



FedFTN: Personalized Federated Learning with Deep Feature Transformation Network for Multi-institutional Low-count PET Denoising

Bo Zhou^{a,*}, Huidong Xie^a, Qiong Liu^a, Xiongchao Chen^a, Xueqi Guo^a, Zhicheng Feng^d, Jun Hou^e, S. Kevin Zhou^f, Biao Li^g, Axel Rominger^h, Kuangyu Shi^{h,i}, James S. Duncan^{a,b,c}, Chi Liu^{a,b,*}

^aDepartment of Biomedical Engineering, Yale University, New Haven, CT, USA.

^bDepartment of Radiology and Biomedical Imaging, Yale School of Medicine, New Haven, CT, USA

^cDepartment of Electrical Engineering, Yale University, New Haven, CT, USA.

^dDepartment of Biomedical Engineering, University of Southern California, Los Angeles, CA, USA

^eDepartment of Computer Science, University of California Irvine, Irvine, CA, USA

^fSchool of Biomedical Engineering & Suzhou Institute for Advanced Research, University of Science and Technology of China, Suzhou, China

^gDepartment of Nuclear Medicine, Ruijin Hospital, Shanghai Jiao Tong University School of Medicine, Shanghai, China

^hDepartment of Nuclear Medicine, Inselspital, Bern University Hospital, University of Bern, Bern, Switzerland

ⁱComputer Aided Medical Procedures and Augmented Reality, Institute of Informatics I16, Technical University of Munich, Munich, Germany

ARTICLE INFO

Article history:

2000 MSC: 41A05, 41A10, 65D05, 65D17

Keywords: low-count PET, Deep Reconstruction, Personalized Federated Learning, Denoising

ABSTRACT

Low-count PET is an efficient way to reduce radiation exposure and acquisition time, but the reconstructed images often suffer from low signal-to-noise ratio (SNR), thus affecting diagnosis and other downstream tasks. Recent advances in deep learning have shown great potential in improving low-count PET image quality, but acquiring a large, centralized, and diverse dataset from multiple institutions for training a robust model is difficult due to privacy and security concerns of patient data. Moreover, low-count PET data at different institutions may have different data distribution, thus requiring personalized models. While previous federated learning (FL) algorithms enable multi-institution collaborative training without the need of aggregating local data, addressing the large domain shift in the application of multi-institutional low-count PET denoising remains a challenge and is still highly under-explored. In this work, we propose FedFTN, a personalized federated learning strategy that addresses these challenges. FedFTN uses a local deep feature transformation network (FTN) to modulate the feature outputs of a globally shared denoising network, enabling personalized low-count PET denoising for each institution. During the federated learning process, only the denoising network's weights are communicated and aggregated, while the FTN remains at the local institutions for feature transformation. We evaluated our method using a large-scale dataset of multi-institutional low-count PET imaging data from three medical centers located across three continents, and showed that FedFTN provides high-quality low-count PET images, outperforming previous baseline FL reconstruction methods across all low-count levels at all three institutions.

© 2023 Elsevier B. V. All rights reserved.

1. Introduction

Positron Emission Tomography (PET) is a commonly used functional imaging modality with wide applications in oncol-

*Corresponding author.

e-mail: bo.zhou@yale.edu (Bo Zhou), chi.liu@yale.edu (Chi Liu)

ogy, cardiology, neurology, and biomedical research. To reconstruct high-quality PET, the patient is injected with a customized dose of radioactive tracer which inevitably introduces radiation exposure to both patients and healthcare providers. Adhere to the principle of As Low As Reasonably Achievable (ALARA), minimizing the radiation dose is of great interest to patients (Strauss and Kaste, 2006), particularly for PET applications where serial scans are commonly required. However, reducing the injection dose in PET would result in increased image noise, poor signal-to-noise ratio (SNR) and image artifacts, which would jeopardize the downstream clinical tasks.

To generate high-quality PET from low-count PET, deep learning-based low-count PET imaging methods have been extensively explored (Xiang *et al.*, 2017; Wang *et al.*, 2018; Lu *et al.*, 2019; Kaplan and Zhu, 2019; Hu *et al.*, 2021; Gong *et al.*, 2021b; Zhou *et al.*, 2020b; Ouyang *et al.*, 2019; Chen *et al.*, 2019; Liu *et al.*, 2020, 2021; Song *et al.*, 2021; Liu *et al.*, 2021; Gong *et al.*, 2021a), which have demonstrated superior performance than conventional methods (Dutta *et al.*, 2013; Maggioni *et al.*, 2013; Mejia *et al.*, 2016). While deep learning-based methods achieve promising performance, they often rely on training using diverse and large-scale paired low-count and full-count datasets that are often prohibitively expensive and difficult to collect. Even though we can alleviate this issue through building a centralized large-scale dataset by transferring all institutional data, the concerns of medical data privacy and security, the difficulty of building data transfer and warehouse protocol, and the laborious process make it challenging to implement this solution in practice (Rieke *et al.*, 2020; Roski *et al.*, 2014).

Federated learning (FL) has recently emerged as a solution to address data privacy concerns in training deep models. This approach enables different local clients to collaboratively learn using their own data and computing resources, without sharing any private data. A client-to-cloud platform is established where the cloud server periodically communicates with local clients to collect local models. These models are then aggregated to generate a global model that is redistributed to the local clients for further local updates. Unlike traditional methods where local data is required to be directly transferred for global training, FL only involves the exchange of model parameters or gradients. As a result, FL can potentially solve the data privacy concerns for training a global model. In the context of low-count PET imaging, the different institutions may have different low-count protocols, different PET systems from different vendors, different reconstructions, and different post-processing protocols, thus can lead to significant data heterogeneity and domain shifts. Unfortunately, it is challenging to generalize a global model trained from the classical FL algorithm, such as FedAvg (McMahan *et al.*, 2017), to different institutions due to different data distributions at each local site. Thus, personalized federated learning is desirable to address this issue.

Previous works have attempted to address the domain shift issues, but mainly focus on image classification (Li *et al.*, 2021; Arivazhagan *et al.*, 2019; Collins *et al.*, 2021; Fallah *et al.*, 2020; Shamsian *et al.*, 2021). In the application of accelerated MR reconstruction using FL, Guo *et al.* (2021) first proposed to ad-

dress the domain shift issue by iteratively aligning the latent feature of UNet (Ronneberger *et al.*, 2015) between target and other client sites. However, their cross-site strategy requires the target client to share both the latent feature and the network parameter with other client sites in each communication round, which could result in additional data privacy concerns (Lyu *et al.*, 2022; Huang *et al.*, 2021). Similarly, Feng *et al.* (2023) proposed a UNet with a globally shared encoder for generalized representation learning and a client-specific decoder for domain-specific reconstruction, but the network architecture is limited to UNet or its variants. While achieving promising performance for the MR reconstruction task, the network architecture is limited to UNet or its variants due to the constraint in their FL algorithm designs. Additionally, instead of using simple encoder-decoder structures for improving image quality, state-of-the-art deep learning-based image restoration networks often deploy advanced designs, such as original resolution restoration (Zhang *et al.*, 2018; Zhou *et al.*, 2021b), recurrent restoration (Zhou and Zhou, 2020; Zhou *et al.*, 2023a), and multi-stage restoration (Zamir *et al.*, 2021; Zhou *et al.*, 2022a). Moreover, all the previous works are limited to 2D MRI reconstruction while PET requires 3D processing. In the application of low-count PET denoising using FL, Zhou *et al.* (2023b) proposed the first FL study for low-count PET reconstruction. However, their study was performed only on simulation experiments with heterogeneous low-count data generated from one scanner at one site. In addition, their generation of personalized models relied on local fine-tuning after global training from FedAvg (McMahan *et al.*, 2017), which may not be optimal. Therefore, the development of 3D personalized federated learning framework and conducting studies with real-world multi-institutional low-count PET imaging data is desirable.

To address these challenges, we propose a personalized federated learning method based on deep feature transformation networks (FedFTN) and perform studies on real-world multi-institutional low-count PET data collected from three medical institutions across three continents, with multiple low-count levels contained at each institution. The general idea is to use a Feature Transformation Network (FTN) to modulate the features in the denoising network. During federated training, the FTNs are kept local, while only the denoising networks' parameters at different institutions are shared and aggregated at the central server. The FTNs are only trained locally to modulate the features from the denoising network, thus enabling personalized PET denoising at each institution. The input to the FTN is the low-count level for individual patients, thus also allowing using a single unified model for multiple low-count levels' denoising at each institution. In addition, we propose a Global Weight Constraint (GWC) loss which helps stabilize the local weight updates of the denoising network during the federated reconstruction learning. Our experimental result on the real-world multi-institutional low-count PET datasets demonstrates that our FedFTN can generate superior low-count PET denoising results as compared to previous federated reconstruction learning methods, as well as locally trained models.

