

A Simple and Robust Framework for Cross-Modality Medical Image Segmentation applied to Vision Transformers

Matteo Bastico¹, David Ryckelynck¹, Laurent Corté¹, Yannick Tillier², Etienne Decencière³
Mines Paris, Université PSL

¹Centre des Matériaux (MAT), UMR7633 CNRS, 91003 Evry, France

²Centre de Mise en Forme des Matériaux (CEMEF), UMR7635 CNRS, 06904 Sophia Antipolis, France

³Centre de Morphologie Mathématique (CMM), 77300 Fontainebleau, France

{matteo.bastico, etienne.decenciere}@minesparis.psl.eu

Abstract

When it comes to clinical images, automatic segmentation has a wide variety of applications and a considerable diversity of input domains, such as different types of Magnetic Resonance Images (MRIs) and Computerized Tomography (CT) scans. This heterogeneity is a challenge for cross-modality algorithms that should equally perform independently of the input image type fed to them. Often, segmentation models are trained using a single modality, preventing generalization to other types of input data without resorting to transfer learning techniques. Furthermore, the multi-modal or cross-modality architectures proposed in the literature frequently require registered images, which are not easy to collect in clinical environments, or need additional processing steps, such as synthetic image generation. In this work, we propose a simple framework to achieve fair image segmentation of multiple modalities using a single conditional model that adapts its normalization layers based on the input type, trained with non-registered interleaved mixed data. We show that our framework outperforms other cross-modality segmentation methods, when applied to the same 3D UNet baseline model, on the Multi-Modality Whole Heart Segmentation Challenge. Furthermore, we define the Conditional Vision Transformer (C-ViT) encoder, based on the proposed cross-modality framework, and we show that it brings significant improvements to the resulting segmentation, up to 6.87% of Dice accuracy, with respect to its baseline reference. The code to reproduce our experiments and the trained model weights are available at <https://github.com/matteo-bastico/MI-Seg> and in the Supplementary Code.

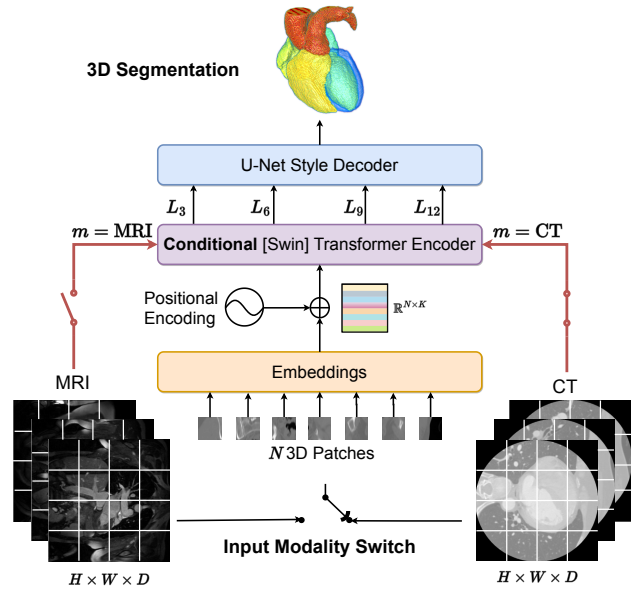


Figure 1: Overview of the proposed cross-modality clinical image segmentation technique applied to the Swin-UNETR segmentation model [16]. The input modality can be arbitrarily switched to obtain the desired segmentation, while keeping the model unchanged. In conditional models, the modality is also used as input to activate the encoder normalizations corresponding to the data type and generate common latent spaces.

1. Introduction

Medical image segmentation is one of the most challenging and studied research fields of deep learning and artificial intelligence in recent years [42, 45]. The precision of segmentation achieved competitive results thanks to recent efforts in developing convolution [32, 38] and transformer-based neural networks [16, 17]. These algorithms are gen-

erally trained on a single medical imaging modality, such as T1- or T2-Weighted Magnetic Resonance Images (MRIs) or Computed Tomography (CT) scans. Therefore, they often suffer from data variability when tested on images different from those seen during training. Data variability is a common issue in the medical field and it is due to several uncontrollable factors that depend on the specific clinic or patient needs [44], e.g., different imaging methods, different scanners, different acquisition settings or patient limitations. In this scenario, the most straightforward approach would be to train a separate model for each domain, requiring a massive amount of different annotated training data and not exploiting inter-domain information.

