arXiv:2110.09260v1 [cs.CV] 18 Oct 2021

A Unified Framework for Generalized Low-Shot Medical Image Segmentation with Scarce Data

Hengji Cui, Dong Wei, Kai Ma, Shi Gu, and Yefeng Zheng, *Senior Member, IEEE*

***Abstract*—Medical image segmentation has achieved remarkable advancements using deep neural networks (DNNs). However, DNNs often need big amounts of data and annotations for training, both of which can be difficult and costly to obtain. In this work, we propose a unified framework for generalized low-shot (one- and few-shot) medical image segmentation based on distance metric learning (DML). Unlike most existing methods which only deal with the lack of annotations while assuming abundance of data, our framework works with extreme scarcity of both, which is ideal for rare diseases. Via DML, the framework learns a multimodal mixture representation for each category, and performs dense predictions based on cosine distances between the pixels' deep embeddings and the category representations. The multimodal representations effectively utilize the inter-subject similarities and intraclass variations to overcome overfitting due to extremely limited data. In addition, we propose adaptive mixing coefficients for the multimodal mixture distributions to adaptively emphasize the modes better suited to the current input. The representations are implicitly embedded as weights of the fc layer, such that the cosine distances can be computed efficiently via forward propagation. In our experiments on brain MRI and abdominal CT datasets, the proposed framework achieves superior performances for low-shot segmentation towards standard DNN-based (3D U-Net) and classical registration-based (ANTs) methods, *e.g.*, achieving mean Dice coefficients of 81%/69% for brain tissue/abdominal multi-organ segmentation using a single training sample, as compared to 52%/31% and 72%/35% by the U-Net and ANTs, respectively.**

***Index Terms*—Semantic segmentation, Low-shot learning, Distance metric learning, Multimodal representation, Adaptive mixing coefficients.**

I. Introduction

SEMANTIC segmentation, one of the fundamental tasks in computer vision, has important applications in medical image analysis, such as risk assessment [1], precise diagnosis [2], therapy planning [3], and clinical research. Benefiting from the rapid development of deep learning techniques, deep neural network (DNN) based methods have advanced the state of the art of medical image segmentation—by large margins over traditional methods—in recent years [4]. To obtain satisfactory performance, large amounts of data with annotations are usually needed for training DNNs. However, accurate annotation of medical images is expertise-demanding, labor-intensive, time-consuming, and error-prone, thus difficult and costly to obtain in large amount. Moreover, in some circumstances, *e.g.*, for rare conditions, even the training data can be very limited.

To cope with the lack of training samples, low-shot (one- and few-shot) learning techniques [5] have been proposed in the natural image domain and achieved encouraging results in both classification [6], [7] and segmentation tasks [8], [9]. In the low-shot natural image domain, it is often assumed (and also thus practiced) that generalizable prior knowledge can be learned on a sufficiently large set of labeled samples of known *base* classes, and then utilized to boost learning of previously unseen *novel* classes given limited samples (*i.e.*, the target task). A few works on low-shot medical image *classification* followed this problem setting of base versus novel classes. For example, Paul *et al.* [10] divided different chest diseases into disjoint base (with large training set) and novel (with a few training samples per class) classes, and proposed a few-shot learning framework based on autoencoder ensemble for disease identification from chest x-ray images. For many *segmentation* tasks of medical images, however, the problem setting of low-shot learning is usually different from the assumption above. Taking one-shot segmentation of brain structures in MRI for example [11]: the concepts of base and novel classes are infeasible here, as all the classes to segment are already presented in the given one-shot example; or in other words, there is no unseen novel class. Consequently, most low-shot segmentation methods proposed for natural images cannot be directly applied to medical images.

Alternatively, most recent works on DNN-based low-shot medical image segmentation [11], [12], [13] relied on image synthesis. These methods utilized unlabeled data of the target task in semi-supervised setting, assuming abundance of data. For example, both Zhao *et al.* [11] and Chaitanya *et al.* [12] proposed to learn and utilize spatial and appearance transformations between the labeled and unlabeled samples to synthesize realistic samples for training the segmentation networks. Although requiring only one or few annotations, these semi-supervised methods [11], [12], [13] require a considerable amount of unlabeled data to work effectively [14], thus inapplicable when the samples are rare, too. The scenario in which both the annotations and data are scarce can be practically relevant, *e.g.*, when developing algorithms for a rare disease, and when the emergence of a novel disease

Manuscript received July 1, 2020; revised October 27, 2020; accepted December 14, 2020. This work was primarily supported by the NSFC General Program 61876032, the Key-Area Research and Development Program of Guangdong Province, China (No. 2018B010111001), the National Key R&D Program of China (2018YFC2000702), and the Scientific and Technical Innovation 2030-"New Generation Artificial Intelligence" Project (No. 2020AAA0104100). *(H. Cui and D. Wei contributed equally to this work.) (Corresponding author: S. Gu.)*

H. Cui and S. Gu are with the School of Computer Science and Engineering, University of Electronic Science and Technology of China (email: chj@std.uestc.edu.cn, gus@uestc.edu.cn). H. Cui contributed to this work as an intern at Tencent Jarvis Lab.

D. Wei, K. Ma, and Y. Zheng are with Tencent Jarvis Lab (email: donwei@tencent.com, kylekma@tencent.com, yefengzheng@tencent.com).

necessitating rapid algorithm development. However, it is more challenging due to the shortage in both data and annotations. Indeed, such scenario is rarely touched in existing literature on DNN-based approaches. In contrast, registration-based segmentation [15]—a classical approach to few-shot segmentation of medical images—can achieve decent performance using only one or few labeled samples (called "atlas"), especially for brains. Nonetheless, this approach can become less effective when large individual variations in appearance and/or structures are present and is usually time-consuming, too.

In this work, we propose a new DNN-based framework for generalized low-shot medical image segmentation using extremely scarce training samples and annotations—a full analogy to humans who can learn reasonably well in such situation, which is able to handle large individual variations and accomplish segmentation in seconds. Unlike most existing DNN-based methods for low-shot medical image segmentation, which relied on image synthesis, the framework employs the distance metric learning (DML) [16] to learn a generalized multimodal representation (comprising one to several prototypes) for each category. The DML of multimodal representation can effectively utilize not only the inter-subject similarities, but also the intraclass variations, for robustness against overfitting in case of scarce data. In the context of medical image data and for a specific imaging modality, the inter-subject similarity refers to the similar appearance of the same tissue, organ, *etc.*, across different individuals, whereas the intraclass variations refer to the intrinsic heterogeneity of a specific organ/tissue. Notably, we propose to adaptively determine the mixing coefficients for the modes of a category based on a self-attention mechanism [17], rather than treating all modes equally or relying on only one of them for prediction [16]. Further, we implement the multimodal representations using the cosine normalization technique [18], by implicitly embedding the representations as weights of the fully connected (fc) layer and computing the cosine similarity by forward propagation. Compared to explicit representation learning [6], [16], the implicit embedding not only simplifies the network structure but also promotes computational efficiency. Accordingly, we name our framework the Multimodal Representation Embedding Network (MRE-Net). Last but not least, the MRE-Net handles both one- and few-shot learning with a unified network structure.

Besides the multimodal representation embedding (MRE), we further employ three strategies targeted at dealing with the unique characteristics of medical image data. First, the online hard example mining (OHEM) [19] technique is adopted in our dense prediction task to deal with the class imbalance problem commonly encountered in medical images. Second, patch-wise spatial information [20] is incorporated for utilization of structural similarities among individuals. Third, the channel-wise attention [17] and atrous spatial pyramid pooling (ASPP) [21] are employed to account for the diversity of the segmentation targets in medical image analysis.

In summary, our contributions are three folds:

- First, we propose the MRE-Net, a unified framework for generalized low-shot (one- and few-shot) medical image segmentation in case of scarcity of both annotations and samples. The MRE-Net employs the DML for multimodal category representation learning to effectively utilize the inter-subject similarities and intraclass variations.
- Second, we propose adaptive mixing coefficients for the modes of a category, a key innovation in multimodal representation learning to ensure that the modes more suitable for the current input are given more emphasis.
- Third, we conduct thorough experiments on two distinct datasets (brain and abdominal multi-organ [22], [23] segmentation) to demonstrate the superiority of the MRE-Net to both classical and modern methods in a variety of low-shot settings. The behaviour of the MRE-Net is studied thoroughly, too, for additional insights.

II. Related Work

1) Semantic Segmentation: Semantic segmentation refers to the process of assigning each pixel in an image a class label. In terms of medical image segmentation, these labels may include organs (*e.g.*, heart, lung, and liver), tissues (*e.g.*, brain structures), bones, just to name a few. Fully convolutional networks (FCNs) [24] are the current state of the art for semantic segmentation, which replaced the fc layers at the end of the network with convolutional layers to output pixel-level dense prediction. Arguably, U-Net and its variants [25], [26] (also belonging to FCNs) may be the most widely used structure for medical image segmentation, which features symmetric downsampling and upsampling paths (the encoder and decoder) with skip connections in between. Although being the *de facto* benchmark for medical image segmentation, U-Nets are data-hungry just like most DNNs. In this work, we use a modified 3D U-Net [26] as the backbone feature extractor, as well as a baseline in the experiments. As will be shown in Section IV, low-shot performance of the U-Net can be substantially improved after enhanced by the proposed MRE and trained in the way of DML.

2) DML and Low-shot Segmentation of Natural Images: DML has been widely used in computer vision applications (see [27] for a survey). Recently, this concept has been revitalized in the realm of deep learning [28] and become one of the central topics in few-shot learning [6], [8], [9], [16]. In DML, a learnable embedding function projects the input into an embedding space, where the projected samples of the same class should be closer to each other than those of different classes. To learn the embedding function, a distance metric (*e.g.*, squared Euclidean [6], [16] and cosine distances [8], [29]) is employed to measure the intra- and inter-class distances, and loss functions are defined accordingly (*e.g.*, the triplet loss [30]) to encourage small intra-class and large inter-class distances. Commonly in low-shot segmentation, masked average pooling is performed [8], [9] to obtain category representations, and segmentation labels of individual pixels are obtained by dense comparison against the representations.

As introduced earlier, it is prohibitive to directly apply low-shot segmentation methods for natural images to medical images in practice. However, we find the idea of DML-based representation learning powerful, and apply it to low-shot medical image segmentation. Notably, our work is most related

to the RepMet [16] and weight imprinting [29]: the former proposed multimodal representations for low-shot detection, and the latter showed that cosine similarity based DML could be equivalently implemented by implicitly embedding category representations as weights of the fc layer. In RepMet, each category was represented by a Gaussian mixture model with multiple modes, whose centers were considered as the representative vectors for the category; however, only one mode was used for a category at the time of classification, and the explicit representations in RepMet unnecessarily complicated the network structure. As to [29], its unimodal weight embedding did not fully exploit the intraclass variations. In this work, we innovatively implement multimodal representation learning by extending the weight embedding technique from single mode to multiple modes, and adaptively emphasize the most informative modes with a self-attention mechanism [17].

3) Low-shot Medical Image Segmentation: Most existing works on DNN-based low-shot medical image segmentation relied on image synthesis. Zhao *et al.* [11] proposed to learn both spatial and appearance transformations from an annotated atlas to a pool of unlabeled images, and synthesized labeled data by combining randomly sampled transformations. The synthesized data were then used as augmented data to train a segmentation DNN. Similarly, Chaitanya *et al.* [12] employed two conditional generative adversarial nets (GANs) [31] to synthesize augmented data using spatial and additive intensity transformations. GAN [32] was also employed in [33] for few-shot segmentation of 3D multi-sequence MR images, using labeled, unlabeled, and fake images for semi-supervised training. Wang *et al.* [13] introduced the forward-backward consistency to the DNN-based registration framework [34], which enabled a variety of extra, cycle-consistency-based supervision signals to make the training process stable, while also boosting the performance of one-shot segmentation. All the works mentioned above belong to the genre of semi-supervised methods, which require a considerable amount of unlabeled data to work effectively, hence inapplicable to the scenario considered in this work.

One exception to the above-reviewed works is the more recent 'channel squeeze & spatial excitation' (sSE) guided method [35], which followed the problem setting of base and novel classes in the natural image domain [36]. Concretely, Roy *et al.* [35] used three abdominal organs as the base classes on which the proposed model was trained, and evaluated the model on the fourth organ (the unseen novel class). A two-arm architecture [36] consisting of a conditioner arm and a segmenter arm was adopted, with the former processing the annotated input, generating task-specific representations, and then passing these representations to the latter for task-specific segmentation guidance. An sSE module was proposed to facilitate the information transfer between the two arms. Encouraging results were reported on contrast-enhanced abdominal CT images. Its principle difference with our work is the problem setting. In [35], a considerable amount of data (65 scans and corresponding annotations) of the base classes were used for training. In contrast, we consider the case in which the data are scarce (one to several examples), *e.g.*, for a rare condition. Another difference is that [35] was intended for binary segmentation of a novel class, whereas our MRE-Net is for multiclass segmentation of all the "base" classes.