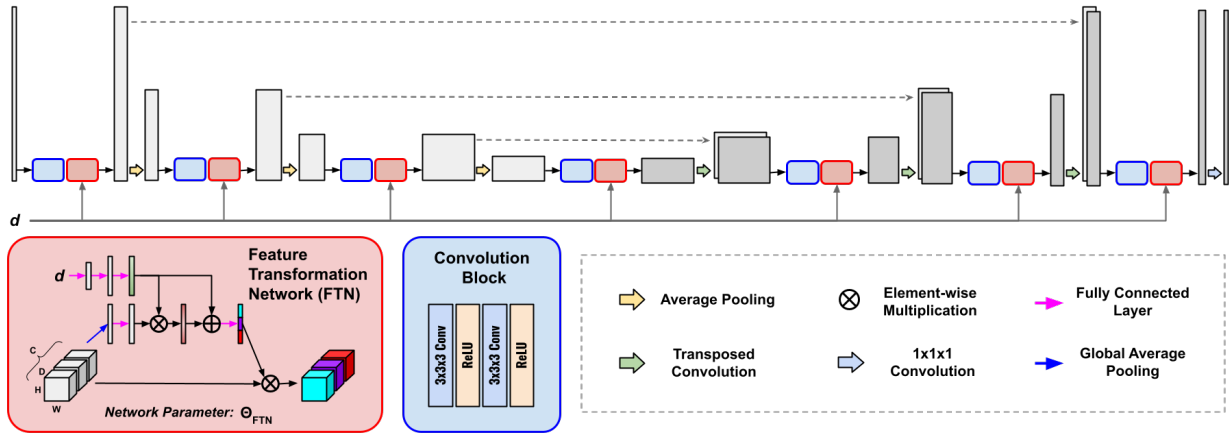


Fig. 1. The architecture of our Feature Transformation Network (FTN)-modulated Denoising Network. Without loss of generality, we deploy a U-Net as our denoising network. The FTNs are appended to the feature outputs at all U-Net resolution levels to transform the features at all levels.

2. Related Work

Low-count PET Denoising. Previous studies on low-count PET denoising can be divided into two categories: conventional post-processing methods (Dutta et al., 2013; Maggioni et al., 2013; Mejia et al., 2016) and deep learning-based methods (Xiang et al., 2017; Wang et al., 2018; Lu et al., 2019; Kaplan and Zhu, 2019; Hu et al., 2021; Gong et al., 2021b; Zhou et al., 2020b; Ouyang et al., 2019; Chen et al., 2019; Liu et al., 2020, 2021; Song et al., 2021). Although conventional methods like Gaussian filtering are standard post-processing techniques for PET reconstruction, they tend to oversmooth the image and have difficulty preserving local structures under low-count conditions with amplified noise. On the other hand, deep learning-based denoising methods have been extensively explored for low-count PET and have demonstrated promising results. For example, Kaplan and Zhu (2019) proposed to use a 2D Generative Adversarial Network (GAN) with UNet as a generator to predict full-count PET images from low-count PET images. Wang et al. (2018) proposed the use of a 3D conditional GAN to directly translate 3D low-count PET images to full-count PET images. Similarly, Gong et al. (2021b) also proposed to use a 3D Wasserstein GAN to stabilize GAN training and to improve the low-count PET denoising performance. Based on previous GAN designs, Ouyang et al. (2019) further proposed to reinforce the denoising performance by incorporating patient-specific information. In parallel to using only the low-count PET images as network input, low-count PET denoising facilitated by other imaging modalities has also been explored. For example, Xiang et al. (2017) developed an auto-context CNN using both low-count PET images and T1 MR images as inputs for full-count PET generation. Similarly, Chen et al. (2019) proposed to input low-count PET images along with multi-contrast MR images into a UNet for ultra-low-count PET denoising. There were also recent developments on unifying the low-count PET denoising model for multiple low-count levels (Xie et al., 2023). In addition to these previous studies on static low-count PET, methods have also been developed for simultaneous motion correction and low-count PET recon-

struction (Zhou et al., 2021a, 2020a, 2023c), and have demonstrated further improved reconstruction quality. Even though deep learning-based methods have shown great promise in improving low-count PET images, these methods have so far been studied only within single institutions, where a single low-count protocol, identical PET system, and reconstruction protocol are assumed. Investigation on how to train generalizable reconstruction models using multi-institutional low-count data with non-identical distribution while addressing data privacy issues is of great significance and remains relatively unexplored.

Federated Learning for Reconstruction. FL with a decentralized learning framework enables multiple local institutions to collaborate in training shared models while maintaining their local data privacy (Li et al., 2020). However, traditional FL algorithms, such as FedAvg (McMahan et al., 2017), do not adequately address the domain shift issue, which arises due to differences in data distribution across clients despite allowing collaborative training without sharing data. To overcome this limitation, recent studies have proposed personalized federated learning strategies in medical image reconstruction tasks. For example, Guo et al. (2021) proposed federated learning with cross-site modeling (FLCM) to address domain shift issues in 2D MR reconstruction by iteratively aligning the latent feature distribution between clients. However, FLCM’s requirement for frequent communication between clients and latent vectors increases the communication cost and risk of potential privacy leakage. On the other hand, Feng et al. (2023) proposed using a client-specific decoder in a UNet to reduce domain shift while only uploading and aggregating the encoder’s parameters. To further enable flexible MRI accelerated imaging operator across sites and improve generalizability, Elmas et al. (2023) proposed FedGIMP that aims to generate site-specific MRI prior which uses a mapper network to produce site-specific latents for the generative network given a site index. A similar strategy has been also adapted to multi-contrast MRI synthesis (Dalmaz et al., 2022). Although these methods show promising performance, they have only been tested on 2D MRI reconstruction tasks, and many rely on specific network architecture, e.g. UNet or Auto-encoder, with latent representation for domain

alignment and personalization. They may not generalize well to applications requiring other network architectures, such as cascade/recurrent/multistage network designs (Zhou and Zhou, 2020; Zhou et al., 2021b, 2022b,a), for better reconstruction quality. For CT reconstruction using FL, strategies have also been developed to adapt to different CT distributions from different sparse-view CT (SVCT) protocols. For example, with simulated multi-institutional 2D CT data, Yang et al. (2022) proposed to extend the PDF framework (Xia et al., 2021) into a hypernet-based FL framework, where the normalization parameter generator sub-networks were kept locally to adapt to different SVCT distributions. For low-count PET denoising, Zhou et al. (2023b) proposed the first FL study for low-count PET imaging. However, their study was limited to simulation experiments with heterogeneous low-count data generated from one scanner at one site. Furthermore, their approach of generating personalized models relied on local fine-tuning after global training from FedAvg(McMahan et al., 2017), which may not be optimal. Therefore, there is a need to develop a 3D personalized federated learning approach and conduct studies with real-world multi-institutional low-count PET data.

3. Methods

Our method included a Feature Transformation Network (FTN)-modulated denoising network (Figure 1) and a customized personalized federated learning framework (Figure 2) for it. Details were elaborated in the following sections.

3.1. Feature Transformation Network

The Feature Transformation Network (FTN) aimed to modulate the global-sharing denoising network to generate personalized denoised images that are specific to the local site and the target low-count level within the site.

As illustrated in Figure 1, given the extracted intermediate features from the convolution blocks of the global-sharing denoising network, the features were inputted into the FTNs for personalized adaptation. Specifically, given a feature input $F = [f_1, f_2, \dots, f_C]$ with $f_n \in \mathbb{R}^{H \times W \times D}$ denoting the individual feature channel, we flattened the feature via global average pooling, generating vector $v \in \mathbb{R}^C$ with its z -th element:

$$v_z = \frac{1}{H \times W \times D} \sum_i^H \sum_j^W \sum_k^D f_z(i, j, k), \quad (1)$$

where vector v embedded the global information of the input feature. Then, v was fed into a fully connected layer with weights of $w_R \in \mathbb{R}^{C \times C}$ and generated $v_R = w_R v$.

In parallel, given the low-count levels $d \in \mathbb{R}^1$ at the local site, we deployed three consecutive fully connected layers with weights of $w_1 \in \mathbb{R}^{\frac{C}{2} \times 1}$, $w_2 \in \mathbb{R}^{C \times \frac{C}{2}}$, and $w_3 \in \mathbb{R}^{C \times C}$, which generated transformation guiding vector:

$$v_d = w_3 \eta(w_2 \eta(w_1 d)), \quad (2)$$

where η is the ReLU activation function. Then, v_R and v_d were fused via:

$$v_{fuse} = \sigma(v_d) v_R + v_d, \quad (3)$$

where σ is the Sigmoid function. Then, we used another fully connected layer with weights of $w_{fuse} \in \mathbb{R}^{C \times C}$ to generate the transformation vector:

$$\hat{v} = w_{fuse} v_{fuse} \quad (4)$$

Finally, the transformation vector was applied to the input feature map using channel-wise multiplication:

$$F_{out} = [f_1 \hat{v}_1, f_2 \hat{v}_2, \dots, f_C \hat{v}_C] \quad (5)$$

Please note the FTNs were deployed at all feature resolution levels in the denoising network (Figure 1), thus allowing progressive feature transformation for personalized denoising.

