

Federated Multi-Organ Segmentation With Inconsistent Labels

Xuanang Xu^{ID}, Hannah H. Deng^{ID}, Jamie Gateno, and Pingkun Yan^{ID}, *Senior Member, IEEE*

Abstract—Federated learning is an emerging paradigm allowing large-scale decentralized learning without sharing data across different data owners, which helps address the concern of data privacy in medical image analysis. However, the requirement for label consistency across clients by the existing methods largely narrows its application scope. In practice, each clinical site may only annotate certain organs of interest with partial or no overlap with other sites. Incorporating such partially labeled data into a unified federation is an unexplored problem with clinical significance and urgency. This work tackles the challenge by using a novel federated multi-encoding U-Net (Fed-MENU) method for multi-organ segmentation. In our method, a multi-encoding U-Net (MENU-Net) is proposed to extract organ-specific features through different encoding sub-networks. Each sub-network can be seen as an expert of a specific organ and trained for that client. Moreover, to encourage the organ-specific features extracted by different sub-networks to be informative and distinctive, we regularize the training of the MENU-Net by designing an auxiliary generic decoder (AGD). Extensive experiments on six public abdominal CT datasets show that our Fed-MENU method can effectively obtain a federated learning model using the partially labeled datasets with superior performance to other models trained by either localized or centralized learning methods. Source code is publicly available at <https://github.com/DIAL-RPI/Fed-MENU>.

Index Terms—Federated learning, deep learning, medical image segmentation, inconsistent labels.

I. INTRODUCTION

CONVOLUTIONAL neural network (CNN) based deep learning (DL) as a data-driven methodology has demonstrated unparalleled performance in various segmentation

Manuscript received 26 March 2023; accepted 17 April 2023. Date of publication 25 April 2023; date of current version 2 October 2023. This work was supported in part by the National Science Foundation (NSF) under CAREER Award OAC 2046708 and in part by the National Institutes of Health (NIH) under Award R01DE022676 and Award R21EB028001. (Corresponding author: Pingkun Yan.)

Xuanang Xu and Pingkun Yan are with the Department of Biomedical Engineering and the Center for Biotechnology and Interdisciplinary Studies, Rensselaer Polytechnic Institute, Troy, NY 12180 USA (e-mail: xux12@rpi.edu; yanp2@rpi.edu).

Hannah H. Deng is with the Department of Oral and Maxillofacial Surgery, Houston Methodist Research Institute, Houston, TX 77030 USA (e-mail: hdeng@houstonmethodist.org).

Jamie Gateno is with the Department of Oral and Maxillofacial Surgery, Houston Methodist Research Institute, Houston, TX 77030 USA, and also with the Department of Surgery (Oral and Maxillofacial Surgery), Weill Medical College, Cornell University, New York, NY 10065 USA (e-mail: jgateno@houstonmethodist.org).

Digital Object Identifier 10.1109/TMI.2023.3270140

tasks, providing that the model can train on large-scale data with sufficient diversity. To suffice the large appetite of CNNs, researchers often collect data from multiple sources to jointly train a model for improved performance. However, in the healthcare domain, such *centralized learning* paradigm is often impractical because the clinical data cannot be easily shared across different institutions due to the regulations, such as Health Insurance Portability and Accountability Act (HIPAA). To overcome the barrier of data privacy and realize large-scale DL on isolated data, federated learning (FL) [1], an emerging *decentralized learning* paradigm, has been adopted in the medical image analysis, solving various clinical problems such as prostate segmentation [2], [3] and COVID-19 diagnosis [4], [5].

FL allows different data-owners to collaboratively train one global DL model without sharing the data. The model training is completed by iterating over a server node and several client nodes. Each client individually trains a copy of the global model using their local data after the server updates the global model by aggregating the locally trained models. By repeating this process, the global model can effectively absorb the knowledge contained in the client datasets without data sharing. As data privacy has become a critical issue concerned by healthcare stakeholders, FL attracts a growing attention from both the research and clinical communities in recent years.

Although the capability of FL for medical image analysis has been demonstrated by the prior studies, it comes with limitations. A critical issue is about the strict requirement for the label consistency. Specifically, in an FL system, all the participating sites need to have identical regions-of-interest (ROIs) annotated on their local data, so the FL model can be optimized following the same objective across different clients. However, in practical scenarios, different clinical sites often have different expertise and thus follow different protocols for data annotation. This leads to inconsistent ROI labels across different sites. The requirement for labeling consistency hinders the FL methods from large-scale deployment in practice. Therefore, a more flexible FL framework supporting the training with inconsistently labeled client data is highly desired. From a technical point of view, this is an FL problem with partial labels since each client merely has a partially labeled dataset with respect to the whole set of ROIs in the federation. To the best of our knowledge, this is a new problem that has not yet been explored before, but at the same time is of great clinical significance and technical urgency to tackle.

In this paper, we propose a novel method, called federated multi-encoding U-Net (Fed-MENU), to address the above challenge. We then demonstrate its performance on a representative task, *i.e.*, multiple abdominal organ segmentation from computed tomography (CT) images. The underlying assumption of our design is that, although the data from different clients show disparities in terms of the labeled ROIs, they share the same or similar image contents and thus can provide complementary information to facilitate the learning of robust features in a unified FL system. Since the client datasets are partially labeled with different ROIs, they can be seen as a set of experts with different expertise. Each expert focuses on learning the features within its expertise (labeled ROIs) while avoiding introducing the unreliable or noisy information from the non-expertise (unlabeled ROIs). To achieve this goal, we design a multi-encoding U-Net (MENU-Net) to decompose the multi-organ feature learning task into a series of individual sub-tasks of organ-specific feature learning. Each of them is bound with a sub-network in the MENU-Net. During the local training stage of FL, the client node can only tune the sub-network associated with the organ on which they have labels (expertise) while keeping the other parts of the network unchanged. Due to this decomposed feature learning strategy, the MENU-Net is encouraged to learn organ-specific features via different sub-networks without interference from other label-absent organs. Moreover, to further encourage the extracted organ-specific features to be informative and distinctive from other organs, we design an auxiliary generic decoder (AGD) to regularize the training of the MENU-Net. As a result, the MENU-Net can extract discriminative features during its encoding stage and thus facilitate the segmentation in the following decoding stage.

To demonstrate the performance of our Fed-MENU method, we conducted extensive experiments using four public abdominal CT datasets, each of which is annotated with a different set of abdominal organs. Our results show that, without sharing the raw data, the proposed Fed-MENU method can effectively utilize the isolated datasets with different partial labels to train a global model for multi-organ segmentation. The performance of the trained model is superior to the localized learning model trained by any single dataset and also the centralized learning model trained by combining all the datasets. The main contributions of this paper are three-fold.

- We addressed a new problem in FL to enable collaboratively training of a global model using isolated datasets with inconsistent labels.
- We proposed a novel Fed-MENU method to deal with this challenging problem. In our method, organ-specific features are individually extracted by different sub-networks of the MENU-Net and further enhanced by the accompanied AGD.
- We evaluated the performance of our Fed-MENU method for multi-organ segmentation using six public abdominal CT datasets. Our experimental results on both in-federation and out-of-federation testing sets demonstrated the effectiveness and superior performance of our design.

The rest of this paper is organized as follows. Section II gives a brief review of the previous literature related to this work. Section III illustrates the details of the proposed method. Section IV presents the experiments and results on four public datasets. Finally, we discuss the limitation of this work and conclude it in Section V.

II. RELATED WORKS

Since our study involves both FL-based medical image analysis and DL-based image segmentation with partially labeled data, we review the related works in these two areas before presenting our proposed method.

A. Background of Federated Learning

FL, first emerged in 2017 [1], is a decentralized learning paradigm designed to address data privacy issue in deep learning. Unlike the conventional centralized deep learning that requires all training data gathered on a server node, FL allows distributed data owners to collaboratively train a model without sharing their data. Federated averaging (FedAvg) [1] was acknowledged as the benchmarking algorithm in FL, which defined the framework for the following studies in this field. In FedAvg, the model is simultaneously trained on several client devices and the clients transmit the trained model (parameters) to the server to update/maintain the global model (parameters). A major challenge in FL is how to deal with the client data that is not independently and identically distributed (non-iid) or heterogeneous [6], [7], [8], [9], [10], [11], [12]. For example, Li et al. [7] derived FedProx from FedAvg by regularizing the local models to be closer to the global model, in which way the method showed more stable and accurate performance when learning over the heterogeneous data. Acar et al. [8] proposed FedDyn by introducing a dynamic regularization term to the local training objective of FedAvg. The resulting model achieved not only better robustness against data heterogeneity but also higher communication efficiency. Li et al. [9] proposed a simple yet effective scheme, FedBN, to address the heterogeneity issue in feature space by keeping the batch normalization layers locally updated. Except for the above FL algorithms, data harmonization [13] can also be considered as a way to address the heterogeneities of multi-center data in FL.

B. Federated Learning for Medical Image Analysis

In recent years, FL has been widely applied in various medical image analysis tasks to address the conflict between large-scale DL model training and healthcare data privacy. For example, Dayan et al. [5] trained an FL model for COVID-19 clinical outcome prediction using chest X-ray images collected from 20 institutes across the globe. The experimental results showed that, by utilizing multi-site data, the trained FL model achieved an average area under the curve (AUC) around 0.920 for predicting the patients oxygen therapy categories after 24- and 72-hour periods from initial admission to the emergency department. This AUC value of the FL model is

16% higher than the average AUC (0.795) of the 20 locally trained models with single-site data. Xia et al. [14] derived an auto-FedAvg method from FedAvg algorithm [1] for medical image segmentation using data that are not independently and identically distributed (non-iid). Their method automatically learns the aggregation weights of each client based on the client data distribution. Roth et al. [2] combined FL with neural architecture search strategy to develop a super-network with a self-learned structure for whole prostate segmentation from multi-institute magnetic resonance imaging (MRI) data. Recently, Yang et al. [4] extended FL to semi-supervised learning paradigm and applied it to COVID region segmentation using chest CT data from three nations. Another work by Park et al. [15] employed the Vision Transformer (ViT) architecture [16] in an FL framework to diagnose the COVID-positive cases from chest X-ray images.

