

Self-Supervised Generative Style Transfer for One-Shot Medical Image Segmentation

Devavrat Tomar¹ Behzad Bozorgtabar^{1,2,3}
 Manana Lortkipanidze Guillaume Vray¹ Mohammad Saeed Rad¹ Jean-Philippe Thiran^{1,2,3}
¹LTS5, EPFL, Switzerland ²CHUV, Switzerland ³CIBM, Switzerland
 {devavrat.tomar, firstname.lastname}@epfl.ch

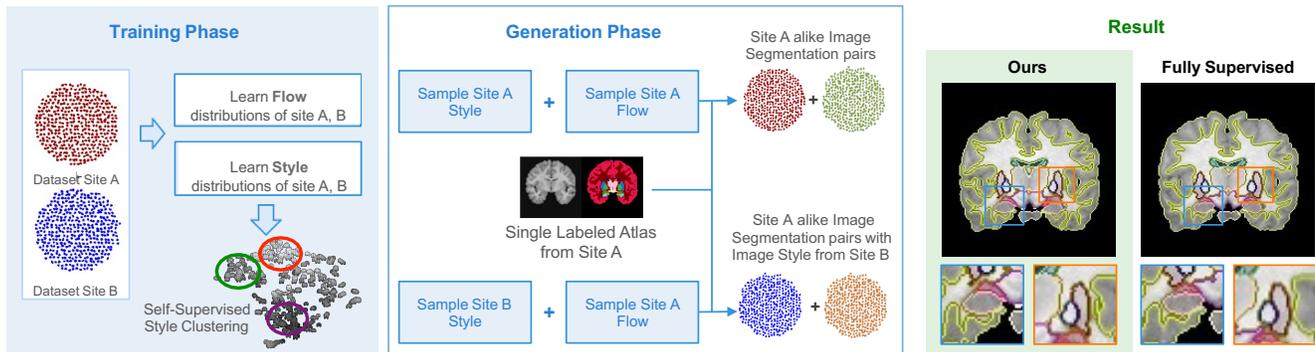


Figure 1: **Cross-site adaptation of the proposed one-shot segmentation vs. Fully Supervised segmentation.** Our method uses volumetric self-supervised learning for style transfer by leveraging unlabeled data. Zoom in for the best view.

Abstract

In medical image segmentation, supervised deep networks’ success comes at the cost of requiring abundant labeled data. While asking domain experts to annotate only one or a few of the cohort’s images is feasible, annotating all available images is impractical. This issue is further exacerbated when pre-trained deep networks are exposed to a new image dataset from an unfamiliar distribution. Using available open-source data for ad-hoc transfer learning or hand-tuned techniques for data augmentation only provides suboptimal solutions. Motivated by atlas-based segmentation, we propose a novel volumetric self-supervised learning for data augmentation capable of synthesizing volumetric image-segmentation pairs via learning transformations from a single labeled atlas to the unlabeled data. Our work’s central tenet benefits from a combined view of one-shot generative learning and the proposed self-supervised training strategy that cluster unlabeled volumetric images with similar styles together. Unlike previous methods, our method does not require input volumes at inference time to synthesize new images. Instead, it can generate diversified volumetric image-segmentation pairs from a prior distribution given a single or multi-site dataset. Augmented data generated by

our method used to train the segmentation network provide significant improvements over state-of-the-art deep one-shot learning methods on the task of brain MRI segmentation. Ablation studies further exemplified that the proposed appearance model and joint training are crucial to synthesize realistic examples compared to existing medical registration methods. The code, data, and models are available at <https://github.com/devavratTomar/SST/>.

1. Introduction

Automated medical image segmentation, for example, to localize anatomical structures, is of great importance for disease diagnosis or treatment planning. Fully supervised deep neural networks (DNNs) [31, 38] achieve state-of-the-art results when trained on large amounts of labeled data. However, acquiring abundant labeled data is often not feasible, as manual labeling is tedious and costly. Using available open-source data for domain adaptation [41, 5, 4] or hand-tuned approaches for augmentation only provides suboptimal solutions. Furthermore, the cross-modality adaptation methods [41, 42, 6] usually rely on fully labeled source datasets. Medical images can vary significantly from institution to

institution in terms of vendor and acquisition protocols [25]. As a result, the pre-trained deep networks often fail to generalize to new test data that are not distributed identically to the training data. Furthermore, although hand-tuned data augmentation (DA) techniques [1, 32] caused improvement in segmentation accuracy [37] partially, manual refinements are not sustainable, and augmented images have limited capacities to cover real variations and complex structural differences found in different images. Therefore, few-shot learning [35, 40] or self-supervised learning [10, 18] based approaches that alleviate the need for large labeled data would be of crucial importance. However, these approaches have not been explored much for low labeled data regime medical image segmentation, e.g., one-shot scenarios. Another classical medical imaging approach often used to reduce the need for labeled data for synthesis and segmentation purposes is an atlas-based approach [7, 2, 26]. In this approach, an atlas is used to register each image and warp one into another (labels undergo the same transformation). One atlas (with its label) is enough for the procedure; however, one can utilize more if available to improve accuracy. Nevertheless, heterogeneity of medical images causes inaccurate warping and, consequently, erroneous segmentation. This heterogeneity issue is even more pronounced for the multi-site medical dataset. Recently, atlas-based approaches empowered by deep convolutional neural networks (CNNs) [50, 14, 49] enable the development of one-shot learning segmentation models.

Among recent one-shot learning methods, two are the most relevant to ours [50, 44]. In the first work, Zhao *et al.* [50] proposed a learning framework to synthesize labeled samples using CNN to warp images based on atlas. However, [50] only recreates exact styles/deformations presented in the dataset **without inducing diversity**. Furthermore, training includes two separate steps for spatial and appearance transformations bringing extra computational overhead. In the second work [44], a one-shot segmentation method has been proposed based on the forward-backward consistency strategy to stabilize the training process. Nonetheless, since the atlas' style does not match the unlabeled image style, this results in imprecise registration and imperfect segmentation. In summary, our main contributions are as follows:

- We propose a novel volumetric self-supervised contrastive learning to learn style representation from unlabeled data facilitating registration and consequently segmentation task in the presence of intra-site and inter-site heterogeneity of MR images (see Fig. 1);
- Unlike current state-of-the-art methods, our method does not require input volumes at inference time to synthesize new images. Instead, it can generate diversified and unlimited image-segmentation pairs by just sampling from a prior distribution;

- Previous methods train the spatial and appearance models separately, resulting in sub-optimal solutions compared to our joint optimization of all modules. Our method achieved state-of-the-art one-shot segmentation performance on two brain T1-weighted MRI datasets and improved the generalization ability for cross-site adaptation scenarios.

