

Modeling the Probabilistic Distribution of Unlabeled Data for One-shot Medical Image Segmentation

Yuhang Ding¹, Xin Yu², Yi Yang^{2*}

¹ Baidu Research, China

² ReLER, University of Technology Sydney, Australia
dyh.ustc.uts@gmail.com, xin.yu@uts.edu.au, yi.yang@uts.edu.au

Abstract

Existing image segmentation networks mainly leverage large-scale labeled datasets to attain high accuracy. However, labeling medical images is very expensive since it requires sophisticated expert knowledge. Thus, it is more desirable to employ only a few labeled data in pursuing high segmentation performance. In this paper, we develop a data augmentation method for one-shot brain magnetic resonance imaging (MRI) image segmentation which exploits only one labeled MRI image (named atlas) and a few unlabeled images. In particular, we propose to learn the probability distributions of deformations (including shapes and intensities) of different unlabeled MRI images with respect to the atlas via 3D variational autoencoders (VAEs). In this manner, our method is able to exploit the learned distributions of image deformations to generate new authentic brain MRI images, and the number of generated samples will be sufficient to train a deep segmentation network. Furthermore, we introduce a new standard segmentation benchmark to evaluate the generalization performance of a segmentation network through a cross-dataset setting (collected from different sources). Extensive experiments demonstrate that our method outperforms the state-of-the-art one-shot medical segmentation methods. Our code has been released at <https://github.com/dyh127/Modeling-the-Probabilistic-Distribution-of-Unlabeled-Data>.

Introduction

Medical image segmentation aims to partition medical images, such as magnetic resonance imaging (MRI) image, into different anatomic regions. It plays an important role in many medical analysis applications, such as computer-assisted diagnosis and treatment planning. In recent years, benefiting from deep convolution neural networks (CNNs), fully supervised medical image segmentation methods (Zhou et al. 2018; Chen et al. 2019) have been extensively studied and achieved promising progress. However, labeling anatomic regions for large-scale 3D images requires a huge amount of time and expert knowledge. Hence, obtaining sufficient labelled data often becomes the bottleneck of fully supervised segmentation methods.

*Corresponding author

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

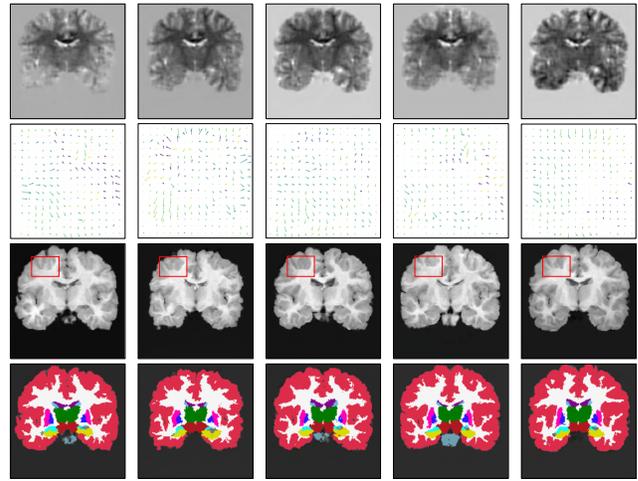


Figure 1: Illustration of our generated diverse deformations. From top to bottom: intensity offsets, shape deformations, synthesized images using the corresponding deformations and segmentation labels. Red frames highlight variations.

One-shot medical image segmentation, also called single atlas-based segmentation, has been proposed to reduce the demand for copious labeled data. Hand-crafted data augmentations (Ronneberger, Fischer, and Brox 2015; Milletari, Navab, and Ahmadi 2016; Roth et al. 2015; Pereira et al. 2016), such as random elastic deformations, generate new labeled images to improve segmentation performance. However, those methods often generate non-realistic images since they do not take the distribution of real images into account. Thus, their learned segmentation networks may not generalize well on real data. Recently, deep learning based data augmentation methods (Zhao et al. 2019; Xu and Niethammer 2019; Chaitanya et al. 2019; Wang et al. 2020; Zhu et al. 2020) have been exploited. Those methods often leverage image registration to obtain profile and intensity differences between the atlas and other MR images, and then combine the profiles and intensities to generate new images for segmentation.

Considering the domain gap and insufficient variations of synthesized data by previous methods, we aim to develop a novel medical image (*i.e.*, MRI) augmentation method to

address one-shot medical image segmentation tasks. To this end, we propose probabilistic data augmentation approach to generate sufficient training images while ensuring them to follow the distribution of real MRI images in terms of brain shapes and MRI intensities, as shown in Fig 1. Thus, our segmentation network trained on our synthesized data will be robustly adapted to real MRI images.

In this work, we firstly employ image registration to obtain the shape deformations and intensity changes between an unlabeled MRI image and the atlas. However, since registration errors might occur in the registration procedure, directly classifying the registered images will lead to erroneous segmentation results. The prior art (Zhao et al. 2019) combines the registered deformation fields and intensity changes to produce new images (with segmentation masks) and exploits them to train a segmentation network, thus mitigating registration errors. However, Zhao et al. (2019) cannot provide new deformation fields and intensity changes. Therefore, the variety of generated images is still limited.

In contrast to prior works, we propose to exploit two variational autoencoders (VAEs) to capture the probabilistic distributions of deformation fields and intensity offsets with respect to the atlas. After that, our VAEs are employed to generate various profile deformations and intensity changes. The generative deformation fields and intensity variations are used to synthesize new MRI images. In this manner, our synthesized training data is not only abundant and diverse but also authentic to the real MRIs. Hence, using our augmented data, we improve the performance of our segmentation network significantly and achieve superior performance compared with the state-of-the-art.

