A Specificity-Preserving Generative Model for Federated MRI Translation

Onat Dalmaz[1,2(✉)], Usama Mirza[1,2], Gökberk Elmas[1,2], Muzaffer Özbey[1,2], Salman U. H. Dar[1,2], and Tolga Çukur[1,2]

[1] Department of Electrical and Electronics Engineering, Bilkent University, Ankara, Turkey
onat@ee.bilkent.edu.tr

[2] National Magnetic Resonance Research Center (UMRAM), Bilkent University, Ankara, Turkey

Abstract. MRI translation models learn a mapping from an acquired source contrast to an unavailable target contrast. Collaboration between institutes is essential to train translation models that can generalize across diverse datasets. That said, aggregating all imaging data and training a centralized model poses privacy problems. Recently, federated learning (FL) has emerged as a collaboration framework that enables decentralized training to avoid sharing of imaging data. However, FL-trained translation models can deteriorate by the inherent heterogeneity in the distribution of MRI data. To improve reliability against domain shifts, here we introduce a novel specificity-preserving FL method for MRI contrast translation. The proposed approach is based on an adversarial model that adaptively normalizes the feature maps across the generator based on site-specific latent variables. Comprehensive FL experiments were conducted on multi-site datasets to show the effectiveness of the proposed approach against prior federated methods in MRI contrast translation.

Keywords: Federated learning · Site-specificity · MRI · Translation · Heterogeneity

1 Introduction

Multi-contrast MRI enables non-invasive diagnostic assessment of anatomy and accumulates complementary information via examination of multiple tissue contrasts [2,26]. Yet, multi-contrast protocols have time and economic costs that can prevent collection of all desired contrasts in an MRI exam [18,34]. This limitation can be addressed by contrast translation, which is the imputation of missing sequences in a protocol from the acquired sequences [15]. Deep models have made remarkable progress in this area, enabling centralized models to significantly improve MRI translation performance. [8,36,39]. Unfortunately, learning generalizable models for medical imaging tasks requires training on diverse datasets. However, compiling such datasets at a central institution would inevitably compromise patient privacy [17].

S. Albarqouni et al. (Eds.): DeCaF 2022/FAIR 2022, LNCS 13573, pp. 79–88, 2022.
https://doi.org/10.1007/978-3-031-18523-6_8

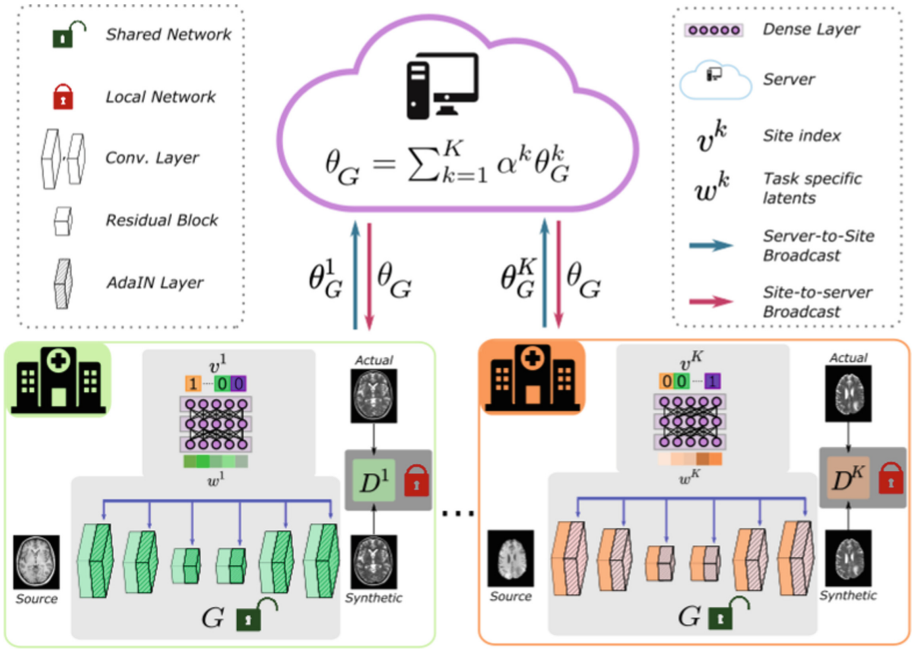

Fig. 1. SPFL-Trans is a decentralized contrast translation method based on federated learning of a conditional adversarial model. AdaIN layers along with site-specific latents produced by a subnetwork effectively modulate feature maps in order to cope with data heterogeneity across different sites.