2. Related Work

Data Augmentation. Various data augmentation techniques have been developed to compensate for the extreme scarcity of labeled data encountered in medical image segmentation [9, 34, 50]. Augmentation approaches range from geometric transformations [31, 33] to data-driven augmentation methods [17, 27]. While former methods are often difficult to tune as they have limited capacities to cover real variations and complex structural differences found in images, latter approaches often learn generative models to synthesize images to combine with real images and train the segmentation model. However, variations in shape and anatomical structures can negatively impact their performance, especially when little training data is available. In this regard, image registration has been effective in approximating deformations between unlabeled images so that the augmented images with plausible deformations can be obtained [50, 48, 8]. The augmented images then allow training deep segmentation models with few labeled examples in a semi-supervised manner. Disappointingly, heterogeneity of medical images often yields inaccurate warping between the moving image and the fixed image and, consequently, inaccurate segmentation. Olut *et al.* [34] proposed an augmentation method to leverage statistical deformation models from unlabeled data via principal component analysis. Shen *et al.* [39] proposed a geometric based image augmentation method that generates realistic images via interpolation from the geodesic subspace to estimate the space of anatomical variability. He *et al.* [19] proposed a Deep Complementary Joint Model (DeepRS) for medical image registration and segmentation.

Image Segmentation in Low Data Regimes. Self-supervised learning and few-shot learning are two facets of the same problem: *training a deep model in low labeled data regime*. These approaches have been used for sparsely annotated images for segmentation but without much success. The former often rely on many training classes to avoid overfitting [43, 40], while the latter requires fine-tuning on sufficient labeled data before testing [16, 45]. Deep atlas-based models [3, 50, 44, 13, 2, 26] using a single atlas or multi-atlas tackled weakly-supervised medical image segmentation. Balakrishnan *et al.* [3] developed VoxelMorph, which aims to estimate pairwise 3D image registration through a learned CNN-based spatial transformation function. In a one-shot learning context, it learns to register a labeled atlas to any

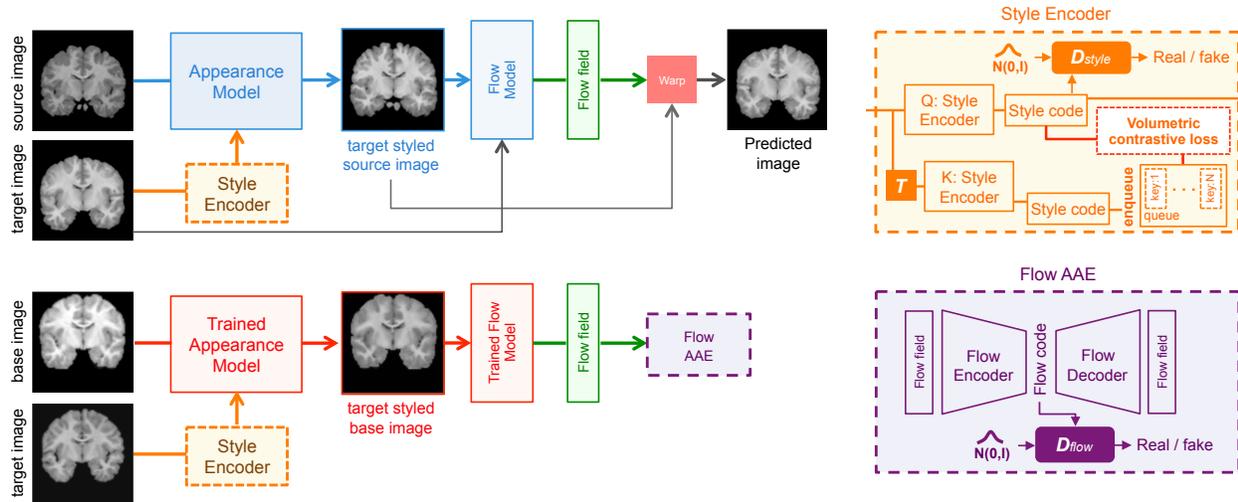


Figure 2: **Schematic description of the training phase.** The **Appearance Model** applies the target image’s style to the source image using its style code predicted by the **Style Encoder**, which is trained in parallel using self-supervised volumetric contrastive loss. Then, the **Flow Model** non-linearly warps the style translated source image to the target image. This allows backpropagation of supervision signals to all three models. Independently, **Flow AAE** maps the flow fields corresponding to the base image generated by using the trained **Flow model** and **Appearance model** to a normally distributed latent space. D_{style} and D_{flow} are trained in an adversarial manner to ensure that the style codes and the flow codes are normally distributed.

other unlabeled volume. This model suffers from variance in voxel intensity confusing the spatial transformer. In this regard, a similar unsupervised method [13] has been proposed that combines a conventional probabilistic atlas-based segmentation with deep learning for MRI scan segmentation. More recently, few deep models [50, 44] explore the one-shot setting for medical image segmentation. Nevertheless, these methods either use samples from a single site (hospital) [44] or aggregate data from multiple sites [50] without cross-dataset transfer learning capability.

3. Method

We first recap the concept of our proposed one-shot atlas-based medical image segmentation. We formulate the synthesis of novel volumetric images and their corresponding segmentation labels as a learned random spatial and style deformation of the given single labeled volumetric atlas image (referred to as **base image**) from learned latent space. We employ a **Style Encoder** (Sec. 3.1) that learns to cluster similar styled images together in a self-supervised manner using volumetric contrastive loss by adapting Momentum Contrast [18] while imposing a normal distribution prior on the latent style codes using adversarial training [28]. The **Appearance Model** (Sec. 3.2) is trained to generate different styles of the base image without changing its spatial structure. For learning the spatial deformation correspondences (referred to as **flow**) between the base image and the target image, we employ **Flow Model** (Sec. 3.3) that is trained on the task of

registering two different image volumes with the same style. This is achieved by first changing the style of moving image (referred to as **source image**) to match the target image’s style (as obtained by **Style Encoder**) using the Appearance Model, followed by morphing it into the target image. As discussed later in Sec 4.1, matching the source image’s style to the target image improves the registration accuracy. We employ an additional adversarial autoencoder [28] that encodes the Flow model’s output for the base image into a Gaussian prior flow latent space to learn the distribution of the spatial deformation fields corresponding to the base image. At test time, we sample flow latent code and style latent code from Gaussian prior to generate new deformation fields and style appearances for the base image, respectively. Fig. 2 and Fig. 3 show an overview of the training procedure and the data generation at test time, respectively. To quantify the images’ quality and their corresponding segmentation labels generated using our approach, we train a separate 3D U-Net [12] on them. We test our model on the real image/segmentation pairs and compare it with the performance of a fully supervised model trained using real data. The loss terms used for training Style Encoder, Appearance Model, and Flow Model, described in the subsequent subsections.