Since different MRI machines (*i.e.*, imaging sources) may lead to different characteristics in MRI images, such as intensity changes and signal-to-noise ratio, we also conduct experiments on unseen MRI sources to evaluate the robustness of our method. Thus, we propose a more challenging benchmark with an additional unseen test set. Benefiting from our generated diverse training data, our segmentation network also performs better than the state-of-the-art on unseen MRI sources, thus demonstrating the superiority of our presented probabilistic augmentation method.

Overall, our contributions are threefold:

- We propose a novel probabilistic data augmentation method based on VAEs to generate diverse and realistic training images for the downstream segmentation task.
- We propose a new challenging segmentation benchmark to evaluate the performance of our proposed method and competing methods. It contains 3D brain MRI images from different sources. Thus, we can also test the generalization ability of the methods on unseen MRI sources.
- Taking advantage of our generated images, our method outperforms the state-of-the-art one-shot segmentation algorithms on both seen and unseen image sources.

Related Work

Atlas-based Segmentation

Atlas-based segmentation methods (Klein et al. 2005; Heckemann et al. 2006) aim to segment target images by explor-

ing knowledge from single or multiple labeled atlas images as well as a few unlabeled training images. Because only a few labeled images are required, atlas-based segmentation methods are more desirable but challenging compared to fully supervised methods.

Single atlas-based segmentation methods (Wang et al. 2020; Dinsdale, Jenkinson, and Namburete 2019) leverage a registration model to learn shape deformations by aligning an atlas to target images, and then transfer the atlas label to the unlabeled ones as target labels. On the other hand, multi-atlas-based segmentation methods mainly focus on atlas selection (Yang et al. 2018) and label fusion (Yang et al. 2018; Ding, Han, and Niethammer 2019). Since our work belongs to the category of single atlas-based methods, we mainly review methods of this category as follows:

Wang et al. (2020) introduce a forward-backward consistency scheme into a registration network to obtain segmentation labels for unlabeled images. However, registration networks may suffer misalignment errors, thus leading to inferior segmentation results. Instead of directly transferring segmentation labels to a target image, some works have been proposed to warp the atlas image and its label to generate new images. Then, the generated data are used to train a segmentation network. For example, Zhao et al. (2019) leverage image registration to learn shape and intensity deformations between target images and the atlas. Then, new images synthesized by the learned deformations are exploited to train their segmentation network. The works (Xu and Niethammer 2019; Zhu et al. 2020) jointly learn image segmentation and registration simultaneously. Since the deformations learned from the unlabeled data are deterministic and only a few, the diversity of generated images is limited.

In contrast, our proposed method is able to arbitrarily generate various shape and intensity deformations that even do not exhibit in any provided images. Moreover, our generated deformations are sampled from the distribution of the deformations between unlabeled images and the atlas via two VAEs. Hence, our synthesized MRI images are not only abundant but also authentic to real ones, thus facilitating the training of our segmentation network.

Medical Image Data Augmentation

Data augmentation is one of the most effective techniques for reducing over-fitting and improving the capability of networks. In medical image analyses, a large number of data augmentation methods have been introduced due to the absence of large-scale labeled training data.

Traditional hand-crafted data augmentation methods are designed to deform medical images in terms of qualities, appearance or shapes. Christ et al. (2016) augment CT images with Gaussian noise while Sirinukunwattana et al. (2017) apply Gaussian blur to augment images for gland segmentation. Dong et al. (2017) enrich the training set by randomly enhancing the brightness of MRI images. The works (Ronneberger, Fischer, and Brox 2015; Çiçek et al. 2016) exploit random elastic deformations to generate annotated images.

Generative adversarial networks (GANs) have been employed to synthesize new images (Mahapatra et al. 2018; Jin et al. 2018; Fu et al. 2018). Fu et al. (2018); Cao et al.

(2020) adopt CycleGANs (Zhu et al. 2017) to generate 3D microscopy images while Mahapatra et al. (2018) use conditional GANs to synthesize realistic chest X-ray images. Chaitanya et al. (2019) exploit conditional GANs to generate deformations of brain MRI images. However, when only one labeled example and only a few unlabeled ones (less than 100 MRI images) are available, GAN based augmentation methods, in this case, might suffer mode collapse, such as outputting all zeros for shape and intensity deformations.

Proposed Method

In this work, we leverage an image registration network and two VAEs to generate diverse and authentic brain MRI training samples. The generative samples are then employed to improve our segmentation network. Here, we introduce the procedure of image registration as well as modeling the probabilistic distributions of those deformations via our shape and intensity 3D VAEs, respectively.

After obtaining the models of the deformations, we randomly sample from the distributions of the deformations and then construct new MRI images with the atlas image. The newly synthesized MRI images with their labels will be used to train our segmentation network.

Learning Deformations from Image Registration

Image registration (Miao et al. 2018; Zitova and Flusser 2003) aims to align an image to a template one, called atlas, by learning shape deformations between them. Most existing registration-based segmentation methods (Wang et al. 2020; Xu and Niethammer 2019; Zhu et al. 2020) only consider the structure differences between two images. However, due to different patients, scan machines and operations, image intensities also vary. Therefore, we model both shape and intensity deformations.