Federated learning (FL) is a powerful framework to address this major limitation based on decentralized model training across multiple institutions [9,21,23,29,30,33]. In this framework, a server aggregates locally optimized models to compute a shared global model [24,35]. Aggregated models can be impaired by the heterogeneity in the data distribution naturally evident for multi-institutional datasets [29,32] due to different scanners, acquisition parameters etc. Previous studies on FL-based medical imaging have introduced several prominent approaches to cope with data heterogeneity in segmentation [5,22,23,27,31,38], classification [3,22,40], and reconstruction [10,11] tasks. However, influence of data heterogeneity on FL-based MRI contrast translation remains understudied.

Here, we introduce a novel Specificity-Preserving Federated Learning method for MRI Translation (SPFL-Trans). In contrast to previous approaches, the proposed method embodies a site-aware architecture that effectively addresses the inherent data heterogeneity in multi-institutional datasets. SPFL-Trans is based on a generator backbone equipped with Adaptive Instance Normalization layers (AdaIN) to adaptively tune the statistics of feature maps for improved generalization across sites. FL experiments conducted on multi-contrast MRI datasets indicate the superiority of the proposed approach against prior FL-based translation methods.

2 Theory

2.1 MRI Translation with Adversarial Models

Conditional generative adversarial networks (cGANs) have emerged as a gold-standard for MRI contrast translation in recent years due to their exceptional recovery for high-frequency textural details in medical images [4,7,8,20]. cGANs perform adversarial learning via a pair of generator (G) and discriminator (D) subnetworks [8]. G predicts a synthetic target-contrast image ($\hat{x}_t$) given as input an acquired source-contrast image (x_s), whereas the D tries to distinguish actual (x_t) and synthetic target-contrast images. To learn image translation, cGANs are typically trained to minimize an aggregate loss function composed of adversarial and pixel-wise terms:

$$\mathcal{L} = \mathbb{E}_{x_s,x_t}[-(D(x_s,x_t)-1)^2 - D(x_s,G(x_s))^2 + \lambda_{pix}||x_t - G(x_s)||_1], \quad (1)$$

where $\mathbb{E}$ denotes expectation, and λ_{pix} is the relative weighing term for the pixel-wise loss.

2.2 Specificity-Preserving Federated Learning of MRI Translation

Network Architecture. The proposed model is a conditional adversarial architecture that takes as input the source image along with site-specifying information (Fig. 1). The first component of the generator is a latent producing block (LPS) to form site-specific latents w^k given one-hot encoding of site index $v^k \in \mathbb{R}^K$:

$$w^k = LPS(v^k) \quad (2)$$

The generator architecture is inspired by the ResNet model [8,13] with a residual bottleneck between a convolutional encoder and a convolutional decoder. To mitigate heterogeneity reflected in the statistics of derived feature maps, an AdaIN layer is inserted after each convolutional layer in the encoder/decoder, and each residual block in the bottleneck. At the output of the ith layer of the generator, the mean and standard deviation of output feature maps $g_i \in \mathbb{R}^{F_i,H_i,W_i}$ are modulated. To do this, site-specific latents w^k are first transformed into scale and bias vectors $\gamma_i, \beta_i \in \mathbb{R}^{F_i}$:

$$\gamma_i = Q_i^{\gamma} w^k + b_i^{\gamma};\ \beta_i = Q_i^{\beta} w^k + b_i^{\beta} \quad (3)$$

where $Q_i^{\gamma,\beta} \in \mathbb{R}^{F_i,J}$ and $b_i^{\gamma,\beta} \in \mathbb{R}^{F_i}$ are learnable linear transformations. The AdaIN layer then modulates the first- and second-order statistics of each channel [14]:

$$g_i^{'} = \mathrm{AdaIN}(g_i,\gamma_i,\beta_i) = \begin{bmatrix} \gamma_i[1]\frac{g_i[1]-\mu(g_i[1])\mathbf{1}}{\sigma(g_j[1])} + \beta_i[1]\mathbf{1} \\ \gamma_i[2]\frac{g_i[2]-\mu(g_i[2])\mathbf{1}}{\sigma(g_j[2])} + \beta_i[2]\mathbf{1} \\ \vdots \end{bmatrix} \quad (4)$$

where $\mathbf{1} \in \mathbb{R}^{H_j,W_j}$ is a matrix of ones, μ, σ compute the mean and standard deviation of individual channels $g_i[j] \in \mathbb{R}^{H_i,W_i}$, and $g_i^{'}$ is the input to the next network layer $i+1$.