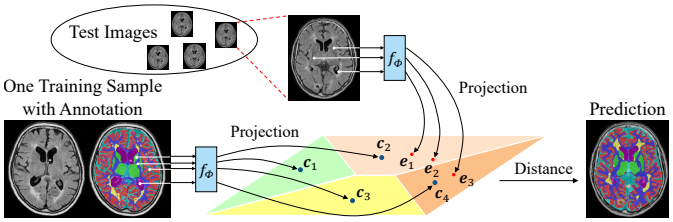

Fig. 1. Conceptual overview of our low-shot segmentation framework. Pixels of test images are projected in the embedding space, where dense predictions are made according to the distances between pixels' embeddings $\boldsymbol{e}_i$ and category prototypes $\boldsymbol{c}_k$.

Registration-based segmentation [15] is a classical approach to low-shot medical image segmentation, and can achieve decent performance using only one or few atlases, especially for brains. Nonetheless, this approach can become less effective when large individual variations in appearance and/or structures are present, and is usually time-consuming. We propose the MRE-Net to address these two issues simultaneously.

4) Data Augmentation: Online data augmentation has become a standard technique for alleviating overfitting by DNNs, since first extensively used by Krizhevsky *et al.* [37]. It applies random transformations to the images during training to artificially increase the diversity of the training data. For medical images, commonly conducted transformations include mirroring, brightness/constrast jittering, rotation, and elastic deformation. In this work, we empirically explore the impact of data augmentation in the low-shot setting.

III. Methods

A. Conceptual Overview: Segmentation via DML-Based Dense Prediction

Given an annotated training set $\mathcal{X}^{\mathrm{tn}} = \{(x^{\mathrm{tn},(l)}, a^{(l)})\}_{l=1}^{L}$, where $x^{\mathrm{tn},(l)}$ is the l^{th} training sample with the annotation $a^{(l)}$, and L is the total number of annotated training samples, we aim to infer the segmentation of any test image x^{ts} based on $\mathcal{X}^{\mathrm{tn}}$. As conceptually illustrated in Fig. 1, our framework first employs an embedding function f_Φ with learnable parameters Φ to project every pixel of the test image into the embedding space: $\boldsymbol{e}_i = f_\Phi(x_i^{\mathrm{ts}}|x^{\mathrm{ts}})$, where i is a pixel index, and $\boldsymbol{e}_i \in \mathbb{R}^{N_{\boldsymbol{e}}\times 1}$ is the resulting embedding vector. Within the embedding space, a representative prototype $\boldsymbol{c}_k$ is learned from $\mathcal{X}^{\mathrm{tn}}$ for each segmentation category $k \in \{1,\ldots,K\}$, where K is the total number of categories (including background). With the notion of DML, the segmentation label for the pixel x_i^{ts} can be inferred non-parametrically by finding the nearest neighbor of $\boldsymbol{e}_i$ among all $\boldsymbol{c}_k$'s (although a more sophisticated mechanism than nearest neighbor is employed in our practical implementation). In this work, we mainly focus on $L \leq 3$, which is generalized low-shot segmentation. Below, we describe our practical implementation of the DML and feature embedding process in detail.

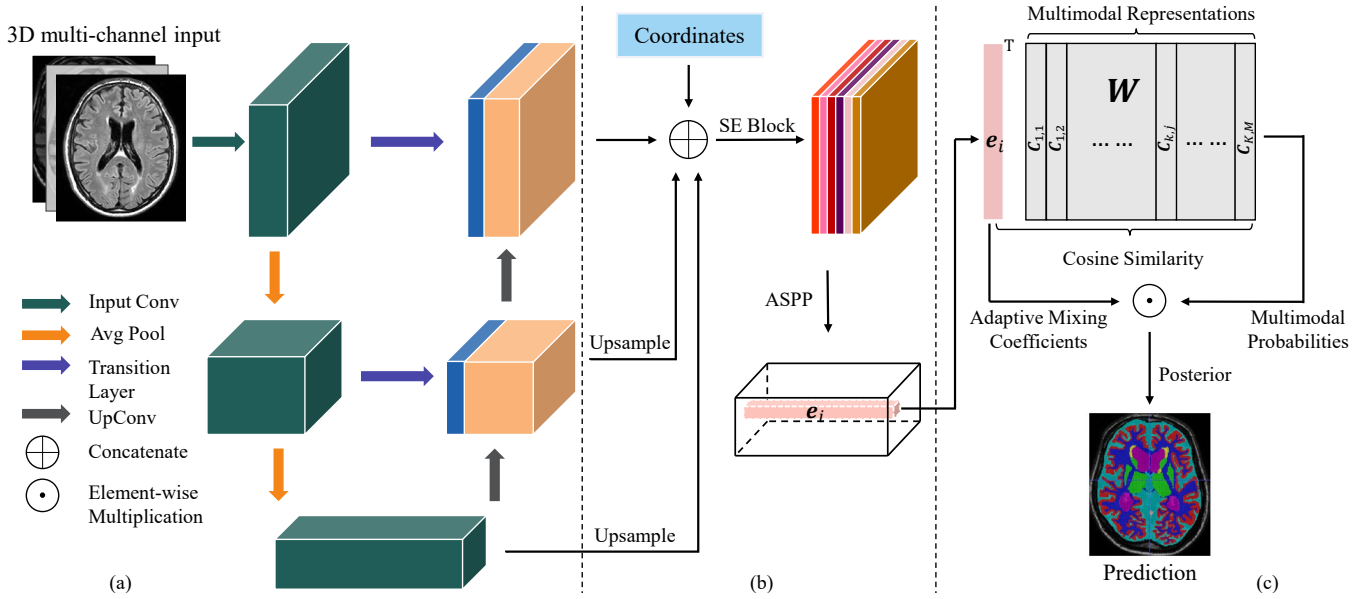

Fig. 2. Overview of the proposed framework: (a) backbone network, (b) dense feature embedding, and (c) cosine distance-based dense prediction. Note that 2D images are used for illustration purpose here, but the real data used in this work are 3D.

B. Backbone Network

U-Net [25] and its variants are widely used for medical image segmentation. Luna and Park [26] used a modified 3D U-Net and won the first place in the Grand Challenge on MR Brain Segmentation at MICCAI 2018 (MRBrainS18)[1]. Most notably, to mitigate the demanding need for a large amount of 3D training data, a structure shallower than the original U-Net was designed; in addition, the skip connection in the original U-Net was enhanced with a transition convolution, which rebalanced the low-level and high-level features by adjusting the number of channels in the former before concatenation. For its outstanding performance on the same dataset as used in this work, we adapt this structure as our backbone feature extractor, and a strong baseline in the experiments.

Structure of the backbone network is illustrated in Fig. 2(a). The input to the network are cropped 3D patches. Multi-sequence data are treated as multichannel input. The encoder comprises three blocks, consisting of three, four, and five convolutions, respectively. The first convolutions in these blocks increase the channels of the feature maps to 64, 128, and 256, respectively. The two average pooling layers between the encoder blocks reduce the size of the feature maps with the strides of [2, 2, 2] and [1, 2, 2] ([depth, height, width]), respectively. Skip connections in the original U-Net [25] are modified to transition layers, by adding channel-reducing convolutions (to 1/4 of the original) to the skip connections. Each of the two blocks in the decoder includes two convolutions, the first of which reduces the channels of the feature maps to 128 and 64, respectively. The transposed convolutions (UpConv) preceding the decoder blocks upsize the feature maps with the strides of [1, 2, 2] and [2, 2, 2], respectively. The upsized feature maps are concatenated with the output of the transition layers, and processed by the decoder blocks. All convolutions are followed by a rectified linear unit (ReLU) and batch normalization (BN). In addition, all convolutions use the kernel size of $3\times3\times3$, except for the transposed convolutions which have kernel sizes equal to those of their strides. The final output is a 64-channel feature map of the same size as the input.

[1]https://mrbrains18.isi.uu.nl/

C. Dense Feature Embedding

From the perspective of FCNs [24], segmentation is actually pixel-wise dense classification. In this sense, we have obtained a 64-dimensional feature vector for each pixel using the U-Net backbone. For normal DNN-based segmentation with abundant training samples and annotations, these features often produce decent results. In case of one-shot segmentation, however, they are not sufficiently descriptive of each pixel, given that the training data not only are limited in amount but also lack diversity in composition (*i.e.*, all examples of a category are from a single/few subjects). To offset this disadvantage, we propose to excavate more information from the backbone. Concretely, we concatenate the feature maps (upsampled if needed) on the right side of the U-Net (Fig. 2(b)) to collect features at different levels. Different from the semantic segmentation in natural images, spatial information has strong implications in medical images as many tissues and organs are location-specific. Hence, we further concatenate Cartesian coordinates as additional channels to the feature maps (see the "Coordinates" rectangle at the top of Fig. 2(b)). Specifically, we calculate a location map for each input image. The location map has the same size as the feature map, and three channels corresponding to the normalized coordinates in the x, y, and z axes, respectively. The maximal and minimal values in each axis are normalized to 1 and -1, respectively.

Considering the diversity in the segmentation targets in medical image analysis, we adopt two countermeasures to make our proposed MRE-Net more robust to the diversity. First, we feed the concatenated feature maps through an SE

attention block [17]. The SE block adaptively recalibrates channel-wise feature responses by explicitly modelling interdependencies between channels. Thus, the channels most informative about the specific segmentation task are emphasized. Second, we employ an atrous spatial pyramid pooling (ASPP) module [21] with five spatial scales for obtaining multi-scale features. The five scales are: one $1\times1\times1$ convolution, three $3\times3\times3$ convolutions with the atrous rates of 6, 12, and 18, respectively, and the image-level features. The outputs of the ASPP module are concatenated and fused by another $1\times1\times1$ convolution to produce the final pixel-wise embedding $\boldsymbol{e}_i$. Together, the concatenation of the (upsampled) multilevel feature maps, SE attention, and ASPP are deemed a holistic enhancement to the original U-Net feature embedding. We abbreviate the enhanced feature embedding as the attentional multi-scale (AMS) embedding, whose effect will be studied later using an ablation experiment.

D. DML of Multimodal Representation via Weight Embedding

To deal with potentially serious overfitting due to the extreme scarcity of training data, we employ DML [16] to learn category-discriminative representations for distance-based pixel-wise dense prediction, which makes good use of the inter-subject similarity. In addition, to fully exploit the intraclass variations, we employ a multimodal representation—instead of a single prototype—for each category, which is expected to further improve model generalization against overfitting. The explicit category representations employed in the previous work [16] caused a cumbersome network structure; in addition, only a single mode was used for classification at a time, thus the intraclass variations captured by the multimodal representations were not fully utilized. In contrast, we present a multimodal representation learning framework with self-adaptive mixing coefficients to simultaneously simplify the network design, enhance computational efficiency, and fully exploit the intraclass variations.

1) DML via Weight Embedding: For simplicity, we first introduce the case of unimodal category representation for DML, before describing our multimodal framework. Following the notation introduced in Section III-A, the distance between a pixel x_i's embedding vector $\boldsymbol{e}_i$ and a category prototype $\boldsymbol{c}_k$ can be computed by a function $d(\boldsymbol{e}_i, \boldsymbol{c}_k)$. Then, the probability of x_i belonging to the k^{th} category is proportional to the exponential of the negative distance: $p_k(x_i) \propto \exp\left[-d(\boldsymbol{e}_i, \boldsymbol{c}_k)\right]$. The discriminative class posterior $\mathbb{P}$ can be computed with a softmax operation over all candidate categories:

$$\mathbb{P}(s_i = k) = \frac{p_k(x_i)}{\sum_{k=1}^{K} p_k(x_i)} \triangleq \frac{\exp\left[-d(\boldsymbol{e}_i, \boldsymbol{c}_k)\right]}{\sum_{k=1}^{K} \exp\left[-d(\boldsymbol{e}_i, \boldsymbol{c}_k)\right]}, \quad (1)$$

where s_i is the segmentation label of x_i. Consider the cosine distance, *i.e.*, $d(\boldsymbol{e}_i, \boldsymbol{c}_k) = 1 - \cos(\theta) = 1 - \left(\frac{\boldsymbol{e}_i}{\|\boldsymbol{e}_i\|}\right)^{\mathrm{T}}\left(\frac{\boldsymbol{c}_k}{\|\boldsymbol{c}_k\|}\right) = 1 - \hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_k$, where θ is the angle between $\boldsymbol{e}_i$ and $\boldsymbol{c}_k$, $\|\cdot\|$ is the L2 norm, and $\hat{\boldsymbol{e}}_i$ and $\hat{\boldsymbol{c}}_k$ are the L2-normalized pixel's embedding vector and category prototype, respectively; it can be easily derived that $\frac{\exp(-1+\hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_k)}{\sum_k \exp(-1+\hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_k)} = \frac{\exp(\hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_k)}{\sum_k \exp(\hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_k)}$. The operation $\hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_k$ can be readily implemented using an fc layer with all biases set to zero, where $\boldsymbol{c}_k$ is the weight vector connected to the k^{th} output node, and each node represents a category. In this way, the cosine distance-based metric learning can be implicitly implemented with the commonly used fc layer, where the category prototypes are embedded as the weights of the fc layer: $\boldsymbol{W} = [\boldsymbol{c}_1, \ldots, \boldsymbol{c}_k] \in \mathbb{R}^{N_e \times K}$, instead of using explicit representations. From a different perspective, Luo *et al.* [18] discussed the benefits of the cosine similarity in terms of preactivation computation, including more flexibility and lower test error, and described the computing process as "cosine normalization"; we refer interested readers to [18] for more details. Closely related to our work, Qi *et al.* [29] used the cosine normalization for unimodal representation learning; in the next section we will describe how we apply the technique to multimodal representation learning and enhance it with self-adaptive mode-mixing coefficients.

