encoder w_e are initialized using pre-trained values. Throughout the training process, the weights in shared encoder w_e remains frozen, ensuring consistent utilization across all tasks. The pre-trained weights can be sourced either from a backbone model trained on publicly available datasets or from client-participated FL tasks. In addition, to reduce the resource consumption caused by training multiple predictors, an adaptive channel-wise model pruning method is proposed for shrinking the footprints of the predictors, which is described in the next subsection in detail.

The local training algorithm based on local parameter sharing of the proposed FedLPS framework is presented in Algorithm 1. In each communication round of FL, the client first prunes each task-specific predictor w_p^t with a pruning ratio ρ , aiming to reduce subsequent training costs (in line 2). Then, the client conducts forward propagation on the shared encoder w_e using local data \mathcal{D}_t specific to each task k . The resulting embedding is denoted as e^t (in line 3). Subsequently, each pruned predictor $w_p^{t'}$ is updated with the embedding e^t (in line 4). Specifically, the update operation can be formulated as follows:

$$w_p^{t'} \leftarrow w_p^{t'} - \eta \nabla \ell_t(w_p^{t'}), \quad (1)$$

where η is the learning rate, ℓ_t is the loss function of task t , and $\nabla \ell_t(w_p^{t'})$ is the local gradients of the predictors $w_p^{t'}$.

¹In this work, we make the fundamental assumption that the model structure remains consistent across all tasks, and this work does not focus on addressing the issue of model heterogeneity.

²For presentation convenience, we omit the superscript i when describing local operation of each client.

Algorithm 1: Local Parameter Sharing-based Training Algorithm of FedLPS.

Input: Task set \mathcal{T} , local data $\mathcal{D}_t (t \in \mathcal{T})$, pre-trained encoder w_e , pruning ratio ρ for pruning predictors

Output: Updated predictors $w_p^{t'}$

- 1: **for** task $t \in \mathcal{T}$ **do**
 - 2: $w_p^{t'} \leftarrow$ Prune the predictor w_p^t with pruning ratio ρ using Eq. 2
 - 3: $e^t \leftarrow$ Forward propagation on shared encoder w_e with local data \mathcal{D}_t
 - 4: $w_p^{t'} \leftarrow$ Update weights in predictor $w_p^{t'}$ using Eq. 1
 - 5: **end for**
 - 6: **return** Updated predictors $w_p^{t'}$
-

Finally, the updated task-specific predictors $w_p^{t'}$ will be sent to the central server for aggregation.

Adaptive Channel-wise Model Pruning

By leveraging the division of the backbone model and the freezing of the pre-trained shared encoder, FedLPS effectively mitigates the training cost associated with the shared encoder. However, the training costs incurred by the predictors continue to pose challenges for clients due to the generally larger footprints associated with these predictors compared to the encoder. Thus, in this subsection, we proposed an adaptive channel-wise model pruning method for FedLPS to reduce the training cost of the task-specific predictors.

In the first communication round of FL, each client prunes the predictors with a punning ratio ρ . Differing from existing model pruning methods in FL that adopt a uniform punning ratio for each client like FedDrop (Caldas et al. 2018), FedLPS prunes the predictors on heterogeneous clients with different punning ratios ρ that are determined by the clients' system capability. Firstly, for each task $t \in \mathcal{T}$ on the client, FedLPS evaluates the importance scores of channels in each layer of the predictor w_p^t by L1-norm (Liu et al. 2017; Li et al. 2016). Then, a fraction ρ of channels corresponding to the smallest importance scores is removed to achieve model pruning. Specifically, the pruning operation can be formulated as follows:

$$w_p^{t'} \leftarrow w_p^t \odot M^t, M^t \in \{0, 1\}^{|w_p^t|}, \quad (2)$$

where \odot denotes the element-wise multiplication, and M^t is a binary mask matrix used to determine the channels to be pruned. In the mask matrix M^t , elements with a value of 0 indicate channels that will be pruned, while elements with a value of 1 indicate channels that will be retained. Finally, the pruned predictors $w_p^{t'}$ will be updated with Eq. 1.