3.1. Style Encoder

Style Encoder aims to incentivize content-invariant image-level representation that brings together similar styled images and pushes apart dissimilar styled images. To do so, we

propose a new volumetric contrastive learning [10, 18, 11] based strategy for training. In particular, we adapt Momentum Contrast [18] for volumetric medical images for clustering task of images’ styles as opposed to the original classification task [18]. More importantly, we use the learned spatial transformer [3] to generate the positive keys (preserving the styles) instead of standard augmentation, e.g., random cropping used in the original formulation. Without loss of generality, we keep a dictionary of keys $\{k_1, k_2, \dots, k_K\}$ that represents different styles. During a training step, we sample a volumetric image \mathbf{q} (called **query**) from the training set and generate its corresponding positive **key** volumetric image \mathbf{k}_+ by warping \mathbf{q} to a randomly selected volumetric image from the training set using a pre-trained spatial transformer \mathcal{T} . This ensures that \mathbf{q} and \mathbf{k}_+ have the same style (with different structural geometries) that is different from the style keys in the dictionary. The dictionary’s style keys are generated by a separate model (**key-Style Encoder**) whose weights are updated as a moving average of the weights of the Style Encoder with momentum $m = 0.99$. The volumetric contrastive loss is computed as below:

$$\mathcal{L}_{\text{vol.cl}} = -\log \frac{\exp(q \cdot k_+ / \tau)}{\sum_{i=0}^K \exp(q \cdot k_i / \tau)} \quad (1)$$

where $k_+ = \mathcal{T}(q)$ and τ is a temperature hyper-parameter [47]. The sum is over one positive and K negative samples.

3.2. Appearance Model

The Appearance Model is responsible for translating the style of the source image (s) to that of the target image (t), given the style latent code as predicted by the Style Encoder. We feed the target image’s style code to the adaptive instance normalization (AdaIN) layers [20] of the model to perform style transfer. Thus, the target styled source image (\tilde{s}) is obtained as:

$$\tilde{s} = \mathcal{A}(s, E_{\text{style}}(t)) \quad (2)$$

where \mathcal{A} denotes the Appearance model, and E_{style} represents the Style Encoder. The Appearance Model loss \mathcal{L}_{app} consists of two components: $\mathcal{L}_{\text{app}} = \mathcal{L}_{\text{cycle}}^{\text{style}} + \mathcal{L}_{\text{id}}^{\text{style}}$, where $\mathcal{L}_{\text{cycle}}^{\text{style}}$ and $\mathcal{L}_{\text{id}}^{\text{style}}$ denote the style consistency loss and style identity loss, respectively that are described below.

Style Consistency Loss. We include a style consistency loss that guides the Appearance model to generate images with the same spatial structure as the source image but with a different style in a consistent cyclic manner. Given the style codes of the source and target images, the following style consistency loss is computed as:

$$\mathcal{L}_{\text{cycle}}^{\text{style}} = \mathcal{L}_{\text{SSIM-L}_1}(s, \mathcal{A}(\tilde{s}, E_{\text{style}}(s))) \quad (3)$$

where $\mathcal{L}_{\text{SSIM-L}_1}$ computes multi-scale structural similarity index [46] and L_1 distance between the two images as:

$$\mathcal{L}_{\text{SSIM-L}_1}(u, v) = \|u - v\|_1 + (1 - \text{SSIM}(u, v)) \quad (4)$$

Style Identity Loss. We also include a regularization loss term, called style identity, that enforces the Appearance Model to generate the same image using its own style.

$$\mathcal{L}_{\text{id}}^{\text{style}} = \mathcal{L}_{\text{SSIM-L}_1}(s, \mathcal{A}(s, E_{\text{style}}(s))) \quad (5)$$

3.3. Flow Model

The Flow Model builds upon a spatial transformer network that warps a *moving image* (I_m) to the *fixed image* (I_f). The spatial transformer (\mathcal{F}) generates a correspondence map δp , referred to as *flow*, which is used to register I_m onto I_f . This warping operation is defined as:

$$\begin{aligned} \delta p &= \mathcal{F}(I_m, I_f) \\ y &= \delta p \odot I_m \end{aligned} \quad (6)$$

where \odot denotes the warping operator, and y is the predicted image. Once we know the correspondence map δp between the base image (b) and the target image (t), we can transfer the segmentation label of the base image (b_{seg}) onto the target image using the same warping operation.

$$t_{\text{seg}} = \mathcal{F}(b, t) \odot b_{\text{seg}} \quad (7)$$

To learn the distribution of deformation fields corresponding to the base image predicted by the Flow Model, we train a separate adversarial autoencoder [28] on its output. This allows us to encode the flow δp in the latent space, which can be used later to generate novel volumetric images and corresponding segmentation labels (Sec. 3.6). The Flow Model loss $\mathcal{L}_{\text{flow}}$ consists of two components: $\mathcal{L}_{\text{flow}} = \mathcal{L}_{\text{recon}}^{\text{flow}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}^{\text{flow}}$, where $\mathcal{L}_{\text{recon}}^{\text{flow}}$ and $\mathcal{L}_{\text{reg}}^{\text{flow}}$ denote the reconstruction loss and the flow regularization loss that are described below.

Reconstruction Loss. In contrast to the Normalized Cross-Correlation loss [3] commonly used for the voxel registration, we include pixel similarity based reconstruction loss between the target image (t) and the spatially warped target styled source image (referred to as **predicted image**) obtained by Flow Model. Using the pixel-wise similarity loss is justified as the Appearance Model changes the source image’s style to match the style of the target image, thus allowing adequate registration of the two images.

$$\mathcal{L}_{\text{recon}}^{\text{flow}} = \mathcal{L}_{\text{SSIM-L}_1}(t, \mathcal{F}(\tilde{s}, t) \odot \tilde{s}) \quad (8)$$

Flow Regularization. We regularize the flow δp by penalizing its spatial gradient, thus ensuring smoothness of the correspondence map generated by our spatial transformer.

$$\mathcal{L}_{\text{reg}}^{\text{flow}} = \|\nabla \mathcal{F}(\tilde{s}, t)\|_1 \quad (9)$$

We prefer the L_1 norm as it helps to stabilize training and results in less noisy flow compared to the L_2 norm.