Algorithm 1: Training of SPFL-Trans

Data: $\{\mathcal{D}^1, \cdots, \mathcal{D}^K\}$ from K sites
Input: P: number of communication rounds
$\alpha^1, \cdots, \alpha^K$: averaging weights for K sites
G: global generator with parameters θ_G
$D^1, \cdots, D^K$: local discriminators with $\theta_{D^1}, \cdots, \theta_{D^K}$
$Opt()$: optimizer for parameter updates
Output: θ_G^* Optimized global generator

Randomly initialize θ_G and $\theta_{D^1}, \cdots, \theta_{D^K}$
for $p = 1$ *to* P **do**
  for $k = 1$ *to* K **do**
    for *one epoch* **do**
      $\theta_G^k \leftarrow \theta_G$ `// Broadcast global generators to the sites`
      Calculate $\nabla_{\theta_G^k} \mathcal{L}^k(\mathcal{D}^k)$ and $\nabla_{\theta_D^k} \mathcal{L}^k(\mathcal{D}^k)$ based on Eq. 1
      $\theta_G^k \leftarrow \theta_G^k - Opt(\nabla_{\theta_G^k} \mathcal{L}^k(\mathcal{D}^k))$; $\theta_D^k \leftarrow \theta_D^k - Opt(\nabla_{\theta_D^k} \mathcal{L}^k(\mathcal{D}^k))$
  $\theta_G = \sum_{k=1}^{K} \alpha^k \theta_G^k$, `// Aggregate locally trained generators`

Federated Training. To train SPFL-Trans, a decentralized learning procedure takes place for a total of P communication rounds between the FL server and individual sites (see Fig. 1, Algorithm 1) [25]. Throughout the procedure, generators are shared across sites, though discriminators are kept unshared for enhanced privacy preservation [12,28]. During local training, models are optimized on the training sets from individual sites according to the aggregate cGAN loss function as expressed in Eq. 1.

3 Methods

3.1 Datasets

Demonstrations were performed on four public datasets taken to represent four different sites in the FL framework: IXI (https://brain-development.org/ixi-dataset/), BRATS [1], MIDAS [6], and OASIS [19]. Multi-contrast brain MRI data including T_1- and T_2-weighted images were analyzed. A total of 53 healthy subjects were selected from the IXI dataset, and data were split into 25 training, 10 validation, 18 test subjects. A total of 55 glioma patients were selected from the BRATS dataset, and data were split into 25 training, 10 validation, 20 test subjects. A total of 66 healthy subjects were selected from the MIDAS dataset, and data were split into 48 training, 5 validation, 13 test subjects. Lastly, a total of 48 healthy subjects were selected from the OASIS dataset, and data were split into 22 training, 9 validation, 17 test subjects. In each subject, nearly 100 axial cross-sections centrally located within the volume were included.

Table 1. Performance of centralized and federated translation models in $T_1 \rightarrow T_2$ and $T_2 \rightarrow T_1$ tasks. Higher PSNR, SSIM scores and lower FID score indicate improved performance. Bold font indicates the top performing federated model for each task.