3.2. FTN-based Personalized Federated Learning

The general pipeline of the FTN-based personalized federated learning is shown in Figure 2. We denoted $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_N$ as the low-count PET datasets from N different institutions. Each institutional dataset \mathcal{D}_n contained pairs of full-count and low-count PET 3D images, where each local dataset has three different low-count level settings. Within each institution, an FTN-modulated denoising network can be trained using the local data with two loss components, including a restoration loss \mathcal{L}_{recon} and a global weight constraint (GWC) loss \mathcal{L}_{gwc} . The first loss can be formulated as:

$$\mathcal{L}_{re} = \sum_{(x,y,d) \in \mathcal{D}_n} \|G_n(x, d | \theta_{R_n}, \theta_{FTN_n}) - y\|_2^2, \quad (6)$$

where G_n is the FTN-modulated denoising network at the n^{th} institution, and was parameterized by θ_{R_n} and θ_{FTN_n} . θ_{R_n} is the denoising network's parameters, and θ_{FTN_n} is the FTN's parameters. x , y , and d are the training pair of the low-count image, full-count image, and low-count level from \mathcal{D}_n . We assume there were Q global training epochs (communication between cloud server & local institutions), and P local training epochs at each institution. During the first two global epochs, the iterative optimization of the local network's parameter only used the reconstruction loss and can be written as:

$$\{\theta_{R_n}^{p+1}, \theta_{FTN_n}^{p+1}\} \leftarrow \{\theta_{R_n}^p, \theta_{FTN_n}^p\} - \alpha \nabla \mathcal{L}_{recon} \quad (7)$$

where α is the learning rate. At the end of each global training epoch, the global-sharing denoising network's parameters θ_{R_n} at each institution can be uploaded to a cloud server for weight aggregation. The cloud server updated the parameters of the global-sharing denoising network by:

$$\theta_{R_A}^q = \frac{1}{N} \sum_{n=1}^N \theta_{R_n}^q \quad (8)$$

where q denotes the q^{th} global epoch. After two global epochs of warm-up, we further added a GWC loss to stabilize the weights update during the local training in the following global epochs. The GWC loss is formulated as:

$$\mathcal{L}_{gwc} = \|\theta_{R_n} - \theta_{R_A}\|_2^2 \quad (9)$$

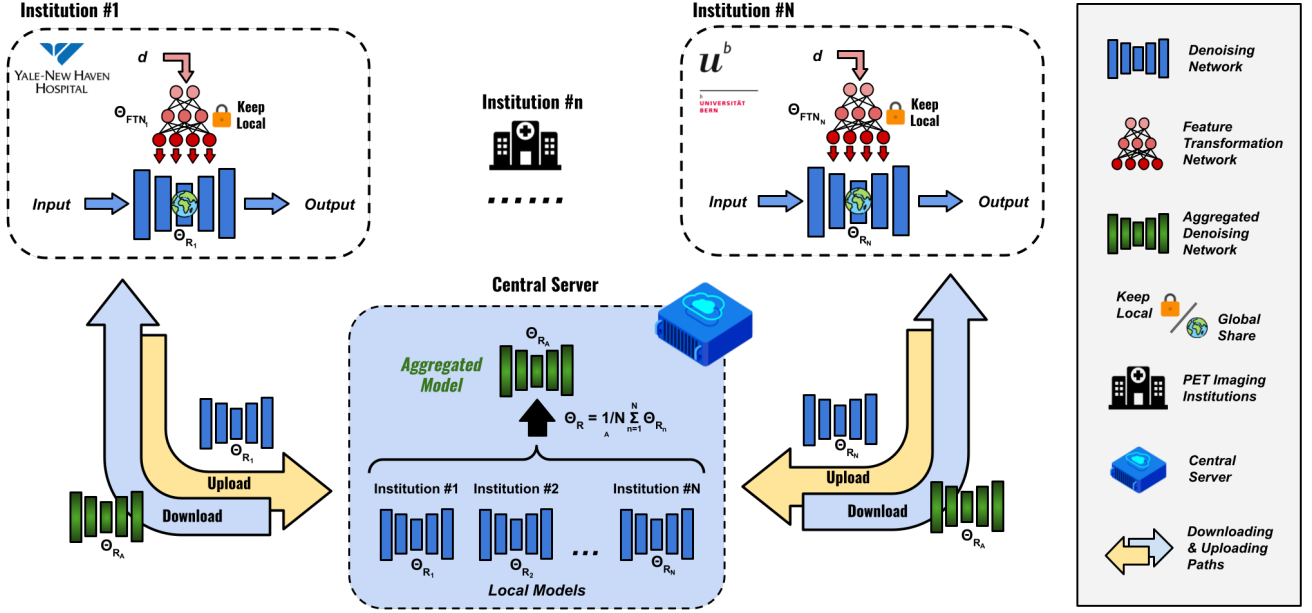


Fig. 2. Personalized federated learning based on feature transformation network (FedFTN). Each institution contains one FTN-modulated denoising network (Figure 1). At each global iteration, each FTN-modulated denoising network at each institution is first trained locally for a fixed number of epochs. During model aggregation, the FTNs are kept local and only the denoising networks' parameters are uploaded to the central cloud server for parameter averaging. The algorithm is summarized in Algorithm 1.

Algorithm 1: Steps of our Personalized Federated Learning Framework.

Input: $\mathcal{D} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_N\}$; ▷ Training datasets from N institutions
Initialize #1: P : the number of local epochs; Q : the number of global epochs
Initialize #2: θ_R : parameters of global denoising network
Initialize #3: $\theta_{FTN_1}, \dots, \theta_{FTN_N}$: parameters of local feature transformation networks
Initialize #4: α : local training learning rate; λ : weight of GWC loss
for $q = 1$ **to** Q **do**
 for $n = 1$ **to** N **in parallel do**
 → deploy averaged global denoising network parameters θ_{R_A} to local institution, if available
 for $p = 1$ **to** P **do**
 if $q < 3$ **then**
 $\mathcal{L}_n = \sum_{(x,y,d) \in \mathcal{D}_n} \|G_n(x, d | \theta_{R_n}, \theta_{FTN_n}) - y\|_2^2$; ▷ Compute local training loss (warm up)
 else
 $\mathcal{L}_n = \sum_{(x,y,d) \in \mathcal{D}_n} \|G_n(x, d | \theta_{R_n}, \theta_{FTN_n}) - y\|_2^2 + \lambda \|\theta_{R_n} - \theta_{R_A}\|_2^2$; ▷ Compute local training loss
 $\{\theta_{R_n}^{p+1}, \theta_{FTN_n}^{p+1}\} \leftarrow \{\theta_{R_n}^p, \theta_{FTN_n}^p\} - \alpha \nabla \mathcal{L}_n$; ▷ Update network parameters
 → upload global denoising network parameters $\theta_{R_n}^q$ to the central server
 $\theta_{R_A}^q = \frac{1}{N} \sum_{n=1}^N \theta_{R_n}^q$; ▷ Update global-sharing denoising network parameter by averaging
Output $\{\theta_{R_A}^Q, \theta_{FTN_1}^Q\}, \{\theta_{R_A}^Q, \theta_{FTN_2}^Q\}, \dots, \{\theta_{R_A}^Q, \theta_{FTN_N}^Q\}$; ▷ Return personalized models for each institution

where the current denoising network's weight θ_{R_n} was constrained to be in proximal to the aggregated weight θ_{R_A} . Starting from the 3rd global epoch, the training loss was the combination of the reconstruction loss and the global weight constraint loss, thus can be written as:

$$\mathcal{L}_{comb} = \mathcal{L}_{re} + \lambda \mathcal{L}_{gwc} \quad (10)$$

where λ is the weight of GWC loss and we set it to 0.001. Then, the iterative optimization of the local network's parameter can be formulated as:

$$\{\theta_{R_n}^{p+1}, \theta_{FTN_n}^{p+1}\} \leftarrow \{\theta_{R_n}^p, \theta_{FTN_n}^p\} - \alpha \nabla \mathcal{L}_{comb} \quad (11)$$

After Q rounds of communication between local institutions and the cloud server, we can obtain collaboratively-trained

FTN-modulated denoising networks with personalized parameters $\{\theta_{R_A}^Q, \theta_{FTN_1}^Q\}, \{\theta_{R_A}^Q, \theta_{FTN_2}^Q\}, \dots, \{\theta_{R_A}^Q, \theta_{FTN_N}^Q\}$. The algorithm is summarized in Algorithm 1.

3.3. Multi-institutional low-count PET Data

We collected multi-institutional low-count PET data from three different medical centers in USA, Switzerland, and China for our study. The first dataset was collected at Yale New Haven Hospital, New Haven, USA. 200 subjects were included in this dataset. The subjects were injected with a ^{18}F -FDG tracer and the whole-body protocol with continuous-bed motion scanning was used. All data were acquired using a Siemens Biograph mCT PET/CT system. The average dose across all patients is $256.3 \pm 16.2\text{MBq}$. We used uniform down-sampling of the PET list-mode data with down-sampling ratios of 5%, 10%, and 20% to generate low-count PET data at three different low-count levels. For both the low-count and full-count images, they were reconstructed using the ordered-subsets expectation maximization (OSEM) algorithm with 2 iterations and 21 subsets, provided by the vendor. A post-reconstruction Gaussian filter with 5mm full width at half maximum (FWHM) was used. The voxel size of the reconstructed image was $2.04 \times 2.04 \times 2.03\text{mm}^3$. The image size was 400×400 in the transverse plane and varied in the axial direction depending on patient height. The 200 subjects were split into 100 subjects for training, 10 subjects for validation, and 90 subjects for evaluation. The second dataset was collected at the Department of Nuclear Medicine, University of Bern, Bern, Switzerland (Xue et al., 2021). 209 subjects with ^{18}F -FDG tracer were included in this dataset. All data were acquired using a Siemens Biograph Vision Quadra whole-body PET/CT system. The average dose across all patients is $264.1 \pm 18.2\text{MBq}$. Here, low-count PET data at 2%, 5%, and 10% low-count levels were generated by down-sampling of the PET list-mode data. For both the low-count and full-count images, they were reconstructed using the OSEM algorithm with 6 iterations and 5 subsets, provided by the vendor. A post-reconstruction Gaussian filter with 5mm FWHM was used. The voxel size of the reconstructed image was $1.65 \times 1.65 \times 1.65\text{mm}^3$. The image size was $440 \times 440 \times 644$. The 209 subjects were split into 109 subjects for training, 10 subjects for validation, and 90 subjects for evaluation. The third dataset was collected at the Ruijin Hospital, Shanghai, China (Xue et al., 2021). 204 subjects with ^{18}F -FDG tracer were included in this dataset. All data were acquired using a United Imaging uExplorer total-body PET/CT system. The average dose across all patients is $260.1 \pm 12.2\text{MBq}$. Here, low-count PET data at 2%, 5%, and 10% low-count levels were generated by down-sampling of the PET list-mode data. For both the low-count and full-count images, they were reconstructed using the OSEM algorithm with 4 iterations and 20 subsets, provided by the vendor. A post-reconstruction Gaussian filter with 5mm FWHM was used. The voxel size of the reconstructed image was $1.66 \times 1.66 \times 2.88\text{mm}^3$. The image size was $360 \times 360 \times 674$. The 204 subjects were split into 104 subjects for training, 10 subjects for validation, and 90 subjects for evaluation.