3.4. Adversarial Loss

We introduce two latent discriminators, called D_{style} and D_{flow} , for enforcing a prior distribution on the latent style codes and the flow codes generated during training, respectively. We use the adversarial loss function described in LS-GAN [29] for training the two discriminators along with Style Encoder and Flow Encoder in an adversarial manner, as shown in Fig. 2.

$$\mathcal{L}_{\text{adv}}^{\text{style}} = \mathbb{E}_{t \sim X_{\text{data}}} [(D_{\text{style}}(E_{\text{style}}(t)) - 1)^2] + \mathbb{E}_{n \sim \mathcal{N}} [D_{\text{style}}(n)^2] \quad (10)$$

where t is sampled from the training images X_{data} , and \mathcal{N} is the normal distribution. Similarly, we train the adversarial autoencoder (AAE) [28] on flow fields generated by Flow Model (Sec. 3.3) using LS-GAN and l_1 reconstruction loss.

3.5. Training Objective

Finally, the proposed training loss $\mathcal{L}_{\text{total}}$ joins the loss terms used to train Style Encoder, Appearance Model, and Flow Model:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{vol.cl}} + \lambda_1 \mathcal{L}_{\text{app}} + \lambda_2 \mathcal{L}_{\text{flow}} + \lambda_3 \mathcal{L}_{\text{adv}}^{\text{style}} \quad (11)$$

where λ 's are the weights of different losses. We observed that pre-training the Style Encoder alone using the loss: $\mathcal{L}_{\text{vol.cl}} + \lambda_3 \mathcal{L}_{\text{adv}}^{\text{style}}$ improves the overall convergence rate and reduces the optimization's complexity. After pre-training the Style Encoder, we jointly optimize it along with Appearance Model and Flow Model by minimizing $\mathcal{L}_{\text{total}}$.

3.6. Generation Phase

As shown in Fig. 3, we can sample new volumetric images and their segmentation maps by mapping Gaussian noise to novel style codes and flow fields using the trained Appearance Model (\mathcal{A}) and Flow Decoder (G_{flow}). The Appearance Model performs style deformation of the base image while the Flow Decoder produces random flow field which is used to warp the styled base image and its segmentation, thus generating new image-segmentation pairs.

$$\begin{aligned} X &= G_{\text{flow}}(n_{\text{flow}}) \odot \mathcal{A}(b, n_{\text{style}}) \\ Y &= G_{\text{flow}}(n_{\text{flow}}) \odot b_{\text{seg}} \end{aligned} \quad (12)$$

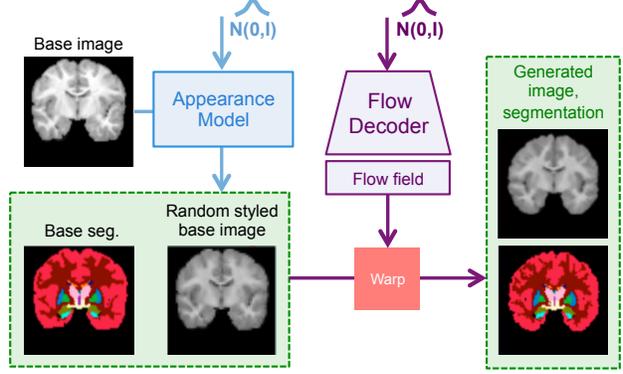


Figure 3: **Data Generation** phase. We use the learned **Appearance Model** and **Flow Decoder** to generate new images and segmentation labels from the normal distribution.

where b and b_{seg} are the base image and base segmentation and $n_{\text{flow}}, n_{\text{style}} \sim \mathcal{N}$. X and Y are the generated image-segmentation pairs.

4. Experiments

This section introduces the implementation details, experimental settings, dataset, results, and ablation studies.

Dataset, Preprocessing and Evaluation Metric. We evaluate our method and other baselines on multi-study datasets from publicly available datasets: CANDI [23] and a large-scale dataset, OASIS [30], each contains 3D T1-weighted MRI volumes. CANDI dataset consists of scans from 103 patients with manual anatomical segmentation labels, whereas the OASIS dataset has 2044 scans with segmentation labels obtained by the FreeSurfer [15] pipeline. As in VoxelMorph [3] and LT-Net [44], 28 anatomical structures are used in our experiments. All the dataset images are first preprocessed by removing the brain's skull region, followed by center cropping the volumes to $128 \times 160 \times 160$. We set aside 20% brain images and their segmentation as a test set, which was untouched during training. The remaining 80% brain images are then used for training and validation, with a split of 90%-10%, in which there is no patient ID overlap among the subsets. For each dataset, there are different acquisition details and health conditions. The most similar image to the anatomical average is selected as the only annotated atlas (base image) used for the training. We only use validation set labels for choosing the best model and hyper-parameters. For the OASIS dataset, the model is trained with a mixture of healthy subjects and diseased patients and is then evaluated on test cases constitutes of both sets. We assess our method's efficiency by training a 3D U-Net-based segmentation model on the generated volumetric image-segmentation pairs and evaluating its perfor-