		IXI		BRATS		MIDAS		OASIS	
		$T_1 \rightarrow T_2$	$T_2 \rightarrow T_1$	$T_1 \rightarrow T_2$	$T_2 \rightarrow T_1$	$T_1 \rightarrow T_2$	$T_2 \rightarrow T_1$	$T_1 \rightarrow T_2$	$T_2 \rightarrow T_1$
Centralized	PSNR	28.6 ± 1.3	27.9 ± 1.1	26.1 ± 0.9	24.5 ± 1.9	28.1 ± 0.5	25.9 ± 1.2	25.2 ± 0.6	21.2 ± 0.8
	SSIM	94.3 ± 1.3	94.4 ± 1.2	93.0 ± 1.1	92.5 ± 1.1	91.9 ± 0.9	87.0 ± 2.1	83.7 ± 2.4	76.8 ± 1.9
	FID	7.4	27.3	24.9	14.2	9.7	11.9	18.1	18.6
SPFL-Trans	PSNR	**28.0 ± 1.4**	**27.6 ± 1.0**	**26.0 ± 0.9**	**24.7 ± 1.6**	**27.9 ± 0.5**	**26.0 ± 1.1**	**24.7 ± 0.5**	**20.9 ± 0.7**
	SSIM	**94.1 ± 1.2**	**94.1 ± 1.2**	**92.8 ± 1.1**	**92.5 ± 1.0**	**91.6 ± 0.9**	**86.4 ± 2.1**	**82.1 ± 2.0**	**75.0 ± 2.6**
	FID	**9.0**	**31.3**	**26.8**	**16.9**	**9.4**	**11.1**	**32.0**	**24.5**
FedGAN	PSNR	26.6 ± 1.1	26.4 ± 0.8	25.2 ± 1.2	22.9 ± 0.7	27.0 ± 0.5	24.6 ± 0.9	21.7 ± 0.5	20.8 ± 1.4
	SSIM	91.8 ± 1.7	92.8 ± 1.1	91.1 ± 1.3	88.8 ± 1.0	89.6 ± 1.4	81.5 ± 2.2	67.7 ± 3.5	74.3 ± 3.5
	FID	14.7	39.2	45.2	47.0	13.3	20.6	43.7	37.7
FedMRI	PSNR	27.5 ± 1.0	27.4 ± 1.1	25.7 ± 0.7	24.7 ± 1.2	27.6 ± 0.6	25.9 ± 1.2	23.6 ± 0.5	20.7 ± 1.0
	SSIM	93.6 ± 1.3	94.0 ± 1.2	92.6 ± 1.1	91.6 ± 0.9	91.3 ± 1.0	85.4 ± 2.1	80.6 ± 2.1	72.1 ± 2.6
	FID	11.2	35.6	34.9	19.4	12.2	12.3	33.2	27.2
FedMedGAN	PSNR	25.6 ± 1.2	24.9 ± 0.7	24.7 ± 1.2	20.9 ± 0.8	26.8 ± 0.5	22.5 ± 0.6	21.5 ± 0.4	20.5 ± 1.3
	SSIM	91.7 ± 1.8	91.6 ± 1.2	90.8 ± 1.3	86.5 ± 0.8	89.9 ± 0.9	82.3 ± 1.9	62.9 ± 3.8	72.6 ± 3.5
	FID	19.8	48.3	36.1	58.6	16.8	84.8	59.1	35.3

3.2 Competing Methods

We demonstrated the proposed approach against a centrally-trained translation model [8], and FL-based translation models including FedGAN [28], FedMRI [11] and FedMedGAN [37]. The centralized model and FedGAN was implemented with matching architecture to the proposed model, except for the AdaIN layers that were excluded. FedMRI was implemented with a U-Net backbone, where encoders were shared while decoders were kept site-specific as originally proposed in [11]. The loss function of the proposed model was adopted for fair comparison. FedMedGAN was implemented with a U-Net backbone as originally proposed in [37]. However, the loss function of the proposed model was used in FedMedGAN as opposed to cycle-consistency loss for fair comparison in the paired translation tasks reported here. FL-based models followed the same federated optimization procedure as the proposed approach. Hyperparameter selection was performed for each model in order to maximize the performance in the validation set. Shared generators across sites and site-specific local discriminators were adopted for all FL models considered here.

3.3 Experiments

SPFL-Trans was implemented with an LPS with 6 dense layers to produce latent variables. The encoder in the generator had three convolutional layers of kernel size $7, 3, 3$. The bottleneck contained 9 residual blocks of kernel size 3. The decoder had three convolutional layers of kernel size $3, 3, 7$. Discriminators for all competing methods were based on the PatchGAN architecture [16]. Network weights were learned via the Adam optimizer run at $\beta_1 = 0.5$ and $\beta_2 = 0.999$. Training was continued for $P = 150$ rounds. A fixed learning rate was selected as