3.4. Evaluation Metrics and Baselines Comparisons

We evaluated the low-count denoised results using Peak Signal-to-Noise Ratio (PSNR) and Normalized Mean Square Error (NMSE) computed against full-count reconstruction ground truth. For baseline comparisons, we first compared our results against previous federated reconstruction algorithms, including Federated Averaging (FedAvg, McMahan et al. (2017)), Federated Learning with Proximal Term (FedProx, Li et al. (2020)), Federated Learning with Local Batch Normalization (FedBN, Li et al. (2021)), Specificity-Preserving Federated Learning (FedSP, Feng et al. (2023)), and Personalized Federated Learning with Hypernetwork (FedHyper, Shamsian et al. (2021)). For a fair comparison, all the methods used an identical reconstruction/denoising network, as shown in Figure 1. We also compared our FedFTN against two types of locally trained models, including local single models and local unified models. For each local single model, an FTN-modulated denoising network was trained using one single low-count level data at one specific institution. On the other hand, the local unified models are also based on FTN-modulated denoising networks, but all three low-count levels' data at a specific institution are used as training data. Furthermore, we also performed ablative studies on federated transfer learning (Zhou et al., 2023b), where FedFTN are further fine-tuned using local data for site adaptation.

3.5. Implementation Details

We implemented our method in Pytorch and performed experiments using an NVIDIA Quadro RTX 8000 GPU with 48GB memory. The Adam solver was used to optimize our models with $lr = 1 \times 10^{-4}$, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. We used a batch size of 3 and trained all models for 300 global epochs. The number of the local epoch was set to 3. To prevent overfitting at each local site, we also implemented 'on-the-fly' data augmentation. During training, we performed $64 \times 64 \times 64$ random cropping, and then randomly flipped the cropped volumes along the x, y, and z-axis. During the site adaptation with fine tuning on local data, we used a reduced learning rate of $2e - 5$ and trained the FedFTN models using the local dataset for 10 epochs, and the batch size was also set to 3 with 'on-the-fly' data augmentation. Our method's training takes about 160 hours to complete. For baseline methods, FedHyper takes about 172 hours, and the rest baselines also take about 160 hours.

4. Experimental Results

Figure 3 shows qualitative comparisons of low-count PET denoised images using locally trained denoising models and our FedFTN method. The lowest low-count level denoised results from each institution were visualized for comparison. As we can see, since only less than 5% count was used, the original images (1st row) suffer from high noise and image artifacts. While the locally trained models can reduce the noise and recover the general structure, the detail recovery was still not ideal. For example, the 5% low-count denoised image from the locally trained model at Institution #1 created additional artifacts at the intersection between the liver and kidney. The 2% low-count

Table 1. Quantitative comparisons of the low-count PET denoised images between our FedFTN and locally trained models. The local single model means FTN-modulated denoising networks trained at one specific low-count level at the specific institution. The local unified model means FTN-modulated denoising networks trained with all three low-count levels within each institution. The performance of FedFTN with Site Adaptation (SA) via further local data fine-tuning is reported in the last row. The best results are marked in bold. “†” indicates that the difference between FedFTN and all compared methods is significant at $p < 0.005$ based on the non-parametric Wilcoxon signed rank test.

Evaluation	Institution #1			Institution #2			Institution #3		
PSNR/NMSE/SSIM	5%	10%	20%	2%	5%	10%	2%	5%	10%
Original	20.46/.129/.931	23.80/.059/.954	27.40/.026/.972	19.90/.076/.923	23.66/.031/.951	26.26/.017/.968	20.61/.121/.933	24.89/.050/.961	27.55/.030/.975
Local Single Models	25.93/.034/.966	27.91/.021/.975	30.22/.012/.983	25.37/.021/.971	27.16/.014/.980	28.64/.010/.985	26.00/.035/.974	28.17/.023/.982	29.90/.018/.986
Local Unified Model	26.02/.032/.968	28.01/.020/.975	30.21/.012/.983	25.42/.020/.970	27.21/.014/.980	28.64/.010/.985	26.02/.035/.975	28.20/.023/.982	29.91/.018/.986
FedFTN	27.24/.025/.999[†]	28.96/.017/.999[†]	30.82/.011/.999[†]	26.12/.018/.999[†]	27.80/.013/.999[†]	29.03/.009/.999[†]	26.83/.031/.999[†]	28.91/.021/.999[†]	30.23/.017/.999[†]

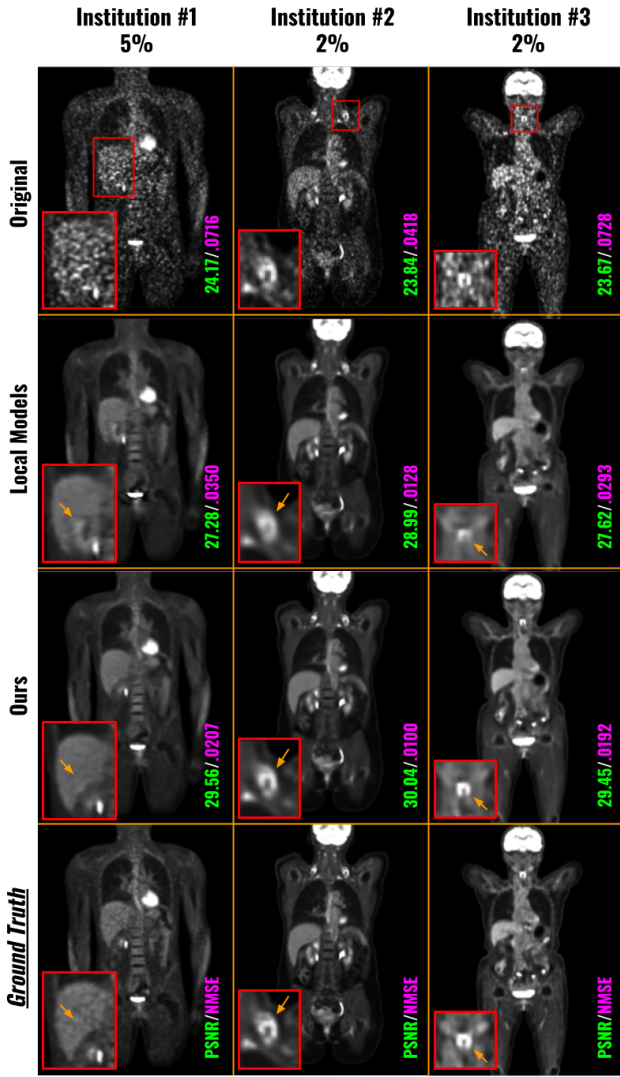


Fig. 3. Visual comparison of the low-count denoised images from locally trained models and from our FedFTN. The local single model means FTN-modulated denoising networks trained at one specific low-count level at the specific institution. The local unified model means FTN-modulated denoising networks trained with all three low-count levels within each institution. The PET images with the lowest count levels at each institution are visualized here.

denoised images from the locally trained models at Institution #2 and Institution #3 suffered from heavy blurring on important regions, such as the hypermetabolic lesions in the zoomed boxes. The quantitative comparisons were summarized in Table 1. We reported quantitative results for all available low-count

levels for the three institutions. Similar to the observation from the visualizations, all the original PET images suffered from low SNR, resulting in low PSNR and NMSE values. Taking Institution #1 as an example, we can see the locally trained unified model, i.e. the FTN-modulated Denoising Network, was able to improve the PSNR from 20.46 to 26.02 for the 5% low-count PET. Please note that the local unified model is the network shown in Figure 1 that is trained locally with the local site’s low-count level as an additional input. Using the FedFTN, we can further improve the PSNR from 26.02 to 27.24 with statistical significance. Similar observations on improvement over the locally trained models can be found for other low-count levels at Institution #1 and other institutions. The average inference time for testing data of institution #1, institution #2, and institution #3 were 8.08 ± 2.83 s, 11.32 ± 1.05 s, and 10.04 ± 1.09 s, respectively.