| | Cerebral-WM | Cerebral-CX | Lateral-Vent | Cerebellum-WM | Cerebellum-CX | Thalamus-Proper | Caudate | Putamen | Pallidum | 3rd-Vent | 4th-Vent | Brain-Stem | Hippocampus | Amygdala | CSF | VentralDC | Mean±std |
|------------------------|-------------|-------------|--------------|---------------|---------------|-----------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|------------------|
| (CANDI) | | | | | | | | | | | | | | | | | |
| MABMIS | 86.3 | 90.5 | 86.1 | 75.9 | 90.2 | 84.2 | 79.2 | 79.5 | 67.6 | 66.7 | 79.8 | 90.8 | 73.4 | 64.8 | 60.4 | 77.0 | 78.3±2.9 |
| VoxelMorph | 81.1 | 87.3 | 83.7 | 69.9 | 82.4 | 85.9 | 82.6 | 81.0 | 74.7 | 64.3 | 73.4 | 89.2 | 66.2 | 66.7 | 59.4 | 81.1 | 76.9±3.2 |
| Bayesian | 89.5 | 84.5 | 85.3 | 84.9 | 82.4 | 82.1 | 83.7 | 83.0 | 77.2 | 57.5 | 75.7 | 84.1 | 74.3 | 72.8 | 48.6 | 75.9 | 77.7±2.3 |
| DataAug | 84.8 | 89.0 | 77.4 | 72.3 | 86.8 | 89.0 | 84.6 | 86.3 | 79.8 | 71.2 | 78.7 | 91.3 | 76.0 | 72.3 | 63.3 | 82.7 | 80.4±3.2 |
| LT-Net** | 85.8 | 90.9 | 83.1 | 80.1 | 91.6 | 87.9 | 85.5 | 88.4 | 80.5 | 68.4 | 79.7 | 92.4 | 71.6 | 71.6 | 67.1 | 82.3 | 81.7±8.0 |
| Ours | 90.9 | 94.3 | 89.2 | 83.4 | 89.5 | 89.2 | 88.3 | 86.7 | 81.1 | 71.3 | 81.9 | 92.2 | 76.0 | 70.9 | 67.6 | 82.2 | 83.5 ±3.0 |
| 3D U-Net* | 94.1 | 97.0 | 93.8 | 89.6 | 96.8 | 91.5 | 90.4 | 90.8 | 83.0 | 74.4 | 87.0 | 94.9 | 84.1 | 79.0 | 79.6 | 88.1 | 88.1 ±1.5 |
| (OASIS) | | | | | | | | | | | | | | | | | |
| MABMIS | 79.4 | 62.1 | 84.4 | 77.1 | 82.6 | 85.3 | 77.5 | 78.9 | 66.5 | 81.5 | 62.3 | 88.4 | 72.5 | 72.9 | 63.6 | 78.3 | 75.8±6.9 |
| VoxelMorph | 75.2 | 63.1 | 85.3 | 77.9 | 83.4 | 85.0 | 79.1 | 81.3 | 72.3 | 81.0 | 58.8 | 90.5 | 70.3 | 72.8 | 67.6 | 79.4 | 76.4±4.3 |
| Bayesian | 90.0 | 45.8 | 64.7 | 82.4 | 72.9 | 84.1 | 70.1 | 70.7 | 40.6 | 33.4 | 18.8 | 87.8 | 48.3 | 56.9 | 33.0 | 62.1 | 59.5±3.0 |
| DataAug | 74.8 | 69.0 | 69.2 | 67.2 | 80.1 | 78.0 | 61.8 | 66.2 | 50.4 | 70.5 | 50.6 | 83.2 | 64.8 | 56.8 | 56.8 | 69.2 | 66.8±5.7 |
| Ours | 89.4 | 89.2 | 89.2 | 86.3 | 91.7 | 84.8 | 80.5 | 80.1 | 65.1 | 82.0 | 70.3 | 91.5 | 74.7 | 69.4 | 65.4 | 77.9 | 80.5 ±3.9 |
| 3D U-Net* | 94.2 | 95.6 | 94.6 | 92.0 | 97.5 | 91.8 | 89.1 | 89.5 | 79.7 | 90.9 | 89.8 | 96.9 | 89.3 | 85.1 | 86.3 | 87.9 | 90.6 ±1.4 |
| (OASIS → CANDI) | | | | | | | | | | | | | | | | | |
| w/o Style Adap. | 84.9 | 88.4 | 58.9 | 75.3 | 90.6 | 78.8 | 55.5 | 47.9 | 40.6 | 55.5 | 65.5 | 83.3 | 55.2 | 44.8 | 44.6 | 60.9 | 64.4±3.5 |
| w/ Style Adap. | 85.0 | 90.0 | 77.0 | 74.7 | 89.8 | 83.3 | 78.1 | 75.8 | 68.6 | 66.3 | 73.6 | 88.1 | 58.8 | 52.3 | 52.4 | 69.2 | 74.0±3.2 |
| (CANDI → OASIS) | | | | | | | | | | | | | | | | | |
| w/o Style Adap. | 80.2 | 83.1 | 58.1 | 32.4 | 71.4 | 37.2 | 39.6 | 20.6 | 6.2 | 3.8 | 7.6 | 12.5 | 12.0 | 17.9 | 2.3 | 12.6 | 31.1±11.7 |
| w/ Style Adap. | 85.4 | 85.4 | 81.6 | 45.3 | 67.8 | 67.6 | 51.9 | 29.0 | 6.6 | 30.5 | 14.2 | 67.7 | 31.8 | 45.0 | 32.2 | 44.2 | 49.1±10.2 |

Table 1: Comparison of segmentation performance (mean *Dice* score in %) of MABMIS (2 atlas) [22], VoxelMorph [3], Bayesian [13], DataAug [50], LT-Net** [44], and **Ours** across various brain structures on CANDI and OASIS datasets. Abbreviations used: white matter (WM), cortex (CX), ventricle (Vent), and cerebrospinal fluid (CSF). ** as reported in the published paper. * fully supervised model. The last four rows (OASIS→CANDI and CANDI→OASIS) indicate the results of our method with (w/) and without (w/o) cross-site style adaptation.

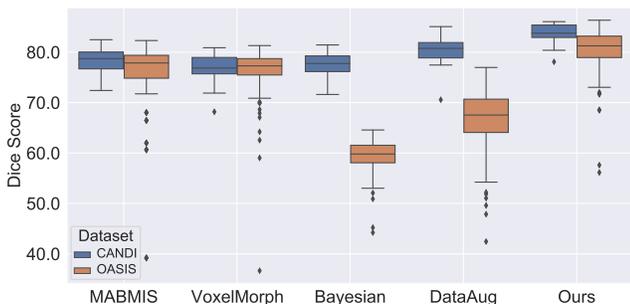


Figure 4: Comparison of volume-wise segmentation accuracy (*Dice* score %) of our method with MABMIS (2 atlas) [22], VoxelMorph [3], Bayesian [13] and DataAug [50]. We outperform the second best baseline by a margin of **3.1%** on CANDI and **4.1%** on OASIS dataset (p-value of 5.8×10^{-4} , 3.2×10^{-16} respectively using paired t-test).

performance on the untouched test data. We use the *Dice* similarity coefficient between the ground truth segmentation and the predicted result in assessing the segmentation accuracy.

Experimental Settings. All our models are based on 3D CNNs [21]. The Appearance model has 3 AdaIN [20] layers for performing style transfer using the style codes. Flow

Model is a lighter version of VoxelMorph [3], while AAE Model has a 3D convolutional encoder-decoder architecture. We implement all the models in PyTorch [36]. For training Style Encoder, Appearance Model, Flow Model, and Flow Autoencoder, we used Adam [24] optimizer with a learning rate of 2×10^{-4} and $(\beta_1 = 0.9, \beta_2 = 0.999)$. We use the same optimizer with the same learning rate for training the latent style code and flow code’s discriminators but with $(\beta_1 = 0.5, \beta_2 = 0.999)$. A hyper-parameter search was conducted to find the optimal values. We choose the loss term weights as $\lambda_1 = 5.0$, $\lambda_2 = 1.0$, $\lambda_3 = 0.1$ and $\lambda_{reg} = 0.1$. The temperature coefficient τ for volumetric contrastive loss is set to 0.7. We use a batch size of 32 for pre-training the Style Encoder and 4 when all the models are trained end-to-end. We utilize the same experimental setup for all baseline experiments for a fair comparison.