Following [29], we add a learnable scaling factor ξ to adjust the exact values of the exponential function, thus we have:

$$\mathbb{P}(s_i = k) = \exp\left(\xi\hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_k\right) \Big/ \sum_{k=1}^{K} \exp\left(\xi\hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_k\right). \quad (2)$$

2) Multimodal Weight Embedding with Adaptive Mixing Coefficients: We now consider the case of multimodal representation learning via the cosine normalization technique (Fig. 2(c)), and more importantly, how to determine the relative contribution of each mode to the posterior $\mathbb{P}$. Specifically, we represent each category by a mixture distribution model, where the centers of the modes are the category prototypes (*i.e.*, a category is represented by a set of prototypes, instead of a single one). Let us denote the prototypes of the k^{th} category by $R_k = \{\boldsymbol{c}_{k,j}\}_{j=1}^{M}$, where $\boldsymbol{c}_{k,j}$ represents the center of the j^{th} mode and M is the number of modes. Same as in the unimodal case, each prototype gives an individual probability $p_{k,j}(x_i) \propto \exp\left[-d(\boldsymbol{e}_i, \boldsymbol{c}_{k,j})\right]$, and the multimodal mixture distribution is defined as $p_k(x_i) = \sum_{j=1}^{M} \alpha_{k,j} p_{k,j}(x_i)$, where $\alpha_{k,j}$ is the mixing coefficient for the j^{th} mode of the k^{th} category, and $\sum_{j=1}^{M} \alpha_{k,j} = 1$. Then, the posterior $\mathbb{P}$ can now be computed using the mixture models:

$$\begin{aligned}\mathbb{P}(s_i = k) &= \frac{p_k(x_i)}{\sum_{k=1}^{K} p_k(x_i)} = \frac{\sum_{j=1}^{M} \alpha_{k,j} p_{k,j}(x_i)}{\sum_{k=1}^{K}\sum_{j=1}^{M} \alpha_{k,j} p_{k,j}(x_i)} \\ &\triangleq \frac{\sum_{j=1}^{M} \alpha_{k,j} \exp\left(\xi\hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_{k,j}\right)}{\sum_{k=1}^{K}\sum_{j=1}^{M} \alpha_{k,j} \exp\left(\xi\hat{\boldsymbol{e}}_i^{\mathrm{T}}\hat{\boldsymbol{c}}_{k,j}\right)}.\end{aligned} \quad (3)$$

Conveniently, the multimodal DML can still be implemented using an fc layer, whose biases and weight matrix are set to zero and $\boldsymbol{W} = [\boldsymbol{c}_{1,1}, \boldsymbol{c}_{1,2}, \ldots, \boldsymbol{c}_{k,j}, \ldots, \boldsymbol{c}_{K,M}] \in \mathbb{R}^{N_e \times KM}$, respectively, where each category is jointly represented by M output nodes and $\boldsymbol{c}_{k,j}$ is the weight vector connected to the j^{th} node of the k^{th} category.

The mixing coefficient $\alpha_{k,j}$ controls the relative contribution of each mode, thus is an important design consideration. Karlinsky *et al.* [16] set $\alpha_{k,j} = 1$ if $j = \arg\max_j p_{k,j}$ and $\alpha_{k,j} = 0$ otherwise, *i.e.*, setting $\mathbb{P}$ to the upper bound on the actual class posterior. In this way, $\mathbb{P}$ was solely determined by the mode which was closest to the pixel's embedding vector for each category, thus the intraclass variations were not fully utilized. One way to make use of all the modes is to set $\alpha_{k,j} = 1/M$, *i.e.*, treating each mode of a category

equally. However, the equal-weight setting implied an underlying assumption that $\boldsymbol{e}_i$ was equally similar to $\boldsymbol{c}_{k,j}$ for all j, which may not always hold. In this work, we propose self-adaptable mixing coefficients for each category's multimodal distribution. It is implemented as a modified "squeeze-and-excitation" (SE) self-attention mechanism [17], which allows end-to-end training and is simple and straightforward. First, an fc layer takes $\boldsymbol{e}_i$ (of dimension $N_{\boldsymbol{e}} \times 1$) as input and outputs a vector of dimension 512×1 (the "squeeze"), where $N_{\boldsymbol{e}} \geq 512$, followed by ReLu. Then, another fc layer projects the squeezed and ReLu'ed vector to the dimension $KM \times 1$ (the "excitation"), followed by a sigmoid activation, producing a vector $\boldsymbol{\beta} = \left[\beta_{k,j} | k = 1, \ldots, K, \text{and } j = 1, \ldots, M\right]$ (we denote the parameter of these two fc layers by θ_{mix}). Lastly, the mixing coefficients are given by a per-category softmax operation:

$$\alpha_{k,j} = \exp(\beta_{k,j}) \Big/ \sum\nolimits_{j=1}^{M} \exp(\beta_{k,j}). \tag{4}$$

As we can see, this process adaptively calculates the mixing coefficients for the modes depending on the actual input $\boldsymbol{e}_i$, thus selectively emphasizing as well as neglecting different modes as needed. It is worth noting that setting $M = 1$ would degrade the framework to be unimodal. Hence, the unimodal and multimodal settings are unified within a single framework.

E. Loss Computation

We use the per-category normalized cross-entropy loss as the dense prediction loss:

$$\mathcal{L}_{\text{CE}} = -\sum\nolimits_{k} \delta_k(a_i) \log \left[\mathbb{P}(s_i = k)\right] / N_k, \tag{5}$$

where $a_i \in \{1, \ldots, K\}$ is the annotation label for pixel i, $\delta_k(a_i)$ is the Dirac delta function, which equals 1 if $k = a_i$ and 0 otherwise, and N_k is the number of pixels of the k^{th} category that are used for loss computation in a mini-batch. The category representations are implicitly embedded as weights of the fc layer, and updated by descending along the stochastic gradient of $\mathcal{L}_{\text{CE}}$. Hence, the framework learns the backbone network parameters, the embedding space, and the representations simultaneously in an end-to-end manner by optimizing $\mathcal{L}_{\text{CE}}$.

Class imbalance is a common problem for medical images, which may result in under-sampling and unsatisfactory performance for categories with considerably fewer examples. OHEM [19] has shown outstanding ability in dealing with foreground/background imbalance in object detection tasks, thus we adapt it to work pixel-wise for our low-shot segmentation task. Specifically, all the categories for segmentation are empirically grouped into a majority group and a minority group, based on the number of pixels in each category. Then, in each mini-batch, we keep all the pixels from the minority group for the loss computation, and rank those from the other according to their loss values. We keep only the most difficult majority-group pixels (identified by greatest losses) and discard the rest. In practice, the number of the majority-group pixels to keep is set to an integer (N_{Lg}) multiple of the total number of minority-group pixels.

Algorithm 1 Training procedure of the proposed MRE-Net.

Input: Annotated training set $\mathcal{X}^{\text{tn}} = \{(x^{\text{tn},(l)}, a^{(l)})\}_{l=1}^{L}$, where $L \leq 3$, and learning rate hyperparameter η
Output: Learned network parameters $\Omega = \{\Phi, \boldsymbol{W}, \theta_{\text{mix}}\}$

```
Randomly initialize Ω
for number of training iterations do
    Sample mini-batch of images (patches) from X^tn
    Evaluate L_CE in Equation (5) using the mini-batch
    Update parameters with gradient descent:
    Ω ← Ω − η∇_Ω L_CE
end for
```

F. Training and Testing Procedures

Training procedures of the proposed MRE-Net is listed in Algorithm 1. Here, the feature embedding function f_Φ includes both Fig. 2(a) (the backbone network) and Fig. 2(b) (the SE block and ASPP module). As we can see, all the learnable parameters (*i.e.*, Φ, $\boldsymbol{W}$, and θ_{mix}) of the MRE-Net are updated together by training the network end-to-end. For testing, we only need to forward the test image through the MRE-Net (following the arrows in Fig. 2), and the semantic label for a pixel can be determined by

$$s_i = \text{argmax}_k \, \mathbb{P}(s_i = k). \tag{6}$$

IV. Experiments

A. Datasets and Preprocessing

We evaluate the proposed MRE-Net on two public datasets: the MRBrainS18 dataset and the Beyond the Cranial Vault (BTCV) Abdomen [22], [23] dataset.

1) MRBrainS18: The task of the challenge is to segment various brain structures on multi-sequence brain MR images with and without (large) pathologies. The publicly available data (*i.e.*, the official training data) of MRBrainS18 comprises seven subjects, with fully annotated multi-sequence MRI scans acquired on a 3T scanner at the UMC Utrecht (the Netherlands), including T1-weighted (repetition time (TR): 7.9 ms and echo time (TE): 4.5 ms), T1-weighted inversion recovery (TR: 4416 ms, TE: 15 ms, and inversion time (TI): 400 ms), and T2-FLAIR (TR: 11000 ms, TE: 125 ms, and TI: 2800 ms) sequences. Officially, the scans were bias field corrected using the N4ITK algorithm [38], and the T1-weighted and T1-weighted inversion recovery images were aligned with the T2-FLAIR image for each subject. As illustrated in Fig. 2(a), multiple sequences of a subject are concatenated as channels of the 3D input. The manual annotations contain 10 labels excluding the background, and officially the evaluation is carried out on eight of them: cortical gray matter (GM), basal ganglia (BG), white matter (WM), white matter hyperintensities (WMH), cerebrospinal fluid (CSF), ventricles (VT), cerebellum (CB), and brain stem (BS). The MRI volumes are of size $240 \times 240 \times 48$ voxels, with the voxel size of $0.958 \times 0.958 \times 3.0$ mm^3. For computational efficiency, we crop the central region of size $216 \times 216 \times 48$ voxels for each scan, which is large enough to contain the whole brain. The z-score standardization is conducted on each MRI volume to

make the pixel intensities have a zero mean and unit standard deviation. No skull stripping is performed.

2) BTCV: The dataset comprises abdominal CT scans acquired at the Vanderbilt University Medical Center from metastatic liver cancer or post-operative ventral hernia patients. The scans were captured during portal venous contrast phase with variable volume sizes ($512 \times 512 \times 85$–$512 \times 512 \times 198$ voxels) and field of views (approx. $280 \times 280 \times 280$ mm^3–$500 \times 500 \times 650$ mm^3). The in-plane resolution varies from 0.54×0.54 mm^2 to 0.98×0.98 mm^2, while the slice thickness ranges from 2.5 to 5.0 mm. A total of 13 abdominal organs were manually annotated; following [39], we report segmentation performances on a subset of eight organs: spleen, right kidney (r.kid.), left kidney (l.kid.), gallbladder, esophagus, liver, stomach, and pancreas. For preprocessing, we crop the CT volumes to focus on the segmentation targets and reduce computational cost, according to the coordinates given in [39]; the cropping coordinates for three scans (out of the 30 of which the annotations are publicly available) turn out problematic, reducing the total number of usable scans to 27. After that, we resize the cropped volumes to $324 \times 396 \times 72$ voxels, and truncate the Hounsfield units of all scans outside the range $[-210, 290]$ to ignore irrelevant image content. Lastly, the image intensities are linearly scaled to $[0, 1]$.

A subset of seven scans is randomly selected for the low-shot experiments, mainly for two reasons. First, this setting is consistent with that for the MRBrainS18 dataset, allowing direct comparison between results on both datasets in a common sense; second, the remaining 20 scans will be used by a semi-supervised method as unlabeled auxiliary training data in a comparative experiment (to be described in Section IV-K2.)