Figure 4 presented qualitative comparisons of low-count PET denoised images using different federated methods. For each institution, we showed one patient example for each low-count level, and all low-count levels are visualized. At Institution #1, all three low-count levels’ denoised images suffered from low SNR, with the noise level increasing as the dose level decreases. While previous FL methods can improve the image quality by suppressing the noise, the detail recovery was still suboptimal. For instance, for the 5% low-count level, the bone marrows in the spine were heavily blurred by the previous methods, whereas FedFTN showed much sharper spine, demonstrating the best consistency with the ground truth full-count reconstruction. Additionally, for the 10% low-count level, a strong false positive signal was visible in the spine from the original low-count reconstruction, and previous methods failed to suppress it, which may lead to misdiagnosis. In contrast, using FedFTN, we can suppress this false positive signal and provided a high-quality image that best matches the ground truth. At Institution #2, for the 2% low-count level, a false positive of cardiac defect was visible from the original low-count reconstruction. While previous methods can suppress the noise, this false positive defect signal cannot be fully removed, particularly for FedSP. In contrast, FedFTN demonstrated significantly better performance in false positive defect removal and the most consistent cardiac shape as compared to the ground truth. Similarly, for the 5% low-count level, FedFTN better recovered the signal of the aorta wall than previous methods. We observe similar results in the patient examples from Institution #3, where FedFTN can provide better image quality as compared to previous methods. The corresponding quantitative comparisons were summarized in Table 2. As mentioned previously, the orig-

inal low-count reconstructions from all three institutions suffered from poor image quality, where the PSNR values were all lower than 21.00 at the lowest low-count levels for all three institutions. By deploying previous federate learning methods, as compared to the locally trained models (Figure 3 and Table 1), we can see that these previous FL methods can already improve the low-count image quality across all institutions at all low-count levels. For example, FedSP was able to increase the PSNR from 26.02 to 26.69 for the 5% low-count reconstruction at Institution #1, from 25.42 to 25.80 for the 2% low-count reconstruction at Institution #2, and from 26.02 to 26.44 for the 2% low-count reconstruction at Institution #3. In the second last row, we showed that our FedFTN can significantly outperform these previous FL baselines, and achieved state-of-the-art performance across all the low-count levels at all three institutions. For example, as compared to the FedHyper with the best performance among previous FL reconstruction algorithms, our FedFTN can further improve the PSNR from 26.88 to 27.24 for the 5% low-count reconstruction at Institution #1, from 25.85 to 26.12 for the 2% low-count reconstruction at Institution #2, and from 26.44 to 26.89 for the 2% low-count reconstruction at Institution #3.

Similar to the process of FTL (Zhou et al., 2023b), we further performed Site Adaptation (SA) through local fine-tuning for our FedFTN. The quantitative results were reported in the last row in Table 2. We can observe further improved image quality metrics of FedFTN through SA for all low-count levels at all three institutions. In addition, we also found the denoising results from FedFTN with SA still significantly outperformed all previous FL baseline methods, in terms of both PSNR and NMSE. Visual comparisons of low-count denoising before and after SA of FedFTN are shown in Figure 5. We can see SA improves the image quality by further suppressing the false positive signal and improving the image resolution across different low-count levels at Institution #1.

We conducted ablative studies on the Global Weight Constraint (GWC) loss that was used in our FedFTN framework to stabilize the training process. The results were summarized in Table 3, which included the image quality analysis for all low-count levels across all three institutions. Adding the GWC component consistently improved image quality compared to using FedFTN without GWC. For instance, at Institution #1, the PSNR for the 5% PET imaging increased from 27.11 to 27.24, from 28.82 to 28.96 for the 10% PET imaging, and from 30.71 to 30.82 for the 20% PET imaging. Figure 7 provided visual comparisons from multiple institutions. We observed that adding GWC helps stabilize the FedFTN training process and helps maintain the subtle true positive signals within the kidneys, as seen in the denoised images from Institutions #1 and #2. Similarly, adding GWC results in better resolution recovery for clustered thin anatomic structures, such as ribs, as demonstrated in the Institution #3 comparison. While the quantitative and qualitative improvements are incremental, as shown in Figure 6, we can observe that FedFTN with GWC can provide faster and more stable training loss convergence, as compared to the one without GWC, demonstrating the GWC's benefits during FL training.

5. Discussion

In this work, we developed a novel personalized federated learning approach, called FedFTN, for multi-institutional low-count PET denoising. Specifically, we proposed to use a deep Feature Transformation Network (FTN) that is kept at the local institutions and takes the low-count level as input, to transform the intermediate feature outputs from the globally shared denoising network. There are several key advantages of this design. First of all, with different FTNs modulating the denoising network features at different local sites, we can personalize the denoising network, adapting to different low-count PET data with different distributions caused by differences in scanners, pre-/post-processing protocols, etc. Second, unlike previous personalized FL modules for image classification (Li et al., 2021), MRI synthesis (Elmas et al., 2023), and 2D SVCT reconstruction (Yang et al., 2022) that use mapper sub-network directly generate scalars for feature map channel-wise multiplication and addition, our FTN first squeezes the feature map into a latent representation and fuse with the site and subject-specific latent vector before re-excitation. In Table 2, we found our FedFTN can provide better reconstruction performance than FedBN (Li et al., 2021) which also uses a locally kept module but generates normalization parameters for direct feature map affine transformation. Third, each institution often has multiple low-count protocols, thus requiring the denoising network to be able to adapt to inputs with different low-count levels. Instead of blindly inputting the original low-count reconstruction into the network without the knowledge of the low-count level, the FTN-modulated denoising network takes the information of the low-count level as additional input, thus enabling dose-level-aware denoising. Please also note that because individual FTN network is kept locally at each institution, the institution can define its own low-count levels. The sites do not need to share the exact same set of levels and no information about the low-count level is shared between institutions. Even though it is possible that better performance could be achieved if different sites share identical and matched low counts, having identical low counts across different sites based on different scanners, different injection protocols, and different processing/reconstruction protocols, may not be easily realizable in real-world FL scenarios. Lastly, we also proposed a Global Weight Constraint (GWC) loss that regularizes the denoising network parameters not to have strong deviation over the aggregated parameter during the training at local sites, helping stabilize the federated learning process, and thus improving the final personalized denoising performance at each site.

We collected large-scale real-world low-count PET data from three different institutions in the U.S.A., Europe, and China, to validate our method. From our experimental results, we demonstrated the feasibility of using our FedFTN for collaborative training without sharing data, while enabling personalized low-count PET denoising at different institutions. First, as we can observe from Table 1 and Figure 3, even though training a denoising model from scratch using local data with limited diversity can generate reasonable denoising performance and potentially avoid domain shift issues, our FedFTN can provide significantly better denoising results. For example, as shown in the