Recently, to tackle this issue, some multi-modal [55] and cross-modality [51, 27, 60, 63] medical image segmentation techniques have been proposed. The first group of methods aims to produce better segmentation by simultaneously exploiting information from multiple sources. In this case, it is good practice to stack several images of different types to generate a combined input for the segmentation model [62]. Furthermore, multi-modality can also be leveraged to produce synthetic images of missing domains at inference time [7]. Here, different modalities are rather seen as separate inputs for the model that generates a common representation and produces synthetic images. However, in both scenarios, registered medical images [19] are needed. Although registration is relatively simple, the acquisition of several images of the same patient is often limited by local resources and time constraints.

The second group of methods, i.e. cross-modality methods, focus on producing segmentation independently of the data type provided in input. More specifically, an *assistant modality*, of which several annotated data are available, is often used to improve the segmentation performance on another modality with fewer annotated data, i.e. *target modality*. Fine-tuning [47] or transfer learning [37] are straightforward ways of adapting deep learning models to target modalities after they have been trained on some specific assistant modality. However, these approaches do not fully exploit shared cross-modality information. One intuitive way to overcome this problem is joint training [51], where a model is trained simultaneously with multiple modalities. Nevertheless, it is relatively hard to directly learn common features when the domain shift is significant, e.g MRIs and CTs. Therefore, to directly learn from multi-modality data with large appearance discrepancy, the use of several feature extractors, i.e., one per modality, has been first proposed in an X-shaped architecture [51], implying a significant overhead on the segmentation model and requiring adjustment to generalize to other clinical tasks. To address this issue, several techniques based on synthetic image generation prior to segmentation have been introduced [6, 26, 63]. In many cases, the assistant modality is used to synthesize

the target, or vice versa, and train the model on a mix of real and synthetic data [27]. The additional pre-processing step required by these methods to make the clinical images segmentation-ready adds computational complexity, which is often critical in real-time applications such as computer-assisted surgery [48].

In this work, we propose a general framework, applicable to any encoder-decoder literature architecture, aiming to produce high-quality cross-modality segmentation without introducing overhead on the segmentation model nor needing registered clinical images for training. With our framework, medical images from different domains can be fed individually directly into a single modality-conditioned model, which generates the desired segmentation by self-adapting its encoder normalization layers. The adaptation is based on Conditional Instance Normalization (CIN), which has been originally proposed for the transfer of arbitrary artistic styles by Dumoulin *et al.* [10]. Thanks to that, and similarly to joint training, our models can be trained in an end-to-end way by randomly mixing different data of several modalities, in what we call *interleaved mixed training* fashion, without prior synthetic image style transfer. Furthermore, based on the proposed framework, we formally define the Conditional Vision Transformer (C-ViT) encoder architecture to build modality-agnostic ViT-based models [9] for image segmentation or classification, as shown in Figure 1. We extensively evaluate our method on the Multi-Modality Whole Heart Segmentation (MM-WHS) 2017 Challenge [64]. We show that C-ViT improves the baseline model and other multi-modal techniques and that our framework outperforms other cross-modality segmentation methods, when applied to the same baseline model. Our contributions are summarized as follows:

- We present a novel framework to enhance the quality of cross-modality segmentation, while removing expensive input pre-processing steps and the need of registered data for training.
- Following the framework, we introduce the Conditional Vision Transformer encoder, which adapts its normalization layers based on the input modality and can be used for several tasks.
- We evaluate our framework on the public whole heart and substructures segmentation challenge [64], showing that it outperforms other cross-modality learning techniques both on UNet and ViT baselines.

2. Related Works

Multi-Modality Learning in Medical Imaging. In recent years, several deep learning architectures have been proposed for image segmentation, achieving extremely good performance [5, 16, 17, 33, 40]. Among them, UNet [38] is surely the most popular and is used as a baseline to create better performing models [42]. More recently,