Comparison with SOTA Methods. Our method surpasses the state-of-the-art methods in most semantic classes and, on average, on both CANDI and OASIS datasets (see Fig. 4). A qualitative comparison of our method with several baselines on coronal brain slices’ segmentation task is shown in Fig. 5. MABMIS [22] and VoxelMorph [3] produce visibly noisy and inaccurate boundaries compared to other methods. Bayesian [13] and DataAug [50] struggle with outer

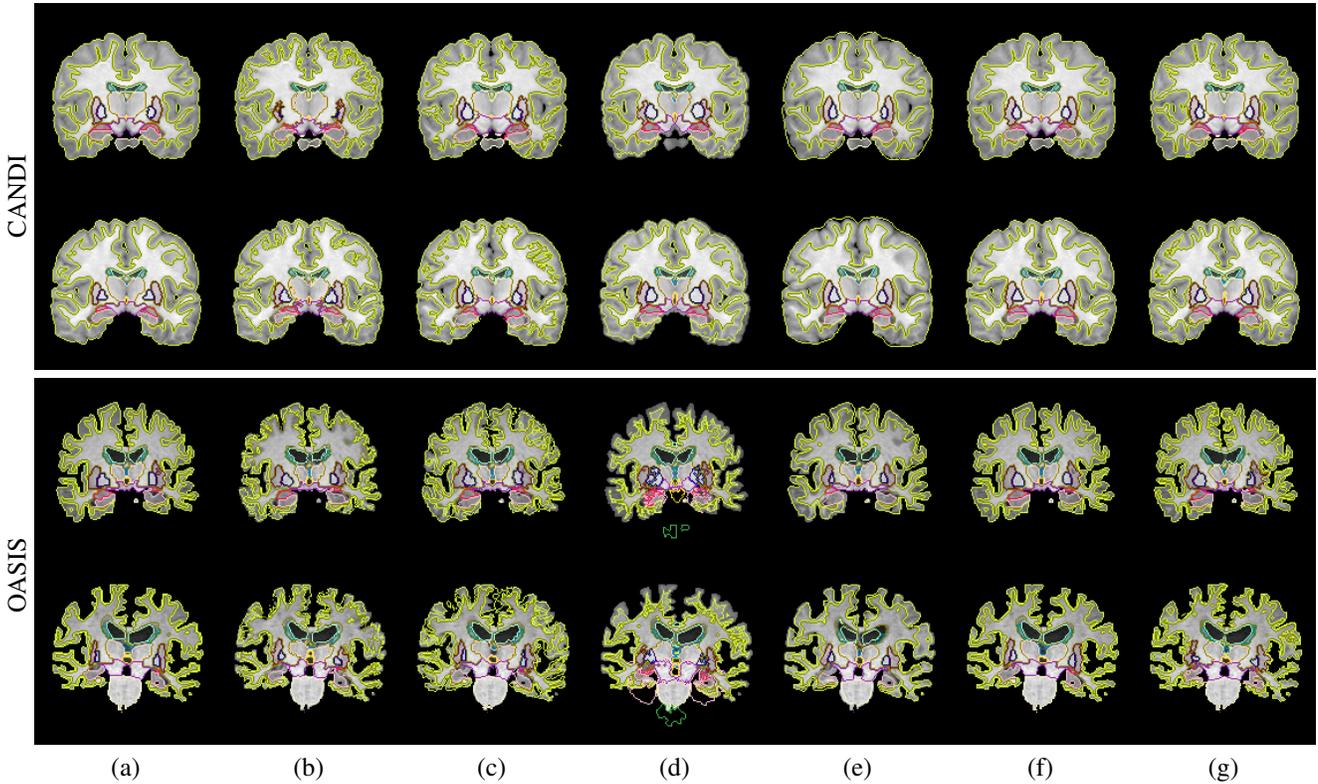


Figure 5: Qualitative comparison of our method with other baselines on the segmentation task of Brain MRI volumetric images from CANDI and OASIS dataset. Left to right: (a) Ground Truth, the segmentation results from (b) MABMIS [22], (c) VoxelMorph [3], (d) Bayesian [13], (e) DataAug [50], (f) Fully Supervised 3D U-Net, (g) **Ours**. Best visualized in color.

contours, where one crops it unnecessarily, and another includes background. Furthermore, VoxelMorph and DataAug encounter difficulties in identifying small segmentation regions, whereas Bayesian overproduces them. Instead, our method handles outer/inner regions and smaller anatomical regions well and compares closely to supervised/ground truth results. The qualitative observations are backed by quantitative metrics (see Table 1).

4.1. Ablations

Ablation Study on Appearance Model. The Appearance Model plays a crucial role in generating diversified styled images and improving the Flow Model’s image registration efficiency during training. Table 2 shows the effect of including-excluding the Appearance Model (\mathcal{L}_{app}) on the registration accuracy of segmentation labels predicted by flow model and supervised segmentation accuracy of the 3D U-Net using the generated images. For both these scenarios, the average *Dice* score improves after the Appearance Model’s inclusion. This is expected as registering two images with similar intensity distributions is easier than registering images with different intensities. We can only generate similar styled images without the Appearance model for the generation phase, which leads to a poor generalization of 3D U-Net

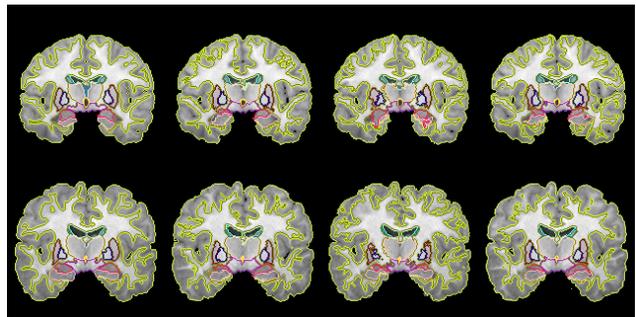


Figure 6: Qualitative comparison of registration accuracy evaluated on segmentation for different methods. *Left to Right*: Ground-truth, VoxelMorph[3], MABMIS[22], Ours.

on the test set. A qualitative evaluation of registration on the CANDI test set is shown in Fig. 6.

Ablation Study on Joint Training. We observed that training our model end-to-end is critical for improved registration accuracy (Table 2). We experimented with the pre-training Appearance model, Flow Model, and Style encoder separately using the losses defined in Sec. 3 followed by fine-tuning and found significant improvement when Appearance Model and Flow Model are trained end-to-end. However,

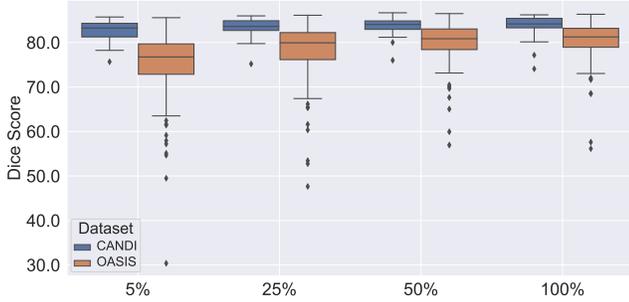


Figure 7: Ablation for segmentation accuracy of 3D U-Net using different sizes of generated data on CANDI and OASIS datasets. 100% implies 1850 generated image-segmentation volumes.