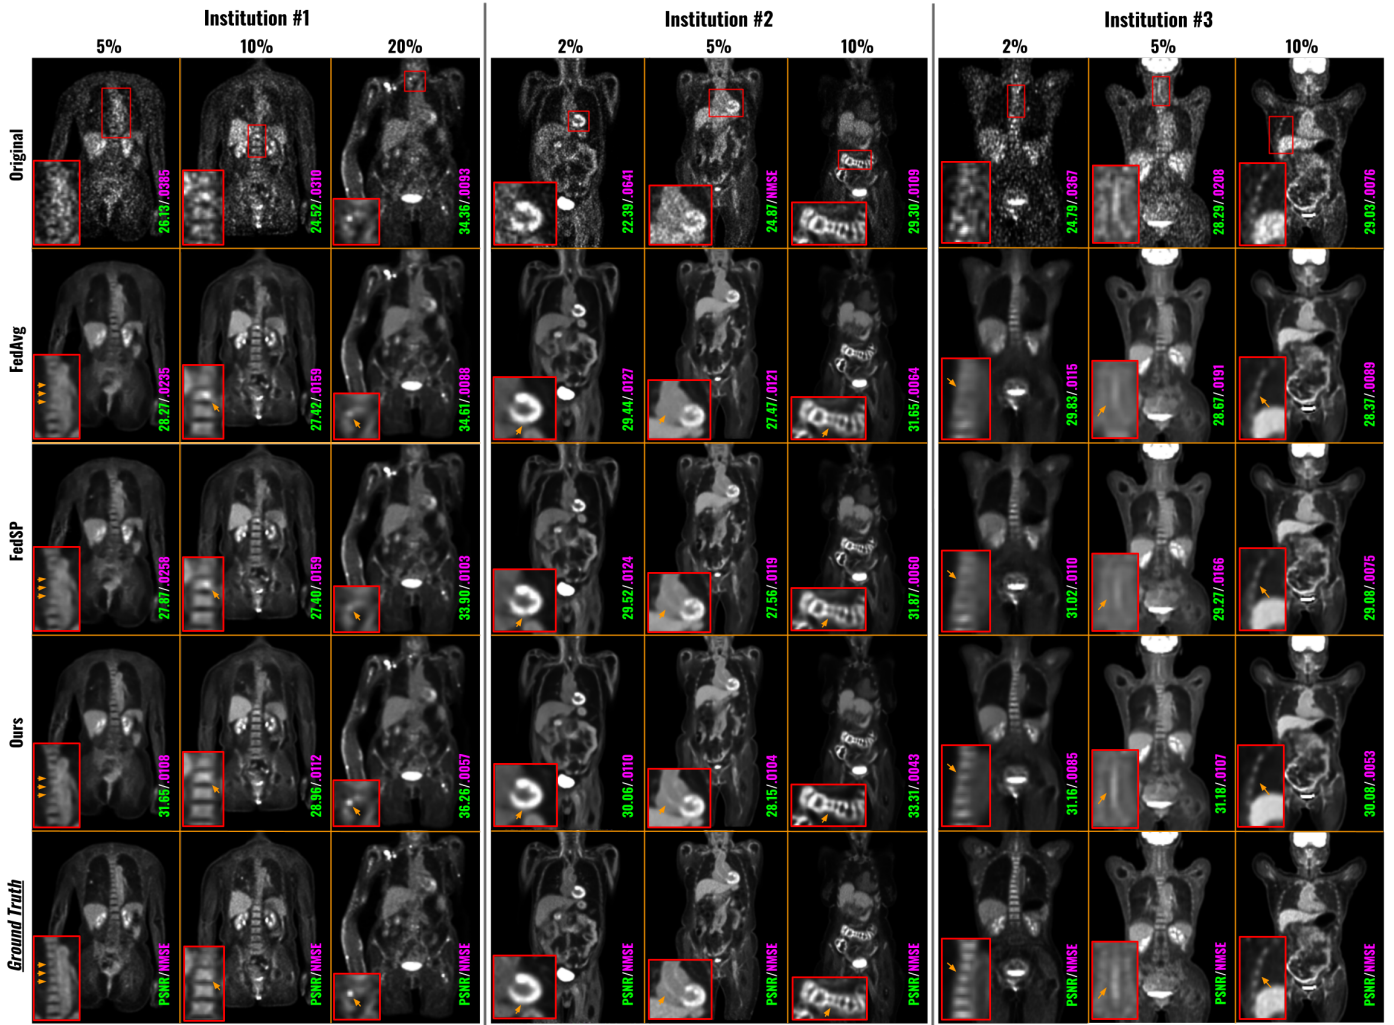


Fig. 4. Qualitative comparison of low-count denoising with different federated learning methods. Each institution contains three-different low-count levels. The original low-count reconstruction and the full-count ground truths are shown in the first row and last row, respectively. The image quality metrics of each image are indicated at the bottom of the images.

Table 2. Quantitative comparisons of low-count PET denoised images using different federated learning methods. Each institution contains three different low-count levels. The best results are marked in bold. “+” indicates that the difference between FedFTN and all compared methods is significant at $p < 0.005$ based on the non-parametric Wilcoxon signed rank test.

Evaluation PSNR/NMSE/SSIM	Institution #1			Institution #2			Institution #3		
	5%	10%	20%	2%	5%	10%	2%	5%	10%
Original	20.46/.129/.931	23.80/.059/.954	27.40/.026/.972	19.90/.076/.923	23.66/.031/.951	26.26/.017/.968	20.61/.121/.933	24.89/.050/.961	27.55/.030/.975
FedAvg	26.62/.029/.969	28.30/.019/.976	29.95/.013/.981	25.70/.020/.971	27.35/.014/.980	28.43/.010/.985	26.07/.035/.973	28.07/.023/.982	29.16/.019/.986
FedBN	26.68/.028/.970	28.35/.019/.978	30.05/.013/.982	25.79/.019/.973	27.47/.013/.981	28.60/.010/.986	26.31/.033/.975	28.34/.022/.983	29.38/.018/.987
FedProx	26.64/.028/.969	28.32/.018/.977	29.99/.013/.981	25.75/.020/.972	27.39/.014/.981	28.50/.010/.985	26.13/.034/.974	28.17/.023/.983	29.23/.019/.986
FedSP	26.69/.028/.969	28.36/.019/.976	30.16/.013/.982	25.80/.019/.973	27.51/.013/.981	28.63/.010/.985	26.23/.034/.974	28.33/.022/.983	29.48/.018/.987
FedHyper	26.88/.027/.971	28.56/.018/.978	30.33/.012/.983	25.85/.019/.974	27.49/.013/.982	28.65/.010/.986	26.44/.033/.976	28.52/.022/.983	29.73/.018/.987
FedFTN	27.24/.025/.979[†]	28.96/.017/.983[†]	30.82/.011/.990[†]	26.12/.018/.980[†]	27.80/.013/.989[†]	29.03/.009/.991[†]	26.83/.031/.980[†]	28.91/.021/.990[†]	30.23/.017/.992[†]
FedFTN + SA	27.32/.024/.980[†]	28.99/.016/.985[†]	30.83/.011/.991[†]	26.21/.017/.981[†]	27.84/.012/.990[†]	29.05/.009/.992[†]	26.89/.031/.980[†]	28.94/.021/.991[†]	30.25/.017/.992[†]

last row of Table 1, our FedFTN demonstrated superior PSNR and NMSE values as compared to the locally trained models across all the low-count levels at all the institutions. This is mainly due to the fact that the FedFTN utilizes all the institutional data with a wider spectrum of data diversity for collaborative learning while using the FTN modulation to mitigate the domain shift issues. Second, as we can see from Table 2 and Figure 4, our method generating personalized FTN-modulated denoising networks for individual institutions can consistently

outperform previous FL reconstruction methods that either only produce one global model (FedAvg) or deploy personalized FL strategies (FedSP and FedHyper). Using an identical backbone reconstruction/denoising network, our FedFTN achieved the best image quality over all the previous FL baselines, in terms of both PSNR and NMSE with statistical significance, as reported in the second last row of Table 2. Further fine-tuning the personalized denoising models from FedFTN with local data slightly boosted the performance. In addition, we

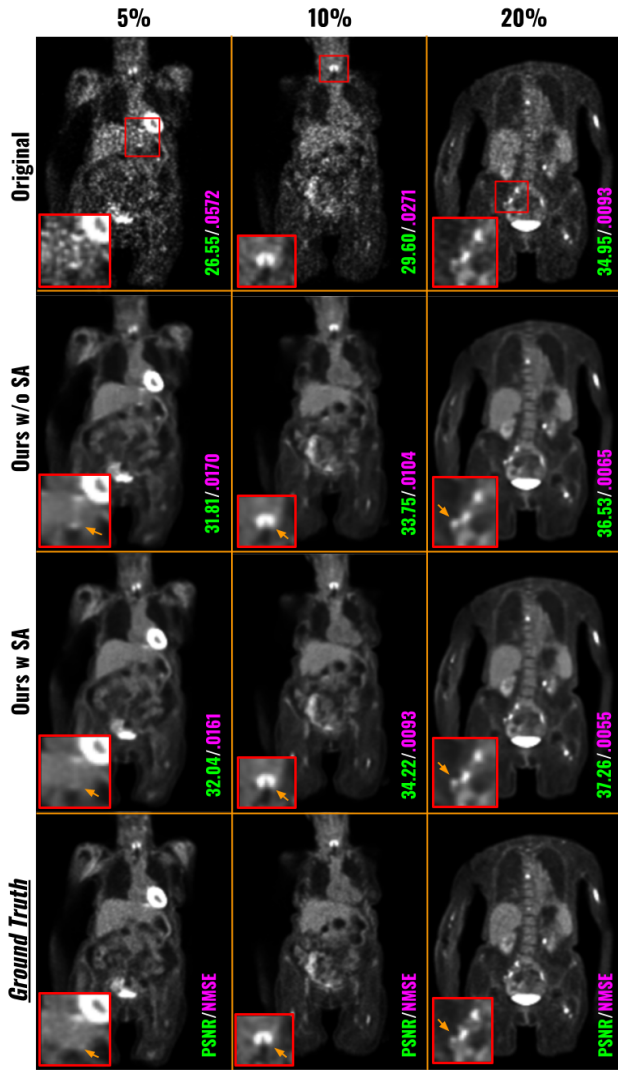


Fig. 5. Visual comparison of low-count PET denoised images from FedFTN before and after Site Adaptation (via) local fine-tuning. The three low-count levels at the Institution #1 are shown here.

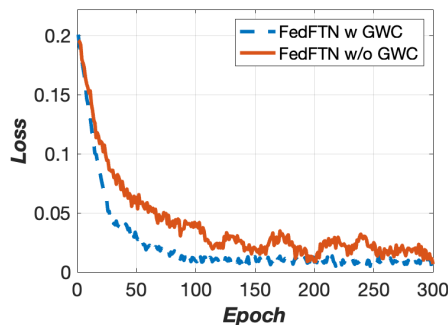


Fig. 6. Institution #1's training curves of FedFTN with and without the GWC components.

found that adding the Global Weight Constraint (GWC) loss helped stabilize the FedFTN and improved the image quality at all low-count levels at all institutions, as demonstrated in Figure 7 and Table 3.