some variation based on vision transformers (ViT) [9] and swin-transformers [30], such as TransUNET [5], UNETR [17] and Swin-UNETR [16], have been shown to outperform previous versions of UNet. In the clinical imaging field, multi-modality learning is a similar, yet challenging, task. Multiple medical image domains have been exploited for synthetic image generation based on Generative Adversarial Networks (GAN) [7, 8, 14, 46, 58] or multi-modal image segmentation [12, 15, 55, 62]. Interestingly, the two tasks have also been merged by Zhang *et al.* [59] in a cross-task feedback fusion GAN that first generates synthetic CT and then performs multi-modal segmentation. The latter uses cross-domain information to achieve better performances, but it is restricted to when registered images are available. To relax the need of aligned medical images and exploit inter-domain features during training, cross-modality segmentation approaches have recently gained attraction. Firstly, Zheng *et al.* [61] investigated the effectiveness of shape priors learned from an assistant modality to improve the segmentation on a target modality by marginal space learning. Valindria *et al.* [51] developed dual-stream encoder-decoder models having different branches for each modality and implementing weights sharing techniques to extract cross-modality features. Similarly to our work, some researcher tried to play with the normalization layers of the models to improve their generalization capacities [11, 35, 39, 41, 63]. Among others, Pan *et al.* [35] introduced the IBN-Net which simultaneously exploiting Instance and Batch Normalization to capture both appearance changes and content information. Segu *et al.* [39] proposed to collect prior domain-dependent statistics by training ad-hoc Batch Normalization (BN) to map the modalities on a shared latent space.

With the aforementioned advances in synthetic images generation, many works tried to tackle the challenge by assisting segmentation models with prior image translation [6, 21, 26, 27, 60, 63]. Among them, Zhang *et al.* [60] tried to improve the segmentation for modalities with limited training samples by easing the appearance gap with other modalities using a GAN. Similarly, Li *et al.* [27] proposed an Image Alignment Module to reduce the appearance gap between assistant and target modality and introduced a Mutual Knowledge Distillation scheme [20] to exploit modality-shared knowledge. More recently, Zhou *et al.* [63] proposed to simulate the possible appearance changes of a target domains by non-linear transformation to augment source-similar and source-dissimilar images.

Different from these methods, which mainly rely on image translation or prior trainings, our framework tries to shrink the model complexity and, at the same time, improve the segmentation accuracy extracting all the available cross-modality features by conditioning the segmentation models and training them jointly with assistant and target modality.




Figure 2: Overview of the proposed conditional framework: the conditional encoder, $E(\cdot, m)$, generates S common latent spaces, $\{C^1, \dots, C^S\}$, one for each encoder-decoder connection, adapting itself to the input modality m and the decoder, $D(\cdot)$, is unique and does not need adaptation.

Moreover, many of the previous techniques do not allow for arbitrary switch of the input modality, but they just exploit assistant modality information to improve the target segmentation.

Normalization Layers. Our work is also related with normalization layers used in neural networks, which may differ based on the spatial dimension to which they are applied. The four main classes of normalization layers are batch, layer, group, and instance normalization. Batch Normalization was the first technique introduced to deal with the variability of layers input during training by normalizing the features across mini-batches [23]. Layer Normalization (LN) [2], used mainly in Natural Language Processing (NLP) and Recurrent Neural Networks (RNNs), was introduced to relax the dependency on the batch, normalizing only across features. Similarly, Group Normalization [57] computes the mean and standard deviations in groups of channels rather than using the entire feature space. Finally, Instance Normalization (IN) has been introduced, in the imaging context, to prevent instance-specific mean and covariance shift, simplifying the learning process [50]. Huang *et al.* [22] proposed a further variation of IN, called Adaptive IN to perform arbitrary image style transfer, that is, synthetic image generation, by adjusting the mean and variance of the content input to match those of the desired style input. The latter does not have learnable parameters, but it computes them directly from the input style. In our framework, we propose to use a similar variation of IN recently proposed by Dumoulin *et al.* [10], called Conditional Instance Normalization (CIN), which learns a different set of parameters γ_m and β_m for each style $m \in [1, M]$. Given a three-dimensional volume with C channels, $\mathbf{z} \in \mathbb{R}^{C \times H \times W \times D}$, of modality m , the CIN is calculated as

$$\text{CIN}(\mathbf{z}, m) = \gamma_m \left(\frac{\mathbf{z} - \mu(\mathbf{z})}{\sigma(\mathbf{z})} \right) + \beta_m \quad (1)$$

where $\mu(\mathbf{z})$ and $\sigma(\mathbf{z})$ are computed across spatial dimensions independently for each channel and each sample [22]. It allows the sharing of all convolutional weights of a style

transfer network across many image styles by only changing the learnable normalization parameters. Inspired by that, in this work, CIN is exploited to generate a shared latent space for different input modalities, which is used to obtain the segmentation output independently of the medical input image domain.