The prior studies [2], [4], [5], [12], [14], [17], [18], [19], [20] have demonstrated the feasibility and effectiveness of FL in solving the problem of large-scale DL model training without data sharing in the healthcare community. However, these successes are built upon a prerequisite that the participating sites have consistently labeled data so that they can contribute to the same training objective of one global model. However, such a strong requirement may not be met in the real scenarios, where different clinical sites often follow different protocols to annotate their local data. Although some prior efforts like FedMix proposed by Wicaksana et al. [21] allowed mixed supervised FL with different levels of label (*e.g.*, pixel-level mask, bounding-box-level annotation, and image-level labels), the gap caused by the label inconsistency (in terms of the target classes) remained unresolved, which largely narrows the application scope of the FL-based methods for medical image analysis and thus motivates our study in this work.

It is worth noting that Dong et al. [22], [23] worked on a similar problem of federated learning-based X-ray image classification with partially labeled data, which is related to our work and thus provided positive implications for us during the revision of this paper. However, there are significant differences between these two works. First of all, Dong et al.'s work is targeting at the image classification task while our work focused on the image segmentation task, which is more general and common in terms of the scope of clinical application. As a consequence, the methods proposed in these two works are fundamentally different and thus cannot be directly compared. Moreover, the method by Dong et al. was only evaluated in an in-federation setting, *i.e.*, the testing data follows the distribution of the training data from one of the clients, while our study also includes an out-of-federation evaluation, *i.e.*, the testing data comes from an unseen domain out of any one of the client datasets, which is more exhaustive and challenging.

C. Medical Image Segmentation With Partial Labels

Due to the high cost of data annotation, medical images are often partially labeled with different ROIs or labels, even though they may share the same imaging field. This created a barrier for the DL-based medical image segmentation

methods [24], [25] when the model is trained on the partially labeled data. To solve this problem, Yan et al. [26] developed a universal lesion detection algorithm trained by CT images labeled with different lesion types. In their method, the missing labels were mined by exploiting clinical prior knowledge and cross-dataset knowledge transfer. Recently, Petit et al. [27] proposed to conduct partially labeled DL training through an iterative confidence relabeling method, in which a self-supervised scheme was employed to iteratively relabel the missing organs by introducing pseudo labels into the training set. For the partial label problem in segmentation tasks, Fang and Yan [28] proposed a pyramid-input-pyramid-output feature abstraction network (PIPO-FAN) for multi-organ segmentation, in which a target adaptive loss is integrated with a unified training strategy to enable image segmentation over multiple partially labeled datasets with a single model. Shi et al. [29] designed two types of loss functions, namely marginal loss and exclusion loss, to train a multi-organ segmentation network using a set of partially labeled datasets. Furthermore, Zhang et al. [30] proposed an ingenious approach, namely dynamic on-demand network (DoDNet), to achieve multi-organ segmentation with partially labeled data. Unlike the conventional deep neural networks with fixed parameters after training, the DoDNet can dynamically generate model parameters to adapt to different organ segmentation tasks. This innovative design largely improved the model efficiency and flexibility.

Although the problem of training DL models using partially labeled data [26], [27], [28], [29], [30], [31] had been studied in the context of centralized learning, it is still an unexplored area for the field of FL. In this paper, we present the Fed-MENU method, which can learn from partially labeled data distributed on different sites.

III. METHOD

Fig. 1 gives an overview of the proposed Fed-MENU method. For a better understanding, it is presented in a specific scenario where the federation contains three client nodes and each node possesses a CT dataset partially labeled with one of the abdominal organs (liver, kidney, and pancreas). We then describe the FL framework adopted by our method in Section III-A. The technical innovation of this work resides in the MENU-Net designed for organ-specific feature extraction (Section III-B) and the AGD designed for organ-specific feature enhancement (Section III-C). The training configurations and other technical details of the proposed method are provided in Section III-D.

A. Federated Learning Framework

Our method inherits the framework of the FedAvg [1] algorithm, which is a commonly adopted benchmark in FL. It consists of one server node and K client nodes. The server takes the responsibility of coordinating the communication and computation across different clients, while the clients focus on model training using their local data and devices. Given the universal label set of M organs, each of the client nodes has a local dataset D_k partially labeled with a subset of

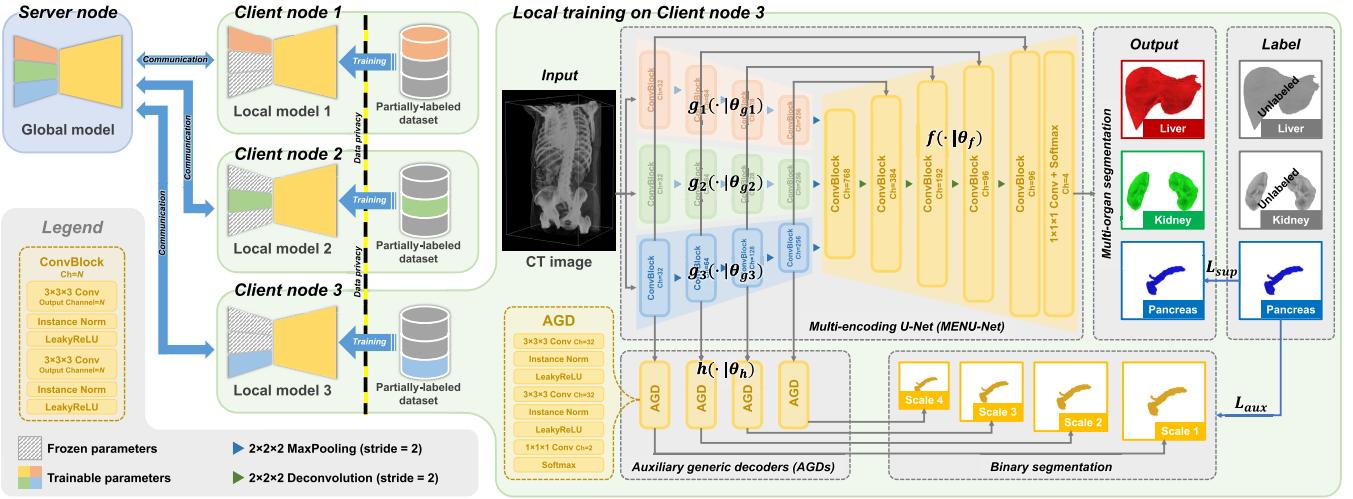


Fig. 1. Scheme of the proposed Fed-MENU method. For better understanding, the method is presented using an example scenario where three clinical sites with partially labeled abdominal CT images collaboratively train an FL model for multi-organ segmentation of the liver, kidney, and pancreas. The left panel of the figure shows the federation structure. The right panel details the local training procedure on Client node 3, where the pancreas is the only labeled organ in this client dataset. For brevity, we present only the connections between the third sub-encoder (colored in blue) and the shared decoder, as well as the AGDs. In the final implementation, all sub-encoders are connected to the shared decoder and AGDs at equivalent positions.

$N_k (\leq M)$ organs. Our goal is to train a M -organ segmentation network $F(\cdot, \theta)$ using the partially labeled datasets $\{D_k\}_{k=1}^K$, which reside in different client nodes and cannot be shared or gathered for centralized training.

The training process of the FL framework consists of T rounds of communication between the server and the client nodes. Specifically, in the t -th communication round, each client node k will first download the parameters of the current segmentation network $F(\cdot, \theta^t)$ on server (denoted as global model), resulting in a local copy of $F(\cdot, \theta^t)$ (denoted as local model). The client node then trains the local model using its local dataset D_k and device for E epochs. After the local training, the server collects the trained local models $F(\cdot, \theta_k^t)$ from all the K client nodes and aggregates them into a new global model through a parameter-wise averaging:

$$\theta^{t+1} = \sum_k^K \frac{|D_k|}{\sum_j^K |D_j|} \cdot \theta_k^t, \quad (1)$$

where $|D_k|$ indicates the sample size of dataset D_k . The whole FL procedure is accomplished by repeating the above process until the global model $F(\cdot, \theta^t)$ converges.

Because the local models are individually trained on the client nodes, the server node only needs to transfer the model parameters instead of the raw data from the clients. That helps the FL model obtain knowledge contained in the isolated client datasets without violating the data privacy. However, since the original FedAvg framework is designed for FL with fully labeled data, there is a technical gap to bridge before we can deploy it to the case of partially labeled data. Therefore, we propose the following MENU-Net with AGD to meet the need.

B. MENU-Net for Organ-Specific Feature Extraction

A straightforward way to train a multi-organ segmentation network using the partial labels is to calculate the training

loss merely using the labeled organs while ignoring the unlabeled ones. However, in the case of FL with partial labels, different organs are labeled in the client datasets, focusing on different expertise for image segmentation. As a consequence, the segmentation performance of the trained local models would be biased to the labeled organs. Intuitively, a uniform aggregation of all local models (see Eq. 1) may weaken the expert client models. A more reasonable way may be to promote the strength of the local models on segmenting the labeled organs and avoid the interference from their weaknesses with the unlabeled organs.