B. Experimental Settings

We compare the performance of our proposed MRE-Net to that of the U-Net [26] (which ranked the first place in the MRBrainS18 challenge) in one-shot setting. Concretely, the model is trained with one sample and tested with the other six, yielding a mean performance averaged over the six test samples for each segmentation category or the cross-category mean metric. However, noting that the choice of the single labeled sample is crucial and can lead to a potential bias in the results, we repeat the experiment seven times by using each of the seven samples for training in turn, and report the mean performance and standard deviation computed over the seven runs. We report the performance both with and without online data augmentation. Also, we compare with the classical registration-based one-shot approach using the registration tool provided in the widely used Advanced Normalization Tools (ANTs) [15], which is one of the state-of-the-art methods for classical intensity-based registration. As to the few-shot settings, we experiment with two- and three-shot learning for our MRE-Net, to gain insights into how it performs in such setting with increasing amount of training data. Concretely, we randomly select pairs/triplets of samples for training and test on the remaining five/four samples. Again, the experiments are repeated seven times and the mean results averaged over the seven runs are reported.

TABLE I
DETAILS ABOUT TRAINING CONFIGURATIONS.

	MRBrainS18		BTCV	
	Ours	U-Net [26]	Ours	U-Net [26]
M	3	-	3	-
$N_{\boldsymbol{e}}$	2048	-	2048	-
N_{Lg}	6	-	4	-
Optimizer	Adam [40]		Adam [40]	
Learning rate (initial)	0.001		0.001	
Learning rate (adjustment)	Multiply by 0.1 every 1,500 iterations			
Input patch size	$36 \times 36 \times 12$		$36 \times 36 \times 12$	
Mini-batch size	24	64	24	64
Training iterations	3,000	20,000	3,000	20,000
GPU device	NVIDIA TITAN Xp (12GB memory)			
No. of GPUs	4×	1×	4×	1×

In addition, we evaluate the leave-one-out performance of the U-Net [26] on each dataset (*i.e.*, each of the seven samples is used for testing in turn while the other six being used for training), which is served as the upper bound of the low-shot settings. The Dice coefficient and 95$^{\text{th}}$ percentile of the Hausdorff distance (HD95) are used as the evaluation metrics. Statistical analysis is conducted to compare the mean performance of our MRE-Net with those of other approaches at the significance level of 5% (with Bonferroni correction when applicable). For reference, we also present results reported in related, non-low-shot-learning works [26], [39] which also involved the two datasets used in this work, along with our experimental results. However, due to the different amounts/compositions/sources of training and test data, backbone networks, and/or testing schemes, direct comparisons are precluded.

C. Implementation

The PyTorch framework is used for experiments. We train the networks from scratch and until convergence. The input are 3D patches randomly sampled from the given training samples. When data augmentation is involved, random left-right mirroring (50% probability; MRBrainS18 only), and random brightness and contrast jittering are implemented. For OHEM, the GM, WM, and VT are considered as the majority group for the MRBrainS18 dataset as they jointly occupy more than 85% of brain tissues according to our statistics, whereas gallbladder and esophagus are considered as the minority group for the BTCV dataset considering their small sizes. Although for both datasets only a subset of all the annotation labels is used for evaluation, we still maintain representations for non-evaluation categories for better performance on target categories. More training details are charted in Table I.

For testing, a straightforward sliding window approach is employed, in which the segmentation labels of consecutive patches of size $28 \times 28 \times 8$ voxels are predicted. Notably, to avoid the unwanted boarder effects, the actual input to the trained network is an expanded patch of size $36 \times 36 \times 12$ voxels and only the prediction results for the central $28 \times 28 \times 8$ voxels are used. Being aware that more sophisticated testing approaches are likely to produce better results (*e.g.*, Luna and Park [26] predicted every voxel eight times in their testing scheme), we choose the simple approach to focus

TABLE II
RESULTS OF THE ABLATION STUDY ON DESIGN COMPONENTS IN ONE-SHOT SETTING. VALUES ARE MEAN (STD.).

	DML	Cartesian coordinate	AMS embedding	OHEM	MRBrainS18		BTCV	
					Dice (%)	HD95 (mm)	Dice (%)	HD95 (mm)
(a)	-	-	-	-	14.86 (19.11)	59.30 (18.21)	25.91 (5.53)	58.29 (16.08)
(b)	✓	-	-	-	65.38 (2.27)	14.47 (2.08)	53.36 (2.75)	35.47 (9.26)
(c)	✓	✓	-	-	69.41 (2.72)	9.36 (1.34)	59.08 (2.47)	26.59 (8.70)
(d)	✓	✓	✓	-	74.47 (1.65)	6.95 (1.77)	67.86 (1.90)	21.18 (8.99)
(e)	✓	✓	✓	✓	**78.39** (1.07)	**6.30** (1.73)	**69.13** (1.43)	**19.02** (8.63)
(f)	-	✓	✓	✓	38.17 (4.18)	24.56 (3.54)	31.23 (4.82)	42.63 (12.69)
(g)	✓	-	✓	✓	70.67 (2.34)	12.76 (2.18)	56.96 (2.67)	22.69 (10.22)

TABLE III
IMPACT OF THE EMBEDDING DIMENSION N_e (WITH $M = 3$). VALUES ARE MEAN (STD.), WHERE APPLICABLE.

N_e	512	1024	2048	4096
Dice (%)	75.77 (1.93)	77.56 (0.95)	**78.39** (1.07)	76.99 (1.26)
HD95 (mm)	6.59 (1.98)	6.62 (1.62)	**6.30** (1.73)	6.47 (1.86)
GPU memory (GB)	∼4×5.90	∼4×7.68	∼4×10.18	∼4×14.44

TABLE IV
IMPACT OF THE NUMBER OF MODES M FOR EACH CATEGORY (WITH $N_e = 2048$). VALUES ARE MEAN (STD.).

M	1	2	3	4	5
One-shot					
Dice (%)	74.91 (1.98)	**78.42** (1.13)	78.39 (**1.07**)	78.05 (1.30)	77.68 (1.62)
HD95 (mm)	7.08 (1.81)	**6.21** (1.55)	6.30 (1.73)	6.54 (1.83)	6.93 (**1.40**)
Three-shot					
Dice (%)	79.21 (1.25)	81.39 (1.43)	**81.98** (**1.11**)	80.99 (1.39)	81.23 (1.91)
HD95 (mm)	5.88 (1.83)	**5.25** (1.79)	5.26 (**1.69**)	5.44 (1.76)	5.50 (1.75)

on demonstrating the effectiveness of our proposed network infrastructure in generalised low-shot medical image segmentation, instead of pursuing extreme performance by adopting complicated implementations.

For a fair comparison, we use the same loss function and testing scheme for the U-Net and the proposed MRE-Net.

D. Important Design Options

Using the MRBrainS18 dataset, we empirically study two important design options of the proposed MRE-Net: dimensionality of the feature embedding N_e and the number of modes M for each category. (i) We vary N_e in the set $\{512, 1024, 2048, 4096\}$, and chart the one-shot performance and GPU memory consumption during training in Table III. As expected, the memory consumption increased markedly with the increase of N_e. Meanwhile, the performance improved noticeably from $N_e = 512$ and peaked at 2048. Combining both factors, we select $N_e = 2048$ for other experiments. (ii) We vary M from 1 to 5 in both one- and three-shot settings, and present the corresponding performance in Table IV (the GPU memory consumption is relatively stable thus not included). As we can see, the trends were similar for both settings: the performance first increased and then fluctuated as M increased. Taking into account both the means and standard deviations of both metrics in both settings, we select $M = 3$ for experiments in this paper.

TABLE V
IMPACT OF MIXING COEFFICIENTS. VALUES ARE MEAN (STD.).

	One-hot [16]	Average	Adapt (proposed)
Dice (%)	77.04 (1.28)	76.08 (1.38)	**78.39** (1.07)
HD95 (mm)	7.59 (1.96)	6.54 (1.78)	**6.30** (1.73)

E. Impact of Mixing Coefficients

Using one-shot learning on the MRBrainS18 dataset, we compare the performance of different approaches to setting the mixing coefficients of the multimodal representations (described in Section III-D2), *i.e.*, setting that of the nearest prototype to 1 and others to 0 ("one-hot") [16], using all equal coefficients ("average"), and adapting the coefficients according to the input ("adapt"). The results are presented in Table V. As we can see, the proposed adaptive approach achieves the best performance in terms of both the Dice and HD95 metrics.

F. Ablation Study on Design Components

Considering inherent characteristics of medical image data, we innovatively incorporate four targeted solutions in our framework: (i) DML of multimodal representations to utilize inter-subject similarity and intra-class variations against overfitting, (ii) incorporation of spatial information (the Cartesian coordinates) to utilize inter-subject structural similarity, (iii) the attentional multi-scale (AMS) embedding for diverse medical images, and (iv) pixel-level OHEM to overcome the class imbalance problem. To verify the effectiveness of the proposed solutions, we conduct an ablation study by incrementally adding them to a baseline network to gradually build up our framework. The results are shown in Table II, from which we notice similar trends on the MRBrainS18 and BTCV datasets. Therefore, in the following we only describe the results on the former for conciseness. As we can see, the introduction of the DML brought improvements over 300% in both Dice and HD95 over the baseline. On top of that, the Cartesian coordinates, AMS embedding, and OHEM brought further conspicuous incremental improvements in both metrics. To summarize, all the proposed components contributed to the good performance of our framework, with the DML being most significant.

To further validate the critical role the DML plays in the MRE-Net, we conduct an additional ablation experiment in which the DML implementation is removed (*i.e.*, producing class posteriors in the way of FCNs) while the other advanced

TABLE VI
RESULTS OF THE ABLATION STUDY ON DISTANCE METRICS IN ONE-SHOT SETTING. VALUES ARE MEAN (STD.).

	MRBrainS18				BTCV			
Distance	Dice (%)	HD95 (mm)	GPU memory	Training time	Dice (%)	HD95 (mm)	GPU memory	Training time
Cosine	78.39 (1.07)	6.30 (1.73)	∼4×10.18 GB	∼150 min	69.13 (1.43)	19.02 (8.63)	∼4×9.81 GB	∼140 min
Euclidean	77.57 (1.01)	6.52 (1.91)	∼8×11.30 GB	∼260 min	67.68 (1.48)	18.95 (9.17)	∼8×10.95 GB	∼250 min

TABLE VII
EXPERIMENTAL RESULTS ON THE MRBRAINS18 DATASET ("AUG." STANDS FOR DATA AUGMENTATION; "-*n*" MEANS n-SHOT LEARNING; "MEAN EXCL." REPRESENTS RESULTS EXCLUDING WMH). VALUES ARE MEAN (STD.). BOLD FACES DENOTE BEST RESULTS PER COLUMN (GROUPED BY ONE- AND FEW-SHOT SETTINGS); ASTERISKS (*) DENOTE STATISTICALLY SIGNIFICANT DIFFERENCES FROM THE BEST RESULTS IN "MEAN" AND "MEAN EXCL." COLUMNS.