Heterogeneous Predictor Aggregation

The adaptive pruning operation on the predictors is capable to shrink their footprint and thereby reducing the resource consumption of the heterogeneous clients. However, the adaptive pruning operation leads to multiple heterogeneous predictors that can not be aggregated by popular FL aggregation algorithms (e.g., FedAvg (McMahan et al.

Algorithm 2: Heterogeneous Predictor Aggregation Algorithm of FedLPS.

Input: Pruned predictors $w_p^{i,t'} (i \in \mathcal{C}, t \in \mathcal{T})$, mask matrices $M^{i,t}$, backbone predictor w_p

Output: Aggregated predictors w_p^{t*}

- 1: **for** task $t \in \mathcal{T}$ **do**
- 2: **for** client $i \in \mathcal{C}$ **do**
- 3: $w_p^{i,t} \leftarrow$ Recover pruned predictor $w_p^{i,t'}$ using backbone predictor w_p and mask matrix $M^{i,t}$ (Eq. 3)
- 4: **end for**
- 5: $w_p^{t*} \leftarrow$ Aggregate the recovered predictors of task t using Eq. 4
- 6: **end for**
- 7: **return** Aggregated predictors w_p^{t*}

2017)). Existing heterogeneous aggregation algorithms such as Hermes (Li et al. 2021) and FedMP (Jiang et al. 2022), aggregate the overlapped parameters of the heterogeneous models. Unfortunately, these algorithms result in a degradation in model performance due to the pruned parameters lacking any significant contribution to the aggregated global model. Moreover, the remained parameters are difficult to learn knowledge from other clients when these parameters are pruned in other clients. Thus, in this subsection, we propose a novel aggregation algorithm that leverages the knowledge within the pre-trained backbone model to aggregate the heterogeneous predictors.

The heterogeneous predictor aggregation algorithm of the proposed FedLPS framework is presented in Algorithm 2. In each FL communication round, the central server receives $|\mathcal{T}|$ task-specific predictors from each selected client. FedLPS first recovers the pruned parameters in each task-specific predictor $w_p^{i,t'}$ using the backbone predictor w_p extracted from the backbone model and the mask matrix $M^{i,t}$ on task t of client i (in line 3). The recovery operation can be formulated as follows:

$$w_p^{i,t} \leftarrow w_p^{i,t'} + [w_p - (w_p \odot M^{i,t})], \quad (3)$$

where $w_p^{i,t}$ represents the recovered predictor containing the updated weights on task t of client i .³ Subsequently, FedLPS can aggregate these recovered predictors of task t by weighted averaging (in line 5). Formally, the aggregation operation can be represented as follows:

$$w_p^{t*} \leftarrow \sum_{i \in \mathcal{C}} \frac{|\mathcal{D}_{i,t}|}{|\mathcal{D}_t|} w_p^{i,t}, \quad (4)$$

where \mathcal{D}_t denotes the entire data of task t across all clients and $\mathcal{D}_{i,t}$ denotes the local data of task t on the client i . For each task $t \in \mathcal{T}$, FedLPS aggregates the predictor $w_p^{i,t}$ using Eq. 4 until all predictors have been aggregated.

Finally, the central server sends the aggregated global predictors $w_p^{i,t}$ back to the selected clients for the further round of local training.

³It is noteworthy that these updated weights differ from the initial predictor, despite using the same symbol.

Experimental Evaluation

In this section, we first implement the proposed FedLPS framework on a real federated learning platform FedML (He et al. 2020) and conduct a comprehensive performance comparison of FedLPS against five state-of-the-art (SOTA) frameworks in multiple tasks-enabled FL environments. Next, we evaluate the effect of varying pruning ratios on the learning performance of the FedLPS. Finally, we evaluate the effect of the layer number n of the shared encoder on the learning performance of the FedLPS.