First, as shown in Fig. 2, we leverage a Unet-based (Ronneberger, Fischer, and Brox 2015) registration network (named shape registration network) to learn 3D shape deformations. Denote an atlas image and its segmentation mask as (x^a, y^a) and N unlabeled images as $\{x_1^u, x_2^u, \dots, x_N^u\}$. Taking the atlas image x^a and an unlabeled training image x_i^u as the input, the registration network is trained to propagate the atlas image x^a to an unlabeled image x_i^u by estimating a shape deformation S_i . In other words, S_i is optimized to warp x^a to x_i^u : $x_i^u \leftarrow x^a \circ S_i$, where \circ represents a warping operation implemented by a differentiable bilinear interpolation-based spatial transformer layer (Balakrishnan et al. 2018). Following the work (Balakrishnan et al. 2019), we employ a local cross-correlation (CC) loss \mathcal{L}_{CC} and a deformation smoothness regularization \mathcal{L}_S^{reg} to train our shape registration network in an unsupervised manner and its objective \mathcal{L}_{srn} is formulated as:

$$\begin{aligned} \mathcal{L}_{CC} &= \sum_i \sum_{p \in \Omega} \frac{g(x_i^u, [x^a \circ S_i], p)^2}{g(x_i^u, x_i^u, p)g([x^a \circ S_i], [x^a \circ S_i], p)}, \\ \mathcal{L}_S^{reg} &= \sum_i \|\nabla S_i\|_2, \\ \mathcal{L}_{srn} &= -\mathcal{L}_{CC} + \mathcal{L}_S^{reg}, \end{aligned} \quad (1)$$

where $g(a, b, p)$ denotes the correlation between local patches a and b on voxel p : $g(a, b, p) = \sum_{p_j} (a(p_j) - \bar{a}(p))(b(p_j) - \bar{b}(p))$, and $\bar{a}(p)$ indicates the mean of local patch intensities on p : $\bar{a}(p) = \frac{1}{\|p\|_1} \sum_{p_j} a(p_j)$. p represents a n^3 cube in a 3D image Ω and p_j denotes the pixels in the cube. We set n to 9 similar to prior methods (Balakrishnan et al. 2019). \mathcal{L}_{CC} encourages the structure similarities between two images regardless of the intensity variations while \mathcal{L}_S^{reg} aims to constrain shape deformations to be smooth. ∇S_i denotes the spatial gradients of the shape variations.

Similar to learning shape deformations, we also use a Unet-based network, called intensity alignment network, to align 3D intensity deformations. As visible in Fig. 2, the network takes the atlas image x^a and the inverse-warped image \hat{x}_i^u as input to measure the intensity deformations I_i . \hat{x}_i^u is generated by aligning x_i^u to x^a , and thus \hat{x}_i^u and x^a share similar profile structure. Similar to (Zhao et al. 2019), we exploit a pixel-wise reconstruction loss \mathcal{L}_{sim} between x^a and x_i^u and an intensity smoothness regularization \mathcal{L}_I^{reg} to train our intensity alignment network. The objective function \mathcal{L}_{irn} is expressed as:

$$\begin{aligned} \mathcal{L}_{sim} &= \sum_i \|(x^a + I_i) \circ S_i - x_i^u\|_2, \\ \mathcal{L}_I^{reg} &= \sum_i \sum_{q_j} (1 - c^a(p_j)) |\nabla I_i(p_j)|, \\ \mathcal{L}_{irn} &= \mathcal{L}_{sim} + \lambda \mathcal{L}_I^{reg}. \end{aligned} \quad (2)$$

Here, \mathcal{L}_I^{reg} is designed to prevent dramatic changes of the I_i in the same brain area. $\nabla I_i(p_j)$ denotes the gradients of I_i at p_j . c^a denotes the mask of contours across different areas. λ is a trade-off weight and set to 0.02, following the work (Zhao et al. 2019).

Diverse Image Generation via VAEs

After image registration, we obtain N shape deformations and N intensity changes from the atlas and N unlabeled images. In the work (Zhao et al. 2019), these variations are directly combined to generate new labeled training images for segmentation. However, only N kinds of shape and intensity transformations are involved during training, and the diversity of the samples is not rich enough to train an accurate segmentation network. Chaitanya et al. (2019) employ GANs to generate new deformations but their method requires a large number of unlabeled data to train GANs. However, we only have less than 100 unlabeled images and their method will suffer mode collapse and is not applicable in our case.

Different from previous methods, we adopt a 3D shape VAE and a 3D intensity VAE to learn the probabilistic distributions of the variations with respect to the atlas separately, since VAE does not suffer mode collapse. Furthermore, inspired by beta-VAE (Higgins et al. 2017; Burgess et al. 2018), we reduce impacts of the Kullback-Leibler (KL) divergence in a conventional VAE to increase the diversity of generated samples. Doing so is also driven by the insufficiency of the training samples. After training, we sample

