

Learning Federated Visual Prompt in Null Space for MRI Reconstruction

Chun-Mei Feng¹ Bangjun Li² Xinxing Xu^{1,*} Yong Liu^{1,*} Huazhu Fu¹ Wangmeng Zuo³

¹Institute of High Performance Computing (IHPC),

Agency for Science, Technology and Research (A*STAR), Singapore

²School of Information Science and Engineering, Shangdong University, China

³Harbin Institute of Technology, Harbin, China

strawberry.feng0304@gmail.com, {xuxinx, liuyong}@ihpc.a-star.edu.sg

<https://github.com/chunmeifeng/FedPR>

Abstract

Federated Magnetic Resonance Imaging (MRI) reconstruction enables multiple hospitals to collaborate distributedly without aggregating local data, thereby protecting patient privacy. However, the data heterogeneity caused by different MRI protocols, insufficient local training data, and limited communication bandwidth inevitably impair global model convergence and updating. In this paper, we propose a new algorithm, FedPR, to learn federated visual prompts in the null space of global prompt for MRI reconstruction. FedPR is a new federated paradigm that adopts a powerful pre-trained model while only learning and communicating the prompts with few learnable parameters, thereby significantly reducing communication costs and achieving competitive performance on limited local data. Moreover, to deal with catastrophic forgetting caused by data heterogeneity, FedPR also updates efficient federated visual prompts that project the local prompts into an approximate null space of the global prompt, thereby suppressing the interference of gradients on the server performance. Extensive experiments on federated MRI show that FedPR significantly outperforms state-of-the-art FL algorithms with $< 6\%$ of communication costs when given the limited amount of local training data.

1. Introduction

Federated Magnetic Resonance Imaging (MRI) reconstruction enables multiple hospitals to train a powerful global model in a distributed manner without sharing private data [7, 9, 13]. In federated MRI, each client (*i.e.*, hospital) uses its local computing power, memory, and private data to train local models independently, while the server aggregates all the local models in each communica-

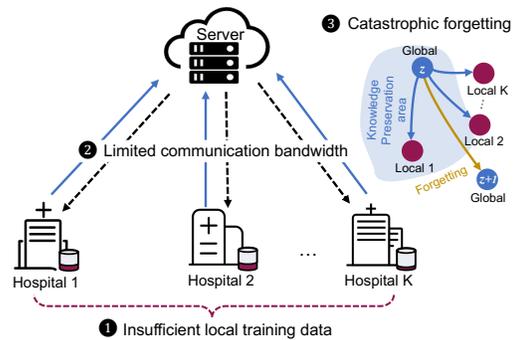


Figure 1. **Illustration** of the three **key issues** in federated MRI.

tion round and distributes the global model to each client again [26].

Existing federated MRI reconstruction techniques usually improve federated learning (FL) by *enhancing the aggregation process* [9] and *reducing the local parameter's variance* [5, 13]. Such techniques require a large communication bandwidth and sufficient local training data. However, federated MRI often faces two issues, **1** *insufficient amount of local training data* due to the difficulty of acquiring the ground-truth of MRI reconstruction [10, 11, 12] and **2** *limited communication bandwidth* due to unbalanced regional development (see Fig. 1). To cope with the issue **1**, *pre-trained* models have exhibited superior performance, and have shown to *close* the gap between federated and centralized performance [3, 28]. However, since the model parameters need to be shared between the client and the server for updating, the large number of parameters of pre-trained models will result in a *huge communication cost*. On contrary, prompt tuning has recently been suggested as a new fine-tuning paradigm by freezing the models and only learning a small number of learnable parameters in the input space [15, 31, 32]. Benefited from pre-trained models, only parameter-efficient prompts are required in learning and communication, and prompt tuning can be conducted with a limited number of samples, making it very appealing in tackling the above two issues (*i.e.*,

*Corresponding author.

❶ and ❷) in federated MRI.

Besides, there is another critical issue for federated MRI, *i.e.*, ❸ *catastrophic forgetting*, caused by data heterogeneity due to the different imaging protocols of MRI scanners adopted by different clients [9, 39] (see Fig. 1). In the local update, the global model from the prior round tends to be overfitted to the local training data, leading to catastrophic forgetting. Analogous to continual learning, several techniques have been presented by introducing proper regularization terms in local models [4, 33, 39] or retaining previously acquired knowledge by knowledge distillation [14, 35, 38]. However, these strategies require seeking a balance between multiple losses and relying on proxy datasets. Instead, Adam-NSC mitigates catastrophic forgetting in continual learning with a new network parameter update rule, *i.e.*, forces the network update to lie in the approximate null space of the input features of previous tasks at each layer [36]. However, the null space at feature level is built upon large local training data, which generally cannot be satisfied in federated MRI, *i.e.*, the issue ❶. Taking the issues ❶ and ❷ into account, we suggest to optimize local prompts in the approximate null space of global prompts instead of input features, thereby preventing the loss of previously gained global knowledge.

In a nutshell, we explore a new FL paradigm, *i.e.*, FedPR, to learn federated visual prompts for MRI reconstruction. To begin with, we pre-train the model on a large-scale public dataset. Given limited amount of local training data, visual prompt tuning is adopted to learn local prompts in a *distributed* manner. For the issues ❶ and ❷, *federated visual prompts* are introduced to learn a strong global model, where only local and global prompts are learnable and communicated. As for the issue ❸, we perform singular value decomposition (SVD) on the uncentered covariance matrix of global prompts to obtain the approximate *null space*. FedPR tunes only the local parameters in the null space of global prompts, thereby well preserving the *prior knowledge* of previous rounds and resulting in low *communication costs*. In particular, FedPR achieves a > 4.5 dB gain in PSNR with less than 6% of communication costs. To sum up, our contributions are as follows:

- We propose a *federated visual prompt* algorithm, FedPR, to solve the three key issues in federated MRI. By leveraging powerful pre-trained models and freezing backbone networks in FL, only a small amount of parameters in the input space are trainable, thereby reducing communication costs.
- We explore how to alleviate *catastrophic forgetting* in FL while reducing communication costs. By optimizing local parameters only in the null space of global prompts, FedPR well preserves the previously acquired global knowledge in each round, maintaining competitive performance with only a few local data.

- We evaluate the performance of FedPR for federated MRI reconstruction. In comparison to the state-of-the-art FL methods, FedPR achieves superior performance in complex scenarios, *e.g.*, less local data, lower communication costs, and faster convergence.