Experimental Setting

FL environments. In our experiments, we simulate 10 heterogeneous clients and deploy 5 classification tasks on each client. Both the ResNet18 model (He et al. 2016) and the ShuffleNetV2 model (Zhang et al. 2018) are adopted to conduct these classification tasks. The heterogeneous client set \mathcal{C} are uniformly divided into 5 levels according to their system capabilities, denoted as $\mathcal{C} = \{\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \mathcal{C}_4, \mathcal{C}_5\}$. The system capabilities are decreasing from \mathcal{C}_1 to \mathcal{C}_5 linearly. Only the clients in \mathcal{C}_1 can conduct these 5 classification tasks without resource optimizing techniques (i.e., model pruning). All the heterogeneous clients are selected to perform FL training. The experiments are conducted on a GPU server with 2 NVIDIA RTX 3080Ti GPUs, and each experiment is executed three times for calculating average metrics.

Datasets and data partition. In our experiments, we adopt five widely recognized datasets: MNIST (LeCun et al. 1998), FashionMNIST (Xiao, Rasul, and Vollgraf 2017), SVHN (Netzer et al. 2011), CIFAR10 (Krizhevsky, Hinton et al. 2009), and CIFAR100 (Krizhevsky, Hinton et al. 2009) datasets to simulate the classification tasks. In the IID setting, each dataset is equally assigned to the clients. In the non-IID setting, we use the Latent Dirichlet Allocation (LDA) (Luo et al. 2021; Wang et al. 2020a) method to build the non-IID data. In LDA, a concentration parameter α is used to control the data heterogeneity. We adopt the conventional setting of $\alpha = 0.5$ in our experiments to construct the non-IID data.

Comparison frameworks. We compare the proposed FedLPS framework with FedAvg (McMahan et al. 2017), FedDrop (Caldas et al. 2018), FedProx (Li et al. 2020), Hermes (Li et al. 2021), and FedMP (Jiang et al. 2022). FedAvg is the classical FL framework that needs each client to update the entire model, thus can involve only the clients in \mathcal{C}_1 to perform FL training in the heterogeneous FL environment. FedDrop leverages model pruning techniques to generate a compact global model that can be updated by the clients in \mathcal{C}_5 , thus involving all clients to participate in the FL training. FedProx encourages local models to maintain similarity with the global model by introducing a regularization term meanwhile allowing the clients in $\{\mathcal{C}_2, \mathcal{C}_3, \mathcal{C}_4, \mathcal{C}_5\}$ to perform fewer local updates. Hermes extracts a tailored sub-model for each client by structured pruning and aggregates only the intersection of the local models. FedMP prunes each local model with a dynamic pruning ratio in each round thereby enabling all clients to participate in the FL training.

Data partition	FL frameworks	MNIST	Fashion-MNIST	SVHN	CIFAR10	CIFAR100	Average
IID partition	FedAvg(McMahan et al. 2017)	97.07	86.35	89.32	64.15	25.27	72.43
	FedDrop(Caldas et al. 2018)	77.92	73.26	37.28	56.77	23.97	53.84
	FedProx(Li et al. 2020)	96.63	86.10	89.41	73.88	38.65	76.93
	Hermes(Li et al. 2021)	97.92	87.73	91.24	76.32	38.92	78.43
	FedMP(Jiang et al. 2022)	97.08	87.29	88.45	74.20	38.78	77.16
	FedLPS (Ours)	97.99	87.94	91.39	78.71	49.66	81.14
Non-IID partition	FedAvg(McMahan et al. 2017)	56.03	59.36	79.23	26.39	20.04	48.21
	FedDrop(Caldas et al. 2018)	77.55	61.81	36.37	41.91	23.06	48.14
	FedProx(Li et al. 2020)	93.77	83.16	82.96	62.14	37.17	71.84
	Hermes(Li et al. 2021)	95.32	82.53	86.29	60.09	38.26	72.50
	FedMP(Jiang et al. 2022)	95.60	84.21	82.34	61.59	37.19	72.19
	FedLPS (Ours)	96.67	83.39	88.00	71.15	47.70	77.38

Table 1: Comparison of model accuracy (%) on the ShuffleNetV2 model with both IID and non-IID data.