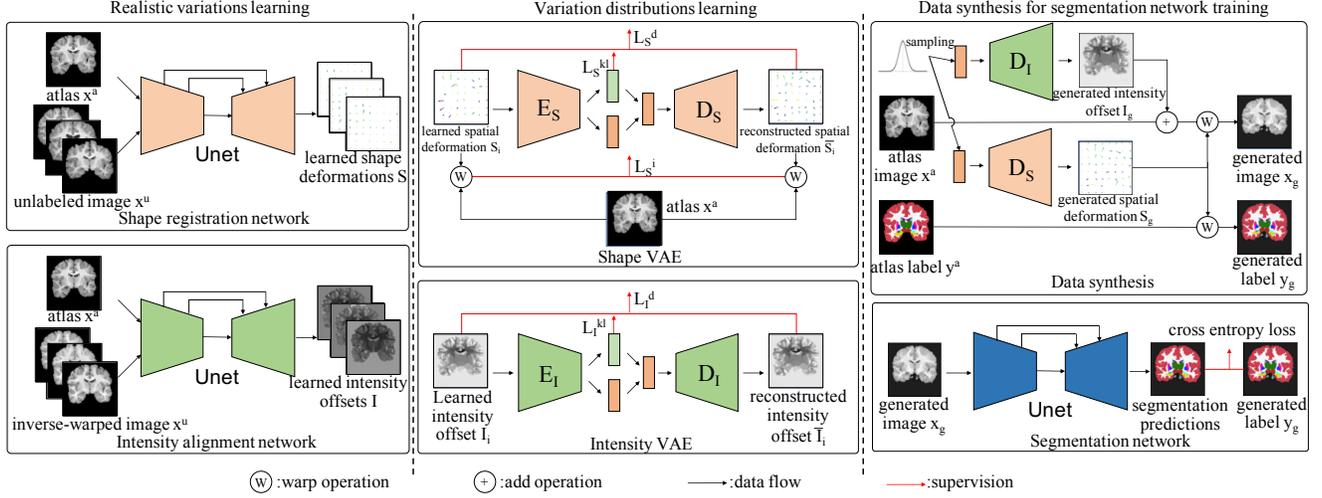


Figure 2: The framework of our proposed method: (i) image deformations are obtained by two Unet-based registration networks; (ii) our shape and intensity VAEs are proposed to learn the variation distributions and generate new deformations; (iii) new training samples are synthesized by applying the generated deformations to the atlas image and our segmentation network is trained on these samples.

deformations from our shape and intensity VAEs, and then generate a large number of various training images.

As illustrated in Fig. 2, our shape VAE first uses an encoder to project an input shape deformation into a latent vector $z = \mathbf{E}_S(S_i)$ and then decodes z to the image domain, *i.e.*, a reconstructed shape deformation $\bar{S}_i = \mathbf{D}_S(z)$. During training, three objectives, including KL divergence \mathcal{L}_S^{kl} and pixel-wise reconstruction losses on the deformations \mathcal{L}_S^d and image intensities \mathcal{L}_S^i , are employed to train our shape VAE, written as:

$$\begin{aligned}
 \mathcal{L}_S^{kl} &= \sum_i D_{kl}(q(z|S_i)||p(z)), \\
 \mathcal{L}_S^d &= \sum_i \|S_i - \bar{S}_i\|_2, \\
 \mathcal{L}_S^i &= \sum_i \|(x^a \circ S_i) - (x^a \circ \bar{S}_i)\|_2, \\
 \mathcal{L}_S &= (\mathcal{L}_S^d + \mathcal{L}_S^i) + \beta \mathcal{L}_S^{kl},
 \end{aligned} \tag{3}$$

where \mathcal{L}_S^{kl} forces the distribution of latent vector z to be a standard normal distribution, (*i.e.*, $z \sim \mathcal{N}(0, 1)$), $q(z|\cdot)$ denotes the posterior distribution, $p(z)$ denotes the Gaussian prior distribution modeled by a standard normal distribution, and β is a hyper-parameter controlling rigidity of the distributions of the latent variable z and the quality of reconstruction. Here, we not only compare the decoded shape deformations with the input ones but also measure the differences between the warped images by the input shape deformations and reconstructed ones.

Smaller β indicates less attention is paid to the KL divergence loss during training and will result in a larger KL divergence between the posterior and prior distributions. As suggested by Burgess et al. (2018), larger KL divergence allows a latent vector to reside in a large space. In other words, smaller β allows our VAE to preserve variations of input images especially when the training samples are scarce. There-

fore, using a small β is more preferable when the number of training samples is limited. Moreover, since the latent space has been enlarged, more variations can be generated from this latent vector space via our decoder in the testing phase. Therefore, we set β to a small value (*i.e.*, 0.1) for all the experiments.

It is worth noting that we employ both \mathcal{L}_S^d and \mathcal{L}_S^i as the reconstruction loss for our shape VAE instead of only reconstructing network inputs by \mathcal{L}_S^d as in the original VAE. When \mathcal{L}_S^d is only employed, image structure information is neglected. In particular, shape deformations should pay attention to the consistency of image contour movements. However, \mathcal{L}_S^d treats the movement of each pixel individually and thus may not perform consistent movements along the contour regions. On the contrary, the reconstruction loss \mathcal{L}_S^i is sensitive to the movements of image contours because image intensities around contours change dramatically. In other words, small reconstruction errors in the deformations of the contours will lead to large intensity differences between two warped images. On the other hand, since \mathcal{L}_S^i only measures intensity similarities, it may not preserve boundary information when two areas have similar intensities. Therefore, we leverage both \mathcal{L}_S^i and \mathcal{L}_S^d as the reconstruction loss in learning our shape VAE.

Similar to our shape VAE, we employ a VAE to model the distribution of the intensity variations with respect to the atlas. Here, we adopt the standard KL divergence loss and a pixel-wise reconstruction loss to train our intensity deformation VAE, expressed as:

$$\begin{aligned}
 \mathcal{L}_I^{kl} &= \sum_i D_{kl}(q(z|I_i)||p(z)), \\
 \mathcal{L}_I^d &= \sum_i \|I_i - \bar{I}_i\|_2, \\
 \mathcal{L}_I &= \mathcal{L}_I^d + \beta \mathcal{L}_I^{kl},
 \end{aligned} \tag{4}$$

where \bar{I}_i is the intensity deformation reconstructed from I_i .

After modeling the deformation distributions, our shape and intensity VAEs are exploited to generate diverse variations by random sampling. Specifically, in the process of the generation, the decoders \mathbf{D}_S and \mathbf{D}_I take random latent vectors sampled from a Gaussian distribution $\mathcal{N}(0, \sigma)$ as input and output various shape deformations S_g and intensity changes I_g , respectively. Then, our synthesized labeled training images are constructed as:

$$x_g = (x^a + I_g) \circ S_g, \quad y_g = y^a \circ S_g, \quad (5)$$

where x_g and y_g represent the synthesized images and their corresponding segmentation masks. Note that, different from MRI images, segmentation masks are warped by a nearest-neighbor interpolation-based 3D spatial transformer layer (Balakrishnan et al. 2018).