To achieve the above goal, we designed the MENU-Net to explicitly decompose multi-organ feature extraction into several individual processes of organ-specific feature extraction. As illustrated in the right panel of Fig. 1, the MENU-Net consists of M sub-encoders $\{g_m(\cdot | \theta_{g_m})\}_{m=1}^M$ followed by a shared decoder $f(\cdot | \theta_f)$ as well as the skip connections between them. Fig. 1 shows an example with $M=3$. Each sub-encoder $g_m(\cdot | \theta_{g_m})$ serves as an organ-specific feature extractor for the m -th organ. An input image x is fed to all the sub-encoders to get the features of all the organs in parallel. The extracted features are then concatenated into one stream and fed to the shared decoder $f(\cdot | \theta_f)$, which is used to interpret the organ-specific features into the multi-organ segmentation masks. This process is formally defined as

$$F(x) = f(g_1(x) \oplus g_2(x) \dots \oplus g_M(x)) = f\left(\bigoplus_{m=1}^M g_m(x)\right), \quad (2)$$

where \oplus denotes channel-wise concatenation.

Given a client node with partial labels on the m -th organ, the local model training only tunes the parameters in the sub-encoder $g_m(\cdot | \theta_{g_m})$ and the shared decoder $f(\cdot | \theta_f)$, which

can be expressed as:

$$\arg \min_{\theta_f, \theta_{g_m} \subset \theta} L_{sup} [F_m(x | \theta), \hat{y}_m], \quad (3)$$

where $F_m(x)$ and \hat{y}_m are the predicted and ground-truth segmentation of the m -th organ, respectively. L_{sup} denotes the supervised loss measuring the similarity between the prediction and ground truth, which can be in any form of pixel-wise loss functions (e.g., Cross-entropy loss and Dice loss [32]) or their combinations.

C. AGD for Organ-Specific Feature Enhancement

Benefiting from the decomposed feature extraction in MENU-Net, organ-specific knowledge from the labeled clients (experts) can be individually learned by different sub-encoders with less interference from other unlabeled clients (non-experts). However, since each local model can only see the images from one client dataset, the organ-specific features learned by the sub-encoders may also include some domain-specific information. This is unfavorable for the subsequent shared decoder. Ideally, the sub-encoder should focus on the structural information of the corresponding organ, which is invariant across domains. In addition, the extracted organ-specific features should be distinctive enough from that of other organs, making the subsequent shared decoder easier to interpret them into different organs' segmentation.

Motivated by the above idea, we design the AGD shared across different sub-encoders to help regularize the training of our MENU-Net, aiming to enhance the extracted organ-specific features. Specifically, as illustrated in the mid-bottom of Fig. 1, given the organ-specific features $g_m(x)$ extracted by an arbitrary sub-encoder $g_m(\cdot | \theta_{g_m})$, we feed them to a set of AGDs to perform an organ-agnostic (binary) segmentation of the m -th organ:

$$G_m(x) = h(g_m(x)). \quad (4)$$

Here, we denote the collection of all AGDs as $h(\cdot | \theta_h)$ parameterized by θ_h . The AGD has a lite structure consisting of three convolutional layers. The first two convolutional layers contain $3 \times 3 \times 3$ kernels followed by instance normalization [34] and LeakyReLU activation [35]. The last convolutional layer has $1 \times 1 \times 1$ kernels with two output channels followed by a softmax activation. Due to the shallow structure of AGD, it has limited representation ability and thus enforces the preceding sub-encoder layers to extract more discriminative features to approach the organ segmentation. On the other hand, since the AGD is working as a generic decoder for different organs when combined with different sub-encoders, the sub-encoders are encouraged to extract features distinctive enough from each other so that the following AGD can interpret them into the corresponding organ's segmentation without extra information.

During the local training stage, the AGDs are tuned together with different sub-encoders over all client nodes by minimizing the segmentation error between the output $G_m(x)$ and the corresponding ground truth \hat{y}_m . This process can be

expressed as:

$$\arg \min_{\theta_h, \theta_{g_m}} L_{aux} [G_m(x | \theta_h, \theta_{g_m}), \hat{y}_m], \quad (5)$$

where L_{aux} denotes the auxiliary loss quantifying the binary segmentation error between the AGDs' output and the corresponding ground truth. We calculate L_{aux} as a sum of the Cross-entropy loss and Dice loss [32]. Note that, since the AGDs are connected to the multi-scale levels of the sub-encoder, they have multiple outputs in different scales. All these outputs will be resampled to the original size of the ground-truth segmentation \hat{y}_m and counted in the auxiliary loss in Eq. 5. After the local training, the tuned parameters of AGDs from all clients are aggregated through the FedAvg algorithm as shown in Eq. 1, which is the same as that of the shared decoder $f(\cdot | \theta_f)$. The AGDs are used to regularize the training of the proposed MENU-Net. During the inference stage, we only utilize the trained MENU-Net to infer the multi-organ segmentation masks.

D. Implementation Details

The proposed method is implemented for 3D segmentation using PyTorch. Model parameters in the segmentation network are initialized using the Xavier algorithm [36] and optimized by an SGD optimizer. The learning rate is initialized to be 0.01 and decayed throughout the training following a poly learning rate policy [37] with a momentum factor of 0.99. We train the model for $T=400$ rounds of communication ($E=1$ epoch of local training per round) and evaluate its performance on the validation set every epoch using Dice similarity coefficient (DSC) as the metric. The model achieving the highest DSC on the validation set is selected as the final model to be evaluated on the testing set. The training batch size is set to 4 on four NVIDIA A100 GPUs. The CT images are resampled to a uniform resolution of $1.0 \times 1.0 \times 1.5 \text{ mm}^3$ (pixel spacing of $1.0 \times 1.0 \text{ mm}^2$ and slice thickness of 1.5 mm) before training. Image patches with a fixed pixel size of $256 \times 256 \times 32$ are randomly cropped from the resampled CT images as the input of the segmentation network. The image intensities are normalized from [-200.0, 400.0] Hounsfield units (HU) to [0.0, 1.0] for a good soft-tissue contrast. Random translation ([−20, 20] mm) and rotation ([−0.1, 0.1] rad) are used to augment the training samples. In the inference time, an unseen CT image is first resampled to the resolution of $1.0 \times 1.0 \times 1.5 \text{ mm}^3$ and divided into a series of patches, which are then fed to the trained segmentation network. The predicted patch-wise segmentation will be assembled to get the final segmentation mask, which is then resampled back to their original CT resolution for quantitative evaluation. Unless otherwise noted, all the competing methods and ablation models in our experiments are trained and evaluated using the same configuration as our method. For better reproducibility, the source code is released at <https://github.com/DIAL-RPI/Fed-MENU>.

IV. EXPERIMENTS

A. Datasets

We conducted experiments using six public abdominal CT image datasets, including 1) liver tumor segmentation

TABLE I
DETAILED INFORMATION OF CLIENT DATASETS INVOLVED IN THIS STUDY. (L: LIVER, K: KIDNEY, P: PANCREAS, S: SPLEEN, G: GALLBLADDER, ✓: LABELED ORGAN, ✕: UNLABELED ORGAN, OoF: OUT-OF-FEDERATION)

Datasets	Image information					Labeled organ					Imaging phase
	# of CT images (train/val./test)	Slice size [in pixel]	# of slice [per image]	Spacing [mm]	Thickness [mm]	L	K	P	S	G	
Client #1: Liver	131 (79/13/39)	512	74~987	0.56~1.00	0.70~5.00	✓	✗	✗	✗	✗	Portal-venous phase
Client #2: Kidney	210 (126/21/63)	{512,796}	29~1059	0.44~1.04	0.50~5.00	✗	✓	✗	✗	✗	Late arterial contrast phase
Client #3: Pancreas	281 (169/28/84)	512	37~751	0.61~0.98	0.70~7.50	✗	✗	✓	✗	✗	Portal-venous phase
Client #4: Spleen	41 (24/5/12)	512	31~168	0.61~0.98	1.50~8.00	✗	✗	✗	✓	✗	Portal-venous phase
Client #5: AMOS	200 (120/20/60)	60~768	64~512	0.45~3.00	0.82~5.00	✓	✓	✓	✓	✓	See Appendix B.1 in [33]
OoF client: BTCV	30 (0/0/30)	512	85~198	0.59~0.98	2.50~5.00	✓	✓	✓	✓	✓	Portal-venous phase
Total	893(518/87/288)	60~796	29~1059	0.44~3.00	0.50~8.00	-	-	-	-	-	-

challenge (LiTS) dataset [38],¹ 2) kidney tumor segmentation challenge (KiTS) dataset [39], [40],² 3) pancreas and 4) spleen segmentation datasets in medical segmentation decathlon challenge [41], [42]³ (Task #7 and #9, respectively), 5) multi-modality abdominal multi-organ segmentation challenge (AMOS) dataset [33],⁴ and 6) multi-atlas labeling beyond the cranial vault challenge (BTCV) dataset [43].⁵ For brevity, we denote the six datasets as the liver, kidney, pancreas, spleen, AMOS, and BTCV, respectively, in the rest of the paper. Table I shows the detailed information of these datasets.

Liver dataset contains 131 images, whose sizes range between [74~987]×[512]×[512] (in $D \times H \times W$ pixels). The in-plane spacing of these CT slices varies from 0.56mm to 1.00mm and the slice thickness varies from 0.70mm to 5.00mm. Each CT image has a pixel-wise annotation of liver and tumor segmentation stored in the same size as the corresponding image. We treat the liver tumor regions as part of the liver in our experiments.

Kidney dataset contains 210 images, whose sizes range between [29~1059]×[512]×{512,796} (in $D \times H \times W$ pixels). The in-plane spacing of CT slices varies from 0.44mm to 1.04mm and the slice thickness varies from 0.50mm to 5.00mm. Each CT image has a pixel-wise annotation of kidney and tumor segmentation stored in the same size as the corresponding image. We treat the kidney tumor regions as part of the kidney in our experiments.

Pancreas dataset contains 281 images, whose sizes range between [37~751]×[512]×[512] (in $D \times H \times W$ pixels). The in-plane spacing of these CT slices varies from 0.61mm to 0.98mm and the slice thickness varies from 0.70mm to 7.50mm. Each CT image has a pixel-wise annotation of pancreas and tumor segmentation stored in the same size as the corresponding image. We treat the pancreas tumor regions as part of the pancreas in our experiments.