Data partition	FL frameworks	MNIST	Fashion-MNIST	SVHN	CIFAR10	CIFAR100	Average
IID partition	FedAvg(McMahan et al. 2017)	98.38	88.82	93.74	76.74	35.06	78.55
	FedDrop(Caldas et al. 2018)	88.59	83.19	69.37	58.39	23.41	64.59
	FedProx(Li et al. 2020)	98.60	90.30	95.05	85.31	54.37	84.73
	Hermes(Li et al. 2021)	98.67	91.32	95.46	81.58	54.54	84.31
	FedMP(Jiang et al. 2022)	97.67	89.25	91.25	79.29	46.10	80.71
	FedLPS (Ours)	98.85	92.15	95.48	86.98	57.87	86.27
Non-IID partition	FedAvg(McMahan et al. 2017)	65.82	66.38	84.58	31.52	27.59	55.18
	FedDrop(Caldas et al. 2018)	82.19	77.24	69.05	42.55	19.73	58.15
	FedProx(Li et al. 2020)	95.87	84.07	91.05	67.60	53.18	78.35
	Hermes(Li et al. 2021)	97.05	84.00	92.12	71.81	53.53	79.70
	FedMP(Jiang et al. 2022)	96.60	83.78	87.66	67.42	42.53	75.60
	FedLPS (Ours)	96.98	86.45	91.98	74.17	55.56	81.08

Table 2: Comparison of model accuracy (%) on the ResNet18 model with both IID and non-IID data.

To conduct fair comparisons, we adopt the same training hyper-parameters for FedLPS and the comparison frameworks in our experiments. The training hyper-parameters are provided in the supplemental materials in detail.

Learning Performance across Multiple Tasks

We compare the model accuracy of the FedLPS with FedAvg, FedDrop, FedProx, Hermes, and FedMP on both IID and non-IID data. In this experiment, the layer number of the shared encoder in FedLPS is set as $n = \lceil 0.25N \rceil$, i.e., the first $\lceil 0.25N \rceil$ layers within the backbone model are used to build the shared encoder. For FedLPS, Hermes, and FedMP frameworks, the pruning ratios ρ adopted by the 5 levels of clients are $\{0, 0.2, 0.4, 0.6, 0.8\}$, respectively. The backbone model is pre-trained on ImageNet (Russakovsky et al. 2015).

Table 1 shows the accuracy on the ShuffleNetV2 model with both IID and non-IID data, while Table 2 shows the accuracy on the ResNet18 model. Additionally, we provide a comparative analysis of communication overhead

in the supplemental materials, assessed through the footprints of the transmitted models. On both the ShuffleNetV2 and ResNet18 models, FedLPS outperforms the comparison frameworks in terms of average model accuracy. The superiority of FedLPS can be attributed to three reasons: firstly, the pre-trained shareable encoder is sophisticated to extract the low-level features and produce general embedding from the local data. Secondly, the task-specific predictors are pruned elaborately to satisfy the resource constraints, and the pruned predictors are trained on each task separately therefore suitable for performing the specific tasks. Thirdly, the heterogeneous predictor aggregation algorithm used in FedLPS leverages the knowledge in the backbone model to assist the aggregation of the local predictors, thus making the predictors learn from other clients better.