Segmentation Network

Once augmented training samples are obtained, we can train our segmentation network on those samples. In order to conduct fair comparisons to the state-of-the-art (Zhao et al. 2019), we employ the same 2D Unet with a five-layer encoder and a five-layer decoder to segment each slice of 3D images individually. In the encoder and decoder, we use 3x3 2D convolutional operations followed by LeakyReLU layers. 2x2 Max-pooling layers are used to reduce the feature resolution while upsampling layers are used to increase resolution by a factor of 2.

In each training iteration, we randomly sample slices from 3D images to construct a batch. The standard cross-entropy loss is employed as follows:

$$\mathcal{L}_{CE} = - \sum_{i=1}^W \sum_{j=1}^H \frac{1}{H \cdot W} \log \frac{\exp(y_p[i, j, y_g(i, j)])}{\sum_{k=1}^K \exp(y_p[i, j, k])}, \quad (6)$$

where y_p is the predicted mask from our segmentation network g (*i.e.*, $y_p = g(x_g; \theta)$) and θ denotes the parameters of the segmentation network. W and H denote the width and height of a 2D slice, respectively. K indicates the number of anatomical components in an MRI image. Similar to the training process, every 3D image is split into 2D slices and segmented in a slice-wise fashion in the testing phase.

Although we incorporate two VAEs to generate labeled data, they are only used in the training phase. During testing, only our segmentation network is exploited. Therefore, our method does not increase the network parameters and FLOPs during inference and thus can be deployed as easily as previous works.

Implementation Details

We adopt the same network architecture for our shape and intensity VAEs, and the VAEs are 3D VAEs since deformations should be consistent in 3D space. More details of the network architecture are described in the supplementary material. In the 3D VAE networks, group normalization (Wu and He 2020) is employed. For the activation function, we use LeakyReLU and ReLU for the encoder and the decoder, respectively. The dimension of the latent vector is set to 512.

During training, Adam (Kingma and Ba 2015) optimizer is used to train our VAEs, where β_1 and β_2 are set to 0.5 and 0.999, respectively. The batch size is set to 1 due to the GPU memory limit. The learning rate is fixed to $1e^{-4}$ for the whole 40k training iterations. The hyper-parameter β in both two VAEs is set to 0.1. In generating deformations, the shape VAE and the intensity VAE take latent vectors sampled from $\mathcal{N}(0, 10)$ as input in order to achieve more diverse data.

For other networks (*i.e.*, shape registration, intensity alignment and segmentation networks), a default Adam with $1e^{-4}$ learning rate is employed. For the shape registration and intensity alignment networks, the batch size is set to 1 and the networks are trained for 500 epochs. For the segmentation network, the batch size is set to 16 and the network is trained for 40k iterations. Our method is trained and tested on an Nvidia Tesla V100 GPU and achieves similar results on Keras with a TensorFlow backend and PaddlePaddle.

Note that, in training the 3D VAEs and segmentation networks, images are generated on-the-fly, and thus we train these networks in terms of iterations. In training registration and alignment networks, only 82 MRI images will be transformed to the atlas, and thus we train the networks in terms of epochs.

Experiments

In this section, we first compare our proposed method with state-of-the-art one-shot based methods and then analyse the contributions of each component in our method. For fair comparisons, we conduct our experiments on the same dataset as previous works (Balakrishnan et al. 2019; Zhao et al. 2019; Wang et al. 2020). Moreover, we propose a more challenging MRI benchmark to evaluate the generalization performance of state-of-the-art one-shot based methods on unseen MRI data.

Dataset and Evaluation Metric

Dataset: CANDI dataset (Kennedy et al. 2011) consists 103 T1-weighted brain MRI images from 57 males and 46 females. In this dataset, four types of diagnostic groups are considered including healthy controls, schizophrenia spectrum, bipolar disorder with psychosis, and bipolar disorder without psychosis. In the experiments, we use the same train and test splits as in (Wang et al. 2020). To be specific, 20, 82 and 1 images are employed as the test set, unlabeled training set and atlas, respectively. Following the work (Wang et al. 2020), we crop a $160 \times 160 \times 128$ volume from the center of an original MRI image. For segmentation, similar to (Wang et al. 2020), we consider 28 primary brain anatomical areas.

Evaluation Metric: Dice coefficient (Dice 1945) is used to measure the segmentation performance, written by:

$$\text{Dice}(M_{y_p}^k, M_{y_{gth}}^k) = 2 \cdot \frac{M_{y_p}^k \cap M_{y_{gth}}^k}{|M_{y_p}^k| + |M_{y_{gth}}^k|}, \quad (7)$$

where $M_{y_p}^k$ and $M_{y_{gth}}^k$ denote segmentation masks of the anatomical region k with predicted labels y_p and its corresponding ground-truth y_{gth} .

Larger Dice scores indicate more overlaps between predictions and ground-truth labels, and thus represent better

| Method | Shape | Intensity | VAE | Mean(std) | Min | Max |
|-------------------------------------|-------|-----------|-----|--------------------|-------------|-------------|
| Registration based (VoxelMorph) | | | | 76.0(9.7) | 61.7 | 80.1 |
| Segmentation with data augmentation | ✓ | | | 81.7(5.6) | 65.4 | 87.4 |
| | ✓ | | ✓ | 83.5(4.2) | 71.1 | 87.8 |
| | ✓ | ✓ | | 84.2(1.7) | 79.7 | 86.5 |
| | ✓ | ✓ | ✓ | 85.1 (1.9) | 80.2 | 87.8 |

Table 1: Ablation study on different types of data augmentation. Shape and Intensity denote that the shape and intensity deformations are from registration. VAE indicates that the deformations are generated from our VAEs.

| Method | Mean(std) | Min | Max |
|--------------------------------|------------------|-------------|-------------|
| Supervised learning | 88.3(1.7) | 83.5 | 90.3 |
| VoxelMorph (Balakrishnan 2019) | 76.0(9.7) | 61.7 | 80.1 |
| DataAug (Zhao 2019) | 80.4(4.3) | 73.8 | 84.0 |
| LT-Net (Wang 2020) | 82.3(2.5) | 75.6 | 84.2 |
| Ours | 85.1(1.9) | 80.2 | 87.8 |

Table 2: Comparison with the state-of-the-art on CANDI. The performance of fully supervised segmentation is also provided as an upper bound. The Dice score (%) is employed, and Mean(std) denotes the standard deviations. Min and Max denote the minimum and maximum Dice scores in the test set, respectively.

segmentation performance. To better demonstrate the performance of each method, we report not only a mean Dice score but also its corresponding standard variance, minimum Dice score and maximum Dice score on the test set.