2. Related Work

Federated MRI Reconstruction. MRI reconstruction refers to reconstructing images without aliasing artifacts from undersampled k -space data [6, 8, 12, 37, 41]. However, existing MRI reconstruction techniques are based on large-scale paired data, which is not only labor-intensive but also violates patient privacy [9]. Driven by these realistic problems, a few FL-based methods for MRI reconstruction have been proposed [5, 9, 13]. Guo *et al.* reduced data heterogeneity by iteratively aligning the data distribution between source and target clients [13]. Feng *et al.* proposed a personalized FL reconstruction scheme and introduced a weighted contrast regularization term to correct the update direction of global generalization [9]. However, these schemes either require frequent communication or are built on a large amount of local data. As a result, these techniques will produce sub-optimal solutions for MRI reconstruction when there is *limited bandwidth* and *insufficient local training data*. In contrast to these works, we build a federated MRI model upon a strong *pre-trained* model by updating and communicating only a *small* portion of parameters for each client using *limited* local data, achieving competitive results.

Catastrophic Forgetting in FL. Due to the heterogeneous data distributions of different tasks in continual learning, when the network is fitted to the current task, the model parameters typically deviate from the areas where the previous knowledge is expected to be preserved [17]. Therefore, catastrophic forgetting is a fundamental challenge of continual learning. Similarly, catastrophic forgetting also affects FL [14, 19, 39]. When the client receives the global model and continues to update locally, it will cause the global model to forget the knowledge of the previous round due to the data heterogeneity across clients [19]. As a result, existing FL techniques that address catastrophic forgetting are designed by referring to continual learning. For example, knowledge distillation-based FL methods maintain previously gained knowledge, but this method strongly relies on proxy data [14, 35, 38]; regularization-based FL methods regularize locally trained parameters, but they require finding a compromise between loss terms [4, 33, 39]. All these studies are built on classification tasks and cannot be directly extended to MRI reconstruction. Inspired by the success of null space in continual learning [18, 23, 36], we try to update the *local* prompt in the approximate *null space* of global prompts, thereby preserving the knowledge acquired by the global model in the previous round and avoid-

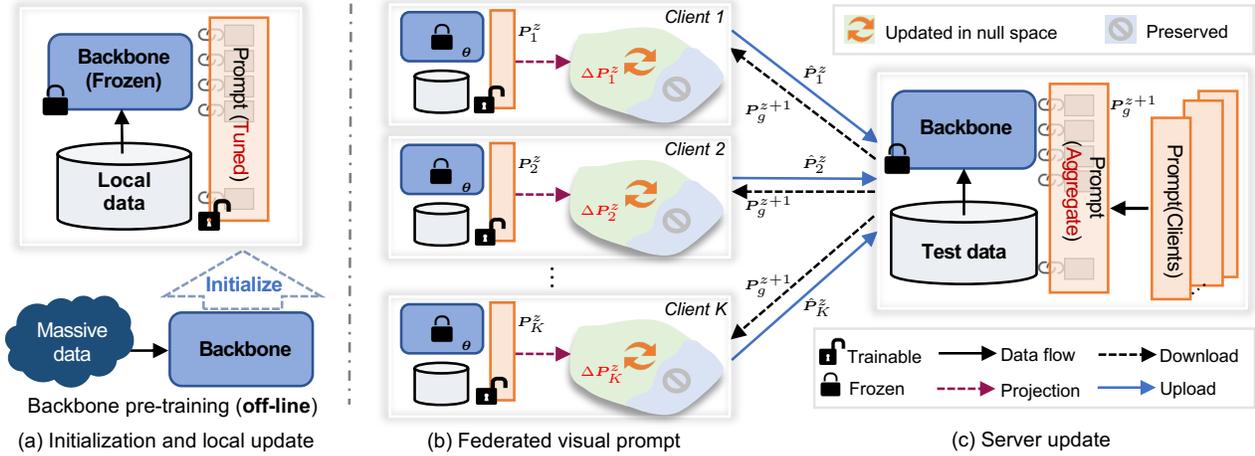


Figure 2. **Illustration** of our FedPR method. (a) For each client, the backbone model is pre-trained on massive data and fine-tuned via prompt (see Sec. 3.2). (b) *Federated visual prompt* is executed by updating the local prompt of each client only in the approximate null space of global prompts while preserving the *previously* acquired global knowledge (see Sec. 3.3). (c) Server update by aggregating local prompts.

ing *catastrophic forgetting*.

Prompt Learning. Prompt tuning was initially proposed to enable pre-trained language models to “understand” downstream tasks by adding trainable tokens to the input text [24]. For example, Jia *et al.* achieved good performance on downstream tasks by introducing a small number of learnable tokens into the input space and keeping the backbone frozen [15]. Hyojin *et al.* suggested to modify pixel space for the frozen visual models [2]. Nie *et al.* proposed a unified prompt tuning framework that performs on different CNN and transformer-based architectures by training only a few additional parameters [29]. Although these works have made progress on various visual tasks, prompt is still limited to centralized systems, and the effectiveness of prompt in distributed framework remains uninvestigated. Therefore, this work focuses on how to learn a *federated visual prompt* to effectively tackle the three key problems in federated MRI reconstruction, *i.e.*, the issues ❶–❸.

3. Methodology

3.1. Federated MRI Reconstruction

MRI reconstruction is an inverse problem of recovering an artifact-free image \mathbf{y} from its undersampled observation \mathbf{x} that can reduce the online collection time and improve the patient experience [34, 37, 40]. Formally, such undersampling process can be expressed as

$$\mathbf{x} = \mathcal{F}^{-1}(\mathcal{M} \odot \mathcal{F}(\mathbf{y}) + \epsilon), \quad (1)$$

where \mathcal{F} is multi-dimensional Fourier transform, ϵ is the measurement noise, and \mathcal{M} is the binary mask operator to undersample the data points in the Fourier space.

The centralized approach violates patient privacy protection roles because it requires collecting large amounts of training data from different hospitals. As a remedy, Federated MRI has been suggested by deploying MRI reconstruction in a distributed manner [9, 13]. Suppose there are K

hospitals (local clients). The private data of all hospitals can be expressed as $\mathcal{D} = \{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^K\}$, where each contains limited pairs of undersampled samples \mathbf{x}_i and fully-sampled images \mathbf{y}_i . Federated MRI aims to learn a global model from the whole dataset \mathcal{D} in a distributed manner, which can be described as

$$\arg \min_{\mathbf{w}} \mathcal{L}(\mathbf{w}) = \sum_{k=1}^K \frac{|\mathcal{D}^k|}{|\mathcal{D}|} \mathcal{L}_k(\mathbf{w}), \quad (2)$$

where $|\mathcal{D}|$ denotes the number of samples in \mathcal{D} , and $\mathcal{L}_k(\mathbf{w})$ is the empirical loss of client k ,

$$\mathcal{L}_k(\mathbf{w}) = \mathbf{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}^k} \ell_k(f(\mathbf{x}; \mathbf{w})), \quad (3)$$

where ℓ_k represents the local loss for MRI reconstruction, *e.g.*, L_1 loss. After training, the reconstructed image $\hat{\mathbf{y}}$ can be produced by $f(\mathbf{x}; \mathbf{w})$.