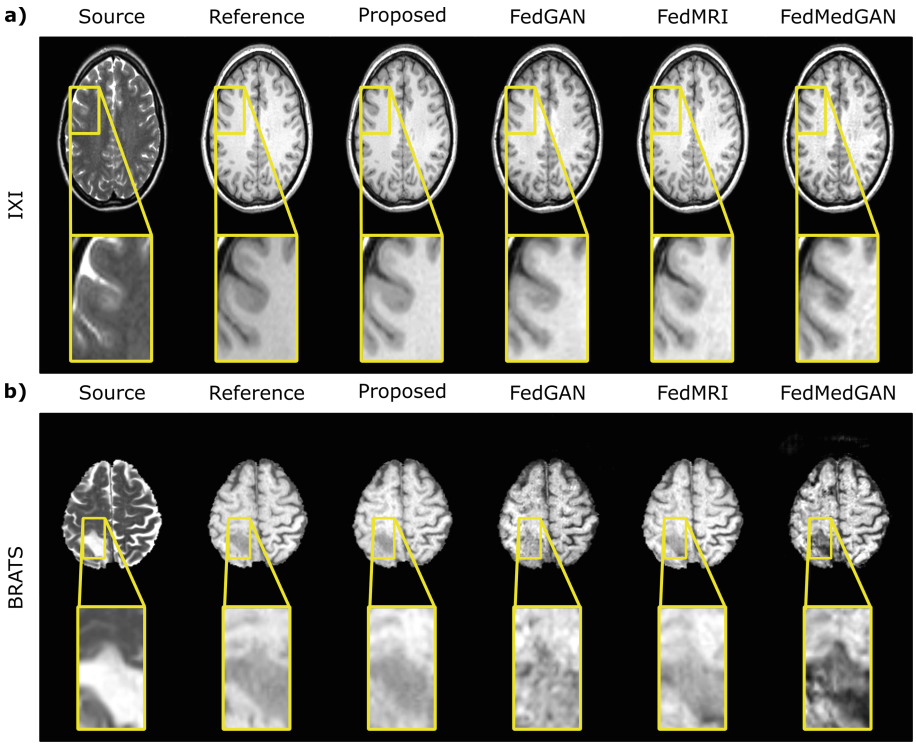

Fig. 2. Synthesized images in IXI and BRATS for the $T_2 \rightarrow T_1$ task. Results are shown for all competing methods, and they are displayed along with the source images and the reference target images.

0.0002 during the first 75 rounds, and it was linearly decayed to 0 during the last 75 rounds. The relative weight of the pixel-wise loss was selected as $\lambda_{pix} = 100$. Modeling was performed via Pytorch framework.

We considered learning two individual one-to-one translation tasks ($T_1 \rightarrow T_2$, $T_2 \rightarrow T_1$ where the mapping is denoted as $source \rightarrow target$) in an FL setup with 4 sites. Translation performance was evaluated via PSNR, SSIM, and Fretchet Inception Distance (FID) metrics. PSNR and SSIM were measured between synthetic and reference target-contrast images for individual cross-sections, and averaged across the volume. Results were reported as mean and standard deviation across test subjects within each individual dataset.