Comparison with State-of-the-art Methods

We mainly compare two state-of-the-art one-shot atlas based method, namely DataAug (Zhao et al. 2019) and TL-Net (Wang et al. 2020). In addition, one unsupervised registration method *i.e.*, VoxelMorph (Balakrishnan et al. 2018) is applied to one-shot medical image segmentation for comparison. VoxelMorph and TL-Net leverage a registration network to align the input MRI images to the atlas and then transfer the segmentation mask of the atlas to the input images as the segmentation results. DataAug employs image registration to achieve shape and intensity transformation, and then augment the atlas image with the attained transformation to train a segmentation network. Note that these state-of-the-art methods do not generate new deformations while our method does.

As seen in Table 2, we demonstrate the segmentation performance of our method is superior to that of the state-of-the-art. As indicated in Table 2, our method achieves superior segmentation performance compared to the state-of-the-art. In particular, our method improves the segmentation performance by 2.8% on the Dice score in comparison to the second best method LT-Net (Wang et al. 2020). Moreover, our method also obtains the smallest variance, demonstrating that our method is more robust.

Ablation Study

To demonstrate the effectiveness of our VAEs, we compare four different types of data augmentation in Table 1. As sim-

| Method | Mean(std) | Min | Max |
|-------------------------------------|---------------------|-------------|-------------|
| \mathcal{L}_S^d | 81.3 (2.8) | 74.4 | 85.0 |
| \mathcal{L}_S^i | 82.3(6.2) | 63.9 | 87.7 |
| $\mathcal{L}_S^d + \mathcal{L}_S^i$ | 83.5 (4.2) | 71.1 | 87.8 |

Table 3: Ablation study on different reconstruction losses in the shape VAE.

ply applying intensity offsets to the atlas does not change the segmentation mask, synthesized images will have the same segmentation labels, thus leading to a trivial segmentation solution.

Effectiveness of our VAEs As indicated in Table 1, compared with direct registration, data augmentation based segmentation methods achieve better segmentation accuracy. Note that all the augmentation methods learn the shape deformations similar to VoxelMorph. Compared with the data augmentation methods using deformations from image registration, our VAEs can generate richer data for training a segmentation network, thus leading to better performance. Moreover, we observe that intensity deformations make great contributions to segmentation performance and various intensity changes facilitate the generalization of our segmentation network. In Table 1, we also notice that our network employing registered shape and intensity deformations achieves better performance than DataAug. This is because DataAug pre-trains a segmentation network with an l2 loss and does not employ the atlas in training the segmentation network. Thus, using the atlas for training segmentation networks is important.

Effectiveness of the Combined Reconstruction Loss To demonstrate the effectiveness of our combined reconstruction loss *i.e.*, $\mathcal{L}_S^d + \mathcal{L}_S^i$, we train the shape VAEs with \mathcal{L}_S^d , \mathcal{L}_S^i and $\mathcal{L}_S^d + \mathcal{L}_S^i$, respectively, and then apply them to augment data. To avoid the influence of the intensity augmentation, we do not use intensity augmentation and the segmentation results are reported in Table 3. As indicated by Table 3, our combined reconstruction loss is more suitable for the shape deformation learning and generation.

Hyper-parameter β in Eq. (3) and Eq. (4), and σ for sampling latent codes As aforementioned, a small β introduces more diversity into the generated deformations, thus improving the segmentation performance. Figure 3 manifests that using a small β , we achieve better segmentation accuracy. Thus, in all the experiments, β is set to 0.1. Fur-

| Method | Seen | | | Unseen | | |
|---------------------|------------------|-------------|-------------|------------------|-------------|-------------|
| | Mean(std) | Min | Max | Mean(std) | Min | Max |
| Supervised learning | 87.6(2.7) | 79.3 | 91.1 | 85.9(1.7) | 81.3 | 87.5 |
| VoxelMorph | 70.3(11.6) | 33.1 | 82.5 | 62.9(13.2) | 32.3 | 79.6 |
| DataAug | 69.6(9.02) | 39.7 | 80.4 | 64.3(9.9) | 35.0 | 77.2 |
| Ours | 76.7(7.4) | 53.2 | 86.5 | 74.8(6.6) | 54.1 | 83.3 |

Table 4: Comparison with the state-of-the-art methods on our newly proposed ABIDE benchmark.

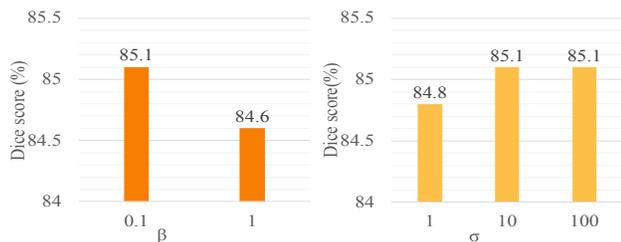


Figure 3: Ablation study on hyper-parameter β and σ . β controls the weight of the KL divergence and σ is the standard deviation of a prior Gaussian distribution $\mathcal{N}(0, \sigma)$ in VAEs.

thermore, as illustrated in Fig. 3, the segmentation performance degrades when the standard deviation σ for sampling latent codes is set to 1. This is because we employ a small β to enforce the KL divergence during training and the latent vector space would deviate from the standard normal distribution. Thus, we use a larger σ to sample latent codes. Figure 3 shows the segmentation accuracy is similar when σ is set to 10 and 100. Thus, σ is set to 10 for all the experiments.