3.2. Learning Federated Prompt

As mentioned in Sec. 1, there are three key issues, *i.e.*, ❶–❸, with the existing federated MRI algorithms. To tackle the issue ❷, we freeze the backbone of the pre-trained model while only tuning and communicating a few learnable parameters for the clients and server. More importantly, it enables the model to achieve very competitive results on limited local data, thereby providing an effective solution to issue ❶.

As shown in Fig. 2, given a pre-trained model with parameters θ , we introduce a set of continuous embeddings $\mathbf{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_l\}$ as the prompts in the input space of each layer [15], where l is the number of prompts. Thus, the overall parameters can be expressed as $\mathbf{w} = \{\mathbf{P}, \theta\}$. Suppose there are Z rounds of communication, each round contains T local updates. During the local update, for the k -th client, only the client-specific prompts \mathbf{P}_k are learnable and communicated while the backbone network is frozen. Therefore, Eq. (2) can be rewritten as:

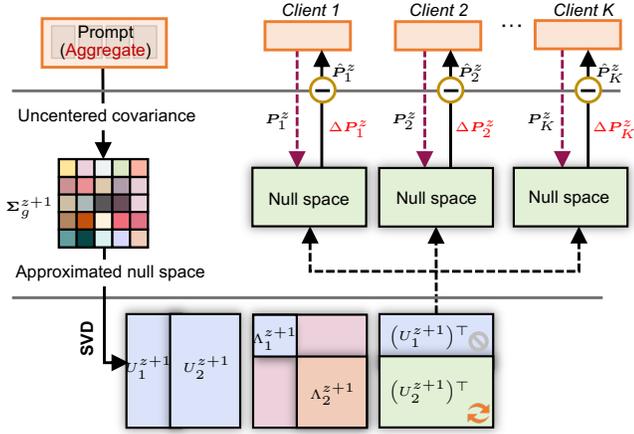


Figure 3. **Illustration of federated visual prompt** for updating local prompts in the null space of global prompts.

$$P_k = \arg \min_P \mathcal{L}(P) = \sum_{k=1}^K \frac{|\mathcal{D}^k|}{|\mathcal{D}|} \mathcal{L}_k(P), \quad (4)$$

where the empirical loss for prompt tuning of client k can be expressed as

$$\mathcal{L}_k(P) = \mathbf{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}^k} \ell_k(f(\mathbf{x}; P, \theta)). \quad (5)$$

Benefited from pre-trained models, federated prompt tuning can be both parameter- and communication-efficient and effective in achieving competitive results with a small amount of local data, thereby serving as a favorable solution to the issues ① and ②.

Local Update Step: In each communication round $z = \{1, 2, \dots, Z\}$, the clients are optimized using the following update rules with a learning rate of η_k :

$$P_k^{z,t+1} \leftarrow P_k^{z,t} - \eta_k \nabla \ell_k(\mathbf{x}^k; P_k^{z,t}), \quad (6)$$

where t denotes the t -th update of the local clients.

Server Update Step: After a round of local updates, all participating clients send their updated prompts P_k^z to the server performing aggregation. Such process can be expressed as follows

$$P_g^{z+1} \leftarrow \sum_{k=1}^K \frac{|\mathcal{D}^k|}{|\mathcal{D}|} P_k^z, \quad (7)$$

where P_g^{z+1} denotes the global prompts of round $z + 1$. To sum up, the issue ② can be largely mitigated because there are only a small number of learnable parameters P communicated between the server and local clients. After Z rounds of communication, we can get a robust global model parameterized by P_g without sharing local private data.

3.3. Local Updating in Prompting Null Space

The local update mechanism in Eq. (6), however, suffers from the issue ③, *i.e.*, catastrophic forgetting, due to data heterogeneity of different clients. Inspired by the continual learning method Adam-NSC [36] that sequentially optimizes network parameters in the null space at feature level, we suggest to optimize the local model in the null space of

Algorithm 1: FedPR

Input: Private datasets from K clients: $\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^K$, local updates T , communication rounds Z , pre-trained model parameters θ , prompt embeddings P , learning rate η , hyperparameter γ ;

```

1 // ServerExecution:
2 Initialize global prompt  $P_g$  with parameters  $\theta$ .
3 for each communication round  $z \in \{1, 2, \dots, Z\}$  do
4   for each client  $k \in \{1, 2, \dots, K\}$  in parallel do
5      $P_k^z \leftarrow P_g^z$ ;
6      $\hat{P}_k^z \leftarrow \text{LocalUpdate}(k, P_k^z)$ ;
7   end
8    $P_g^{z+1} \leftarrow \sum_{k=1}^K \frac{|\mathcal{D}^k|}{|\mathcal{D}|} \hat{P}_k^z$ ;
9   Compute the uncentered covariance matrix:
10   $\Sigma_g^{z+1} = (P_g^{z+1})^\top P_g^{z+1}$ ;
11  Approximate the null space of  $\Sigma_g^{z+1}$ :
12   $\Sigma_g^{z+1} = U^{z+1} \Lambda^{z+1} (U^{z+1})^\top$ ;
13  Select the smallest diagonal singular values of
14   $\Lambda_2^{z+1}$  with the ratio of  $\gamma$ ;
15  Get  $U_2^{z+1}$  corresponding to  $\Lambda_2^{z+1}$ .
16 end
17 return  $P_g^{z+1}$ 
18 // LocalUpdate ( $k, P_k$ ):
19 for each local epoch  $t \in \{1, 2, \dots, T\}$  do
20   Get the updated parameters  $\Delta P_k^z$ :
21    $\Delta P_k^z = U_2^{z+1} (U_2^{z+1})^\top P_k^z$ ;
22    $\hat{P}_k^{z,t+1} \leftarrow P_k^{z,t} - \eta_k \Delta P_k^{z,t}$ ;
23 end
24 return  $\hat{P}_k^z$ 

```

global prompts. Specifically, we train the local model in the approximate null space of the global prompts in a distributed way, preventing knowledge outside the local distribution from being overwritten and resulting in catastrophic forgetting. In comparison to [36], approximating the null space of global prompts instead of features does not require a large amount of local data and not increase the communication cost, thus offering a feasible way to tackle the three issues ①-③ in federated MRI.