	GM	BG	WM	WMH	CSF	VT	CB	BS	Mean (std.)	Mean excl. (std.)
Dice Similarity Coefficient (%)										
ANTs-1 [15]	57.06 (3.18)	**77.27** (1.80)	63.98 (4.16)	26.36 (3.01)	57.29 (2.69)	85.13 (3.49)	84.07 (2.79)	**76.24** (11.25)	*65.93 (1.39)	*71.58 (1.29)
U-Net-1 [26]	14.06 (25.56)	6.44 (12.58)	22.22 (23.88)	0.18 (0.27)	24.63 (20.72)	18.77 (23.22)	16.35 (20.56)	16.21 (13.79)	*14.86 (19.11)	*16.96 (18.83)
Aug. + U-Net-1 [26]	58.84 (13.40)	41.54 (9.46)	56.56 (15.93)	26.38 (7.22)	63.13 (11.97)	56.53 (15.33)	55.50 (11.27)	34.72 (13.43)	*49.15 (11.91)	*52.40 (12.28)
MRE-Net-1 (ours)	**80.84** (1.61)	74.88 (2.71)	**83.70** (1.29)	58.42 (2.03)	**79.50** (1.58)	**90.57** (4.43)	87.57 (1.76)	71.66 (6.22)	**78.39** (1.07)	**81.25** (1.01)
Aug. + MRE-Net-1	80.60 (1.83)	75.16 (1.95)	83.21 (1.30)	**58.82** (1.39)	79.25 (1.63)	90.32 (3.06)	**87.60** (2.13)	71.61 (5.49)	78.32 (0.93)	81.11 (1.04)
MRE-Net-2 (ours)	82.34 (1.03)	78.94 (1.39)	84.78 (0.97)	68.67(14.26)	76.77 (0.66)	**92.23** (3.77)	88.61 (0.84)	72.06 (2.42)	80.55 (1.31)	82.25 (1.23)
Aug. + MRE-Net-2	82.46 (1.24)	78.82 (1.30)	84.92 (1.06)	68.38 (13.99)	76.67(0.63)	91.92 (3.72)	88.46 (0.75)	71.78 (2.72)	80.43 (1.78)	82.15 (1.61)
MRE-Net-3 (ours)	82.85 (1.00)	78.53 (1.20)	**85.21** (1.13)	72.64 (4.43)	**81.56** (0.24)	91.22 (0.75)	**90.51** (0.74)	**73.32** (3.01)	**81.98** (1.11)	**83.31** (1.15)
Aug. + MRE-Net-3	82.42 (0.88)	78.67 (1.32)	84.86 (1.01)	73.42 (4.17)	81.14 (0.22)	91.07 (0.79)	90.11 (0.75)	73.03 (2.52)	81.84 (1.28)	83.04 (1.33)
U-Net-6 [26]	81.26 (2.17)	78.49 (2.69)	81.98 (3.48)	74.10 (10.21)	79.20 (1.54)	86.80 (2.57)	88.53 (2.09)	70.39 (5.30)	80.09 (2.76)	80.95 (3.17)
Aug. + U-Net-6 [26]	**84.20** (1.94)	**78.96** (2.60)	84.80 (2.74)	**76.09** (10.25)	78.73 (1.39)	89.18 (3.77)	88.95 (1.78)	68.48 (5.69)	81.17 (2.60)	81.90 (2.80)
Non-low-shot [26][a]	86.0 (-)	83.4 (-)	88.2 (-)	65.2 (-)	83.7 (-)	93.1 (-)	93.9 (-)	90.5 (-)	-	-
95% Hausdorff Distance (mm)										
ANTs-1 [15]	3.04 (0.75)	**4.39** (0.40)	3.93 (0.98)	15.88 (0.86)	3.86 (0.51)	6.78 (6.59)	**5.01** (1.15)	**11.62** (9.86)	6.81 (1.09)	**5.52** (1.17)
U-Net-1 [26]	42.16 (28.74)	75.39 (29.41)	40.69 (22.27)	78.74 (23.61)	33.78 (27.65)	80.36 (20.30)	71.33 (17.21)	51.93 (4.55)	*59.30 (18.21)	*56.52 (17.97)
Aug. + U-Net-1 [26]	25.37 (20.17)	28.11 (18.52)	15.64 (16.96)	52.90 (19.16)	22.35 (11.33)	40.77 (13.36)	31.26 (9.09)	29.48 (10.34)	*30.74 (12.71)	*27.57 (11.96)
MRE-Net-1 (ours)	2.28 (0.12)	5.36 (1.58)	2.74 (0.19)	**10.33** (0.76)	**2.64** (0.22)	**6.31** (1.23)	7.56 (3.75)	13.16 (9.93)	**6.30** (1.73)	5.72 (1.76)
Aug. + MRE-Net-1	**2.10** (0.15)	5.32 (1.46)	**1.51** (0.21)	10.51 (1.51)	2.73 (0.18)	6.50 (1.45)	7.76 (3.78)	13.23 (8.80)	6.34 (1.87)	5.74 (1.79)
MRE-Net-2 (ours)	1.75 (0.13)	4.36 (0.32)	2.57 (0.20)	10.02 (2.50)	2.43 (0.14)	4.78 (0.53)	6.93 (2.45)	12.04 (2.01)	5.61 (1.06)	4.98 (1.24)
Aug. + MRE-Net-2	1.69 (0.09)	4.26 (0.38)	2.83 (0.21)	9.77 (1.46)	2.54 (0.12)	4.67 (0.79)	6.64 (2.93)	12.26 (4.90)	5.58 (1.62)	4.99 (1.57)
MRE-Net-3 (ours)	1.60 (0.11)	**4.03** (0.44)	2.49 (0.11)	9.87 (1.37)	2.31 (0.11)	3.26 (1.35)	6.17 (1.99)	12.32 (2.37)	5.26 (1.69)	4.60 (1.66)
Aug. + MRE-Net-3	**1.54** (0.09)	4.16 (0.30)	2.64 (0.15)	**9.61** (1.29)	**2.27** (0.14)	**3.00** (0.83)	**6.12** (1.96)	12.24 (2.00)	**5.20** (1.72)	**4.57** (1.75)
U-Net-6 [26]	1.86 (0.30)	4.42 (0.57)	2.91 (0.35)	17.25 (4.13)	2.95 (0.18)	8.24 (4.74)	7.16 (5.90)	15.31 (6.94)	7.51 (3.37)	6.12 (3.17)
Aug. + U-Net-6 [26]	1.65 (0.10)	4.12 (0.44)	**2.37** (0.19)	14.31 (5.32)	2.55 (0.09)	7.64 (3.91)	7.08 (4.77)	**10.32** (6.06)	6.26 (2.60)	5.10 (2.80)
Non-low-shot [26][a]	1.27 (-)	3.64 (-)	2.12 (-)	10.25 (-)	2.22 (-)	2.81 (-)	3.74 (-)	3.92 (-)	-	-

[a] As reported in [26] (missing numbers are represented by '-'). Note that different datasets and testing schemes preclude direction comparisons.

components listed in Table II are kept. The results are shown in row (f) of Table II. As we can see, the results were unusually low compared to those in rows (b)–(e), emphasizing the indispensable role of the DML. In addition, we similarly conduct an exclusion experiment for the Cartesian coordinates (row (g)). After removing the Cartesian coordinates, the performance dropped substantially compared to the full model (row (e)). The results double-confirmed the benefit of utilizing inter-subject structural similarity by incorporating spatial information.

G. Ablation Study on Distance Metrics

As introduced in Section III-D1, we use the cosine distance for the function $d(\boldsymbol{e}, \boldsymbol{c})$ to facilitate implementation of the DML with the fc layer. In this section, we evaluate another distance—the squared Euclidean distance [16]—in comparison with the cosine distance. Specifically, it is defined as $d_{\mathrm{Euc}}(\boldsymbol{e}, \boldsymbol{c}) = \left\| \boldsymbol{e}/\|\boldsymbol{e}\| - \boldsymbol{c}/\|\boldsymbol{c}\| \right\|^2 = \|\hat{\boldsymbol{e}} - \hat{\boldsymbol{c}}\|^2$. Implementation-wise, to use d_{Euc} we must represent the category prototypes with dedicated memory explicitly—which adds complexity to the framework, and can no longer compute the distances by forwarding through the fc layer as using the cosine distance. Instead, $d_{\mathrm{Euc}}(\boldsymbol{e}, \boldsymbol{c})$ is densely computed on GPUs for exhaustive combinations of the pixel-wise embedding $\boldsymbol{e}$ and category prototype $\boldsymbol{c}$. We compare the performance obtained using d_{Euc} on both the MRBrainS18 and BTCV datasets to that using the cosine distance; besides, the computational resource consumption for training is compared, too. The results are presented in Table VI. On one hand, the performance of d_{Euc} is slightly lower than that of the cosine distance, but still comparable, *e.g.*, the mean Dice coefficients on MRBrainS18 are 77.57% (Euclidean) versus 78.39% (cosine), while the corresponding HD95 are 6.52 mm (Euclidean) versus 6.30 mm (cosine). On the other hand, using d_{Euc} consumes substantially more resources than using the cosine distance: it takes more than 2 times GPU memory and ∼1.7–1.8 times duration for training. It is apparent that our choice of the distance metric (the cosine distance) and corresponding implementation (the cosine normalization technique [18]) are beneficial.

H. Evaluation Results on MRBrainS18

The results are presented in Table VII. The WMH category is very challenging in that these are WM abnormalities without regular shape, size, or location. Hence, we additionally present the mean evaluation metrics excluding WMH. As we can see, the U-Net [26] trained with one sample (U-Net-1) without data augmentation performed poorly (mean Dice 14.86% and HD95 59.30 mm). With data augmentation, the performance

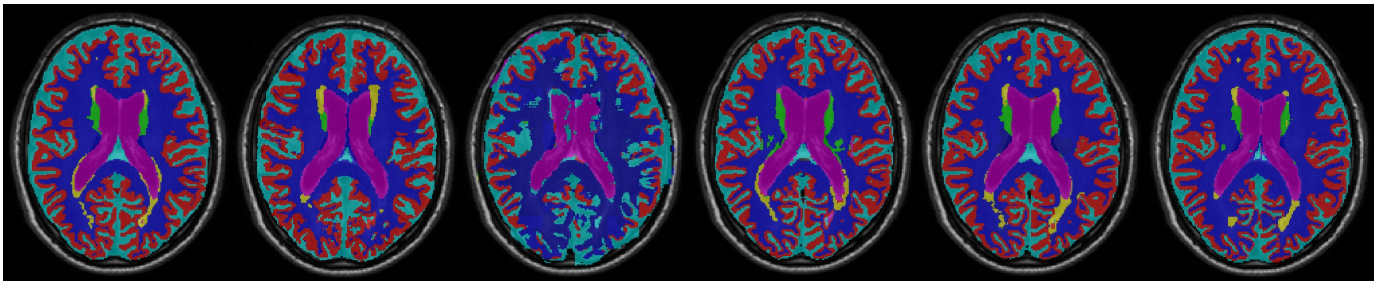

Fig. 3. Segmentation results of different methods on the MRBrainS18 dataset. Best viewed in color and zoomed-in.

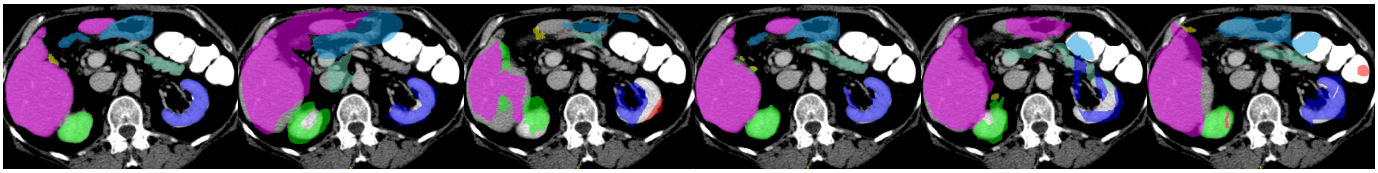

Fig. 4. Segmentation results of different methods on the BTCV dataset. Best viewed in color and zoomed-in.

of the U-Net improved markedly (mean Dice 49.15% and HD95 30.74 mm), yet still unsatisfactory. (Interestingly, the data augmentation in general improved the performance of the U-Net, yet did not cause much difference for the proposed MRE-Net. Accordingly, we only describe the performance of the U-Net with the data augmentation, and that of the MRE-Net without it below.) In contrast, the MRE-Net achieved quite reasonable results in the one-shot setting (mean Dice 78.39% and HD95 6.30 mm), significantly outperforming those of U-Net-1. When compared with U-Net-6, MRE-Net-1 achieved ∼97% of the mean Dice of the upper bound (78.39±1.07% vs. 81.17±2.60%) and a comparable mean HD95 with that of the upper bound (6.30±1.73 mm vs. 6.26±2.60 mm). Meanwhile, ANTs [15] achieved decent one-shot performance: mean Dice 65.93% and HD95 6.81 mm. In comparison, the mean Dice and HD95 by MRE-Net-1 were higher by 12.46% and smaller by 0.51 mm, respectively.

Regarding the few-shot setting, we can observe apparent improvements in both Dice and HD95 by increasing the number of training samples from one to three. Notably, using two training samples, our MRE-Net already exceeded U-Net-6 in terms of mean HD95; when using three training samples, the MRE-Net achieved a mean Dice higher than that of U-Net-6. In summary, the MRE-Net achieved the best performance in both the one- and few-shot settings, demonstrating competence in low-shot brain MRI segmentation. Example segmentation results of different methods are shown in Fig. 3.

When compared with the non-low-shot approach proposed by Luna and Park [26], we observe close performances of our MRE-Net and U-Net-6 in few-shot settings, although the non-low-shot approach yielded an overall superior performance. In fact, all these three approaches employed the same 3D U-Net backbone, and the main differences were three-fold. (i) Our MRE-Net incorporated countermeasures specific for low-shot learning—most notably the multimodal DML—on top of the backbone. (ii) The non-low-shot approach employed a more sophisticated testing scheme for extreme performance in the challenge competition (and actually won the first place). (iii) In addition, it was trained with all the seven subjects used in this work, and tested on 23 test subjects which are not publicly available. With due caution, we try to explain the comparison as following. First, the close performances validated our choice of the backbone. Second, the overall superior performance of the non-low-shot approach to our MRE-Net might be partially attributed to the more complicated testing scheme, and the fact that more training data and different test data were used. Third, we notice that for the WMH our evaluation results were actually better than those reported in [26], which might be caused by different distributions of the test data, urging caution in the interpretation, again.