Our Proposed ABIDE Benchmark

Since the MRI images in CANDI are collected from only one source, the variances (including shape and intensity) mainly come from different individuals. However, different MRI machines and operations may also lead to variations. Therefore, to validate the robustness of our method, we propose a new standard segmentation benchmark, called ABIDE benchmark, as visible in Fig. 4.

We sample T1-weighted MRI images from the autism brain imaging data exchange (ABIDE) database (Di Martino et al. 2014), which are collected from 17 international sites. We sample 190 images from ten imaging sources and split them into 100, 30, 60 volumes for training, validation and testing, respectively. These testing images form a *seen* test set. As suggested by Balakrishnan et al. (2019), the most similar image to the average volume is selected as the atlas. We also sample 60 images from the rest imaging sources as an *unseen* test set. All the volumes are resampled into a $256 \times 256 \times 256$ with 1mm isotropic voxels and then cropped to $160 \times 160 \times 192$. 28 anatomical regions are annotated by FreeSurfer (Fischl 2012).

As our benchmark contains images from multiple sites and includes an unseen test setting, it is more challenging and is also able to evaluate the robustness of a method.

We compare our method with VoxelMorph(Balakrishnan

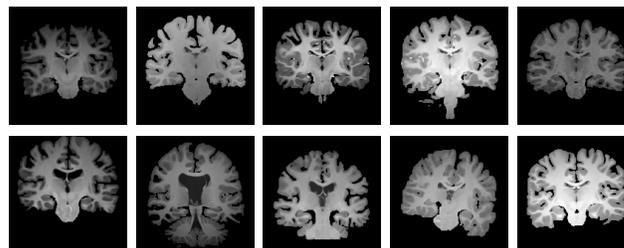


Figure 4: Illustration of significant variances in our ABIDE benchmark. The 96-th slices of ten 3D MRI images are shown. (Top row: images from seen datasets; Bottom row: images from unseen datasets.) More images are shown in supplementary materials.

et al. 2019) and DataAug(Zhao et al. 2019) in Table 4. The performance of the segmentation network trained with full supervision is also reported. Compared with the other two methods, we achieve superior performance on the seen and unseen datasets, demonstrating the effectiveness of our data augmentation method. In addition, our performance only degrades 1.9% on the unseen test dataset while the performance of the competing methods decreases more than 5%. This demonstrates that our method achieves a better generalization ability with the help of our generated various deformations.

Conclusion

In this paper, we propose a 3D VAE based data augmentation scheme to generate realistic and diverse training samples for one-shot medical image segmentation. We present a shape deformation VAE and an intensity deformation VAE to learn the distributions of the deformations of unlabeled real images with respect to an atlas one. With the help of our learned VAEs, we can generate various deformations rather than solely combining existing deformations from unlabeled data, thus significantly enriching training data for segmentation. To evaluate the segmentation performance on unseen imaging sources, we collect and annotate MRI data from different sources and construct a new benchmark. It provides us a standard public testbed for one-shot or few-shot based medical image segmentation methods. Extensive experiments demonstrate that our method outperforms the state-of-the-art on seen and unseen datasets.

References

Balakrishnan, G.; Zhao, A.; Sabuncu, M. R.; Guttag, J. V.; and Dalca, A. V. 2018. An Unsupervised Learning Model