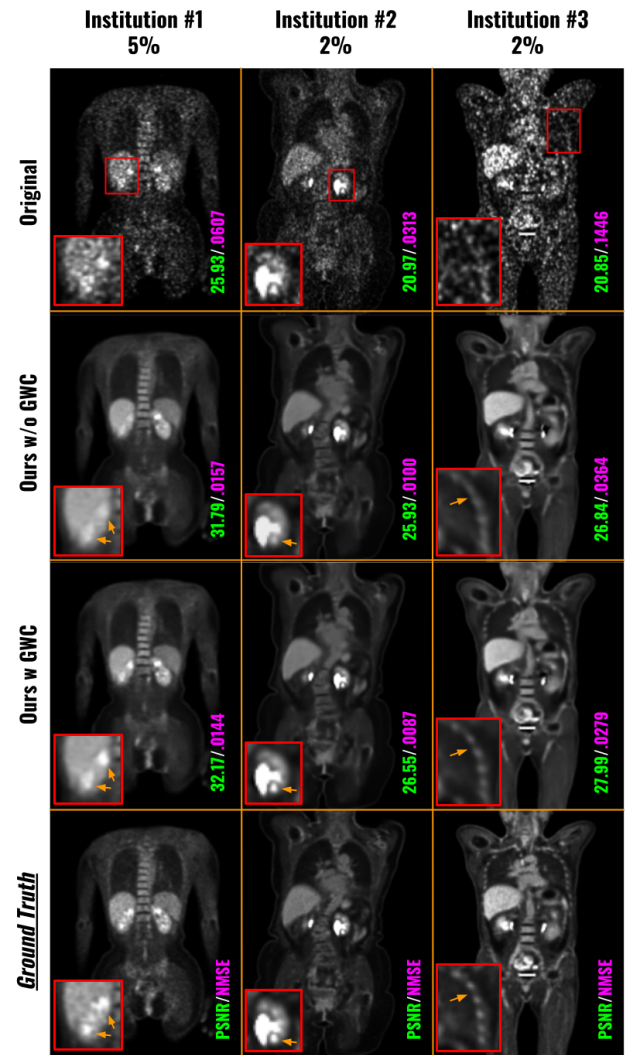


Fig. 7. Visual comparison of low-count denoised images from FedFTN with and without the Global Weight Constraint (GWC) loss. The PET data with the lowest dose levels at each institution are visualized here.

Table 3. Ablation study on Global Weight Constraint (GWC) loss. Quantitative comparisons of FedFTN with and without GWC loss during training are shown here.

Institution #1	5%	10%	20%
FedFTN w/o GWC	27.11/.0262/.978	28.82/.0178/.982	30.71/.0117/.990
FedFTN w GWC	27.24/.0254/.979	28.96/.0172/.983	30.82/.0113/.990
Institution #2	2%	5%	10%
FedFTN w/o GWC	25.91/.0194/.979	27.59/.0137/.988	28.82/.0103/.989
FedFTN w GWC	26.12/.0185/.980	27.80/.0130/.989	29.03/.0098/.991
Institution #3	2%	5%	10%
FedFTN w/o GWC	26.83/.0323/.980	28.87/.0223/.988	30.14/.0179/.991
FedFTN w GWC	26.83/.0319/.980	28.91/.0219/.990	30.23/.0176/.992

The presented work also has limitations with several potential improvements that are the subjects of our future studies. First, our study only considered three institutions and all with ^{18}F -FDG tracer. While there are other PET tracers that could be used for specific applications, ^{18}F -FDG is still the most commonly used PET tracer in clinical practice and thus is the primary focus of our study. Furthermore, our FedFTN framework

can be flexibly adjusted to different numbers of institutions, and potentially adapted to multi-tracer PET scenarios. Specifically, we could incorporate the tracer type, as well as other dose and patient information, as additional inputs to the FTN, which would also allow the FTN to transform the features in the denoising network depending on the input tracer type. In fact, we believe that including more diverse tracer types with expanded training data with more diverse data representation would potentially further improve our performance. In this work, we have already shown that we can use FTN to adapt and unify those different low-count distributions. However, we believe expanding to include multi-tracer multi-institutional data is an important future direction to validate any conclusion. Second, we only evaluated the overall image quality based on image quality metrics, i.e. PSNR/NMSE/SSIM, which use full-count PET as ground truth. The difference between our FL method and locally trained baselines is in the order of 1-2 dB PSNR which implies a significant image quality improvement with our FL strategy as compared to the local training strategy, which can be observed from Table 1 and Figure 3. While the image quality metrics improvements from our FedFTN as compared to the prior SOTA FL methods is in smaller magnitude, the image quality improvements are reflected on more detailed regions, as shown in Figure 4. We believe this kind of improvement could potentially lead to more accurate disease quantification, e.g. lesion radiomic, cardiac function, etc. However, how will such improvement over the prior SOTA FL method be reflected in clinically relevant tasks is an important direction for our following clinical investigation. Given PET has extensive clinical applications in oncology, cardiology, and neurology, our future work also includes evaluations of how the denoised image impacts the downstream tasks, such as the impacts on staging and therapy response, and human experts' evaluations on these clinical tasks. Lastly, for a fair comparison and to demonstrate the idea, we performed all our experiments using a simple UNet as the backbone denoising network. However, our method could be adapted with more advanced restoration network structures. For instance, we could use more advanced network designs, such as cascade-based, transformer-based, and multistage-based reconstruction networks (Zhou et al., 2021b, 2022b; Zhou and Zhou, 2020; Zhou et al., 2022a; Shan et al., 2019), in our FedFTN. Specifically, we can use FTN to transform and modulate the intermediate features in these networks, thus enabling personalized FL with these networks. Deploying these networks in our FedFTN could potentially further improve our performance and will also be an important direction for our future studies.

6. Conclusion

Our work proposes an innovative personalized federated learning method, named FedFTN, for multi-institutional low-count PET denoising. Our method utilizes a deep feature transformation network to generate a personalized low-count PET denoising model for each institution by modulating a globally shared denoising network. During the federated learning process, the FTN remains at the local institutions and is used to

transform intermediate feature outputs from the shared denoising network, thus enabling personalized FL denoising. We utilized a large-scale dataset of multi-institutional low-count PET data from three medical centers located across three continents to validate our method. Our experimental results showed that the FedFTN provides high-quality low-count PET denoised images, outperforming previous baseline FL methods across all low-count levels at all three institutions. We believe our proposed methods could potentially be adapted to other deep learning-based medical imaging challenges where collaborative training without data sharing is needed to improve the medical image quality.

Acknowledgments

This work was supported by the National Institutes of Health (NIH) grant R01EB025468 and grant R01CA275188. Parts of the data used in the preparation of this article were obtained from the University of Bern, Department of Nuclear Medicine and School of Medicine, Ruijin Hospital. As such, the investigators contributed to the design and implementation of DATA and/or provided data but did not participate in the analysis or writing of this report. A complete listing of investigators can be found at: "<https://ultra-low-dose-pet.grand-challenge.org/Description/>"

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Credit authorship contribution statement

Bo Zhou: Conceptualization, Methodology, Software, Visualization, Validation, Formal analysis, Writing original draft. **Huidong Xie:** Data preparation, Methodology, Writing - review and editing. **Qiong Liu:** Data preparation, Writing - review and editing. **Xiongchao Chen:** Data preparation, Writing - review and editing. **Xueqi Guo:** Data preparation, Writing - review and editing. **Zhicheng Feng:** Results analysis, Writing - review and editing. **Jun Hou:** Results analysis, Writing - review and editing. **S. Kevin Zhou:** Writing - review and editing. **Biao Li:** Data preparation, Writing - review and editing. **Axel Rominger:** Data preparation, Writing - review and editing. **Kuangyu Shi:** Data preparation, Writing - review and editing. **James S. Duncan:** Writing - review and editing, Supervision. **Chi Liu:** Conceptualization, Writing - review and editing, Supervision.

References

- Arivazhagan, M.G., Aggarwal, V., Singh, A.K., Choudhary, S., 2019. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*.
- Chen, K.T., Gong, E., de Carvalho Macruz, F.B., Xu, J., Boumis, A., Khalighi, M., Poston, K.L., Sha, S.J., Greicius, M.D., Mormino, E., et al., 2019. Ultra-low-dose 18f-florbetaben amyloid PET imaging using deep learning with multi-contrast MRI inputs. *Radiology* 290, 649–656.