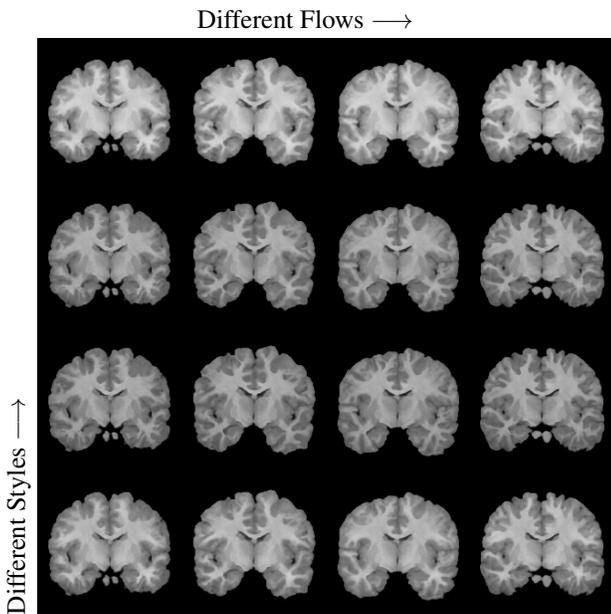


Figure 8: **Manipulating style and flow codes.** *Left to right:* Images generated using the same style code with different flow codes. *Bottom to top:* Images generated using the same flow code and distinct style.

a pre-training Style encoder using volumetric contrastive loss improves the overall model convergence ($10\times$ faster) without qualifying the performance as compared to training it jointly.

Cross-Site One-Shot Adaptation. To evaluate the Appearance model’s efficacy in generating the unlabeled target dataset styles, we experimented with replacing the Appearance Model trained on CANDI (OASIS) with the one trained on OASIS (CANDI) to generate images in the style of the target domain. The volumetric image-segmentation pairs generated are then used to train 3D U-Net, which is further evaluated on the target domain’s test samples. As shown in

| Ablation Type | Method | Mean \pm std | |
|------------------------------------|----------|--------------------------------|--------------------------------|
| | | CANDI | OASIS |
| w/o \mathcal{L}_{app} | Reg. (a) | 71.6 \pm 3.2 | 68.5 \pm 3.8 |
| w/o Joint opt. | Reg. (a) | 73.6 \pm 3.3 | 69.9 \pm 3.7 |
| w/ Joint opt., \mathcal{L}_{app} | Reg. (a) | 78.9\pm2.4 | 73.3\pm3.1 |
| w/o \mathcal{L}_{app} | Sup. (b) | 54.9 \pm 20.4 | 41.2 \pm 22.7 |
| w/ \mathcal{L}_{app} | Sup. (b) | 83.5\pm3.0 | 80.5\pm3.9 |

Table 2: **Effect of style transfer and joint training.** Mean *Dice* score of the segmentation task with and without Appearance Model and joint training using the method as (a) registration of the base image to the target image (b) 3D U-Net trained on generated images-segmentation pairs.

Table 1, the domain gap between the two data sites would severely impede the generalization ability of the trained supervised model on the source data site (w/o style adaptation) while performing style adaptation between OASIS and CANDI dataset provides substantial improvements. We use the same test samples from CANDI and OASIS datasets for evaluation.

Exploring Diversity of the Generated Data. We evaluate 3D U-Net’s performance on the segmentation task using different data sizes generated by our approach. For the best performing 3D U-Net, we report the accuracy obtained by training it on 1850 generated image-segmentation pairs. Fig. 7 shows a box-plot of the *Dice* score (%) of the segmentation accuracy of 3D U-Net trained on different sizes of samples generated using our proposed method trained on CANDI and OASIS. This suggests that our method has learned to generate more diverse and realistic data. Besides, the effects of different flow codes and styles on generated images can be observed in Fig. 8.

5. Conclusion and Future Work

We proposed the novel volumetric contrastive loss used for style transfer by leveraging unlabeled data for one-shot medical image segmentation. We presented a generic method adapted for cross-site one-shot segmentation scenarios to generate arbitrarily diversified volumetric image-segmentation pairs using trained appearance models from one data site and a flow model from another data site. We demonstrated state-of-the-art one-shot segmentation performance on two T1-weighted brain MRI datasets under various settings and ablations. We shed light on our method’s efficacy in closing the gap with a fully-supervised segmentation model in the extreme case of only one labeled atlas. Our method uses neither tissue nor modality-specific information and can be adjusted to other modalities or anatomy. As future work, our method can be easily extended for the few-shot scenario using several atlases.