Spleen dataset contains 41 images, whose sizes range between [31~168]×[512]×[512] (in $D \times H \times W$ pixels). The in-plane spacing of these CT slices varies from 0.61mm

to 0.98mm and the slice thickness varies from 1.50mm to 8.00mm. Each CT image has a pixel-wise annotation of spleen segmentation stored in the same size as the corresponding image.

AMOS dataset contains 200 images, whose sizes range between [64~512]×[60~768]×[192~768] (in $D \times H \times W$ pixels). The in-plane spacing of these CT slices varies from 0.45mm to 3.00mm and the slice thickness varies from 0.82mm to 5.00mm. Each CT image has a pixel-wise annotation of liver, kidney, pancreas, spleen, and gallbladder segmentation stored in the same size as the corresponding image.

BTCV dataset contains 30 images, whose sizes range between [85~198]×[512]×[512] (in $D \times H \times W$ pixels). The in-plane spacing of these CT slices varies from 0.59mm to 0.98mm and the slice thickness varies from 2.50mm to 5.00mm. Each CT image has a pixel-wise annotation of liver, kidney, pancreas, spleen, and gallbladder segmentation stored in the same size as the corresponding image.

For the liver, kidney, pancreas, spleen, and AMOS datasets, we randomly split each of them into training/validation/testing sets with a fixed ratio of 60%:10%:30%, respectively. The experimental results on the five testing sets are used for *in-federation* evaluation, which indicates the model performance when the testing data follows the same distribution as the training and validation data. For the BTCV dataset, we reserve it as an *out-of-federation* testing set, which is completely unseen to the model during training and validation. The performance on the BTCV dataset gives a good indication of the model's generalization ability.

B. Metrics

We calculated the mean and standard deviation (SD) of Dice similarity coefficient (DSC) and average symmetric surface distance (ASD) for each organ to quantitatively evaluate the segmentation results yielded by different methods. For each labeled organ in a certain client dataset, we first calculate the mean value Q_c^k over all cases to get the result for organ c in this client k . Then, the mean value of the C_k labeled organs $Q^k = 1/C_k \sum_{c=1}^{C_k} Q_c^k$ is calculated as the performance index of client k . Finally, the mean value of the K clients $Q = 1/K \sum_{k=1}^K Q^k$ is considered as the global index indicating the

¹<https://competitions.codalab.org/competitions/17094>

²<https://kits19.grand-challenge.org/home>

³<http://medicaldecathlon.com>

⁴<https://amos22.grand-challenge.org>

⁵<https://www.synapse.org/#!Synapse:syn3193805/wiki/89480>

model accuracy. Instead of performing an overall case-wise averaging, this strategy was to avoid the bias from unbalanced sample numbers among the datasets, ensuring the global DSC and ASD calculated from each dataset with different sizes play an equal role. We also conducted paired *t*-tests on the above metrics to check the statistical significance between different groups of results.

C. Comparison With Benchmarks

1) *Benchmarks*: We compared our method with three benchmarks to demonstrate its effectiveness. The benchmarks include:

a) *Localized learning mode*: Four single-organ segmentation networks and a five-organ segmentation network were individually trained using the liver, kidney, pancreas, spleen, and AMOS datasets, respectively. This benchmark simulated the scenario where the clinical sites cannot share their data with each other and no techniques are adopted to deal with the data privacy issue during the model training.

b) *Centralized learning mode*: A five-organ segmentation network was trained using the combination of the liver, kidney, pancreas, spleen, and AMOS datasets. This benchmark simulated an ideal scenario where the clinical sites can freely share their data without any concern on the data privacy during the model training.

c) *Federated learning mode*: A five-organ segmentation network was trained using the FedAvg [1] algorithm with the liver, kidney, pancreas, spleen, and AMOS datasets distributed on five client nodes. This benchmark simulated a practical scenario where the clinical sites cannot share their data with each other and a naive FL solution is adopted to deal with the data privacy issue during the model training.

We adopted U-Net [44] as the segmentation network for the above models. To ensure the performance gain comes from the design choice rather than the additional network parameters, we extended the base channel number of the U-Net from 32 to 96 to make it have a slightly larger number of parameters than our method. In the leftmost column of Table II and Table III, we specified the trainable parameters in each model. The training configuration of all the methods followed the recommended settings by nnU-Net [37]. For the U-Net in federated learning mode, we employed the marginal loss and exclusion loss [29] as the training objective to adapt it to the partially labeled client data. The core idea of the marginal/exclusion loss is to combine the unlabeled organs with the background category and try to maximize/minimize the overlap with the merged background/foreground ground truth, respectively, which can be formulated as follows:

$$L_{margin} = L_{ce}(p', 1 - \hat{y}') + L_{dice}(p', 1 - \hat{y}') \quad (6)$$

$$L_{exclusion} = -L_{ce}(p', \hat{y}') - L_{dice}(p', \hat{y}') \quad (7)$$

where p' and \hat{y}' denote the predicted segmentation of the merged background category and ground-truth segmentation of the merged foreground category, respectively. We utilized the sum of the Cross-entropy loss and Dice loss [32] to quantify the overlapping regions, which is consistent with the original nnU-Net [37] training settings. For the sake of fairness, the

same objective was also used as the supervised loss L_{sup} in our method.

2) *In-Federation Performance*: Table II shows the quantitative results of the in-federation evaluation. It can be observed that, among the three learning modes, the centralized learning models generally exhibited the upper-bound performance. The FL models significantly outperformed the localized learning models, which justified the feasibility of using different partially labeled datasets to collaboratively train a multi-organ segmentation model through FL. When compared with the baseline U-Net model, our method achieved higher global accuracy in both centralized learning mode and FL mode, which demonstrated the superiority of our designs. Although not all the improvement margins are statistically significant, the standard deviation of our results is much lower than those of the U-Net baseline model, which indicated our method had a more stable performance on the multi-center data. We also observed that, while the region-based DSC metric showed relatively small differences, our model performed much better than the U-Net in terms of the distance-based ASD metric. This phenomenon suggested that there are fewer false positives in our results. It is worth noting that the trainable parameter number of our method is ~ 78 million, which is 24% fewer than that of the U-Net (~ 97 million parameters). This result suggests that the performance gain of our method was coming from the design choice rather than the additional parameters. More importantly, the smaller model size can also benefit the FL due to the lower communication and computation costs.

3) *Out-of-Federation Performance*: Table III presents the quantitative results of the out-of-federation evaluation. Similar to the in-federation scenario, the centralized learning models showed superior accuracy over the FL and localized learning models. However, there was a global performance degradation from the in-federation results to the out-of-federation results due to the data distribution shift from the training domain (the liver, kidney, pancreas, spleen, and AMOS datasets) to an unseen testing domain (the BTCV dataset). Furthermore, we observed that the performance gap between the centralized learning models and the FL models decreased for the out-of-federation results, as compared to the in-federation results. These results suggest that out-of-federation segmentation is a more challenging task than in-federation segmentation, and our method can effectively handle it with better performance over the baseline U-Net model.

4) *Result Visualization*: Figs. 2 and 3 visualize the segmentation results yielded by different methods in 2D (axial) and 3D views, respectively. Each column shows a case from one client dataset. In the 2D views, the ground-truth contours are superimposed on the segmentation results as yellow dashed lines for better comparison. In the 3D views, the segmentation results of the unlabeled organs are represented by transparent meshes. It can be observed that our proposed method produced fewer false positive results in either in-federation or out-of-federation results.

D. Ablation Study

In this section, we conducted ablation studies on the proposed method to evaluate the effectiveness of the two

TABLE II

QUANTITATIVE PERFORMANCE EVALUATION OF DIFFERENT METHODS WHEN TESTED ON THE IN-FEDERATION DATA. THE BEST RESULTS ARE MARKED IN **BOLD**. THE UNDERLINED RESULTS INDICATE A STATISTICALLY SIGNIFICANT DIFFERENCE FROM OUR RESULT ($p < 0.05$)