As shown in Fig. 3, we first compute the uncentered covariance matrix of $z+1$ round global prompt P_g^{z+1} by

$$\Sigma_g^{z+1} = (P_g^{z+1})^\top P_g^{z+1}. \quad (8)$$

Then, we find the approximate null space of Σ_g^{z+1} by apply SVD to it

$$\Sigma_g^{z+1} = U^{z+1} \Lambda^{z+1} (U^{z+1})^\top, \quad (9)$$

where

$$\begin{aligned} U^{z+1} &= [U_1^{z+1}, U_2^{z+1}], \\ \Lambda^{z+1} &= \begin{bmatrix} \Lambda_1^{z+1} & 0 \\ 0 & \Lambda_2^{z+1} \end{bmatrix}. \end{aligned} \quad (10)$$

Here U_1^{z+1} are the singular vectors corresponding to the large singular values in Λ_1^{z+1} [36]. According to principal component analysis (PCA), U_1^{z+1} can be considered as principal components, which contain the *global knowledge* that has been captured in the previous round. Thus, we have $\Sigma_g^{z+1} \approx U_1^{z+1} \Lambda_1^{z+1} (U_1^{z+1})^\top$. And the range space of U_2^{z+1} can be denoted as the approximate null space of Σ_g^{z+1} [36]. In our implementation, we introduce a $\gamma\%$ to control U_2^{z+1} by selecting the last $\gamma\%$ of singular values in Λ_2^{z+1} to constitute U_2^{z+1} .

To tackle the issue $\textcircled{9}$, P_k^z can be projected into the approximate null space of Σ_g^{z+1} to get the updated parameters, while the principal components U_1^{z+1} containing the *global knowledge* from the previous round are *preserved*. Thus, the updated parameters ΔP_k^z can be optimized in the null space of Σ_g^{z+1} by

$$\Delta P_k^z = U_2^{z+1} (U_2^{z+1})^\top P_k^z, \quad (11)$$

where $U_2^{z+1} (U_2^{z+1})^\top$ is the projection operator [27]. Formally, the local update step of Eq. (6) can be rewritten as

$$\hat{P}_k^{z,t+1} \leftarrow P_k^{z,t} - \eta_k \Delta P_k^{z,t}. \quad (12)$$

The detailed FedPR algorithm is given in Algorithm 1.

4. Experiments

4.1. Experimental Setup

Implementation Details. Our method is trained by Pytorch with one NVIDIA Tesla V100 GPU and 32GB of memory. We use Adam as the optimizer, with a momentum of 0.9 and weight decay of 0.0005. The models are trained with 50 communication rounds and 10 local epochs for each round. We set the batch size and initial learning rate to 8 and 1×10^{-1} , respectively. The hyperparameter γ is empirically set to 80%. The prompt embeddings are added with a size of $8 \times 20 \times 256$.

Datasets. We use **fastMRI**¹, the largest publicly available MRI dataset, to train the pre-training model, with Swin Transformers and a standard ConvNet head serving as the backbone of our experiments [25]. We note that the mean and std of BN layer in the head are also shared with the server to update the original statistical properties. For the clients, we employ four datasets with a size of 320×320 , including **FeTS**² [30], **IXI**³ [1], and two clinical datasets,

¹<https://fastmri.org/>.

²<https://www.synapse.org/#!/Synapse:syn28546456/wiki/>.

³<https://brain-development.org/>

to comprise our local data. In our experiments, we divide them into 15 clients according to the acquiring institutions, *i.e.*, **FeTS** is divided into 10 clients with only 120, 304, 304, 240, 160, 160, 280, 224, 264, and 184 images, respectively; **IXI** is divided into 3 clients with only 360, 328, and 296 images, respectively; two clinical datasets collected from the United Imaging Healthcare uMR 790 scanner and 3T Siemens Magnetom Skyra system form two clients with 432 and 360 images, respectively. The input data of each client employs a *1D random* sampling pattern with $3 \times$ acceleration. We note that our clients have only a small number of images. In comparison, FedMRI [9] and FL-MRCM [13] require thousands to tens of thousands of images. For In-Federation evaluation, we divide each dataset with a ratio of 7 : 3 for local train/test. For Out-of-Federation evaluation, we use the test set of **FeTS** as our test set because it comes from a separate institution with a quite different distribution.

Baselines. To demonstrate the effectiveness of our proposed method, we compare it with three categories of methods, including: **a)** SingleSet, a single model that each client is trained with their local data without FL; **b)** Centralized, a single model that is trained with the combination of all the local data, which serves as the upper-bound of FL models; and **c)** eight state-of-the-art FL algorithms, including: (1) FedAvg [26], a classical FL algorithm that is trained by averaging parameters of all the participating clients; (2) FedBN [22], a FL algorithm that alleviates the client-shift by using batch normalization on each local client; (3) FedProx [21], a FL algorithm that applies a proximal term to the local objective function; (4) SCAFFOLD [16], a FL algorithm that uses control variates to correct the client-drift; (5) MOON [20], a FL algorithm that utilizes the similarity between model representations to correct the local update; (6) FedReg [39], a FL algorithm that regularizes locally trained parameters with the loss on generated pseudo data to alleviate knowledge forgetting; (7) FL-MRCM [13], a federated MRI algorithm that aligns the latent features between the source and target clients; and (8) FedMRI [9], a federated MRI algorithm that preserves the client-specific properties to improve the FL performance. For a fair comparison, we adopt Swin Transformers and a standard ConvNet head as the reconstruction network for each client among all the baselines.

4.2. Comparison with State-of-the-arts

In-Federation Performance. The first subtable of Table 1 provides a comprehensive evaluation of the In-Federation setting with regards to various FL algorithms, where *SingleSet* indicates that each client is trained to use their local data without FL, and *Centralized* indicates that all the local datasets are gathered, which serves as the upper-bound of FL models. **Ours** (*w. Pro*) and **Ours** (*Full*) are the variants of our method that only with prompts and

Table 1. **Quantitative comparison** of state-of-the-art FL methods with regard to In-Federation and Out-of-Federation scenarios, where # **Com.cost** is the communication cost, Ub indicates the upper-bound of FL algorithms, \uparrow and \downarrow indicate increments and decrements compared with FedAvg (w. FFt). Detailed analyses are provided in Sec. 4.2.