- for Deformable Medical Image Registration. In *CVPR*, 9252–9260.
- Balakrishnan, G.; Zhao, A.; Sabuncu, M. R.; Guttag, J. V.; and Dalca, A. V. 2019. VoxelMorph: A Learning Framework for Deformable Medical Image Registration. *IEEE Trans. Medical Imaging* 38(8): 1788–1800.
- Burgess, C. P.; Higgins, I.; Pal, A.; Matthey, L.; Watters, N.; Desjardins, G.; and Lerchner, A. 2018. Understanding disentangling in β -VAE. *Arxiv*.
- Cao, B.; Zhang, H.; Wang, N.; Gao, X.; and Shen, D. 2020. Auto-GAN: Self-Supervised Collaborative Learning for Medical Image Synthesis. In *AAAI*, 10486–10493.
- Chaitanya, K.; Karani, N.; Baumgartner, C. F.; Becker, A. S.; Donati, O.; and Konukoglu, E. 2019. Semi-supervised and Task-Driven Data Augmentation. In *IPMI*, 29–41.
- Chen, X.; Williams, B. M.; Vallabhaneni, S. R.; Czanner, G.; Williams, R.; and Zheng, Y. 2019. Learning Active Contour Models for Medical Image Segmentation. In *CVPR*, 11632–11640.
- Christ, P. F.; Elshaer, M. E. A.; Ettliger, F.; Tatavarty, S.; Bickel, M.; Bilic, P.; Rempfler, M.; Armbruster, M.; Hofmann, F.; D’Anastasi, M.; Sommer, W. H.; Ahmadi, S.; and Menze, B. H. 2016. Automatic Liver and Lesion Segmentation in CT Using Cascaded Fully Convolutional Neural Networks and 3D Conditional Random Fields. In *MICCAI*, 415–423.
- Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S. S.; Brox, T.; and Ronneberger, O. 2016. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In *MICCAI*, 424–432.
- Di Martino, A.; Yan, C.-G.; Li, Q.; Denio, E.; Castellanos, F. X.; Alaerts, K.; Anderson, J. S.; Assaf, M.; Bookheimer, S. Y.; Dapretto, M.; et al. 2014. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular psychiatry* 19(6): 659–667.
- Dice, L. R. 1945. Measures of the amount of ecologic association between species. *Ecology* 26(3): 297–302.
- Ding, Z.; Han, X.; and Niethammer, M. 2019. VoteNet: A Deep Learning Label Fusion Method for Multi-atlas Segmentation. In *MICCAI*, 202–210.
- Dinsdale, N. K.; Jenkinson, M.; and Namburete, A. I. 2019. Spatial warping network for 3d segmentation of the hippocampus in mr images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 284–291.
- Dong, H.; Yang, G.; Liu, F.; Mo, Y.; and Guo, Y. 2017. Automatic Brain Tumor Detection and Segmentation Using U-Net Based Fully Convolutional Networks. In *MIUA*, 506–517.
- Fischl, B. 2012. FreeSurfer. *Neuroimage* 62(2): 774–781.
- Fu, C.; Lee, S.; Ho, D. J.; Han, S.; Salama, P.; Dunn, K. W.; and Delp, E. J. 2018. Three Dimensional Fluorescence Microscopy Image Synthesis and Segmentation. In *CVPRW*, 2221–2229.
- Heckemann, R. A.; Hajnal, J. V.; Aljabar, P.; Rueckert, D.; and Hammers, A. 2006. Automatic anatomical brain MRI segmentation combining label propagation and decision fusion. *NeuroImage* 33(1): 115–126.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *ICLR*.
- Jin, D.; Xu, Z.; Tang, Y.; Harrison, A. P.; and Mollura, D. J. 2018. CT-Realistic Lung Nodule Simulation from 3D Conditional Generative Adversarial Networks for Robust Lung Segmentation. In *MICCAI*, 732–740.
- Kennedy, D. N.; Haselgrove, C.; Hodge, S. M.; Rane, P. S.; Makris, N.; and Frazier, J. A. 2011. CANDIShare: a resource for pediatric neuroimaging data. *Neuroinformatics* 10(3): 319–322.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- Klein, A.; Mensh, B.; Ghosh, S. S.; Tourville, J. A.; and Hirsch, J. 2005. Mindboggle: Automated brain labeling with multiple atlases. *BMC Medical Imaging* 5: 7.
- Mahapatra, D.; Bozorgtabar, B.; Thiran, J.; and Reyes, M. 2018. Efficient Active Learning for Image Classification and Segmentation Using a Sample Selection and Conditional Generative Adversarial Network. In *MICCAI*, 580–588.
- Miao, S.; Piat, S.; Fischer, P. W.; Tuysuzoglu, A.; Mewes, P. W.; Mansi, T.; and Liao, R. 2018. Dilated FCN for Multi-Agent 2D/3D Medical Image Registration. In *AAAI*, 4694–4701.
- Milletari, F.; Navab, N.; and Ahmadi, S. 2016. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In *3DV*, 565–571.
- Pereira, S.; Pinto, A.; Alves, V.; and Silva, C. A. 2016. Brain Tumor Segmentation Using Convolutional Neural Networks in MRI Images. *TMI* 35(5): 1240–1251.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI*, 234–241.
- Roth, H. R.; Lu, L.; Farag, A.; Shin, H.; Liu, J.; Turkbey, E. B.; and Summers, R. M. 2015. DeepOrgan: Multi-level Deep Convolutional Networks for Automated Pancreas Segmentation. In *MICCAI*, 556–564.
- Sirinukunwattana, K.; Pluim, J. P. W.; Chen, H.; Qi, X.; Heng, P.; Guo, Y. B.; Wang, L. Y.; Matuszewski, B. J.; Bruni, E.; Sanchez, U.; Böhm, A.; Ronneberger, O.; Cheikh, B. B.; Racoceanu, D.; Kainz, P.; Pfeiffer, M.; Urschler, M.; Snead, D. R. J.; and Rajpoot, N. M. 2017. Gland segmentation in colon histology images: The glas challenge contest. *Medical Image Anal.* 35: 489–502.
- Wang, S.; Cao, S.; Wei, D.; Wang, R.; Ma, K.; Wang, L.; Meng, D.; and Zheng, Y. 2020. LT-Net: Label Transfer by Learning Reversible Voxel-Wise Correspondence for One-Shot Medical Image Segmentation. In *CVPR*, 9159–9168.

- Wu, Y.; and He, K. 2020. Group Normalization. *Int. J. Comput. Vis.* 128(3): 742–755.
- Xu, Z.; and Niethammer, M. 2019. DeepAtlas: Joint Semi-supervised Learning of Image Registration and Segmentation. In *MICCAI*, 420–429.
- Yang, H.; Sun, J.; Li, H.; Wang, L.; and Xu, Z. 2018. Neural multi-atlas label fusion: Application to cardiac MR images. *Medical Image Anal.* 49: 60–75.
- Zhao, A.; Balakrishnan, G.; Durand, F.; Guttag, J. V.; and Dalca, A. V. 2019. Data Augmentation Using Learned Transformations for One-Shot Medical Image Segmentation. In *CVPR*, 8543–8553.
- Zhou, Z.; Siddiquee, M. M. R.; Tajbakhsh, N.; and Liang, J. 2018. UNet++: A Nested U-Net Architecture for Medical Image Segmentation. In *MICCAI*, 3–11.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.
- Zhu, W.; Myronenko, A.; Xu, Z.; Li, W.; Roth, H.; Huang, Y.; Milletari, F.; and Xu, D. 2020. NeurReg: Neural Registration and Its Application to Image Segmentation. In *WACV*, 3606–3615.
- Zitova, B.; and Flusser, J. 2003. Image registration methods: a survey. *Image and vision computing* 21(11): 977–1000.