Models (# of param.)	DSC [Mean(SD) %]										Global		
	Client #1		Client #2		Client #3		Client #4		Client #5				
	Liver	Kidney	Pancreas	Spleen	Liver	Kidney	Pancreas	Spleen	Gallbladder				
Localized													
- Client #1	93.46(3.20)	-	-	-	-	-	-	-	-				
- Client #2	-	<u>91.17</u> (7.60)	-	-	-	-	-	-	-				
- Client #3	-	-	<u>76.82</u> (13.01)	-	-	-	-	-	-		88.50		
- Client #4	-	-	-	<u>91.98</u> (3.75)	-	-	-	-	-				
- Client #5	-	-	-	-	95.93(2.00)	94.80(2.81)	78.91(11.06)	95.47(2.78)	80.29(19.40)				
Centralized													
- U-Net (97M)	95.22(2.68)	95.52(3.30)	80.50(11.82)	95.52(1.82)	96.69(1.28)	93.27(4.48)	82.25(9.37)	96.17(2.95)	84.76(17.77)		91.48		
- MENU-Net (78M)	<u>95.27</u> (2.81)	95.52(6.28)	<u>81.48</u> (9.09)	96.32 (1.58)	<u>96.99</u> (1.00)	<u>95.00</u> (2.33)	<u>83.88</u> (8.67)	96.56 (2.35)	85.20(17.98)		92.02		
Federated													
- U-Net (97M)	94.48(2.28)	<u>94.54</u> (6.18)	79.38(12.69)	93.75(4.71)	96.76(1.55)	<u>92.40</u> (4.56)	80.79(11.89)	95.72(2.97)	<u>83.38</u> (15.62)		90.39		
- MENU-Net (78M)	94.07(3.09)	95.94 (3.75)	80.05(12.02)	94.65(2.98)	96.70(1.15)	93.74(2.23)	82.14(9.45)	96.34(2.09)	<u>86.08</u> (14.48)		91.14		
Models (# of param.)	ASD [Mean(SD) mm]										Global		
	Client #1		Client #2		Client #3		Client #4		Client #5				
	Liver	Kidney	Pancreas	Spleen	Liver	Kidney	Pancreas	Spleen	Gallbladder				
Localized													
- Client #1	<u>4.00</u> (2.88)	-	-	-	-	-	-	-	-				
- Client #2	-	<u>2.40</u> (4.26)	-	-	-	-	-	-	-				
- Client #3	-	-	<u>3.01</u> (3.82)	-	-	-	-	-	-		2.34		
- Client #4	-	-	-	<u>1.11</u> (0.62)	-	-	-	-	-				
- Client #5	-	-	-	-	0.95(1.08)	0.87(3.31)	1.48(1.03)	<u>0.67</u> (1.01)	<u>1.99</u> (3.33)				
Centralized													
- U-Net (97M)	1.95(1.63)	<u>3.26</u> (6.31)	2.04(2.80)	1.17(1.62)	1.02(1.82)	<u>4.87</u> (9.15)	<u>1.30</u> (1.26)	0.64(1.93)	1.52(3.26)		2.06		
- MENU-Net (78M)	<u>1.86</u> (2.07)	0.88(1.98)	1.85 (1.49)	<u>0.30</u> (0.13)	0.61 (0.47)	<u>0.72</u> (1.40)	<u>1.19</u> (1.28)	0.36 (0.62)	1.24 (2.22)		1.14		
Federated													
- U-Net (97M)	3.75(4.08)	<u>3.11</u> (6.11)	2.36(4.01)	1.54(2.72)	<u>1.40</u> (1.86)	<u>2.82</u> (6.02)	1.93(3.87)	0.45(0.60)	1.56(2.89)		2.48		
- MENU-Net (78M)	2.56(2.47)	0.79 (1.03)	2.10(2.45)	0.54(0.42)	0.94(1.59)	0.80(0.89)	1.53(1.54)	0.41(0.52)	1.34(2.70)		1.40		

TABLE III

QUANTITATIVE PERFORMANCE EVALUATION OF DIFFERENT METHODS WHEN TESTED ON THE OUT-OF-FEDERATION DATA. THE BEST RESULTS ARE MARKED IN **BOLD**. THE UNDERLINED RESULTS INDICATE A STATISTICALLY SIGNIFICANT DIFFERENCE FROM OUR RESULT ($p < 0.05$)

Models (# of param.)	DSC [Mean(SD) %]					Global	ASD [Mean(SD) mm]					Global
	Liver	Kidney	Pancreas	Spleen	Gallbladder		Liver	Kidney	Pancreas	Spleen	Gallbladder	
Localized												
- Client #1	<u>92.63</u> (3.46)	-	-	-	-		<u>3.70</u> (4.93)	-	-	-	-	
- Client #2	-	<u>85.07</u> (16.53)	-	-	-		-	<u>5.28</u> (17.98)	-	-	-	
- Client #3	-	-	<u>73.78</u> (11.47)	-	-	76.13	-	-	<u>6.02</u> (7.19)	-	-	4.57
- Client #4	-	-	-	<u>85.01</u> (15.57)	-		-	-	-	<u>3.64</u> (9.40)	-	
- Client #5	<u>92.78</u> (4.97)	89.14 (16.38)	<u>57.13</u> (26.67)	88.18(15.80)	<u>53.45</u> (33.93)		<u>3.15</u> (5.41)	4.19 (11.60)	6.73(7.94)	<u>3.49</u> (5.60)	<u>5.29</u> (9.17)	
Centralized												
- U-Net (97M)	95.35(4.07)	<u>85.50</u> (16.97)	77.44(11.05)	91.46(10.59)	68.02(32.17)	83.55	2.06(4.68)	<u>9.09</u> (15.33)	2.91(3.97)	2.49(5.23)	2.08(3.43)	3.73
- MENU-Net (78M)	<u>95.84</u> (3.35)	88.70(15.30)	78.65 (9.98)	92.57 (8.89)	67.86(31.26)	84.72	<u>1.85</u> (4.57)	4.26(14.35)	2.48 (4.15)	1.58(3.38)	3.79(13.25)	2.79
Federated												
- U-Net (97M)	95.00(5.41)	88.03(15.76)	76.53(10.52)	92.37(12.44)	64.35(30.94)	83.26	2.76(6.66)	5.97(15.98)	2.59(2.71)	1.13 (2.39)	1.70 (1.87)	2.83
- MENU-Net (78M)	95.21(4.11)	88.12(16.88)	77.92(10.91)	90.97(13.29)	69.45 (31.61)	84.33	2.39(5.74)	4.56(14.78)	3.44(9.85)	1.38(2.64)	1.79(3.28)	2.71

key designs, including 1) the MENU-Net for organ-specific feature extraction and 2) the AGD for organ-specific feature enhancement. A U-Net [44] trained using the FedAvg algorithm [1] was used as the baseline method, on which the MENU-Net and AGD components are sequentially imposed. In addition, we also tried to replace the AGD with a non-generic version, *i.e.*, the auxiliary decoders are locally trained on the client nodes without parameter fusion (see Eq. 1)

through the server node. We denote it as *auxiliary localized decoder* (ALD) in the following contents. This ablation model is used to justify the necessity of parameter sharing in learning organ-specific features.

The quantitative results of the ablation studies are listed in Table IV. The observations are three-fold. 1) Compared with the baseline, the design of the MENU-Net (Table IV, “+MENU-Net”) can effectively improve the global

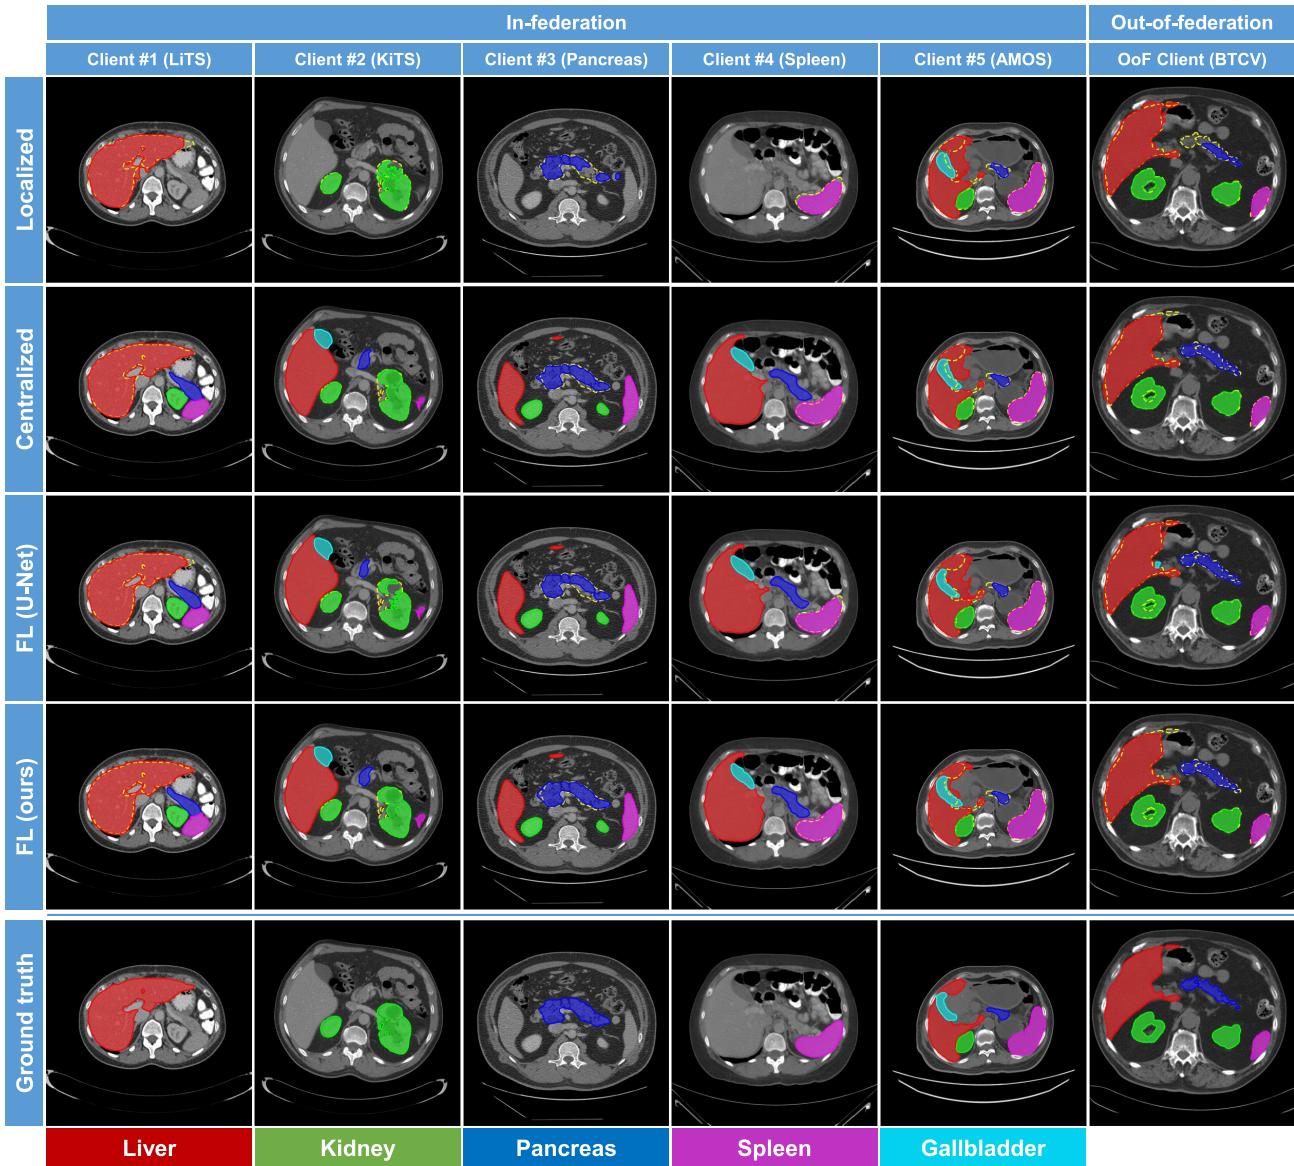


Fig. 2. 2D Visualization of segmentation results yielded by the competing methods. Each column shows a case from one client. For better comparison, the ground-truth contours are also superimposed on the segmentation results as yellow dashed lines.