Method	# Com.cost	In-Federation			Out-of-Federation		
		PSNR	SSIM	NMSE	PSNR	SSIM	NMSE
SingleSet	18.43 M	28.79(3.35) \downarrow	0.805(0.099) \downarrow	0.021(0.011) \downarrow	28.92(2.15) \downarrow	0.808(0.091) \downarrow	0.032(0.011) \downarrow
Centralized (Ub)	18.43 M	36.71(4.57) \uparrow	0.947(0.044) \uparrow	0.009(0.005) \uparrow	35.73(4.66) \uparrow	0.940(0.041) \uparrow	0.008(0.013) \uparrow
FedAvg (w. FFt) [26]	18.43 M	32.14(0.00)	0.903(0.000)	0.010(0.000)	31.07(0.00)	0.899(0.000)	0.021(0.000)
FedBN [22]	18.40 M	28.31(3.83) \downarrow	0.817(0.086) \downarrow	0.044(0.034) \downarrow	25.65(5.42) \downarrow	0.731(0.186) \downarrow	0.080(0.059) \downarrow
FedProx [21]	18.43 M	32.84(0.70) \uparrow	0.901(0.002) \uparrow	0.010(0.000)	31.98(0.91) \uparrow	0.905(0.006) \uparrow	0.018(0.003) \uparrow
SCAFFOLD [16]	18.43 M	33.08(0.94) \uparrow	0.914(0.010) \uparrow	0.009(0.001) \uparrow	32.20(1.13) \uparrow	0.915(0.016) \uparrow	0.018(0.003) \uparrow
MOON [20]	18.43 M	34.06(1.92) \uparrow	0.927(0.023) \uparrow	0.008(0.002) \uparrow	31.16(0.09) \uparrow	0.907(0.008) \uparrow	0.023(0.002) \downarrow
FedReg [39]	18.43 M	33.29(1.15) \uparrow	0.890(0.013) \uparrow	0.009(0.001) \uparrow	32.41(1.34) \uparrow	0.907(0.008) \uparrow	0.017(0.004) \uparrow
FL-MRCM [13]	18.43 M	33.60(1.46) \uparrow	0.922(0.019) \uparrow	0.013(0.003) \downarrow	32.72(1.65) \uparrow	0.911(0.012) \uparrow	0.016(0.005) \uparrow
FedMRI [9]	17.46 M	33.35(1.21) \uparrow	0.923(0.020) \uparrow	0.014(0.004) \downarrow	32.00(0.93) \uparrow	0.914(0.015) \uparrow	0.019(0.002) \uparrow
Ours (w. Pro)	0.11 M	35.29(3.15) \uparrow	0.927(0.024) \uparrow	0.009(0.001) \uparrow	34.65(3.58) \uparrow	0.921(0.022) \uparrow	0.010(0.011) \uparrow
Ours (Full)	0.11 M	36.43 (4.30) \uparrow	0.945 (0.042) \uparrow	0.007 (0.003) \uparrow	35.60 (4.53) \uparrow	0.939 (0.040) \uparrow	0.008 (0.013) \uparrow

our full method that with both prompts and the null space mechanism. FedAvg (w. FFt) [26] indicates the FedAvg algorithm while employing the pretrained model on the local side and full fine-tuning with all the participating clients. For a fair comparison, we retrain them on the two different scenarios with their default parameters and report their optimal results. All the competing methods are pretrained with Swin Transformers on **fastMRI** and fully fine-tuned with 50 communication rounds and 10 local epochs because they all converge before round 50. To assess the difference between a pair of methods, we use the paired Student’s t -test to demonstrate that all results were at the 0.001 level of statistical significance.

From Table 1, **Ours** (w. Pro) is superior to the competing FL methods in terms of PSNR, SSIM, and NMSE. In comparison to FedAvg (w. FFt), **Ours** (w. Pro) increases the PSNR results from 32.14 dB to 35.29 dB and the SSIM results from 0.903 to 0.927, because the FedAvg (w. FFt) is prone to overfitting when the local data is limited (*i.e.*, issue ❶) even though FedAvg (w. FFt) also employ the pretrained model on the local side. However, the method of freezing backbone networks that only fine-tune prompts is consistent with the mechanism of personalized FL, *e.g.*, FedMRI, which divides the local network into a shared global generalized representation and a locally client-specific representation. More importantly, the number of training parameters and communication cost of **Ours** (w. Pro) is only 6% of FedAvg (w. FFt). Moreover, even compared to FedReg [39], which aims to address catastrophic forgetting in FL, our method still achieves 9.43% improvement in terms of PSNR. Furthermore, the reconstruction results of our proposed method are almost on par with the *Centralized* results (upper bound), *i.e.*, PSNR: 36.71 dB vs. **36.41** dB, and SSIM: 0.947 vs. **0.945**.

In contrast, SingleSet has the lowest results, even though

it is also fine-tuned on the pre-trained model, as it cannot solve non-i.i.d. problems. For the classical federated MRI methods, FL-MRCM [13] and FedMRI [9], they lead to poor reconstruction performance when the number of local clients is large and the amount of local data is insufficient. Furthermore, other FL algorithms are either designed for data heterogeneity problems or for catastrophic forgetting problems. They have not considered the clinical issues of federated MRI, *i.e.*, issues ❶-❸, causing it to produce less accurate reconstructions. These findings confirm our core idea that the three issues of federated MRI can be relieved by learning a federated visual prompt with our federated MRI paradigm. In addition, we visualize the reconstructed images and the corresponding error maps for all the competing methods in the first two rows of Fig. 4. The fewer textures in the error map, the better the reconstruction. It is obvious from the error map in Fig. 4 that our method can significantly reduce the reconstruction error.