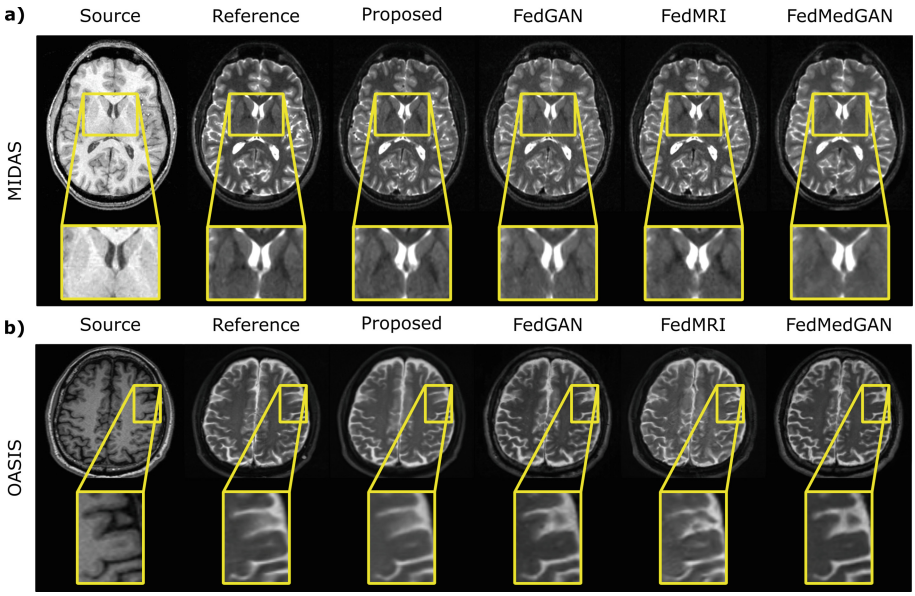

Fig. 3. Synthesized images in MIDAS and OASIS for the $T_1 \rightarrow T_2$ task. Results are shown for all competing methods, and they are displayed along with the source images and the reference target images.

4 Results

To demonstrate the effectiveness of the proposed approach in federated learning of multi-contrast MRI translation, we compared it against state-of-the-art FL-based translation models and a centrally-trained model as a performance baseline. Quantitative performance metrics for the competing methods in each dataset and in each task are listed in Table 1. The proposed approach yields the highest performance across tasks and across sites. On average, SPFL-Trans achieves 1.3 dB higher PSNR, 3.3 % higher SSIM, and 13.5 point lower FID over competing methods. Representative target images from all competing methods are shown in Fig. 3 for the $T_1 \rightarrow T_2$ task in MIDAS and OASIS datasets. Representative results are shown in Fig. 2 for the $T_2 \rightarrow T_1$ task in IXI and BRATS datasets. SPFL-Trans yields superior translation performance in regions where competing models have inaccurate tissue depiction, especially near gray matter and pathology. Overall, SPFL-Trans generates images with fewer artifacts and lower noise levels compared to baselines. These quantitative and qualitative assessments indicate that site-specific modulation of feature statistics in SPFL-Trans enhances translation performance compared to competing federated models.

We also conducted an ablation study to investigate the benefits of the statistical modulation mechanism achieved by AdaIN layers in federated learning of MRI synthesis. To do this, we compared SPFL-Trans with an ablated variant

where LPS and AdaIN layers were jointly removed (w/o LPS and AdaIN) from the generator. Quantitative performance metrics for SPFL-Trans and w/o LPS and AdaIN models are listed in Table 2. The proposed method outperforms the ablated variant across all sites and tasks except for SSIM. These results signals the performance gain brought by the LPS subnetwork and AdaIN layers in SPFL-Trans.

Table 2. Performance of SPFL-Trans and a variant ablated of LPS and AdaIN layers in $T_1 \rightarrow T_2$ and $T_2 \rightarrow T_1$ tasks. Higher PSNR, SSIM scores and lower FID score indicate improved performance. Bold font indicates the top performing model for each task.

		IXI		BRATS		MIDAS		OASIS	
		$T_1 \rightarrow T_2$	$T_2 \rightarrow T_1$	$T_1 \rightarrow T_2$	$T_2 \rightarrow T_1$	$T_1 \rightarrow T_2$	$T_2 \rightarrow T_1$	$T_1 \rightarrow T_2$	$T_2 \rightarrow T_1$
SPFL-Trans	PSNR	**28.0 ± 1.4**	**27.6 ± 1.0**	**26.0 ± 0.9**	**24.7 ± 1.6**	**27.9 ± 0.5**	**26.0 ± 1.1**	**24.7 ± 0.5**	**20.9 ± 0.7**
	SSIM	**94.1 ± 1.2**	**94.1 ± 1.2**	**92.8 ± 1.1**	**92.5 ± 1.0**	**91.6 ± 0.9**	**86.4 ± 2.1**	**82.1 ± 2.0**	75.0 ± 2.6
	FID	**9.0**	**31.3**	**26.8**	**16.9**	**9.4**	**11.1**	**32.0**	**24.5**
w.o. LPS and AdaIN	PSNR	26.4 ± 1.2	26.5 ± 0.8	25.6 ± 0.9	23.2 ± 0.9	27.1 ± 0.4	24.3 ± 0.8	22.0 ± 0.5	20.8 ± 1.3
	SSIM	90.9 ± 2.1	93.1 ± 1.3	90.1 ± 1.6	89.5 ± 0.9	89.6 ± 1.3	80.9 ± 2.4	65.3 ± 3.9	**77.2 ± 3.0**
	FID	14.8	43.2	42.4	48.0	12.9	19.8	43.6	39.3