TABLE IV
EXPERIMENTAL RESULTS OF ABLATION STUDIES ON THE PROPOSED
METHOD. THE BEST RESULTS ARE MARKED IN BOLD

Ablation models	In-federation		Out-of-federation	
	DSC [%]	ASD [mm]	DSC [%]	ASD [mm]
baseline	90.39	2.48	83.26	2.83
+MENU-Net	90.95	1.76	83.86	3.39
+NEMU-Net+ALD	90.78	1.62	83.36	3.65
+NEMU-Net+AGD	91.14	1.40	84.33	2.71

segmentation indices except for the ASD metric in out-of-federation evaluation. This result indicated a positive effect of the task decomposition design (*i.e.*, the multi-encoder network architecture) in learning domain-invariant features. 2) Based on the MENU-Net, the AGD (Table IV, “+MENU-

Net+AGD”) can further boost the segmentation accuracy for both in-federation and out-of-federation evaluation. This result suggested that the auxiliary supervisions on the intermediate layers can facilitate the learning of more discriminative features, and thus, benefit the subsequent shared decoder for accurate segmentation. 3) By replacing the AGD with the ALD, the accuracy of our method (Table IV, “+MENU-Net+ALD”) dropped in both the in-federation setting and the out-of-federation setting. This result demonstrated the importance of parameter sharing in our AGD since the auxiliary decoder shared across the clients enforced the preceding sub-encoders to extract discriminative features for different organs.

Fig. 4 visualized the gradient-weighted class activation maps (Grad-CAMs) [45] produced by our MENU-Net (extracted from the third convolutional block in the sub-encoders) when it is trained without AGD, with ALD, and with AGD,

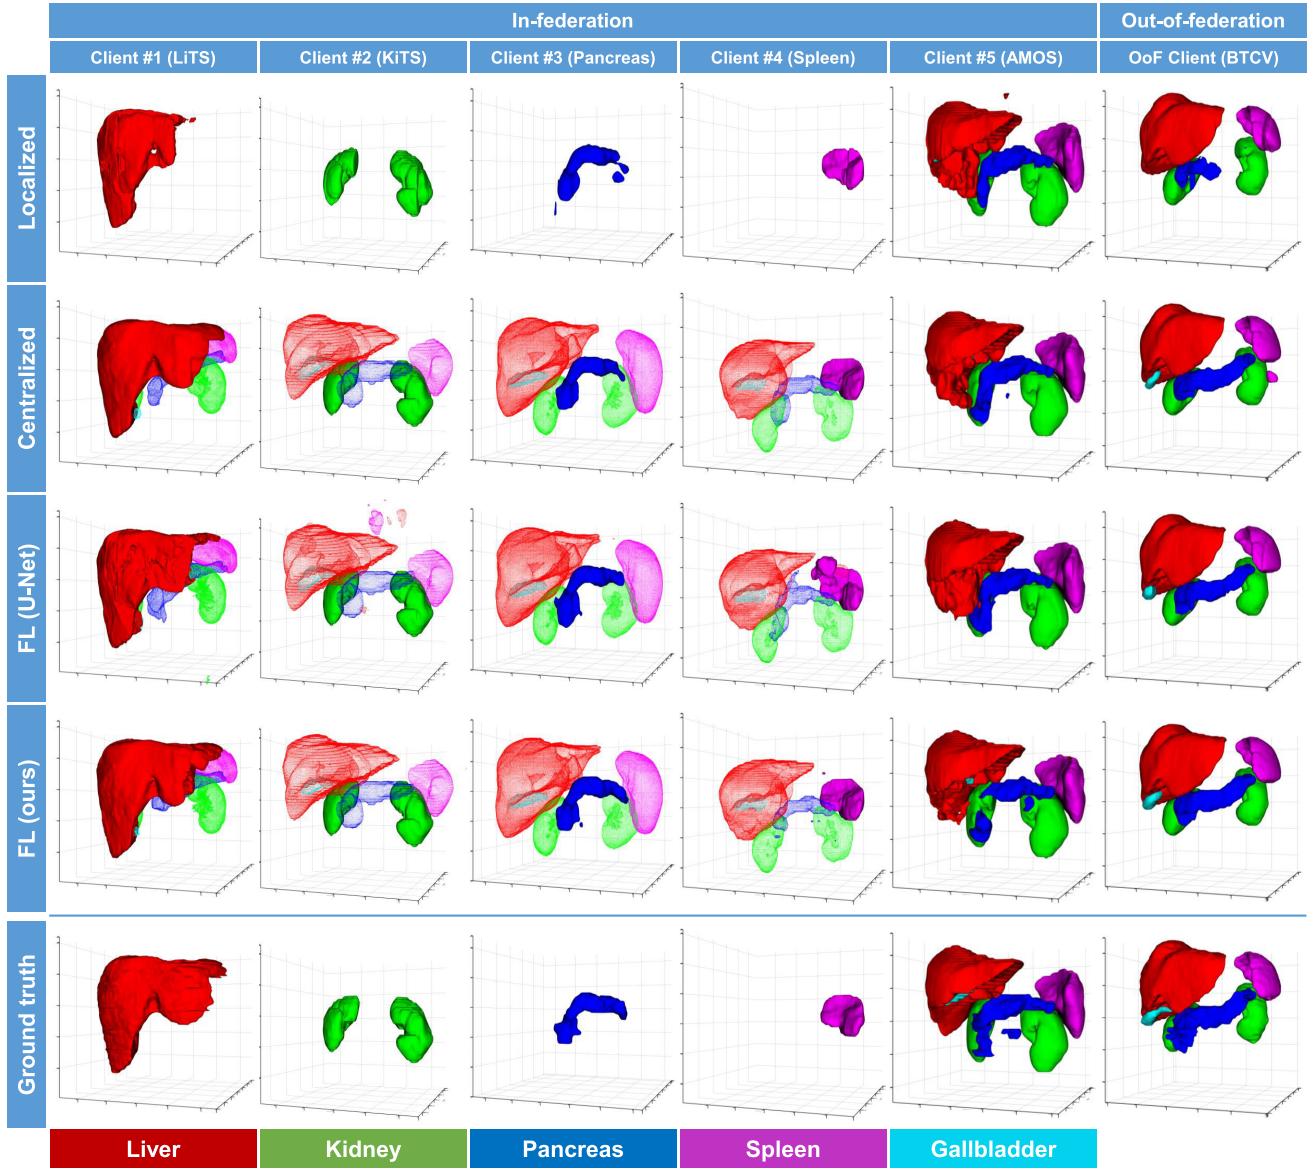


Fig. 3. 3D Visualization of segmentation results by the competing methods. Each column shows a case from one client. The segmentation of the unlabeled organs is represented by transparent meshes.

TABLE V

QUANTITATIVE PERFORMANCE EVALUATION OF THE PROPOSED METHOD TRAINED WITH DIFFERENT COMMUNICATION FREQUENCIES $T \times E$. T IS THE NUMBER OF COMMUNICATION ROUNDS AND E IS THE NUMBER OF LOCAL TRAINING EPOCHS. THE BEST RESULTS ARE MARKED IN BOLD

Communication frequency $T \times E$	In-federation		Out-of-federation	
	DSC [%]	ASD [mm]	DSC [%]	ASD [mm]
25×16	90.63	1.78	83.10	3.47
50×8	90.64	1.92	83.45	2.95
100×4	90.72	1.89	83.58	2.98
200×2	90.95	1.57	83.81	3.27
400×1	91.14	1.40	84.33	2.71

respectively. The model trained with AGD generated activation maps with more accurate shapes and fewer false positive

regions than the other two models, which indicated that the AGD can effectively enhance the organ-specific features extracted by the sub-encoders with better explainability [46].

E. Effects of Communication Frequency

Communication frequency is a key factor affecting the performance of the FL methods in practice. In the proposed method, the communication frequency is jointly controlled by the number of communication rounds T and the number of local training epochs E . We successively trained our method with different combinations of $T \times E$ to investigate the effects of communication frequency. For the sake of fairness, the product of these two parameters are fixed to 400, which means all models are optimized with the same number of training batches (iterations).