Out-of-Federation Performance. The second sub-table of Table 1 provides quantitative evaluations under the Out-of-Federation scenario of various FL algorithms. Since the testing distribution in the Out-of-Federation scenario is unseen by the training distribution, the overall results of the Out-of-Federation scenario are slightly lower than those of In-Federation. However, our method still achieves the best performance in terms of the reconstruction results on the three metrics, *i.e.*, **Ours** (Full) improves PSNR values from 31.07 dB to **35.60** dB, raises SSIM values from 0.899 to **0.939**, and decreases NMSE values from 0.021 to **0.008**. Notably, our method even achieves comparable results to the *Centralized* (Upper-bound) method, *i.e.*, PSNR: 35.73 dB vs. **35.60** dB, SSIM: 0.940 vs. **0.939** and NMSE: 0.008 vs. **0.008**. That is, our method suffers the least from the issue ❹ among all competing

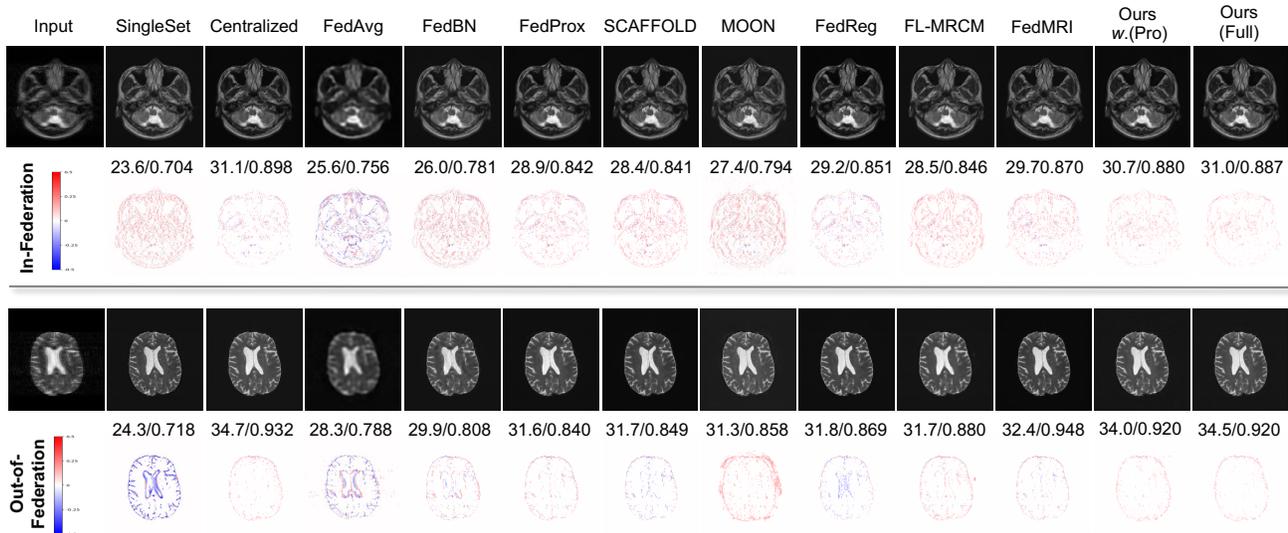


Figure 4. **Qualitative comparison** of different algorithms in terms of reconstruction images and error maps with corresponding quantitative measurements in PSNR/SSIM under In-Federation and Out-of-Federation scenarios. The less texture in the error map, the better reconstruction quality.

algorithms. These results support our core idea that learning federated visual prompts in the null space of global prompts can alleviate the three key issues in federated MRI while achieving competitive performance on limited local data. In addition, we show the qualitative evaluation results for the Out-of-Federation scenario in the second rows of Fig. 4, which visualizes the high-precision results of our method. As can be seen from this figure, our method provides comparable effects to the *Centralized* one, with clear details of tissue texture and boundaries.

4.3. Ablation Studies

Catastrophic Forgetting vs. Local Updates. The catastrophic forgetting (*i.e.*, issue ⑤) of federated MRI is caused by data heterogeneity due to different imaging protocols. As a result, more local updates will cause the local model to deviate from the global optimal solution. Here, we explore whether more local epochs cause the model to forget previously acquired global knowledge. In other words, it needs to be explored whether our proposed method can force local mitigation of this forgetting. Thus, we record the reconstruction results in terms of PSNR, SSIM, and NMSE for different numbers of local epochs in Fig. 5 (a). From this figure, in comparison to FedAvg (*w.* FFT) [26], increasing the local epoch does not greatly affect the accuracy of our method, *i.e.*, **Ours** (*w.* Pro), and **Ours** (Full). In particular, as the number of local epochs increases, the performance of our method gradually stabilizes. In contrast, when the number of local epochs is too large, the reconstruction accuracy of FedAvg (*w.* FFT) will be reduced because it will gradually deviate from the global knowledge area. As a result, the result indicates that our method can learn without forgetting global knowledge over more local epochs.

Analysis of Ratio $\gamma\%$. As we mentioned in Sec. 3.3, the ratio $\gamma\%$ of the smallest singular values in Λ_2^{z+1} controls

the size of the approximate null space in global prompts. That is, the larger the $\gamma\%$, the smaller the area of knowledge preservation. However, the area of knowledge preservation contains the global knowledge from the previous round, which is expected to be unchanged. To analyze the knowledge preservation capability that a good federated MRI framework should have, we discuss the reconstruction accuracy for different $\gamma\%$ values with regard to two scenarios in Fig. 6. It can be seen from this figure that, with the increase in $\gamma\%$ values, the reconstruction accuracy gradually increases. However, when $\gamma\% = 100\%$, the reconstruction accuracy drops rapidly. This is because when the knowledge preservation area is reduced to zero, fitting on the local distribution leads to forgetting the out-of-local distribution, *i.e.*, the issue ⑥. Additionally, we can observe from Fig. 6 that roughly 40% of the area comprises previously acquired global knowledge because our method obtains the highest results at this point. Following [36], we use the proportion of the sum of singular values of Λ_2^{z+1} in the sum of singular values of Λ^{z+1} to verify the rationality of the approximation. We find that the proportion in each layer is smaller than 10^{-7} . That is, the sum of the smallest singular values U_2^{z+1} can be ignored, thereby making it reasonable to approximate the null space through the spatial range of U_2^{z+1} . Interestingly, our method only performs local updates in the approximate null space of global prompts, which greatly improves the stability of the model and thus prevents catastrophic forgetting, *i.e.*, the issue ⑥.