Effect of Pruning Ratios

In the model pruning-enabled FL frameworks, the resource consumption of the training local model is significantly re-




Figure 2: Model accuracy of Hermes, FedMP, and FedLPS on the ResNet18 model with different pruning ratios ρ on the non-IID setting of the CIFAR10 and CIFAR100 datasets.

duced, whereas the model accuracy will decrease when the pruning ratio exceeds a threshold. Thus, we investigate the effect of the pruning ratio on model accuracy in this subsection. In this experiment, the proposed FedLPS framework is compared with two SOTA model pruning-enabled frameworks Hermes and FedMP, and the pruning ratio ρ is set as 0.2, 0.4, 0.6, and 0.8. Fig. 2(a) and Fig. 2(b) show the model accuracy of the ResNet18 model on the non-IID setting of the CIFAR10 dataset and the CIFAR100 dataset during 100 FL rounds, respectively. FedLPS outperforms Hermes and FedMP when the pruning ratio ρ ranges from 0.2 to 0.8 in both the CIFAR10 dataset and the CIFAR100 dataset, although all of them adopt model pruning techniques to shrink the model footprint. The detailed footprints of the predictors within FedLPS are provided in the supplemental materials.

Effect of Local Parameter Sharing

In FedLPS, the larger the layer number n assigned to the shared encoder, the more training costs can be saved. In this experiment, we investigate the effect of the layer number n of the shared encoder on the training cost and model accuracy. The layer number n is set as $\lceil 0.25N \rceil$, $\lceil 0.5N \rceil$, and $\lceil 0.75N \rceil$, respectively. The floating-point operations (FLOPs) of the model are measured to indicate the training cost. In this experiment, the ResNet18 model containing 34.36M FLOPs is adopted. Fig. 3 shows the model accuracy of MNIST, FashionMNIST, SVHN, CIFAR10, and CIFAR100 datasets while Table 3 reports the FLOPs of the sum of the shared encoder and all task-specific predictors (before adaptive pruning). In FedLPS, the FLOPs linearly decrease as the layer number n increases, resulting in a substantial reduction in training costs. Thus, FedLPS are more suitable for clients that are deployed with multiple tasks than existing FL frameworks as the training cost of multiple tasks can be easily reduced by building a shared encoder with n -layer. Take the clients in \mathcal{C}_1 that deployed with 5 tasks as an example, in Fig. 3 and Table 3, FedLPS achieves 86.27% and 81.08% of average accuracy with 21.3% fewer FLOPs (i.e., $n = \lceil 0.25N \rceil$) in the IID setting and non-IID setting, respectively. Whereas the SOTA frameworks only realize 84.31% and 79.70% of average accuracy (in Table 2) with 171.8M FLOPs ($5 \times$ FLOPs of the original ResNet18 model).

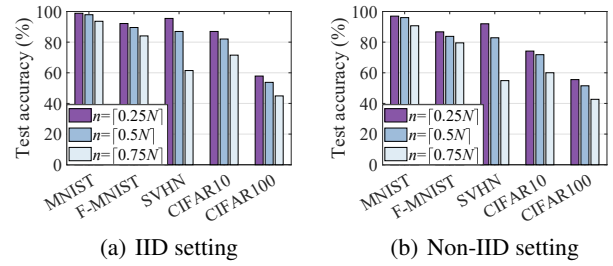


Figure 3: Model accuracy of FedLPS with different layer numbers n on the IID setting and the non-IID setting. Where F-MNIST represents FashionMNIST dataset.

Layer number n	FLOPs	Reduction on FLOPs	Average accuracy(%)
$n = \lceil 0.25N \rceil$	135.2M	-21.3%	86.27(81.08)
$n = \lceil 0.5N \rceil$	101.5M	-40.9%	82.04(77.20)
$n = \lceil 0.75N \rceil$	67.9M	-60.5%	71.12(65.57)

Table 3: Total FLOPs of FedLPS with different layer numbers n and corresponding model average accuracy on both IID data and non-IID data. (The accuracy of IID data is outside the parentheses and the accuracy of non-IID data is inside the parentheses.)