I. Evaluation Results on BTCV

The proposed MRE-Net achieved competitive performance on the MRBrainS18 dataset. To validate the capability of the MRE-Net to generalize for data of other modalities as well as other parts of the body, we also apply it to the BTCV dataset. However, low-shot multi-organ segmentation in abdominal CT is more challenging, mainly for two reasons: (i) the contrast of soft tissues in CT is not as good as that in MRI; and (ii) the inter-subject variations in both appearance and structure—such as shape, size, and location—can be much larger in abdominal CT than in brain MRI. From the preceding experiment on the MRBrainS18 dataset, we have learned that the data augmentation generally works for the U-Net but not for our MRE-Net. Hence, we only experiment with the corresponding settings on the BTCV dataset. The evaluation results are presented in Table VIII. As we can see, our MRE-Net-1 achieved the best performances for all organ-wise and cross-organ evaluation metrics with large margins, substantially outperforming other competing methods including U-Net-6. It is worth noting that neither ANTs nor the U-Net managed to produce reasonable results for this challenging task. On the contrary, our framework was able to overcome the large inter-subject variations using only one sample for

TABLE VIII
EXPERIMENTAL RESULTS ON THE BTCV DATASET ("AUG." STANDS FOR DATA AUGMENTATION, "-n" MEANS n-SHOT LEARNING). VALUES ARE MEAN (STD.). BOLD FACES DENOTE BEST RESULTS PER COLUMN; ASTERISKS (*) DENOTE STATISTICALLY SIGNIFICANT DIFFERENCES FROM THE MRE-NET-1 IN THE "MEAN" COLUMN.

	Spleen	R.kid.	L.kid.	Gallbladder	Esophagus	Liver	Stomach	Pancreas	Mean (std.)
Dice Similarity Coefficient (%)									
ANTs-1 [15]	51.62 (2.97)	33.60 (4.63)	17.65 (5.57)	15.25 (4.25)	56.30 (3.71)	76.48 (1.55)	24.60 (1.61)	4.31 (0.64)	*34.98 (3.73)
Aug. + U-Net-1 [26]	31.48 (3.75)	44.84 (4.11)	34.03 (4.62)	26.38 (2.98)	17.24 (2.80)	61.68 (1.41)	11.56 (0.59)	19.48 (0.82)	*30.84 (5.22)
MRE-Net-1 (ours)	**78.08** (1.92)	**75.92** (1.70)	**77.43** (2.98)	**52.15** (1.76)	**56.83** (1.62)	**91.35** (0.81)	**59.65** (1.53)	**61.26** (2.21)	**69.13** (1.43)
Aug. + U-Net-6 [26]	63.49 (1.56)	66.26 (1.89)	66.29 (0.80)	33.74 (2.32)	30.99 (2.45)	87.27 (0.88)	33.19 (0.89)	29.88 (1.02)	*51.39 (2.98)
Aug. + U-Net-20 [26]	64.81 (1.33)	72.73 (1.77)	69.26 (0.95)	28.63 (1.12)	25.75 (1.59)	85.51 (1.04)	45.95 (1.94)	25.80 (0.43)	*52.30 (2.57)
Non-low-shot[a]									
Zhou *et al.* [41]	92 (-)	-	91 (-)	65 (-)	43 (-)	95 (-)	60 (-)	62 (-)	-
DenseVNet [39]	95 (-)	-	93 (-)	73 (-)	71 (-)	95 (-)	87 (-)	75 (-)	-
95% Hausdorff Distance (mm)									
ANTs-1 [15]	63.30 (10.57)	42.44 (13.59)	105.26 (29.81)	41.56 (16.76)	43.78 (10.11)	31.72 (7.70)	33.09 (5.89)	70.33 (7.43)	*53.93 (15.47)
Aug. + U-Net-1 [26]	47.14 (8.52)	52.93 (10.71)	51.77 (8.09)	44.88 (7.98)	40.39 (10.30)	34.01 (5.27)	32.91 (5.25)	46.15 (6.17)	*43.77 (13.18)
MRE-Net-1 (ours)	**22.93** (3.37)	**18.54** (3.20)	**19.47** (5.66)	**27.61** (7.23)	**15.87** (6.78)	**10.62** (2.78)	**12.98** (2.36)	**24.10** (5.16)	**19.02** (8.63)
Aug. + U-Net-6 [26]	43.05 (3.82)	44.22 (5.00)	47.33 (6.57)	36.16 (9.39)	35.06 (10.44)	27.28 (7.17)	28.45 (10.57)	36.26 (6.69)	*37.23 (9.57)
Aug. + U-Net-20 [26]	42.57 (2.10)	39.08 (4.31)	46.65 (7.95)	32.12 (5.12)	34.87 (8.24)	21.35 (7.48)	20.74 (3.16)	31.17 (6.19)	*33.57 (6.91)

[a] As reported in [39] (missing numbers are represented by '-'); mean values for HD95 were not reported in [39]. Different datasets and testing schemes preclude direction comparisons.

training. To further evaluate the performance of MRE-Net-1 when compared to the U-Net trained with even more data, we train the U-Net with all the 20 scans not selected for the low-shot experiment and test it on the selected seven scans. Results showed that U-Net-20 outperformed U-Net-6 in both mean Dice and HD95, as expected; however, our MRE-Net-1 actually performed even better than U-Net-20. We conjecture that the relatively small improvements of U-Net-20 over U-Net-6 might be partially attributed to the difficulty of the problem itself, and the possibility that these 20 scans did not present enough diversity for well training of the U-Net, which was not tailored for learning with limited data. These results demonstrate superiority of our proposed MRE-Net in dealing with low-shot segmentation of data with higher extents of inter-subject variations. Fig. 4 shows example segmentations produced by the different methods in Table VIII.

The performance of non-low-shot approaches to abdominal multi-organ segmentation reported in the literature [39] was generally better than the one-shot performance of our MRE-Net, with the exceptions that the Dice values of MRE-Net-1 on esophagus, stomach, and pancreas were higher or on a par with those of [41]. Again, different training and testing data were used: [39] used 90 CT scans (27 of which overlapped with the BTCV dataset used in this work) with a 9-fold cross-validation, and [41] used a totally different dataset of 240 CT scans with a 230:10 (train:test) split. Therefore, the comparison results must be interpreted with due caution. The superior evaluation results reported in [39], [41] to those obtained in this work were reasonable and expected, considering that many times more data were used for training. Meanwhile, the superior performance of [39] to that of [41] might be partially attributed to the more advanced network structures of higher capacity employed in [39] than in [41], although differences between the two datasets might also have played a role.

J. Time Efficiency

Table IX compares the time efficiency of the various methods in Table VII, including training time and average inference time, on the MRBrainS18 dataset. Apparently, ANTs—despite yielding decent one-shot performance compared to the U-Net—was orders of magnitudes slower than the DNN-based approaches, kind of representative of the classical registration methods. Meanwhile, despite 30-min longer training time, the proposed MRE-Net presents no practically perceivable difference in the inference time from the U-Net.

TABLE IX
COMPARISON OF THE TIME EFFICIENCY OF DIFFERENT METHODS.

	ANTs [15]	Aug. + U-Net [26]	MRE-Net (ours)
Training (one-shot)	-	∼120 min	∼150 min
Inference (per sample)	∼28 min[a]	∼3s	∼4s

[a] Using 64 threads on an Intel® Xeon® E5-2690 v4 @2.60GHz processor.

K. Comparison with Existing Low-Shot Approaches

As introduced earlier, existing approaches on few-shot medical image segmentation relied on the accessibility to auxiliary training data [11], [35]. We are aware of the different usage scenarios intended, but also believe that it is useful to compare with some state-of-the-art (SOTA) methods among existing approaches. In this part, we compare our proposed MRE-Net to the channel squeeze & spatial excitation (sSE) guided method [35] and data augmentation (DataAug) method [11], with necessary adaption. These two methods were both recently proposed, and representative of the one-shot medical image segmentation methods relying on (i) auxiliary data with annotations for base classes and (ii) unlabeled training data, respectively. It is worth noting that although both methods have the potential for few-shot settings, the optimal way of extension is neither presented by the authors nor obvious. Hence, the comparison is made in one-shot settings.

1) Comparison with the sSE Method: The sSE method [35] was originally developed for the scenario in which a separate dataset with sufficient annotations of base classes is provided, to facilitate segmentation of a novel class versus the background (binary segmentation) given an example volume. We

TABLE X
COMPARISON WITH THE SSE METHOD [35] IN ONE-SHOT SETTING ("AUG." STANDS FOR DATA AUGMENTATION). VALUES ARE MEAN (STD.). BOLD FACES DENOTE BEST RESULTS PER COLUMN, AND ASTERISKS (*) DENOTE STATISTICALLY SIGNIFICANT DIFFERENCES FROM THE MRE-NET.

	MRBrainS18		BTCV	
	Dice (%)	HD95 (mm)	Dice (%)	HD95 (mm)
MRE-Net (ours)	**78.39** (1.07)	**6.30** (1.73)	**69.13** (1.43)	**19.02** (8.63)
Aug. + U-Net [26]	*49.15 (11.91)	*30.74 (12.71)	*30.84 (5.22)	*43.77 (13.18)
Aug. + sSE (k = 10)	*28.20 (2.26)	*31.58 (0.74)	*23.04 (4.05)	*61.53 (13.14)
Aug. + sSE (k = max.)	*22.66 (1.33)	*36.54 (0.90)	*25.37 (1.75)	*63.87 (3.31)

adapt it for the one-shot scenario considered in this work, *i.e.*, one-shot multiclass segmentation in which a single example with all target classes annotated is given but without any auxiliary data. Concretely, we follow the training scheme of sSE to train the model until convergence, by sampling both the support and query sets from the only annotated volume. For testing, the annotated volume is reused as the support volume to infer the segmentation of the query volumes, class by class. We experiment with varying hyperparameters of sSE for optimal performance; especially, we experiment with two different budgets of slice-wise annotations (the hyperparameter k): $k = 10$ (the optimal tradeoff suggested in [35]) and $k =$ the total number of slices which contain a specific organ/tissue (the maximum value allowed, providing more guidance in theory); the latter is enabled in our scenario as the support volume is fully annotated. The same one-shot setting as described in Section IV-B is used.

The results are charted in Table X. We notice that increasing the number of slice-wise annotations from 10 to the maximum values allowed did not result in substantial differences, which was consistent with the finding of Roy *et al.* [35]. However, our results also indicated that the sSE method with the straightforward adaption did not work for the target scenario considered in this work, yielding compromised performance. We conjecture that the unsatisfactory performance was caused by severe overfitting, as sSE was originally designed to handle extreme scarcity of annotations rather than of data.

2) Comparison with the DataAug Method: The intended usage scenario of DataAug partially overlaps with ours in that a single fully annotated volume (the "atlas") for multiclass classification is provided. However, in order to learn the spatial and appearance transformations for synthesizing extra training data, a pool of unlabeled data are additionally needed, making DataAug a semi-supervised method. This difference makes it infeasible to compare DataAug with our proposed MRE-Net in exactly the same setting (in contrast to the comparison with the sSE method); instead, we formulate a comparison setting as described below.

Noticing that there are 20 scans in the BTCV dataset that are not used in the low-shot experiments and thus can be used as unlabeled data, we compare with DataAug on the BTCV dataset. In addition, it can be prohibitively time-consuming to train and test DataAug following the one-shot setting described in Section IV-B (*i.e.*, repeating the one-shot experiment by using each of the seven scans selected for low-shot experiments as the atlas in turn), as it necessitates training of three DNN models (spatial transformation, appearance transformation, and segmentation) for just one of these experiments. Instead, we pick the scan that is closest to the anatomical average of the seven scans as the atlas—as done in [11]—and report the performance evaluated on the remaining six scans. To evaluate the varying benefits of using different amounts of unlabeled data, we experiment with 5, 10, and 20 unlabeled scans for DataAug. Further, we experiment with two strategies proposed in [11] for generating training data with the learned transformation models: (i) data augmentation using single-atlas segmentation (SAS-aug), where the segmentation mask of the atlas is warped to the unlabeled data, which are then used as additional training data for the segmenter; and (ii) data augmentation using independent sampling (IDS-aug), where the spatial and appearance transformations are independently sampled and combined to synthesize training data. The generated training data are used to train the U-Net [26] along with online data augmentation.

The experimental results are presented in Table XI, from which we summarized three findings. First, increasing the number of unlabeled scans from 5 to 20 consistently improved the performances for both variants of DataAug. Second, using the U-Net trained with the only annotated sample as baseline, extra training data produced by DataAug helped the segmentation of certain larger organs (spleen and liver) but hurt that of others, ultimately leading to slightly worse cross-organ mean results. Third, our MRE-Net yielded the best performance for all organs and both metrics. We conjecture that the conflicting performance of DataAug was caused by the unsatisfactory registration of the abdominal CT data, as it was originally proposed for registering and synthesizing brain MRI data [11], which intrinsically present subtler individual variations. Consequently, the unsatisfactory registration resulted in synthesized data/warped masks of low quality, which did not help or even interfered the training of the segmentation networks. A side evidence is that for larger organs that were registered more reasonably, we did observe apparent performance improvement over the U-Net baseline. We thus believe that if a spatial transformation model better suited to the abdominal organs could be integrated into the DataAug pipeline, better performance can be expected.