Communication Efficiency Analysis. Our primal interest is to learn a federated visual prompt that enables FL to perform well with less local data, lower communication cost, and faster convergence. Here, we investigate the communication efficiency of our proposed method in terms of communication cost and different communication rounds,

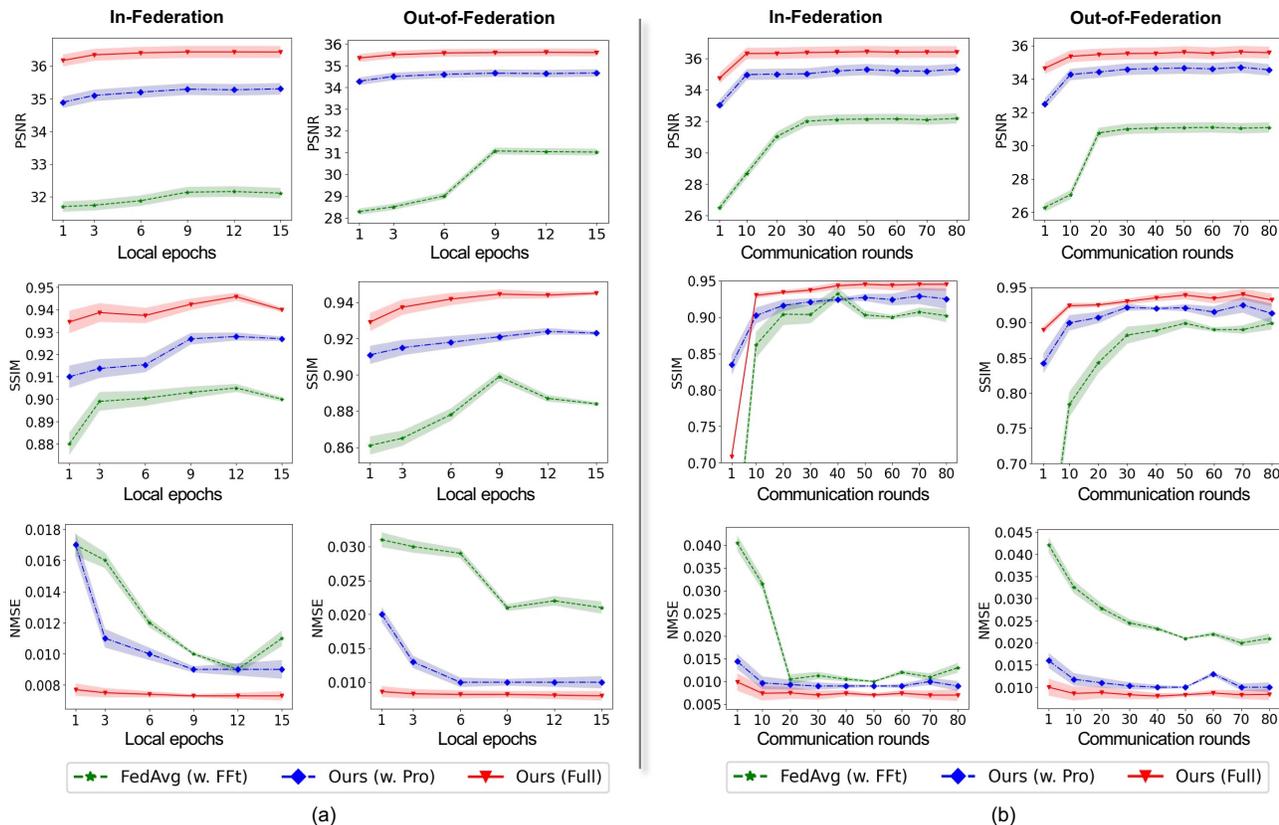


Figure 5. **Reconstruction accuracy** of FedAvg (w. FFt), **Ours** (w. Pro), and **Ours** (Full) versus (a) local epochs and (b) communication rounds under In-Federation and Out-of-Federation scenarios.

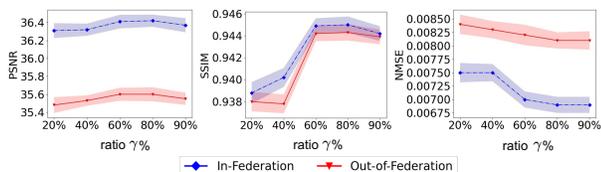


Figure 6. Analysis of the **ratio** γ % of the approximate null space in terms of PSNR, SSIM, and NMSE.

respectively. As shown in Table 1, the proposed method only requires 0.11 M of communication, which is 6% of the others, thereby addressing the issue ②. However, although the classical federated MRI algorithms FL-MRCM [13] and FedMRI [9] can also achieve good accuracy, they ignored this clinical issue of federated MRI. Additionally, the reconstruction accuracy of FedAvg (w. FFt) [26], **Ours** (w. Pro), and **Ours** (Full) under different communication rounds is recorded in Fig. 5 (b), where the number of the local epochs is fixed at 10 for each method. As can be seen from the figure, our method converges within 10 rounds, while FedAvg (w. FFt) [26] converges after 30 rounds. Especially under the mechanism of updating the local prompts in the approximate null space of global prompts, our method provides higher reconstruction results. This is because our federated prompt learning mechanism prevents overfitting caused by the limited amount of local training data and offers a personalized scheme for each client. The results in Fig. 5 (b)

confirm that our method provides the highest communication efficiency even on limited local data, thus relieving the issues ① and ②.

5. Conclusion

This paper presented a federated visual prompt learning method, FedPR, to tackle the three issues in federated MRI reconstruction, *i.e.*, limited communication bandwidth, insufficient local training data, and catastrophic forgetting. Benefiting from a powerful pre-trained model, our FedPR only learns prompts with a small number of learnable parameters. Additionally, FedPR greatly mitigates catastrophic forgetting by updating local prompts only in the approximate null space of global prompts, thereby preventing knowledge outside the local distribution from being overwritten. Experiments under In-Federation and Out-of-Federation scenarios demonstrate the superiority of FedPR in relieving the three issues of federated MRI reconstruction.

Acknowledgements: This work was supported by the Agency for Science, Technology and Research (A*STAR) AME Programmatic Funds (Grant Number: A20H4b0141). This work was also supported by the AI Singapore Tech Challenge (Open-Theme) Funding Scheme (Award No.: AISG2-TC-2021-003).