5 Discussion and Conclusion

Federated MRI translation involves multi-site imaging data collected under different settings, so it has to operate reliably under distributional heterogeneity [22]. In this context, the proposed approach offers a site specificity-preserving global MRI translation model for multi-institutional collaborations. Experiments on public multi-site brain MRI data demonstrate that SPFL-Trans offers competitive performance to a centralized baseline model, while significantly outperforming alternative federated baselines both visually and quantitatively. Our results suggest that SPFL-Trans can improve generalizability and flexibility in multi-site collaborations by enabling training on imaging data from diverse sites and protocols. Improved generalization against domain shifts in the distribution of MRI data renders SPFL-Trans a promising candidate for multi-site training of MRI contrast translation models. In the future, the proposed approach might also be adopted for cross-modal image translation tasks.

References

1. Bakas, S., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. arXiv:1811.02629 (2019)
2. Bakas, S., et al.: Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. Sci. Data **4**, 1–13 (2017)
3. Bdair, T., Navab, N., Albarqouni, S.: FedPerl: semi-supervised peer learning for skin lesion classification. In: de Bruijne, M., et al. (eds.) MICCAI 2021, Part III. LNCS, vol. 12903, pp. 336–346. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87199-4_32

4. Beers, A., et al.: High-resolution medical image synthesis using progressively grown generative adversarial networks. arXiv:1805.03144 (2018)
5. Bercea, C.I., Wiestler, B., Rueckert, D., Albarqouni, S.: FedDis: disentangled federated learning for unsupervised brain pathology segmentation (2021). https://doi.org/10.48550/ARXIV.2103.03705. https://arxiv.org/abs/2103.03705
6. Bullitt, E., et al.: Vessel tortuosity and brain tumor malignancy. Acad. Radiol. **12**, 1232–40 (2005). https://doi.org/10.1016/j.acra.2005.05.027
7. Dalmaz, O., Yurt, M., Çukur, T.: ResViT: residual vision transformers for multimodal medical image synthesis. IEEE Trans. Med. Imaging 1 (2022). https://doi.org/10.1109/TMI.2022.3167808
8. Dar, S.U., Yurt, M., Karacan, L., Erdem, A., Erdem, E., Çukur, T.: Image synthesis in multi-contrast MRI with conditional generative adversarial networks. IEEE Trans. Med. Imaging **38**(10), 2375–2388 (2019). https://doi.org/10.1109/TMI.2019.2901750
9. Dayan, I., et al.: Federated learning for predicting clinical outcomes in patients with COVID-19. Nat. Med. **27**, 1–9 (2021). https://doi.org/10.1038/s41591-021-01506-3
10. Elmas, G., et al.: Federated learning of generative image priors for MRI reconstruction. arXiv:2202.04175 (2022)
11. Feng, C.M., Yan, Y., Fu, H., Xu, Y., Shao, L.: Specificity-preserving federated learning for MR image reconstruction. arXiv:2112.05752 (2021)
12. Han, T., et al.: Breaking medical data sharing boundaries by using synthesized radiographs. Sci. Adv. **6**(49), eabb7973 (2020)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Comput. Vis. Pattern Recognit, pp. 770–778 (2016)
14. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: ICCV (2017)
15. Iglesias, J.E., Konukoglu, E., Zikic, D., Glocker, B., Van Leemput, K., Fischl, B.: Is synthesizing MRI contrast useful for inter-modality analysis? In: Medical Image Computing and Computer-Assisted Intervention, pp. 631–638 (2013)
16. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Computer Vision and Pattern Recognition, pp. 1125–1134 (2017)
17. Kaissis, G.A., Makowski, M.R., Rüeckert, D., Braren, R.F.: Secure, privacy-preserving and federated machine learning in medical imaging. Nat. Mach. Intell. **2**(6), 305–311 (2020)
18. Krupa, K., Bekiesińska-Figatowska, M.: Artifacts in magnetic resonance imaging **80**, 93–106 (2015)
19. LaMontagne, P.J., et al.: OASIS-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and Alzheimer disease. medRxiv (2019). https://doi.org/10.1101/2019.12.13.19014902. https://www.medrxiv.org/content/early/2019/12/15/2019.12.13.19014902
20. Lee, D., Kim, J., Moon, W.J., Ye, J.C.: CollaGAN: collaborative GAN for missing image data imputation. In: Computer Vision and Pattern Recognition, pp. 2487–2496 (2019)
21. Li, W., et al.: Privacy-preserving federated brain tumour segmentation. In: Suk, H.-I., Liu, M., Yan, P., Lian, C. (eds.) MLMI 2019. LNCS, vol. 11861, pp. 133–141. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32692-0_16
22. Li, X., Jiang, M., Zhang, X., Kamp, M., Dou, Q.: FedBN: federated learning on non-IID features via local batch normalization. In: International Conference on Learning Representations (2021). https://openreview.net/pdf?id=6YEQUn0QICG