V. DISCUSSION AND CONCLUSION

In this work, we proposed a new DNN-based framework, the MRE-Net, for generalized low-shot medical image segmentation in the scenario of extremely scarce training samples and annotations based on DML. The framework addressed difficulties in this challenging task (most notably data scarcity and overfitting, followed by class imbalance and great diversity) by fully utilizing inter-subject similarities and intraclass variations via multimodal representation learning, as well as incorporating other targeted solutions including structural similarity, online hard example mining, and attentional multi-scale feature embedding. Experimental results on the MRBrainS18 and BTCV datasets showed that our framework was superior to competing methods including registration in both one-shot and few-shot settings. As far as the authors are aware of, this

TABLE XI
COMPARISON WITH THE SEMI-SUPERVISED DATAAUG METHOD [11] IN ONE-SHOT SETTING ON THE BTCV DATASET ("AUG." STANDS FOR ONLINE DATA AUGMENTATION; "-n" MEANS n UNLABELED SCANS WERE USED). VALUES ARE MEAN (STD.). BOLD FACES DENOTE BEST RESULTS PER COLUMN; ASTERISKS (*) DENOTE STATISTICALLY SIGNIFICANT DIFFERENCES FROM THE MRE-NET IN THE "MEAN" COLUMN.

	Spleen	R.kid.	L.kid.	Gallbladder	Esophagus	Liver	Stomach	Pancreas	Mean (std.)
Dice Similarity Coefficient (%)									
Aug. + U-Net [26]	32.32 (6.38)	47.13 (18.40)	32.40 (8.54)	14.81 (4.14)	19.43 (3.69)	57.61 (11.90)	54.25 (12.08)	19.55 (5.88)	*34.69 (3.84)
MRE-Net (ours)	**78.76** (3.46)	**73.54** (4.42)	**78.36** (3.26)	**60.43** (9.77)	**55.44** (14.27)	**89.79** (3.08)	**62.43** (13.21)	**59.71** (15.94)	**69.81** (4.58)
Semi-supervised [11]									
SAS-aug-5	29.51 (3.44)	8.10 (4.57)	12.58 (4.91)	6.85 (1.98)	4.93 (3.08)	60.78 (7.51)	14.46 (9.68)	18.28 (9.19)	*19.44 (3.15)
IDS-aug-5	36.26 (6.51)	11.66 (6.33)	15.58 (11.16)	2.93 (4.49)	13.59 (5.01)	64.02 (9.29)	23.49 (10.17)	20.23 (8.23)	*23.47 (2.33)
SAS-aug-10	38.76 (6.20)	13.40 (3.93)	13.30 (4.08)	6.12 (3.72)	9.51 (4.29)	74.74 (3.12)	24.08 (3.15)	18.16 (8.87)	*24.76 (2.46)
IDS-aug-10	37.10 (8.11)	12.85 (6.78)	15.58 (11.94)	3.46 (4.17)	14.10 (5.65)	64.61 (9.23)	23.90 (9.83)	21.84 (9.29)	*24.18 (2.61)
SAS-aug-20	42.89 (8.13)	16.51 (13.79)	16.74 (8.27)	6.68 (5.56)	13.06 (2.98)	79.53 (3.79)	30.92 (3.88)	22.34 (8.33)	*28.58 (3.74)
IDS-aug-20	38.85 (8.18)	15.41 (6.86)	16.97 (11.80)	3.55 (2.86)	17.63 (5.05)	68.77 (7.59)	28.98 (10.25)	22.92 (10.20)	*26.64 (2.11)
95% Hausdorff Distance (mm)									
Aug. + U-Net [26]	42.32 (15.49)	51.76 (14.31)	47.37 (24.62)	46.31 (6.96)	33.73 (11.83)	34.15 (11.78)	34.02 (9.10)	39.26 (18.33)	*41.12 (7.43)
MRE-Net (ours)	**19.89** (9.73)	**16.57** (4.36)	**18.49** (5.63)	**21.02** (5.75)	**13.42** (6.11)	**12.31** (7.87)	**9.30** (5.05)	**15.37** (5.28)	**15.80** (3.46)
Semi-supervised [11]									
SAS-aug-5	43.08 (10.20)	54.97 (7.92)	57.66 (12.72)	47.30 (13.55)	37.58 (7.73)	35.79 (7.21)	39.78 (8.01)	55.27 (13.45)	*46.43 (4.68)
IDS-aug-5	43.12 (7.36)	45.40 (12.10)	51.68 (5.49)	47.19 (6.57)	44.86 (7.08)	31.65 (7.25)	34.96 (12.72)	43.20 (13.37)	*42.76 (3.74)
SAS-aug-10	36.98 (9.83)	49.37 (7.20)	51.09 (12.16)	42.15 (13.24)	31.21 (7.94)	30.34 (5.95)	32.95 (7.75)	47.98 (18.94)	*40.26 (3.67)
IDS-aug-10	42.40 (6.60)	44.51 (12.52)	49.07 (5.96)	46.36 (9.74)	44.03 (7.86)	29.71 (6.56)	33.04 (12.69)	42.98 (13.82)	*41.51 (3.86)
SAS-aug-20	34.86 (15.31)	45.52 (13.15)	47.00 (17.51)	43.96 (6.13)	29.15 (12.97)	30.05 (14.22)	27.32 (10.09)	45.57 (24.33)	*37.93 (10.41)
IDS-aug-20	40.13 (6.12)	47.96 (7.93)	46.52 (6.71)	43.71 (12.65)	41.38 (8.77)	29.06 (7.63)	30.88 (11.74)	42.63 (13.63)	*40.28 (3.08)

work is the first attempt at addressing the extreme scarcity of both annotation and data for DNN-based segmentation, which is expected to exert a significant impact in applicable clinical and research scenarios.

1) Clinical Impact: The scenario considered in this work has great clinical value for rare diseases. There are about 7,000 known rare diseases [42] which collectively affect about 400 million people in the world [43]. It is generally difficult to obtain data for these conditions which are usually overlooked by the medical imaging community. For example, Li *et al.* considered as few as one to several samples per class for rare disease diagnosis [44]. Although in this work common brain MRI and abdominal CT data were used for experiments, for a valid evaluation, we deliberately used extremely limited data to simulate the real situation—just like Li *et al.* did [44]—and achieved promising results in our experiments. Hence, we expect the proposed MRE-Net to work effectively in the target scenario of rare diseases. In addition, in case of emergency, the MRE-Net can be employed as a rapid yet effective response prior to the accumulation of sufficient data for fully or semi-supervised approaches.

2) Performance and Comparison: ANTs [15] and U-Net [25], [26] are generally considered as benchmarks of classical intensity-based registration and generic DNN-based medical image segmentation, respectively. In our experiments on MRBrainS18, ANTs achieved decent one-shot results, and the U-Net [26] achieved competitive six-shot results. However, they are known to be either time-consuming or data-hungry. In this work, we aimed at solving these two issues simultaneously. On one hand, the proposed MRE-Net outperformed both ANTs and the U-Net (including U-Net trained with times more data) in the low-shot setting. On the other hand, MRE-Net was orders of magnitudes faster than ANTs (a few seconds versus 28 minutes). Moreover, on the BTCV dataset, neither ANTs nor the U-Net managed to produce meaningful results (Table VIII). In contrast, our MRE-Net produced quite reasonable results even with only one training sample. We attribute the superior performance of the MRE-Net to the employment of the DML for low-shot learning as well as the proposal of several targeted solutions considering characteristics of medical images (Table II). It is noticed that the MRE-Net requires more GPU resources for training than the U-Net [26] (last row of Table I), mainly because the DML necessitates dense distance computations for exhaustive pair-wise combinations of the pixel-wise embeddings and category prototypes. However, considering the substantial performance gains (*e.g.*, up to 124% increase in mean Dice and up to 80% reduction in mean HD95 in one-shot settings, and the superior performance of the MRE-Net to the U-Net in low-shot settings), we believe the trade-off between computational resources and performance is practically acceptable in the realm of low-shot medical image segmentation.

The scenario considered in this work differed from the conventional settings in the literature, making direct comparison with existing methods difficult. We focused on the scenario in which both the annotation and data are scarce, and deliberately used small amounts of data for experiments. We are aware that existing approaches to low-shot medical image segmentation require auxiliary training data; however, we think it would still be informative to compare our proposed MRE-Net with some SOTA approaches in controlled settings. Specifically, we compared with the sSE method [35] in the same one-shot setting as ours, and the DataAug method [11] trained in the semi-supervised one-shot setting. The comparison results emphasized the need for a tailored solution to the specific problem we consider, as both approaches—simply adapting sSE, and applying DataAug to abdominal CT given small to modest amounts of unlabeled data—yield compromised results. In contrast, our MRE-Net was specifically designed for the extremely low-data regime, and a prominent advantage over existing approaches is the avoidance of large amounts of data. Another advantage worth mentioning is that the MRE-

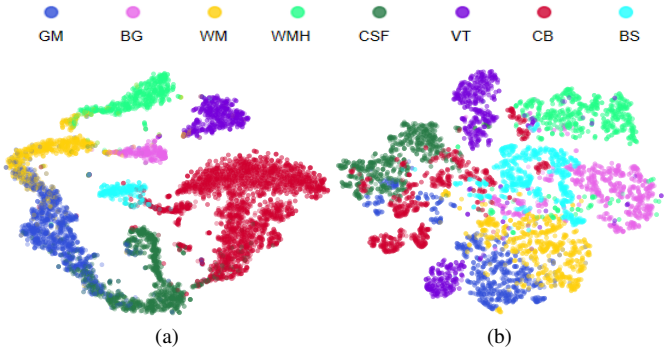

Fig. 5. Test sample visualization (t-SNE) of the feature spaces learned by (a) the proposed MRE-Net and (b) the 3D U-Net [26] (with online data augmentation) on the MRBrainS18 dataset in one-shot setting. Different colors represent different tissues (best viewed digitally).

Net naturally unifies both one- and few-shot learning with a single framework; in contrast, although extension to few-shot setting is possible for the two compared methods [11], [35], the optimal way of extension was neither presented nor obvious.

Compared to the non-low-shot performance reported in two SOTA related works [26], [39], apparent gaps were observed for the low-shot performance of the MRE-Net on the BTCV dataset (although MRE-Net-1 achieved comparable or even superior results to those of [41] for some organs). Such differences were expected, considering that 80 and 230 times data were used for training in [39] and [41], respectively. Note that considering the extremely limited training data in low-shot settings, we chose a relatively lightweight model as compared to the DenseVNet [39], to avoid potential overfitting. Similarly, wide performance gaps between low-shot and non-low-shot results were also reported in [35]. On one hand, if access to many fully annotated data is provided, it is always recommended to design more powerful models of larger capacity to make full use of the data and annotations. On the other hand, our work was focused on the data scarcity problem—instead of trying to beat non-low-shot models trained with abundant data and annotations, and our experimental results demonstrated effectiveness of the proposed MRE-Net in producing reasonable segmentation in the target scenario.

3) Why Does Our MRE-Net Work?: To effectively utilize the inter-subject similarity and model the intraclass variations from the extremely limited data, our proposed MRE-Net employed the DML to learn a multimodal representation for each category and address the key issue of overfitting. The inter-subject similarity was utilized by the MRE-Net in the sense that, the DML essentially determined the semantic label of a pixel based on its similarity (distance) to the category prototypes, which were learned from the training data in the embedding space (see Fig. 1 and Section III-A). This non-parametric approach basically helped prevent overfitting [6]. The critical role of the DML was demonstrated by an ablation experiment, in which the performance went straight down after the DML was removed (row (f) of Table II). For an intuitive interpretation, we employ the t-SNE [45] to visualize the feature embeddings learned by the MRE-Net, in comparison to those by the 3D U-Net [26]. In Fig. 5(a), data points of different categories formed clear and compact clusters for the MRE-Net, whereas those in Fig. 5(b) were more difficult to separate for the 3D U-Net. This visualization demonstrated that the DML helped our MRE-Net learn a better feature space where the embeddings of different categories were more separable (recall that the goal of DML is to make the projected samples of the same class closer to each other than those of different classes in the embedding space). Meanwhile, our MRE-Net captured the intraclass variations with a mixture of multiple modes for each category, such that the diversity contained in the limited training data was fully exploited. As empirically demonstrated in Table IV, the multimodal framework (*i.e.*, $M = 2$ and 3) gained considerable advantages in performance over the unimodal counterpart ($M = 1$), in both one- and three-shot settings.

The proposed framework implemented the concept of multimodal representation learning using the cosine normalization technique [18], by embedding the representations as weights of the fc layer instead of allocating specialized memory for them, and computing the cosine similarities via forward prorogation [29]. Correspondingly, we named it the Multimodal Representation Embedding Network (MRE-Net). In this way, the network design was simplified, and by utilizing the heavily optimized built-in implementation of the fc layer in commonly used deep learning platforms, the DML can become more efficient (Table VI). In addition, the MRE design allowed a unified structure to be applied in both one- and few-shot settings; that is, the intraclass variations, no matter of a single or few training samples, can be modeled the same way by the MRE. Along with the use of multimodal representations is how to set the mixing coefficients for the modes. Instead of the one-hot [16] or all-equal coefficients, we innovatively proposed an input-dependent mode-weighting strategy based on the self-attention mechanism [17]. As experimentally validated, this strategy improved performance upon the one-hot and all-equal coefficients with adaptable ones (Table V).