3. Method

The overview of the proposed cross-modality segmentation framework is shown in Figure 2. The segmentation model, which can be chosen among those available in the literature, is conditioned by the input data modality to adapt its normalization layers to the fed clinical image. Therefore, the aim is to produce the desired segmentation independently of the input data domain.

In the following, the general framework for creating conditional models is introduced and it is applied to the transformer encoder, to obtain the C-ViT encoder. Finally, more details about the framework training procedure, including the loss function, are reported.

Conditional Segmentation Models. Let any state-of-the-art encoder-decoder segmentation architecture be considered as baseline for the proposed framework. In a straightforward case, for a given data modality $m \in [1, M]$, let $E_m(\cdot) : \mathcal{I}_m \rightarrow \mathcal{H}_m$ and $D_m(\cdot) : \mathcal{H}_m \rightarrow \mathcal{O}$ be the two functions that describe the encoder and decoder modules of the chosen architecture, respectively. The role of $E_m(\cdot)$ is to transform the input, $\mathbf{x}_m \in \mathcal{I}_m$, into a latent feature space, \mathcal{H}_m , which can better represent the information carried by raw data. On the other hand, the purpose of $D_m(\cdot)$ is to transform this representation into the final segmentation $\mathbf{y} \in \mathcal{O}$. In general, the space \mathcal{H}_m is learned ad-hoc by the model, based on the type of data used during training.

To make the model independent of the input modality, we propose to create a general conditional encoding module, $E(\cdot, m) : \mathcal{I}_m \rightarrow \mathcal{C}$, which transforms any input $\mathbf{x}_m \in \mathcal{I}_m$, $1 \leq m \leq M$, into a common shared latent feature space \mathcal{C} . In this way, a single decoder $D(\cdot) : \mathcal{C} \rightarrow \mathcal{O}$ is needed to generate the segmentation of all the modalities. The relation between any input \mathbf{x}_m and the output \mathbf{y} in a conditional model can then be expressed as

$$\mathbf{y} = D(E(\mathbf{x}_m, m)), \quad 1 \leq m \leq M \quad (2)$$

This concept is easily extended to architectures with skip connections, e.g. UNet [38], in which the encoder may be seen as a sequence of S sub-modules, each of them generating a common latent space \mathcal{C}^s , $0 \leq s \leq S$, where S is also the number of connections between encoder and decoder. Therefore, the decoder will simply gather the data from all the shared latent spaces to generate a modality-agnostic segmentation output.

In practice, we propose to generate the conditional encoder $E(\cdot, m)$, starting from any simple encoder $E_m(\cdot)$,




Figure 3: Conditional Vision Transformer Encoder. The multi-head self-attention (MSA) mechanism is followed by a MLP and each of them is preceded by CIN. The latter consists of a switch mechanism to activate the IN corresponding the to modality flag in input. A focus on the different normalizing direction of IN and LN is also shown in the zoom.

by adding the data modality as input (often available from meta-data, without manual intervention) and replacing all its normalization layers with CIN. In such way, the only overhead on the segmentation model is given by the extra normalization learnable parameters γ_m and β_m , $1 \leq m \leq M$, of Equation 1, which are negligible with respect to the rest of the architecture. Furthermore, we avoid any pre-processing phase, such as synthetic image generation through GANs, which can heavily impact on the network performances, and we let the model directly learn a common shared latent space for all the input modalities, trying to fully exploit the inter-modality information.

Conditional Vision Transformer. The proposed C-ViT encoder is shown in Figure 3. As originally proposed by Vaswani *et al.* [52], it has two sub-layers, the first is a multi-head self-attention mechanism (MSA), and the second is a simple, position-wise fully connected Multi-Layer Perceptron (MLP). Both sub-layers are preceded by a normalization [9, 53] and a residual connection is employed around each of them. Here, we replace the Layer Normalization (LN) [2], which has historically been used in transformers,

Table 1: Quantitative comparison with other methods for cross-modality medical image segmentation on the target modality (CT). All the techniques have the same UNet baseline [51], are trained using 20 MRIs and 10 CTs and are evaluated on the test set of MM-WHS dataset. The mean DICE score is reported, as well as the ones for all the heart substructures.