Conclusion

In this paper, we have proposed FedLPS for multiple-tasks-enabled heterogeneous FL environments, aiming to reduce the resource consumption of the clients during the FL training process while maintaining satisfying model accuracy. FedLPS realizes local parameter sharing by dividing the local model into a shareable encoder and multiple task-specific predictors, thus achieving the reduction of training costs across multiple tasks on individual clients. To tackle the system heterogeneity problem, an adaptive channel-wise model pruning method is proposed for FedLPS to allow the heterogeneous clients to participate in the FL training with heterogeneous task-specific predictors. Furthermore, a novel aggregation algorithm is proposed for FedLPS to efficiently aggregate the heterogeneous predictors with the assistance of the knowledge within the pre-trained backbone model. The comparison results on five popular datasets and two modern DNN models demonstrated the superiority of the FedLPS in terms of both average model accuracy and resource consumption.

Limitations and Prospects for Future Research. The hyper-parameter n employed in FedLPS currently lacks the flexibility for dynamic adjustments during FL training, a promising future direction involves delving into the nuanced interactions among the layer number n , the pruning ratio ρ , and the ensuing model accuracy. Moreover, the scope of this study is confined to utilizing FedLPS exclusively for classification tasks. To extend the applicability of FedLPS, an exciting direction involves comprehensively exploring its performance across diverse task domains.

Acknowledgments

This research is supported part by the National Key Research and Development Program of China No.2020YFB1707601, National Natural Science Foundation of China No.92267104.

References

- Caldas, S.; Konečný, J.; McMahan, H. B.; and Talwalkar, A. 2018. Expanding the reach of federated learning by reducing client resource requirements. *arXiv preprint arXiv:1812.07210*.
- Di, S.; Zhang, H.; Li, C.-G.; Mei, X.; Prokhorov, D.; and Ling, H. 2017. Cross-domain traffic scene understanding: A dense correspondence-based transfer learning approach. *IEEE transactions on intelligent transportation systems*, 19(3): 745–757.
- Ezzeldin, Y. H.; Yan, S.; He, C.; Ferrara, E.; and Avestimehr, A. S. 2023. Fairfed: Enabling group fairness in federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 7494–7502.
- Fu, C.; Huang, H.; Chen, X.; Tian, Y.; and Zhao, J. 2021. Learn-to-share: A hardware-friendly transfer learning framework exploiting computation and parameter sharing. In *International Conference on Machine Learning*, 3469–3479. PMLR.
- Gao, D.; Yao, X.; and Yang, Q. 2022. A Survey on Heterogeneous Federated Learning. *arXiv preprint arXiv:2210.04505*.
- He, C.; Li, S.; So, J.; Zhang, M.; Wang, H.; Wang, X.; Vepakomma, P.; Singh, A.; Qiu, H.; Shen, L.; Zhao, P.; Kang, Y.; Liu, Y.; Raskar, R.; Yang, Q.; Annavaram, M.; and Avestimehr, S. 2020. FedML: A Research Library and Benchmark for Federated Machine Learning. *arXiv preprint arXiv:2007.13518*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- He, Y.; Zhang, X.; and Sun, J. 2017. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision*, 1389–1397.
- Jiang, Z.; Xu, Y.; Xu, H.; Wang, Z.; Qiao, C.; and Zhao, Y. 2022. FedMP: Federated Learning through Adaptive Model Pruning in Heterogeneous Edge Computing. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, 767–779.
- 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.
- Khan, L. U.; Pandey, S. R.; Tran, N. H.; Saad, W.; Han, Z.; Nguyen, M. N.; and Hong, C. S. 2020. Federated learning for edge networks: Resource optimization and incentive mechanism. *IEEE Communications Magazine*, 58(10): 88–93.
- Khan, L. U.; Saad, W.; Han, Z.; Hossain, E.; and Hong, C. S. 2021. Federated Learning for Internet of Things: Recent Advances, Taxonomy, and Open Challenges. *IEEE Communications Surveys & Tutorials*, 23(3): 1759–1799.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images. *Technical report*.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Li, A.; Sun, J.; Li, P.; Pu, Y.; Li, H.; and Chen, Y. 2021. Hermes: an efficient federated learning framework for heterogeneous mobile clients. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, 420–437.
- Li, H.; Kadav, A.; Durdanovic, I.; Samet, H.; and Graf, H. P. 2016. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2: 429–450.
- Liu, Z.; Li, J.; Shen, Z.; Huang, G.; Yan, S.; and Zhang, C. 2017. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE international conference on computer vision*, 2736–2744.
- Long, M.; Zhu, H.; Wang, J.; and Jordan, M. I. 2017. Deep transfer learning with joint adaptation networks. In *International conference on machine learning*, 2208–2217. PMLR.
- Luo, M.; Chen, F.; Hu, D.; Zhang, Y.; Liang, J.; and Feng, J. 2021. No Fear of Heterogeneity: Classifier Calibration for Federated Learning with Non-IID Data. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 5972–5984. Curran Associates, Inc.
- Ma, J.; Zhao, Z.; Chen, J.; Li, A.; Hong, L.; and Chi, E. H. 2019. Snr: Sub-network routing for flexible parameter sharing in multi-task learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 216–223.
- 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.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; and Ng, A. Y. 2011. Reading Digits in Natural Images with Unsupervised Feature Learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; Berg, A. C.; and Li, F. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3): 211–252.
- Tan, Y.; Liu, Y.; Long, G.; Jiang, J.; Lu, Q.; and Zhang, C. 2023. Federated learning on non-iid graphs via structural knowledge sharing. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 9953–9961.