The experimental results are shown in Table V and Fig. 5. It can be seen that higher communication frequency generally

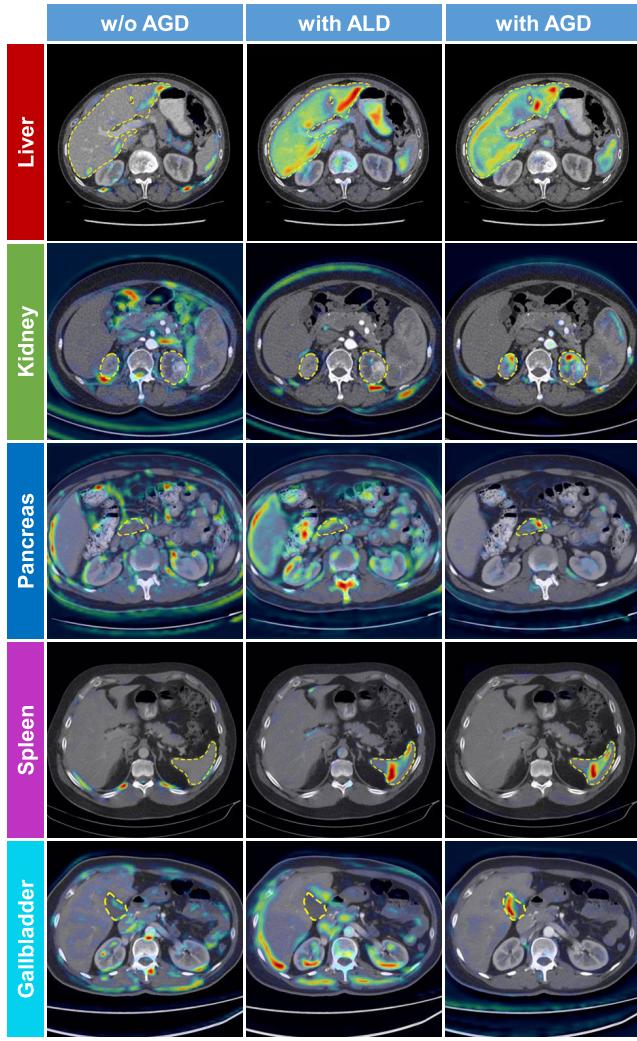


Fig. 4. Gradient-weighted class activation maps (Grad-CAMs) [45] generated by the MENU-Net trained without AGD (left), with ALD (middle), and with AGD (right). Yellow dashed lines indicate ground-truth contours. We use the activation (output) of the third convolutional block in the sub-encoders to calculate the Grad-CAMs.

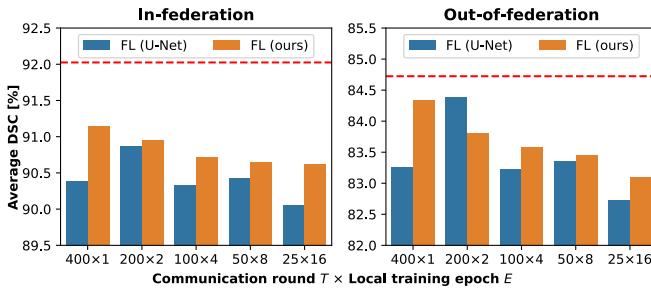


Fig. 5. Average DSC of FL models trained with different communication frequencies. The red dashed line indicates the average DSC achieved by centralized learning model.

brought better performance to the trained model, which is in line with observations in other FL-based methods [1], [4], [9], [14], and our method consistently outperformed the baseline FL U-Net model. Considering our method is designed for the cross-silo FL [47] scenario, where the clients (clinical sites) have stable internet connections with sufficient bandwidth,

TABLE VI
QUANTITATIVE PERFORMANCE EVALUATION OF THE PROPOSED METHOD WHEN COMBINED WITH DIFFERENT FEDERATED LEARNING STRATEGIES. THE BEST RESULTS ARE MARKED IN BOLD

Models	In-federation		Out-of-federation	
	DSC [%]	ASD [mm]	DSC [%]	ASD [mm]
FedAvg [1]	91.14	1.40	84.33	2.71
FedAvgM [6]	90.28	1.98	82.32	3.29
FedProx [7]	90.95	1.62	83.73	2.97
FedDyn [8]	91.33	1.56	84.60	3.23
FedSM [11]	91.15	1.41	84.38	2.70
FedNorm [10]	90.92	1.75	84.21	3.12
Fed-DoDNet [30]	91.09	1.75	83.82	2.74

TABLE VII
QUANTITATIVE PERFORMANCE EVALUATION OF THE PROPOSED METHOD WHEN THE LESION REGIONS WERE TREATED AS BACKGROUND/A PART OF THE TARGET ORGANS. THE BEST RESULTS ARE MARKED IN BOLD

Models	In-federation		Out-of-federation	
	DSC [%]	ASD [mm]	DSC [%]	ASD [mm]
Lesion as background	89.76	1.65	84.00	2.81
Lesion as part of organ	91.14	1.40	84.33	2.71

we finally choose $T=400$ and $E=1$ in our method to achieve higher accuracy.

F. Comparison With Different FL Strategies

Another important factor relating to the FL systems performance is the choice of the learning strategy, which determines how the locally trained models are optimized and aggregated to update the server model. In the prior experiments, to focus on our own designs, we consistently used the most basic FL strategy, *i.e.*, the FedAvg algorithm [1], for this choice. Our method is also flexible to be combined with other FL strategies. To demonstrate this property, we tried to deploy our method with six different FL strategies, including FedAvg [1], FedAvgM [6], FedProx [7], FedDyn [8], FedSM [11], and FedNorm [10]. The experimental results of these hybrid models are shown in Table VI. It can be observed that, by combining with more advanced FL strategies, the segmentation accuracy of our method can be further improved. The highest DSC was achieved when our method was combined with the FedDyn [8] algorithm. This experiment demonstrates the flexibility of our proposed method. While the existing FL strategies mainly focus on resolving the problem of image heterogeneities, our method is designed to deal with label heterogeneities. Thus, the proposed method and FL strategies can complement each other. We also compared our method with another multi-organ segmentation method for partially labeled data, *i.e.*, DoDNet [30], which was originally designed for centralized learning scenario and extended in this experiment by combining with the FedAvg [1] algorithm (denoted as “Fed-DoDNet” in Table VI). The segmentation accuracy of this competing model is slightly lower than our

model, which demonstrated the superiority of our method designs.

V. DISCUSSION AND CONCLUSION

In our experiments, the FL model was trained with five client datasets, each containing varying numbers of images that were labeled with different organs. This setup was designed to simulate real scenarios in clinical settings. Our experimental results revealed that our proposed method is flexible and capable of handling such complex scenarios. However, it is important to note that our approach has a minimum criterion for the number of labeled organs required. Specifically, each client dataset must be labeled with at least one of the target organs. For more complicated scenarios involving unlabeled images, it would require some additional semi-supervised or unsupervised learning techniques, such as the teacher-student consistency learning strategy [48], [49], to deal with. Fortunately, our Fed-MENU framework is highly adaptable and can seamlessly integrate these methods, making it possible to extend the framework to these situations.

Most of the CT images used in this study contained tumors, and we treated the lesion regions as a part of the surrounding organ in the experiments. There could be a potential influence of these abnormalities on the performance of the proposed method. To investigate this impact, we compared the performance of our method when the lesion regions in the CT images were treated 1) as background and 2) as part of the foreground organs, respectively. The quantitative results yielded from these two settings are reported in Table VII. It can be seen that, when the lesion regions were treated as the background category, the resulting performance of the proposed method degraded. The reason can be attributed to the fact that, although the lesion regions exhibited different textures or intensities than the healthy tissues, they may not substantially change the contour shape of the target organs. Thus, it would be relatively easy to segment these regions as a part of the foreground. Otherwise, if the lesions were treated as the background, it may significantly change the structure of the target organs, which would eventually affect the segmentation accuracy. Since the proposed method is designed to deal with the inconsistently labeled data for FL, we can also consider the tumor labels as a new class inconsistently distributed in different client nodes and train our method to segment the tumors. This could lead us to a potential application of disease diagnosis. Naturally, to be adaptive to the task of diagnosis (classification), the output branch of the proposed MENU-Net may be replaced with a global pooling layer followed by fully-connected layers. As this study mainly focused on the application of multi-organ segmentation, this direction will be explored in our future research.

The experiments conducted in this study are done in a virtual environment, in which we focus on the training problem of inconsistently labeled data. In practical scenarios, there are some security issues that should be considered. For example, encryptions are required to secure the communications between server and client nodes. Unencrypted communication would raise a risk of information leakage and interference

from malicious attacks [50]. In addition, challenges of data memorization [51] as well as adversarial inference [52] also raise privacy concerns to the FL methods. It has been demonstrated that an FL model can be maliciously manipulated to reconstruct the private training data if no protective measures are taken [53]. To address this issue, differential privacy (DP) [54] techniques can be used to enhance the FL methods by introducing noise to their input data, computation results, or optimization process [55]. Since the goal of this study is to develop an FL framework to enable collaborative learning on inconsistently labeled data, we did not draw extra efforts to the security issue. Fortunately, our method is compatible to the above-mentioned techniques and thus can be combined with them to achieve secure FL.

One potential limitation of the proposed method could be its scalability. In this study, we designed our method for a five-organ segmentation task. However, when more organs are desired in the segmentation task, more encoders are required in the MENU-Net. As a consequence, the concatenation of the output feature maps from the multi-encoders would increase model complexity and GPU memory consumption. One solution to this issue could be using one single shared encoder combined with conditional input to replace the multi-encoder design. This would result in a variant of our method sharing the DoDNet structure [30]. In Table VI, we reported the experimental results of this variant model (denoted as “**Fed-DoDNet**”), which demonstrated the feasibility of this shared-encoder+conditional input design.

In this paper, we revealed and defined a new problem of FL with partially labeled data in the context of medical image segmentation, which is of great clinical significance and technical urgency to solve. Subsequently, a novel Fed-MENU method was proposed to tackle this challenging problem. Compared with the conventional FL framework that worked on the fully-labeled data, our Fed-MENU method had two key designs to adapt to the partially labeled client datasets, 1) the MENU-Net for organ-specific feature extraction and 2) the AGD for organ-specific feature enhancement. Extensive experiments were conducted using six public abdominal CT image datasets. The experimental results comprehensively demonstrated the feasibility and effectiveness of our method in solving the partial label problem in the context of FL.

REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Y. Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Proc. 20th Int. Conf. Artif. Intell. Statist.*, 2017, pp. 1273–1282.
- [2] H. R. Roth et al., “Federated whole prostate segmentation in MRI with personalized neural architectures,” in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2021, pp. 357–366.
- [3] K. V. Sarma et al., “Federated learning improves site performance in multicenter deep learning without data sharing,” *J. Amer. Med. Inform. Assoc.*, vol. 28, no. 6, pp. 1259–1264, Jun. 2021.
- [4] D. Yang et al., “Federated semi-supervised learning for COVID region segmentation in chest CT using multi-national data from China, Italy, Japan,” *Med. Image Anal.*, vol. 70, May 2021, Art. no. 101992.
- [5] I. Dayan et al., “Federated learning for predicting clinical outcomes in patients with COVID-19,” *Nature Med.*, vol. 27, no. 10, pp. 1735–1743, 2021.

[6] T.-M. Harry Hsu, H. Qi, and M. Brown, "Measuring the effects of non-identical data distribution for federated visual classification," 2019, *arXiv:1909.06335*.

[7] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in *Proc. Mach. Learn. Syst.*, vol. 2, 2020, pp. 429–450.

[8] D. A. E. Acar, Y. Zhao, R. M. Navarro, M. Mattina, P. N. Whatmough, and V. Saligrama, "Federated learning based on dynamic regularization," 2021, *arXiv:2111.04263*.

[9] X. Li, M. Jiang, X. Zhang, M. Kamp, and Q. Dou, "FedBN: Federated learning on non-IID features via local batch normalization," 2021, *arXiv:2102.07623*.

[10] T. Bernecker et al., "FedNorm: Modality-based normalization in federated learning for multi-modal liver segmentation," 2022, *arXiv:2205.11096*.

[11] A. Xu et al., "Closing the generalization gap of cross-silo federated medical image segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 20866–20875.

[12] X. Xu et al., "Federated cross learning for medical image segmentation," in *Proc. Med. Imag. Deep Learn.*, 2023, pp. 1–10.

[13] Y. Nan et al., "Data harmonization for information fusion in digital healthcare: A state-of-the-art systematic review, meta-analysis and future research directions," *Inf. Fusion*, vol. 82, pp. 99–122, Jun. 2022.

[14] Y. Xia et al., "Auto-FedAvg: Learnable federated averaging for multi-institutional medical image segmentation," 2021, *arXiv:2104.10195*.

[15] S. Park, G. Kim, J. Kim, B. Kim, and J. C. Ye, "Federated split task-agnostic vision transformer for COVID-19 CXR diagnosis," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 24617–24630.

[16] A. Dosovitskiy et al., "An image is worth 16 × 16 words: Transformers for image recognition at scale," 2020, *arXiv:2010.11929*.

[17] H. R. Roth et al., "Federated learning for breast density classification: A real-world implementation," in *Domain Adaptation and Representation Transfer, and Distributed and Collaborative Learning*. Lima, Peru: Springer, 2020, pp. 181–191.

[18] Q. Liu, C. Chen, J. Qin, Q. Dou, and P.-A. Heng, "FedDG: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 1013–1023.

[19] M. Jiang, H. Yang, C. Cheng, and Q. Dou, "IOP-FL: Inside-outside personalization for federated medical image segmentation," 2022, *arXiv:2204.08467*.

[20] P. Guo et al., "Auto-FedRL: Federated hyperparameter optimization for multi-institutional medical image segmentation," 2022, *arXiv:2203.06338*.

[21] J. Wicaksana et al., "FedMix: Mixed supervised federated learning for medical image segmentation," *IEEE Trans. Med. Imag.*, early access, Dec. 30, 2022, doi: [10.1109/TMI.2022.3233405](https://doi.org/10.1109/TMI.2022.3233405).

[22] N. Dong, M. Kampffmeyer, and I. Voiculescu, "Learning underrepresented classes from decentralized partially labeled medical images," in *Proc. Med. Image Comput. Comput. Assist. Intervent. 25th Int. Conf. (MICCAI)*. Singapore: Springer, Sep. 2022, pp. 67–76.

[23] N. Dong, M. Kampffmeyer, I. Voiculescu, and E. Xing, "Federated partially supervised learning with limited decentralized medical images," *IEEE Trans. Med. Imag.*, early access, Dec. 20, 2022, doi: [10.1109/TMI.2022.3231017](https://doi.org/10.1109/TMI.2022.3231017).

[24] F. Shi et al., "Deep learning empowered volume delineation of whole-body organs-at-risk for accelerated radiotherapy," *Nature Commun.*, vol. 13, no. 1, p. 6566, Nov. 2022.

[25] X. Ye et al., "Comprehensive and clinically accurate head and neck cancer organs-at-risk delineation on a multi-institutional study," *Nature Commun.*, vol. 13, no. 1, p. 6137, Oct. 2022.

[26] K. Yan et al., "Learning from multiple datasets with heterogeneous and partial labels for universal lesion detection in CT," *IEEE Trans. Med. Imag.*, vol. 40, no. 10, pp. 2759–2770, Oct. 2021.

[27] O. Petit, N. Thome, and L. Soler, "Iterative confidence relabeling with deep ConvNets for organ segmentation with partial labels," *Computerized Med. Imag. Graph.*, vol. 91, Jul. 2021, Art. no. 101938.

[28] X. Fang and P. Yan, "Multi-organ segmentation over partially labeled datasets with multi-scale feature abstraction," *IEEE Trans. Med. Imag.*, vol. 39, no. 11, pp. 3619–3629, Nov. 2020.

[29] G. Shi, L. Xiao, Y. Chen, and S. K. Zhou, "Marginal loss and exclusion loss for partially supervised multi-organ segmentation," *Med. Image Anal.*, vol. 70, May 2021, Art. no. 101979.

[30] J. Zhang, Y. Xie, Y. Xia, and C. Shen, "DoDNet: Learning to segment multi-organ and tumors from multiple partially labeled datasets," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 1195–1204.

[31] P. Liu, L. Xiao, and S. K. Zhou, "Incremental learning for multi-organ segmentation with partially labeled datasets," 2021, *arXiv:2103.04526*.

[32] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-Net: Fully convolutional neural networks for volumetric medical image segmentation," in *Proc. 4th Int. Conf. 3D Vis. (3DV)*, Oct. 2016, pp. 565–571.

[33] Y. Ji et al., "AMOS: A large-scale abdominal multi-organ benchmark for versatile medical image segmentation," in *Proc. 36th Conf. Neural Inf. Process. Syst. Datasets Benchmarks Track*, 2022, pp. 36722–36732.

[34] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," 2016, *arXiv:1607.08022*.

[35] A. L. Maas et al., "Rectifier nonlinearities improve neural network acoustic models," in *Proc. ICML*, vol. 30, no. 1. Princeton, NJ, USA: Citeseer, 2013, p. 3.

[36] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. 13th Int. Conf. Artif. Intell. and Statist.*, 2010, pp. 249–256.

[37] F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation," *Nature Methods*, vol. 18, no. 2, pp. 203–211, Dec. 2020.

[38] P. Bilic et al., "The liver tumor segmentation benchmark (LiTS)," 2019, *arXiv:1901.04056*.

[39] N. Heller et al., "The KiTS19 challenge data: 300 kidney tumor cases with clinical context, CT semantic segmentations, and surgical outcomes," 2019, *arXiv:1904.00445*.

[40] N. Heller et al., "The state of the art in kidney and kidney tumor segmentation in contrast-enhanced CT imaging: Results of the KiTS19 challenge," *Med. Image Anal.*, vol. 67, Jan. 2020, Art. no. 101821.

[41] A. L. Simpson et al., "A large annotated medical image dataset for the development and evaluation of segmentation algorithms," 2019, *arXiv:1902.09063*.

[42] M. Antonelli, A. Reinke, and S. Bakas, "The medical segmentation decathlon," *Nature Commun.*, vol. 13, p. 4128, Jul. 2022.

[43] B. Landman, Z. Xu, J. Iglesias, M. Styner, T. Langerak, and A. Klein, "MICCAI multi-atlas labeling beyond the cranial vault–workshop and challenge," in *Proc. Multi-Atlas Labeling Beyond Cranial Vault Workshop Challenge (MICCAI)*, vol. 5, 2015, p. 12.

[44] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2015, pp. 234–241.

[45] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 618–626.

[46] G. Yang, Q. Ye, and J. Xia, "Unbox the black-box for the medical explainable AI via multi-modal and multi-centre data fusion: A mini-review, two showcases and beyond," *Inf. Fusion*, vol. 77, pp. 29–52, Jan. 2022.

[47] P. Kairouz et al., "Advances and open problems in federated learning," *Found. Trends Mach. Learn.*, vol. 14, nos. 1–2, pp. 1–210, 2019.

[48] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 1–10.

[49] X. Xu, T. Sanford, B. Turkbey, S. Xu, B. J. Wood, and P. Yan, "Shadow-consistent semi-supervised learning for prostate ultrasound segmentation," *IEEE Trans. Med. Imag.*, vol. 41, no. 6, pp. 1331–1345, Jun. 2022.

[50] G. A. Kaassis, M. R. Makowski, D. Rückert, and R. F. Braren, "Secure, privacy-preserving and federated machine learning in medical imaging," *Nature Mach. Intell.*, vol. 2, no. 6, pp. 305–311, Jun. 2020.

[51] C. Song, T. Ristenpart, and V. Shmatikov, "Machine learning models that remember too much," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Oct. 2017, pp. 587–601.

[52] D. Usynin et al., "Adversarial interference and its mitigations in privacy-preserving collaborative machine learning," *Nature Mach. Intell.*, vol. 3, no. 9, pp. 749–758, Sep. 2021.

[53] G. Kaassis et al., "End-to-end privacy preserving deep learning on multi-institutional medical imaging," *Nature Mach. Intell.*, vol. 3, pp. 473–484, Jun. 2021.

[54] C. Dwork et al., "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, nos. 3–4, pp. 211–407, 2014.

[55] A. Rajkumar and S. Agarwal, "A differentially private stochastic gradient descent algorithm for multiparty classification," in *Proc. 15th Int. Conf. Artif. Intell. Statist.*, 2012, pp. 933–941.