Model	Avg Dice	Dice of substructures of heart						
		MYO	LA	LV	RA	RV	AA	PA
Baseline	87.06	87.02	89.22	90.86	83.86	84.60	92.52	81.34
Fine-tune	87.69	87.16	90.40	90.79	84.43	85.26	92.74	83.05
Joint-training	87.43	86.65	90.76	91.23	82.78	84.92	93.02	82.66
X-shape [51]	87.67	87.19	89.79	90.94	85.51	84.44	93.43	82.40
Zhang <i>et al.</i> [60]	88.50	87.81	91.12	91.34	85.14	86.31	94.30	83.42
Li <i>et al.</i> [27]	90.12	89.34	91.90	92.67	87.47	88.14	95.95	85.38
Ours	90.77	90.06	92.68	93.77	88.22	90.85	94.70	84.52

with CIN. In other words, a switch mechanism that selects the IN corresponding to the input modality is introduced before each sub-layer.

In a general scenario with M different input modalities, let L , N and K be the number of stacked transformer blocks that make up a transformer encoder, the number of input patches and the dimensions of the embedding space, respectively. For a given input embedding of the i -th transformer block, $\mathbf{z}_{i-1} \in \mathbb{R}^{N \times K}$, we have

$$\hat{\mathbf{z}}_i = \text{MSA}(\text{CIN}(\mathbf{z}_{i-1}, m)) + \mathbf{z}_{i-1}, \quad 1 \leq i \leq L \quad (3)$$

$$\mathbf{z}_i = \text{MLP}(\text{CIN}(\hat{\mathbf{z}}_i, m)) + \hat{\mathbf{z}}_i, \quad 1 \leq i \leq L \quad (4)$$

where $1 \leq m \leq M$ is the modality corresponding to the input and $\text{CIN}(\cdot, \cdot)$ is defined in Equation 1. Note that \mathbf{z}_0 corresponds to the linear projection of the input patches, i.e. non-overlapping sub-patches of dimension $P \times P \times P$ of the input volume, $\{\mathbf{p}_j \in \mathbb{R}^{P^3} : 1 \leq j \leq N\}$, into a K -dimensional embedding space using a projection matrix $\mathbf{E} \in \mathbb{R}^{P^3 \times K}$ [17]. Additionally, learnable positional embeddings, $\mathbf{E}_{\text{pos}} \in \mathbb{R}^{N \times K}$, are included to code spatial information, finally obtaining

$$\mathbf{z}_0 = [\mathbf{p}_1 \mathbf{E}; \dots; \mathbf{p}_J \mathbf{E}] + \mathbf{E}_{\text{pos}}. \quad (5)$$

The rationale for substituting LN, which is inherited from the applications of the first transformers in NLP [56], for CIN, or more generally IN, is also motivated on the zoom of Figure 3. LN computes the average and variance of the single input patches along the K dimensions of the embedding space. Therefore, when it comes to images, the latter produces an independent normalization for each image patch without considering the entire input data. IN addresses this issue by computing the metrics for normalization across all the input patches but separately for each embedding dimension. Intuitively, the model is able to generate more meaningful embeddings when the input volume is considered as a whole for normalization.

The technique presented here can also be generalized to the *Swin-transformer* model by replacing the MSA of Equation 3 with W-MSA or SW-MSA, denoting window-based multi-head self-attention using regular and shifted window partitioning configurations, respectively [30]. Finally, in this case, to have a fully conditional encoder, also the patch merging block, used to generate hierarchical representations of the input patches, must be adapted in the normalizations.

Loss function and Training. To train cross-modality segmentation models using our framework, we used a combination of Dice Loss [25], $\mathcal{L}_{\text{Dice}}$, and Focal Loss [29], $\mathcal{L}_{\text{Focal}}$, as

$$\mathcal{L} = \lambda_D \mathcal{L}_{\text{Dice}} + \lambda_F \mathcal{L}_{\text{Focal}} \quad (6)$$

where λ_D and λ_F are scalar factors greater than 0 to adjust the contribution of each individual loss to the total. We preferred this version of the loss, rather than the standard combination of Dice loss and Cross-Entropy (CE), to tackle the unbalancing of clinical image segmentation, which is often significant between background and foreground. The focal loss, instead, reshapes the standard CE loss so that it down-weights the loss assigned to well-classified examples. In this way, more weight will be given in the loss to hard classified samples, i.e., labels difficult to segment, improving the learning process.