23. Liu, Q., Chen, C., Qin, J., Dou, Q., Heng, P.: FedDG: federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. In: CVPR, pp. 1013–1023 (2021)
24. McMahan, H.B., Moore, E., Ramage, D., Hampson, S., Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: AISTATS (2017)
25. McMahan, H.B., Moore, E., Ramage, D., Hampson, S., Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data (2016)
26. Moraal, B., et al.: Multi-contrast, isotropic, single-slab 3D MR imaging in multiple sclerosis. Neuroradiol. J. **22**, 33–42 (2009)
27. Pati, S., et al.: The federated tumor segmentation (FeTS) challenge (2021). https://doi.org/10.48550/ARXIV.2105.05874. https://arxiv.org/abs/2105.05874
28. Rasouli, M., Sun, T., Rajagopal, R.: FedGAN: federated generative adversarial networks for distributed data. arXiv:2006.07228 (2020)
29. Rieke, N., et al.: The future of digital health with federated learning. NPJ Digit. Med. **3**(1), 119 (2020). https://doi.org/10.1038/s41746-020-00323-1
30. Roth, H.R., et al.: Federated learning for breast density classification: a real-world implementation. In: DART, DCL, pp. 181–191 (2020)
31. Roth, H.R., et al.: Federated whole prostate segmentation in MRI with personalized neural architectures. In: de Bruijne, M., et al. (eds.) MICCAI 2021. LNCS, vol. 12903, pp. 357–366. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87199-4_34
32. Sheller, M., et al.: Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Sci. Rep. **10** (2020). https://doi.org/10.1038/s41598-020-69250-1
33. Sheller, M.J., Reina, G.A., Edwards, B., Martin, J., Bakas, S.: multi-institutional deep learning modeling without sharing patient data: a feasibility study on brain tumor segmentation. In: Crimi, A., Bakas, S., Kuijf, H., Keyvan, F., Reyes, M., van Walsum, T. (eds.) BrainLes 2018. LNCS, vol. 11383, pp. 92–104. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-11723-8_9
34. Thukral, B.: Problems and preferences in pediatric imaging **25**, 359–364 (2015)
35. Wang, J., et al.: A field guide to federated optimization (2021)
36. Wei, W., et al.: Fluid-attenuated inversion recovery MRI synthesis from multi-sequence MRI using three-dimensional fully convolutional networks for multiple sclerosis **6**(1), 014005 (2019)
37. Xie, G., et al.: FedMed-GAN: federated domain translation on unsupervised cross-modality brain image synthesis (2022)
38. Yang, D., et al.: Federated semi-supervised learning for COVID region segmentation in chest CT using multi-national data from China, Italy, Japan. Med. Image Anal. **70**, 101992 (2021). https://doi.org/10.1016/j.media.2021.101992. https://www.sciencedirect.com/science/article/pii/S1361841521000384
39. Yu, B., Zhou, L., Wang, L., Shi, Y., Fripp, J., Bourgeat, P.: Ea-GANs: edge-aware generative adversarial networks for cross-modality MR image synthesis. IEEE Trans. Med. Imaging **38**(7), 1750–1762 (2019). https://doi.org/10.1109/TMI.2019.2895894
40. Zhou, S., Landman, B.A., Huo, Y., Gokhale, A.: Communication-efficient federated learning for multi-institutional medical image classification. In: Deserno, T.M., Park, B.J. (eds.) Medical Imaging 2022: Imaging Informatics for Healthcare, Research, and Applications, vol. 12037, pp. 6–12. International Society for Optics and Photonics. SPIE (2022). https://doi.org/10.1117/12.2611654