- Tu, L.; Ouyang, X.; Zhou, J.; He, Y.; and Xing, G. 2021. Feddl: Federated learning via dynamic layer sharing for human activity recognition. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, 15–28.
- Wallingford, M.; Li, H.; Achille, A.; Ravichandran, A.; Fowlkes, C.; Bhotika, R.; and Soatto, S. 2022. Task adaptive parameter sharing for multi-task learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7561–7570.
- Wang, H.; Yurochkin, M.; Sun, Y.; Papailiopoulos, D. S.; and Khazaeni, Y. 2020a. Federated Learning with Matched Averaging. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Wang, J.; Liu, Q.; Liang, H.; Joshi, G.; and Poor, H. V. 2020b. Tackling the Objective Inconsistency Problem in Heterogeneous Federated Optimization. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 7611–7623. Curran Associates, Inc.
- Weiss, K.; Khoshgoftaar, T. M.; and Wang, D. 2016. A survey of transfer learning. *Journal of Big data*, 3(1): 1–40.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Ye, H.; Zhang, B.; Chen, T.; Fan, J.; and Wang, B. 2023. Performance-aware Approximation of Global Channel Pruning for Multitask CNNs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Yosinski, J.; Clune, J.; Bengio, Y.; and Lipson, H. 2014. How transferable are features in deep neural networks? In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Yu, S.; Nguyen, P.; Abebe, W.; Qian, W.; Anwar, A.; and Jannesari, A. 2022. SPATL: salient parameter aggregation and transfer learning for heterogeneous federated learning. In *2022 SC22: International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 495–508. IEEE Computer Society.
- Zhang, J.; Hua, Y.; Wang, H.; Song, T.; Xue, Z.; Ma, R.; and Guan, H. 2023. Fedala: Adaptive local aggregation for personalized federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11237–11244.
- Zhang, X.; Zhou, X.; Lin, M.; and Sun, J. 2018. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6848–6856.
- Zhuang, F.; Qi, Z.; Duan, K.; Xi, D.; Zhu, Y.; Zhu, H.; Xiong, H.; and He, Q. 2020. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1): 43–76.