Differently from joint training [51], in which training data is divided into batches of different modalities, conditional models are trained using an interleaved mixed data training fashion. This means that different modalities are randomly fed into the model during training, even in the same batch. Therefore, the overall loss will provide an overview of the goodness of the segmentation independently from the input modality, pushing the network to be fair with all different types of data.

4. Experiments

In this section we describe the experiments carried out to evaluate the proposed cross-modality framework with in-

Table 2: Performance comparison on the target modality (CT) of our framework applied to the Swin-UNETR model [16]. All the models are trained using 20 MRI and 10 CT volumes and are evaluated on the test set of MM-WHS dataset. The mean DICE score is reported, as well as the ones for all the heart substructures and Whole Heart (WH) segmentation.

Model	WH	Avg Dice	Dice of substructures of heart						
			MYO	LA	LV	RA	RV	AA	PA
Baseline	83.11	85.32	84.79	88.12	89.80	81.04	84.08	77.58	76.36
Fine-tune	86.33	86.09	82.57	90.03	87.92	83.29	85.51	90.28	83.06
Joint-training	89.03	87.99	86.25	92.07	91.97	85.35	87.98	89.55	82.76
C-ViT [ours]	89.98	89.33	88.13	91.64	92.39	86.36	89.33	93.39	84.05

Table 3: Performance comparison on MRIs, used as assistant modality, of our framework applied to the Swin-UNETR model [16]. The baseline is the model used to do fine-tuning on the target modality.

Model	WH	Mean Dice
Baseline	83.32	81.87
Joint-training	86.65	84.82
C-ViT [ours]	85.57	84.30

terleaved mixed data training. The experiment settings are described, including the dataset, the implementation details, and the algorithms to tune the models hyper-parameters. Finally, the results are shown and discussed.

4.1. Experiment settings

Dataset. We evaluate the proposed method on the Multimodality Whole Heart Segmentation Challenge 2017 (MM-WHS 2017) dataset [64], which contains non-registered 20 MRI and 20 CT volumes for training and the ground-truth (GT) annotations of 7 cardiac substructures including the Left Ventricle blood cavity (LV), the Right Ventricle blood cavity (RV), the Left Atrium blood cavity (LA), the Right Atrium blood cavity (RA), the MYOcardium of the left ventricle (MYO), the Ascending Aorta (AA), and the Pulmonary Artery (PA). Following previous cross-modality segmentation works on this dataset [27], we used MRI as assistant modality and CT as target modality, since MRI has better contrast for soft tissue and ideally provide better information for heart substructures segmentation. We evenly and randomly split the CT data in order to perform a two-fold cross-validation and train the conditional models with 20 MRIs and 10 CTs at a time, to simulate data shortage for the target modality. Furthermore, MRI volumes are pre-processed using the N4 bias field correction algorithm to correct low frequency intensity non-uniformity present in MRI image data known as gain field [49]. Finally, all the volumes are re-sampled to an isotropic space and their values are normalized between 0 and 1.

Implementation Details. To provide a fair comparison,

we first applied our framework to the same baseline used by other cross-modality previous works [27, 51]. Namely, we implemented the network architecture with encoder-decoder structure proposed by Valindria *et al.* [51], consisting of a 3D UNet [38] with residual blocks with pre-activation [18]. After that, we evaluated the efficiency of the proposed C-ViT encoder applied to the Swin-UNETR [16] segmentation model, originally consisting of a Swin-transformer encoder [30] and UNet style decoder, as shown in Figure 1. For all the models, we define an input size of $96 \times 96 \times 96$ pixels and sub-patches, in the case of transformer-based architectures, of $16 \times 16 \times 16$ pixels. Therefore, at each epoch, a fixed number of random crops of input size is extracted, for each volume. Note that, in this way, the effective batch size is given by $N_s \times N_c$, where N_s and N_c are the number of samples and crops, respectively. For the transformer-based models, we follow the original implementation of Hatamizadeh *et al.* [16], with $L = 8$ blocks and hidden size $K = 768$ in the encoder. Furthermore, data augmentation is applied to artificially increase training data and deal with inter-modality variations by randomly shifting and scaling the intensity of the input volumes, along with random flips in all three dimensions. Finally, to validate and test the models, sliding window inference is used with an overlap of 0.5.

The code is implemented in PyTorch, with the help of the MONAI library [4], and trained on a cluster equipped with multiple NVIDIA V100 GPUs. Optimization is performed using Adam with decoupled weight decay [31] and a cosine-decay Learning Rate (LR) scheduler with warm-up, for a maximum of 2500 epochs. All the models are tuned to find the best set of hyper-parameters as described below.