References

- [1] brain-development.org.
- [2] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Exploring visual prompts for adapting large-scale models, 2022.
- [3] Hong-You Chen, Cheng-Hao Tu, Ziwei Li, Han-Wei Shen, and Wei-Lun Chao. On pre-training for federated learning. *arXiv preprint arXiv:2206.11488*, 2022.
- [4] Jiahua Dong, Lixu Wang, Zhen Fang, Gan Sun, Shichao Xu, Xiao Wang, and Qi Zhu. Federated class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10164–10173, 2022.
- [5] Gokberk Elmas, Salman UH Dar, Yilmaz Korkmaz, Emir Ceyani, Burak Susam, Muzaffer Özbey, Salman Avestimehr, and Tolga Çukur. Federated learning of generative image priors for mri reconstruction. *arXiv preprint arXiv:2202.04175*, 2022.
- [6] Chun-Mei Feng, Huazhu Fu, Shuhao Yuan, and Yong Xu. Multi-contrast mri super-resolution via a multi-stage integration network. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 140–149. Springer, 2021.
- [7] Chun-Mei Feng, Yunlu Yan, Geng Chen, Yong Xu, Ying Hu, Ling Shao, and Huazhu Fu. Multi-modal transformer for accelerated mr imaging. *IEEE Transactions on Medical Imaging*, 2022.
- [8] Chun-Mei Feng, Yunlu Yan, Huazhu Fu, Li Chen, and Yong Xu. Task transformer network for joint mri reconstruction and super-resolution. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24*, pages 307–317. Springer, 2021.
- [9] Chun-Mei Feng, Yunlu Yan, Shanshan Wang, Yong Xu, Ling Shao, and Huazhu Fu. Specificity-preserving federated learning for mr image reconstruction. *IEEE Transactions on Medical Imaging*, 2022.
- [10] Chun-Mei Feng, Yunlu Yan, Kai Yu, Yong Xu, Ling Shao, and Huazhu Fu. Exploring separable attention for multi-contrast mr image super-resolution. *arXiv preprint arXiv:2109.01664*, 2021.
- [11] Chun-Mei Feng, Zhanyuan Yang, Geng Chen, Yong Xu, and Ling Shao. Dual-octave convolution for accelerated parallel mr image reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 116–124, 2021.
- [12] Chun-Mei Feng, Zhanyuan Yang, Huazhu Fu, Yong Xu, Jian Yang, and Ling Shao. Donet: dual-octave network for fast mr image reconstruction. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [13] Pengfei Guo, Puyang Wang, Jinyuan Zhou, Shanshan Jiang, and Vishal M Patel. Multi-institutional collaborations for improving deep learning-based magnetic resonance image reconstruction using federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2423–2432, 2021.
- [14] Wenke Huang, Mang Ye, and Bo Du. Learn from others and be yourself in heterogeneous federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10143–10153, 2022.
- [15] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. *arXiv preprint arXiv:2203.12119*, 2022.
- [16] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pages 5132–5143. PMLR, 2020.
- [17] Hyo-Eun Kim, Seungwook Kim, and Jaehwan Lee. Keep and learn: Continual learning by constraining the latent space for knowledge preservation in neural networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 520–528. Springer, 2018.
- [18] Yajing Kong, Liu Liu, Zhen Wang, and Dacheng Tao. Balancing stability and plasticity through advanced null space in continual learning. In *European Conference on Computer Vision*, pages 219–236. Springer, 2022.
- [19] Gihun Lee, Yongjin Shin, Minchan Jeong, and Se-Young Yun. Preservation of the global knowledge by not-true self knowledge distillation in federated learning. *arXiv preprint arXiv:2106.03097*, 2021.
- [20] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10713–10722, 2021.
- [21] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020.
- [22] Xiaoxiao Li, Meirui Jiang, Xiaofei Zhang, Michael Kamp, and Qi Dou. Fedbn: Federated learning on non-iid features via local batch normalization. *arXiv preprint arXiv:2102.07623*, 2021.
- [23] Guoliang Lin, Hanlu Chu, and Hanjiang Lai. Towards better plasticity-stability trade-off in incremental learning: A simple linear connector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 89–98, 2022.
- [24] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*, 2021.
- [25] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021.
- [26] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.
- [27] Carl D Meyer. *Matrix analysis and applied linear algebra*, volume 71. Siam, 2000.

- [28] John Nguyen, Kshitiz Malik, Maziar Sanjabi, and Michael Rabbat. Where to begin? exploring the impact of pre-training and initialization in federated learning. *arXiv preprint arXiv:2206.15387*, 2022.
- [29] Xing Nie, Bolin Ni, Jianlong Chang, Gaomeng Meng, Chunlei Huo, Zhaoxiang Zhang, Shiming Xiang, Qi Tian, and Chunhong Pan. Pro-tuning: Unified prompt tuning for vision tasks. *arXiv preprint arXiv:2207.14381*, 2022.
- [30] Sarthak Pati, Ujjwal Baid, Brandon Edwards, Micah J Sheller, Patrick Foley, G Anthony Reina, Siddhesh P Thakur, Chiharu Sako, Michel Bilello, Christos Davatzikos, et al. The federated tumor segmentation (fets) tool: an open-source solution to further solid tumor research. *Physics in Medicine & Biology*, 2022.
- [31] Timo Schick and Hinrich Schütze. It’s not just size that matters: Small language models are also few-shot learners. *arXiv preprint arXiv:2009.07118*, 2020.
- [32] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*, 2020.
- [33] Neta Shoham, Tomer Avidor, Aviv Keren, Nadav Israel, Daniel Benditkis, Liron Mor-Yosef, and Itai Zeitak. Overcoming forgetting in federated learning on non-iid data. *arXiv preprint arXiv:1910.07796*, 2019.
- [34] Liyan Sun, Zhiwen Fan, Xueyang Fu, Yue Huang, Xinghao Ding, and John Paisley. A deep information sharing network for multi-contrast compressed sensing mri reconstruction. *IEEE Transactions on Image Processing*, 28(12):6141–6153, 2019.
- [35] Anastasiia Usmanova, François Portet, Philippe Lalanda, and German Vega. A distillation-based approach integrating continual learning and federated learning for pervasive services. *arXiv preprint arXiv:2109.04197*, 2021.
- [36] Shipeng Wang, Xiaorong Li, Jian Sun, and Zongben Xu. Training networks in null space of feature covariance for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 184–193, 2021.
- [37] Shanshan Wang, Zhenghang Su, Leslie Ying, Xi Peng, Shun Zhu, Feng Liang, Dagan Feng, and Dong Liang. Accelerating magnetic resonance imaging via deep learning. In *2016 IEEE 13th international symposium on biomedical imaging (ISBI)*, pages 514–517. IEEE, 2016.
- [38] Guoyizhe Wei and Xiu Li. Knowledge lock: Overcoming catastrophic forgetting in federated learning. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 601–612. Springer, 2022.
- [39] Chencheng Xu, Zhiwei Hong, Minlie Huang, and Tao Jiang. Acceleration of federated learning with alleviated forgetting in local training. *arXiv preprint arXiv:2203.02645*, 2022.
- [40] Yan Yang, Jian Sun, Huibin Li, and Zongben Xu. Admm-cnnet: A deep learning approach for image compressive sensing. *IEEE transactions on pattern analysis and machine intelligence*, 42(3):521–538, 2018.
- [41] Zizhao Zhang, Adriana Romero, Matthew J Muckley, Pascal Vincent, Lin Yang, and Michal Drozdal. Reducing uncertainty in undersampled mri reconstruction with active acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2049–2058, 2019.