4) Application to Ultrasound Images: We noticed that our experiments so far were all conducted with volumetric CT and MRI images. In practice, ultrasound imaging also plays an important role in clinics and the segmentation of ultrasound images remains challenging. For example, Huang *et al.* proposed a novel method based on semantic classification of superpixels (SCS) for tumor segmentation in breast ultrasound (BUS) images, which is difficult due to the poor image quality [46]. To gain an understanding of how our MRE-Net would perform on such challenging tasks of vastly different characteristics (2D segmentation of low-quality images), we applied it to the BUS dataset collected in [46]. This dataset included 220 images for training and 100 for testing, with a relatively small, normalized size of 128×128 pixels. For few-shot learning, we used 1%, 5%, and 10% of all the training images to train a 2D version of our MRE-Net (replacing 3D operations in the backbone in Fig. 2(a) with 2D counterparts), and compared its performance with that of the FCN [24]. The FCN was trained with the task-optimized hyperparameters described in [46]. For evaluation, we adopted the same metrics as in [46], including: averaged radial error (ARE), true positive (TP), false positive (FP), and F1-score. ARE evaluates object boundary approximation error, whereas TP and FP are region-

TABLE XII
FEW-SHOT SEGMENTATION RESULTS ON THE BUS DATASET [46] (↑: HIGHER IS BETTER; ↓: LOWER IS BETTER). SCS REPT. AND EXPT. REPRESENT THE NUMBERS ORIGINALLY REPORTED IN [46] AND OBTAINED IN OUR EXPERIMENT, RESPECTIVELY. VALUES ARE MEAN (STD.). BOLD FACES DENOTE THE OVERALL BEST NUMBERS FOR DIFFERENT METRICS, AND ASTERISKS (*) INDICATE STATISTICALLY SIGNIFICANT DIFFERENCES FROM THE RESULTS OF OUR MRE-NET FOR EACH RATIO SETTING.

No. images (ratio)	F1-score↑ (%)	ARE↓ (%)	TP↑ (%)	FP↓ (%)
2 (1%)				
MRE-Net	89.69 (3.55)	15.95 (14.32)	87.52 (5.58)	**7.81** (9.49)
FCN [24], [46]	*79.89 (4.53)	*21.60 (11.45)	*73.07 (6.34)	*9.81 (8.13)
11 (5%)				
MRE-Net	90.97 (5.45)	13.87 (15.18)	93.45 (4.73)	12.73 (14.93)
FCN [24], [46]	*87.69 (4.82)	*19.59 (15.16)	*92.36 (4.51)	*19.00 (14.58)
22 (10%)				
MRE-Net	**91.88** (3.97)	10.49 (9.36)	**94.17** (4.69)	11.13 (10.43)
FCN [24], [46]	*88.24 (3.70)	*18.24 (10.70)	*90.97 (4.41)	*15.49 (9.75)
220 (100%)				
FCN [24], [46]	90.49 (2.80)	12.78 (8.64)	91.93 (4.75)	11.36 (7.71)
SCS rept. [46]	89.97 (4.05)	9.95 (4.42)	91.41 (5.04)	12.22 (9.42)
SCS expt. [46]	91.17 (2.88)	**9.81** (6.28)	93.75 (4.68)	12.06 (8.10)

based and F1-score is a balanced metric between TP and FP (for more details please refer to [46]).

The results are shown in Table XII. As we can see, our MRE-Net yielded significantly better results than those of the FCN for all metrics in all different ratio settings; meanwhile, our few-shot results appeared to be quite practical, and generally improved as the amount of training data increased. In addition, using only 10% of the training images, our MRE-Net achieved superior or comparable performances with the FCN and SCS approaches trained with all available training images. These results demonstrated the wide applicability of our MRE-Net to common clinic image data for generalized low-shot segmentation.

5) Future Work: We identified directions for future work based on current limitations. First, as mentioned in Section IV-H, the WMH is very challenging to segment for its irregular shape, size, or location. Specifically, we examined the ablation study results of WMH before and after adding the spatial information (rows (b) and (c) in Table II). Contrary to what was expected, WMH also benefited from incorporating the coordinates (Dice $+4.47\%$ and HD95 -4.02 mm), suggesting weak spatial preference. In future work we would like to design more effective strategies exclusively for improving segmentation of irregular abnormalities like the WMH. Second, we did not explore the impact of pretraining followed by transfer learning in this work. According to the literature, such a procedure often boosts performance of DNNs in case of limited training data [47]. Conforming to the specific scenario considered (*i.e.*, extremely scarce training data for the target task), we plan to pretrain the MRE-Net on data of a different task and transfer the learned knowledge onto the target task. Lastly, the proposed MRE-Net consumes considerable computational resources. Optimization of the resource consumption is worth future exploration.

REFERENCES

[1] C. Klifa *et al.*, "Magnetic resonance imaging for secondary assessment of breast density in a high-risk cohort," *Magn. Reson. Imaging*, vol. 28, no. 1, pp. 8–15, 2010.

[2] M. Silveira *et al.*, "Comparison of segmentation methods for melanoma diagnosis in dermoscopy images," *IEEE J. Sel. Topics Signal Process.*, vol. 3, no. 1, pp. 35–45, 2009.

[3] P. Jackson, N. Hardcastle, N. Dawe, T. Kron, M. S. Hofman, and R. J. Hicks, "Deep learning renal segmentation for fully automated radiation dose estimation in unsealed source therapy," *Front. Oncol.*, vol. 8, p. 215, 2018.

[4] G. Litjens *et al.*, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, 2017.

[5] F.-F. Li, R. Fergus, and P. Perona, "One-shot learning of object categories," *IEEE Pattern Anal. Mach. Intell.*, vol. 28, no. 4, pp. 594–611, 2006.

[6] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *NeurIPS*, 2017, pp. 4077–4087.

[7] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra, "Matching networks for one shot learning," in *NeurIPS*, 2016, pp. 3630–3638.

[8] K. Wang, J. H. Liew, Y. Zou, D. Zhou, and J. Feng, "PANet: Few-shot image semantic segmentation with prototype alignment," in *ICCV*, 2019, pp. 9197–9206.

[9] C. Zhang, G. Lin, F. Liu, R. Yao, and C. Shen, "CANet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning," in *CVPR*, 2019, pp. 5217–5226.

[10] A. Paul, Y.-X. Tang, and R. M. Summers, "Fast few-shot transfer learning for disease identification from chest x-ray images using autoencoder ensemble," in *Medical Imaging 2020: Computer-Aided Diagnosis*, H. K. Hahn and M. A. Mazurowski, Eds., vol. 11314, International Society for Optics and Photonics. SPIE, 2020, pp. 33 – 38.

[11] A. Zhao, G. Balakrishnan, F. Durand, J. V. Guttag, and A. V. Dalca, "Data augmentation using learned transformations for one-shot medical image segmentation," in *CVPR*, 2019, pp. 8543–8553.

[12] K. Chaitanya, N. Karani, C. F. Baumgartner, A. Becker, O. Donati, and E. Konukoglu, "Semi-supervised and task-driven data augmentation," in *Int. Conf. IPMI*. Springer, 2019, pp. 29–41.

[13] S. Wang *et al.*, "LT-Net: Label transfer by learning reversible voxel-wise correspondence for one-shot medical image segmentation," in *CVPR*, 2020, pp. 9162–9171.

[14] N. Tajbakhsh, L. Jeyaseelan, Q. Li, J. N. Chiang, Z. Wu, and X. Ding, "Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation," *Med. Image Anal.*, p. 101693, 2020.

[15] B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, and J. C. Gee, "A reproducible evaluation of ANTs similarity metric performance in brain image registration," *NeuroImage*, vol. 54, no. 3, pp. 2033–2044, 2011.

[16] L. Karlinsky *et al.*, "RepMet: Representative-based metric learning for classification and few-shot object detection," in *CVPR*, 2019, pp. 5197–5206.

[17] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *CVPR*, 2018, pp. 7132–7141.

[18] C. Luo, J. Zhan, X. Xue, L. Wang, R. Ren, and Q. Yang, "Cosine normalization: Using cosine similarity instead of dot product in neural networks," in *Int. Conf. Artif. Neural Netw.* Springer, 2018, pp. 382–391.

[19] A. Shrivastava, A. Gupta, and R. Girshick, "Training region-based object detectors with online hard example mining," in *CVPR*, 2016, pp. 761–769.

[20] F. Liao, M. Liang, Z. Li, X. Hu, and S. Song, "Evaluate the malignancy of pulmonary nodules using the 3D deep leaky noisy-or network," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 11, pp. 3484–3495, 2019.

[21] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," *arXiv preprint arXiv:1706.05587*, 2017.

[22] B. Landman, Z. Xu, J. Igelsias, M. Styner, T. Langerak, and A. Klein, "MICCAI multi-atlas labeling beyond the cranial vault – workshop and challenge," 2015.

[23] Z. Xu *et al.*, "Evaluation of six registration methods for the human abdomen on clinically acquired CT," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 8, pp. 1563–1572, 2016.

[24] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *CVPR*, 2015, pp. 3431–3440.

[25] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. MICCAI*. Springer, 2015, pp. 234–241.

[26] M. Luna and S. H. Park, "3D patchwise U-Net with transition layers for MR brain segmentation," in *MICCAI Brain-Lesion Workshop*. Springer, 2018, pp. 394–403.

[27] B. Kulis, "Metric learning: A survey," *Found. Trends Mach. Learn.*, vol. 5, no. 4, pp. 287–364, 2012.
[28] M. Kaya and H. Ş. Bilge, "Deep metric learning: A survey," *Symmetry*, vol. 11, no. 9, p. 1066, 2019.
[29] H. Qi, M. Brown, and D. G. Lowe, "Low-shot learning with imprinted weights," in *CVPR*, 2018, pp. 5822–5830.
[30] K. Q. Weinberger and L. K. Saul, "Distance metric learning for large margin nearest neighbor classification." *J. Mach. Learn. Res.*, vol. 10, no. 2, 2009.
[31] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
[32] I. Goodfellow *et al.*, "Generative adversarial nets," in *NeurIPS*, 2014, pp. 2672–2680.
[33] A. K. Mondal, J. Dolz, and C. Desrosiers, "Few-shot 3D multi-modal medical image segmentation using generative adversarial learning," *arXiv preprint arXiv:1810.12241*, 2018.
[34] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "VoxelMorph: A learning framework for deformable medical image registration," *IEEE Trans. Med. Imag.*, vol. 38, no. 8, pp. 1788–1800, 2019.
[35] A. G. Roy, S. Siddiqui, S. Pölsterl, N. Navab, and C. Wachinger, "'Squeeze & excite' guided few-shot segmentation of volumetric images," *Med. Image Anal.*, vol. 59, p. 101587, 2020.
[36] A. Shaban, S. Bansal, Z. Liu, I. Essa, and B. Boots, "One-shot learning for semantic segmentation," *arXiv preprint arXiv:1709.03410*, 2017.
[37] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *NeurIPS*, 2012, pp. 1097–1105.
[38] N. J. Tustison *et al.*, "N4ITK: improved N3 bias correction," *IEEE Trans. Med. Imag.*, vol. 29, no. 6, pp. 1310–1320, 2010.
[39] E. Gibson *et al.*, "Automatic multi-organ segmentation on abdominal CT with dense V-networks," *IEEE Trans. Med. Imag.*, vol. 37, no. 8, pp. 1822–1834, 2018.
[40] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
[41] X. Zhou, T. Ito, R. Takayama, S. Wang, T. Hara, and H. Fujita, "Three-dimensional CT image segmentation by combining 2D fully convolutional network with 3D majority voting," in *Deep Learning and Data Labeling for Medical Applications*. Springer, 2016, pp. 111–120.
[42] J. Jia, R. Wang, Z. An, Y. Guo, X. Ni, and T. Shi, "RDAD: A machine learning system to support phenotype-based rare disease diagnosis," *Front. Genet.*, vol. 9, p. 587, 2018.
[43] C. M. Molster *et al.*, "The evolution of public health genomics: exploring its past, present, and future," *Front. Public Health*, vol. 6, p. 247, 2018.
[44] X. Li, L. Yu, Y. Jin, C.-W. Fu, L. Xing, and P.-A. Heng, "Difficulty-aware meta-learning for rare disease diagnosis," in *Proc. Int. Conf. MICCAI*. Springer, 2020, pp. 357–366.
[45] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, no. Nov, pp. 2579–2605, 2008.
[46] Q. Huang, Y. Huang, Y. Luo, F. Yuan, and X. Li, "Segmentation of breast ultrasound image with semantic classification of superpixels," *Med. Image Anal.*, vol. 61, p. 101657, 2020.
[47] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A survey on deep transfer learning," in *Int. Conf. Artif. Neural Netw.* Springer, 2018, pp. 270–279.