- Collins, L., Hassani, H., Mokhtari, A., Shakkottai, S., 2021. Exploiting shared representations for personalized federated learning, in: International conference on machine learning, PMLR. pp. 2089–2099.
- Dalmaz, O., Mirza, U., Elmas, G., Özbey, M., Dar, S.U., Ceyani, E., Avestimehr, S., Çukur, T., 2022. One model to unite them all: Personalized federated learning of multi-contrast MRI synthesis. *arXiv preprint arXiv:2207.06509*.
- Dutta, J., Leahy, R.M., Li, Q., 2013. Non-local means denoising of dynamic PET images. *PloS one* 8, e81390.
- Elmas, G., Dar, S.U.H., Korkmaz, Y., Ceyani, E., Susam, B., Ozbey, M., Avestimehr, S., Çukur, T., 2023. Federated learning of generative image priors for mri reconstruction. *IEEE Transactions on Medical Imaging* 42, 1996–2009. doi:10.1109/TMI.2022.3220757.
- Fallah, A., Mokhtari, A., Ozdaglar, A., 2020. Personalized federated learning: A meta-learning approach. *arXiv preprint arXiv:2002.07948*.
- Feng, C.M., Yan, Y., Wang, S., Xu, Y., Shao, L., Fu, H., 2023. Specificity-preserving federated learning for mr image reconstruction. *IEEE Transactions on Medical Imaging* 42, 2010–2021. doi:10.1109/TMI.2022.3202106.
- Gong, K., Kim, K., Cui, J., Wu, D., Li, Q., 2021a. The evolution of image reconstruction in PET: From filtered back-projection to artificial intelligence. *PET clinics* 16, 533–542.
- Gong, Y., Shan, H., Teng, Y., Tu, N., Li, M., Liang, G., Wang, G., Wang, S., 2021b. Parameter-transferred wasserstein generative adversarial network (PT-WGAN) for low-dose PET image denoising. *IEEE Transactions on Radiation and Plasma Medical Sciences* 5, 213–223. doi:10.1109/TRPMS.2020.3025071.
- Guo, P., Wang, P., Zhou, J., Jiang, S., Patel, V.M., 2021. 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, pp. 2423–2432.
- Hu, Z., Xue, H., Zhang, Q., Gao, J., Zhang, N., Zou, S., Teng, Y., Liu, X., Yang, Y., Liang, D., Zhu, X., Zheng, H., 2021. DPIR-Net: Direct PET image reconstruction based on the wasserstein generative adversarial network. *IEEE Transactions on Radiation and Plasma Medical Sciences* 5, 35–43. doi:10.1109/TRPMS.2020.2995717.
- Huang, Y., Gupta, S., Song, Z., Li, K., Arora, S., 2021. Evaluating gradient inversion attacks and defenses in federated learning. *Advances in Neural Information Processing Systems* 34.
- Kaplan, S., Zhu, Y.M., 2019. Full-dose PET image estimation from low-dose PET image using deep learning: a pilot study. *Journal of digital imaging* 32, 773–778.
- Li, T., Sahu, A.K., Talwalkar, A., Smith, V., 2020. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine* 37, 50–60.
- Li, X., Jiang, M., Zhang, X., Kamp, M., Dou, Q., 2021. Fed{bn}: Federated learning on non-{iid} features via local batch normalization, in: International Conference on Learning Representations.
- Liu, H., Wu, J., Lu, W., Onofrey, J.A., Liu, Y.H., Liu, C., 2020. Noise reduction with cross-tracer and cross-protocol deep transfer learning for low-dose PET. *Physics in Medicine & Biology* 65, 185006.
- Liu, J., Malekzadeh, M., Mirian, N., Song, T.A., Liu, C., Dutta, J., 2021. Artificial intelligence-based image enhancement in PET imaging: Noise reduction and resolution enhancement. *PET clinics* 16, 553–576.
- Lu, W., Onofrey, J.A., Lu, Y., Shi, L., Ma, T., Liu, Y., Liu, C., 2019. An investigation of quantitative accuracy for deep learning based denoising in oncological PET. *Physics in Medicine & Biology* 64, 165019.
- Lyu, L., Yu, H., Ma, X., Chen, C., Sun, L., Zhao, J., Yang, Q., Yu, P.S., 2022. Privacy and robustness in federated learning: Attacks and defenses. *IEEE Transactions on Neural Networks and Learning Systems*, 1–21doi:10.1109/TNNLS.2022.3216981.
- Maggioni, M., Katkovnik, V., Egiazarian, K., Foi, A., 2013. Nonlocal transform-domain filter for volumetric data denoising and reconstruction. *IEEE Transactions on Image Processing* 22, 119–133. doi:10.1109/TIP.2012.2210725.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A., 2017. Communication-efficient learning of deep networks from decentralized data, in: Artificial Intelligence and Statistics, PMLR. pp. 1273–1282.
- Mejia, J., Mederos, B., Mollineda, R.A., Ortega Maynez, L., 2016. Noise reduction in small animal PET images using a variational non-convex functional. *IEEE Transactions on Nuclear Science* 63, 2577–2585. doi:10.1109/TNS.2016.2589246.
- Ouyang, J., Chen, K.T., Gong, E., Pauly, J., Zaharchuk, G., 2019. Ultra-low-dose PET reconstruction using generative adversarial network with feature matching and task-specific perceptual loss. *Medical physics* 46, 3555–3564.
- Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H.R., Albarqouni, S., Bakas, S., Galtier, M.N., Landman, B.A., Maier-Hein, K., et al., 2020. The future of digital health with federated learning. *NPJ digital medicine* 3, 1–7.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, Springer. pp. 234–241.
- Roski, J., Bo-Linn, G.W., Andrews, T.A., 2014. Creating value in health care through big data: opportunities and policy implications. *Health affairs* 33, 1115–1122.
- Shamsian, A., Navon, A., Fetaya, E., Chechik, G., 2021. Personalized federated learning using hypernetworks, in: International Conference on Machine Learning, PMLR. pp. 9489–9502.
- Shan, H., Padole, A., Homayounieh, F., Kruger, U., Khera, R.D., Nitiwarangkul, C., Kalra, M.K., Wang, G., 2019. Competitive performance of a modularized deep neural network compared to commercial algorithms for low-dose CT image reconstruction. *Nature Machine Intelligence* 1, 269–276.
- Song, T.A., Yang, F., Dutta, J., 2021. Noise2void: unsupervised denoising of PET images. *Physics in Medicine & Biology* 66, 214002.
- Strauss, K.J., Kaste, S.C., 2006. The alara (as low as reasonably achievable) concept in pediatric interventional and fluoroscopic imaging: striving to keep radiation doses as low as possible during fluoroscopy of pediatric patients—a white paper executive summary. *Radiology* 240, 621–622.
- Wang, Y., Yu, B., Wang, L., Zu, C., Lalush, D.S., Lin, W., Wu, X., Zhou, J., Shen, D., Zhou, L., 2018. 3d conditional generative adversarial networks for high-quality PET image estimation at low dose. *Neuroimage* 174, 550–562.
- Xia, W., Lu, Z., Huang, Y., Liu, Y., Chen, H., Zhou, J., Zhang, Y., 2021. CT reconstruction with PDF: Parameter-dependent framework for data from multiple geometries and dose levels. *IEEE Transactions on Medical Imaging* 40, 3065–3076. doi:10.1109/TMI.2021.3085839.
- Xiang, L., Qiao, Y., Nie, D., An, L., Lin, W., Wang, Q., Shen, D., 2017. Deep auto-context convolutional neural networks for standard-dose PET image estimation from low-dose PET/MRI. *Neurocomputing* 267, 406–416.
- Xie, H., Liu, Q., Zhou, B., Chen, X., Guo, X., Liu, C., 2023. Unified noise-aware network for low-count pet denoising. *arXiv preprint arXiv:2304.14900*.
- Xue, S., Guo, R., Bohn, K.P., Matzke, J., Viscione, M., Alberts, I., Meng, H., Sun, C., Zhang, M., Zhang, M., et al., 2021. A cross-scanner and cross-tracer deep learning method for the recovery of standard-dose imaging quality from low-dose PET. *European journal of nuclear medicine and molecular imaging*, 1–14doi:https://doi.org/10.1007/s00259-021-05644-1.
- Yang, Z., Xia, W., Lu, Z., Chen, Y., Li, X., Zhang, Y., 2022. Hypernetwork-based personalized federated learning for multi-institutional ct imaging. *arXiv preprint arXiv:2206.03709*.
- Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L., 2021. Multi-stage progressive image restoration, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14821–14831.
- Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y., 2018. Residual dense network for image super-resolution, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2472–2481.
- Zhou, B., Chen, X., Xie, H., Zhou, S.K., Duncan, J.S., Liu, C., 2022a. Dudoynet: dual-domain under-to-fully-complete progressive restoration network for simultaneous metal artifact reduction and low-dose CT reconstruction. *IEEE Transactions on Medical Imaging* 41, 3587–3599.
- Zhou, B., Chen, X., Zhou, S.K., Duncan, J.S., Liu, C., 2022b. DuDoDR-Net: Dual-domain data consistent recurrent network for simultaneous sparse view and metal artifact reduction in computed tomography. *Medical Image Analysis* 75, 102289.
- Zhou, B., Dey, N., Schlemper, J., Salehi, S.S.M., Liu, C., Duncan, J.S., Sofka, M., 2023a. Dsformer: A dual-domain self-supervised transformer for accelerated multi-contrast mri reconstruction, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 4966–4975.
- Zhou, B., Miao, T., Mirian, N., Chen, X., Xie, H., Feng, Z., Guo, X., Li, X., Zhou, S.K., Duncan, J.S., Liu, C., 2023b. Federated transfer learning for low-dose PET denoising: A pilot study with simulated heterogeneous data. *IEEE Transactions on Radiation and Plasma Medical Sciences* 7, 284–295. doi:10.1109/TRPMS.2022.3194408.

- Zhou, B., Tsai, Y.J., Chen, X., Duncan, J.S., Liu, C., 2021a. MDPET: A unified motion correction and denoising adversarial network for low-dose gated PET. *IEEE Transactions on Medical Imaging* 40, 3154–3164. doi:10.1109/TMI.2021.3076191.
- Zhou, B., Tsai, Y.J., Liu, C., 2020a. Simultaneous denoising and motion estimation for low-dose gated PET using a siamese adversarial network with gate-to-gate consistency learning, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer. pp. 743–752.
- Zhou, B., Tsai, Y.J., Zhang, J., Guo, X., Xie, H., Chen, X., Miao, T., Lu, Y., Duncan, J.S., Liu, C., 2023c. Fast-MC-PET: A novel deep learning-aided motion correction and reconstruction framework for accelerated PET, in: *International Conference on Information Processing in Medical Imaging*, Springer. pp. 523–535.
- Zhou, B., Zhou, S.K., 2020. DuDoRNet: Learning a dual-domain recurrent network for fast MRI reconstruction with deep T1 prior, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 298–313.
- Zhou, B., Zhou, S.K., Duncan, J.S., Liu, C., 2021b. Limited view tomographic reconstruction using a cascaded residual dense spatial-channel attention network with projection data fidelity layer. *IEEE Transactions on Medical Imaging* 40, 1792–1804. doi:10.1109/TMI.2021.3066318.
- Zhou, L., Schaefferkoetter, J.D., Tham, I.W., Huang, G., Yan, J., 2020b. Supervised learning with cycleGAN for low-dose FDG PET image denoising. *Medical Image Analysis* 65, 101770.