References

- [1] Zeynettin Akkus, Alfiia Galimzianova, Assaf Hoogi, Daniel L Rubin, and Bradley J Erickson. Deep learning for brain mri segmentation: state of the art and future directions. *Journal of digital imaging*, 30(4):449–459, 2017.
- [2] Hossein Arabi, Nikolaos Koutsouvelis, Michel Rouzaud, Raymond Miralbell, and Habib Zaidi. Atlas-guided generation of pseudo-ct images for mri-only and hybrid pet–mri-guided radiotherapy treatment planning. *Physics in Medicine & Biology*, 61(17):6531, 2016.
- [3] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019.
- [4] Behzad Bozorgtabar, Mohammad Saeed Rad, Hazim Kemal Ekenel, and Jean-Philippe Thiran. Learn to synthesize and synthesize to learn. *Computer Vision and Image Understanding*, 185:1–11, 2019.
- [5] Behzad Bozorgtabar, Mohammad Saeed Rad, Hazim Kemal Ekenel, and Jean-Philippe Thiran. Using photorealistic face synthesis and domain adaptation to improve facial expression analysis. In *2019 14th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2019)*, pages 1–8. IEEE, 2019.
- [6] Behzad Bozorgtabar, Mohammad Saeed Rad, Dwarikanath Mahapatra, and Jean-Philippe Thiran. Syndemo: Synergistic deep feature alignment for joint learning of depth and ego-motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4210–4219, 2019.
- [7] Ciprian Catana, Andre van der Kouwe, Thomas Benner, Christian J Michel, Michael Hamm, Matthias Fenchel, Bruce Fischl, Bruce Rosen, Matthias Schmand, and A Gregory Sorensen. Toward implementing an mri-based pet attenuation-correction method for neurologic studies on the mr-pet brain prototype. *Journal of Nuclear Medicine*, 51(9):1431–1438, 2010.
- [8] Krishna Chaitanya, Neerav Karani, Christian F Baumgartner, Anton Becker, Olivio Donati, and Ender Konukoglu. Semi-supervised and task-driven data augmentation. In *International conference on information processing in medical imaging*, pages 29–41. Springer, 2019.
- [9] Chen Chen, Chen Qin, Huaqi Qiu, Cheng Ouyang, Shuo Wang, Liang Chen, Giacomo Tarroni, Wenjia Bai, and Daniel Rueckert. Realistic adversarial data augmentation for mr image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 667–677. Springer, 2020.
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [11] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [12] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016.
- [13] Adrian V Dalca, Evan Yu, Polina Golland, Bruce Fischl, Mert R Sabuncu, and Juan Eugenio Iglesias. Unsupervised deep learning for bayesian brain mri segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 356–365. Springer, 2019.
- [14] Mohamed S Elmahdy, Jelmer M Wolterink, Hessam Sokooti, Ivana Išgum, and Marius Staring. Adversarial optimization for joint registration and segmentation in prostate ct radiotherapy. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 366–374. Springer, 2019.
- [15] Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.
- [16] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- [17] Changhee Han, Leonardo Rundo, Ryosuke Araki, Yujiro Furukawa, Giancarlo Mauri, Hideki Nakayama, and Hideaki Hayashi. Infinite brain mr images: Pggan-based data augmentation for tumor detection. In *Neural approaches to dynamics of signal exchanges*, pages 291–303. Springer, 2020.
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.
- [19] Yuting He, Tiantian Li, Guanyu Yang, Youyong Kong, Yang Chen, Huazhong Shu, Jean-Louis Coatrieux, Jean-Louis Dillenseger, and Shuo Li. Deep complementary joint model for complex scene registration and few-shot segmentation on medical images. *arXiv preprint arXiv:2008.00710*, 2020.
- [20] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017.
- [21] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):221–231, 2012.
- [22] Hongjun Jia, Pew-Thian Yap, and Dinggang Shen. Iterative multi-atlas-based multi-image segmentation with tree-based registration. *NeuroImage*, 59(1):422–430, 2012.
- [23] David N Kennedy, Christian Haselgrove, Steven M Hodge, Pallavi S Rane, Nikos Makris, and Jean A Frazier. Candishare: a resource for pediatric neuroimaging data, 2012.
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [25] Kelvin K Leung, Matthew J Clarkson, Jonathan W Bartlett, Shona Clegg, Clifford R Jack Jr, Michael W Weiner, Nick C Fox, Sébastien Ourselin, Alzheimer’s Disease Neuroimaging Initiative, et al. Robust atrophy rate measurement in

- alzheimer’s disease using multi-site serial mri: tissue-specific intensity normalization and parameter selection. *Neuroimage*, 50(2):516–523, 2010.
- [26] Maria Lorenzo-Valdés, Gerardo I Sanchez-Ortiz, Raad Mohiaddin, and Daniel Rueckert. Atlas-based segmentation and tracking of 3d cardiac mr images using non-rigid registration. In *International conference on medical image computing and computer-assisted intervention*, pages 642–650. Springer, 2002.
- [27] Dwarikanath Mahapatra, Behzad Bozorgtabar, and Ling Shao. Pathological retinal region segmentation from oct images using geometric relation based augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9611–9620, 2020.
- [28] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- [29] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2017.
- [30] Daniel S Marcus, Tracy H Wang, Jamie Parker, John G Csernansky, John C Morris, and Randy L Buckner. Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *Journal of cognitive neuroscience*, 19(9):1498–1507, 2007.
- [31] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016.
- [32] Pim Moeskops, Max A Viergever, Adrienne M Mendrik, Linda S De Vries, Manon JNL Benders, and Ivana Išgum. Automatic segmentation of mr brain images with a convolutional neural network. *IEEE transactions on medical imaging*, 35(5):1252–1261, 2016.
- [33] Americo Oliveira, Sérgio Pereira, and Carlos A Silva. Augmenting data when training a cnn for retinal vessel segmentation: How to warp? In *2017 IEEE 5th Portuguese Meeting on Bioengineering (ENBENG)*, pages 1–4. IEEE, 2017.
- [34] Sahin Olut, Zhengyang Shen, Zhenlin Xu, Samuel Gerber, and Marc Niethammer. Adversarial data augmentation via deformation statistics. In *European Conference on Computer Vision*, pages 643–659. Springer, 2020.
- [35] Cheng Ouyang, Carlo Biffi, Chen Chen, Turkey Kart, Huaqi Qiu, and Daniel Rueckert. Self-supervision with superpixels: Training few-shot medical image segmentation without annotation. In *European Conference on Computer Vision*, pages 762–780. Springer, 2020.
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [37] Sérgio Pereira, Adriano Pinto, Victor Alves, and Carlos A Silva. Brain tumor segmentation using convolutional neural networks in mri images. *IEEE transactions on medical imaging*, 35(5):1240–1251, 2016.
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [39] Zhengyang Shen, Zhenlin Xu, Sahin Olut, and Marc Niethammer. Anatomical data augmentation via fluid-based image registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 318–328. Springer, 2020.
- [40] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017.
- [41] Devavrat Tomar, Manana Lortkipanidze, Guillaume Vray, Behzad Bozorgtabar, and Jean-Philippe Thiran. Self-attentive spatial adaptive normalization for cross-modality domain adaptation. *IEEE Transactions on Medical Imaging*, 2021.
- [42] Devavrat Tomar, Lin Zhang, Tiziano Portenier, and Orcun Goksel. Content-preserving unpaired translation from simulated to realistic ultrasound images. *arXiv preprint arXiv:2103.05745*, 2021.
- [43] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9197–9206, 2019.
- [44] Shuxin Wang, Shilei Cao, Dong Wei, Renzhen Wang, Kai Ma, Liansheng Wang, Deyu Meng, and Yefeng Zheng. Lt-net: Label transfer by learning reversible voxel-wise correspondence for one-shot medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9162–9171, 2020.
- [45] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12275–12284, 2020.
- [46] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee, 2003.
- [47] Zhirong Wu, Yuanjun Xiong, X Yu Stella, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [48] Zhenlin Xu and Marc Niethammer. Deepatlas: Joint semi-supervised learning of image registration and segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 420–429. Springer, 2019.

- [49] Heran Yang, Jian Sun, Huibin Li, Lisheng Wang, and Zongben Xu. Neural multi-atlas label fusion: Application to cardiac mr images. *Medical image analysis*, 49:60–75, 2018.
- [50] Amy Zhao, Guha Balakrishnan, Fredo Durand, John V Guttag, and Adrian V Dalca. Data augmentation using learned transformations for one-shot medical image segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8543–8553, 2019.