Hyper-parameters Tuning. To take full advantage of the segmentation models, we tuned them using the Optuna framework [1], following the guidelines from [13]. Therefore, the batch size is not tuned but is kept unchanged among trials as the maximum that can fit our memory availability, that is, 16 three-dimensional crops ($N_s = 4$ and $N_c = 4$). The optimal LR is found for each model using the cyclical learning rate technique [43] and a logarithmic search space centered on that value is defined for tuning.




Figure 4: Qualitative comparison of the segmentation result on the CT target domain of our framework applied to the Swin-UNETR model [16] on the MM-WHS dataset. Our framework is compared with the baseline model, trained on the target modality, fine-tuning and joint-training with interleaved mixed fashion.

The number of LR Warm-Up epochs is chosen among the values of 2, 4 and 6% of the maximum training epochs [24] and the optimizer weight decay is also tuned. Finally, the loss function of Equation 6 is used with $\lambda_D = \lambda_F = 1$.

For each trial, the hyper-parameters are sampled from the search space using the Tree-structured Parzen Estimator (TPE) algorithm [3, 34]. This technique aims to find the best set of parameters by fitting two Gaussian Mixture Models (GMMs) to the good and bad trials and minimizing their ratio. Furthermore, Asynchronous Successive Halving (ASHA) [28] is used to early-stop bad trials and save computational time and resources.

4.2. Experiment results

We report here the results of our experiments using the proposed framework on the MM-WHS dataset [64], quantitative using the Dice similarity coefficient and qualitative by comparing segmentation images. In the following, we first compare our proposal with other techniques and then we evaluate the accuracy of the proposed C-ViT.

Comparison with Other Methods. The comparative quantitative results for heart substructures segmentation are shown in Table 1. In fine-tuning the model is first trained using the assistant modality only and then, the acquired

knowledge is transferred to the target modality with a subsequent training. In joint-training the baseline model is trained with both modalities simultaneously by alternating them on different batches. We also compare the accuracy with the X-shape architecture proposed by Valindria *et al.* [51] and the online training with synthetic image generation approach of Zhang *et al.* [60]. Finally, a knowledge distillation cross-modality segmentation technique is also considered for comparison [27]. As highlighted by the latter, fine-tuning, joint-training and the X-shaped architecture bring marginal improvements on the segmentation accuracy since they do not fully exploit cross-modality information. Online GAN-based synthetization with and without mutual knowledge distillation start to significantly improve the mean accuracy of target modality segmentation, up to 3.06%. Nevertheless, in both cases, significant overhead is introduced on the model, limiting its possible real-time applications. Our method achieves significant improvements in the segmentation accuracy, while reducing the complexity of both training and inference. We observe an enhancement of 0.65% in the mean Dice and up to 2.71% on the single heart substructures, namely in the RV. However, a small decrease in performances is registered for AA and PA, probably due to the fact that the manual segmentation of these




Figure 5: Comparison of the 3D segmentation of heart substructures from CT using different cross-modality adaptations.




Figure 6: Qualitative comparison of the segmentation result, on the MRI assistant domain, of our framework applied to the Swin-UNETR model [16] on the MM-WHS dataset.

structures generally covers beyond the real size of the vessels while the test tool cuts them to a limited length [64].

Evaluation of C-ViT and Ablation Study. We compare the performance of the C-ViT-based segmentation model with its baseline [16] and the result obtained by fine-tuning and joint training the same model (with interleaved mixed fashion). Table 2 summarizes the quantitative results in terms of Dice accuracy of Whole Heart Segmentation (WHS) and single substructures segmentation. Furthermore, Figure 4 shows a qualitative comparison on single slices of the segmentation obtained on CT by using different methods, and Figure 5 directly compares the 3D segmentation of a whole sample. Our framework enhances of 4% the mean Dice of heart substructures, with significant improvements in each of them, and 6.87% the accuracy of WHS, with respect to the baseline model. Qualitative, we can see important improvements when using C-ViT with respect to other methods, that may also result in refinement of the GT (e.g. the segmentation of PA on the second row of Figure 4). In the 3D comparison, we observe the improvement in the segmentation produced by C-ViT which does not contain false positives outside the region of the GT. Although in our cross-validation the ViT-based conditional model performs

better than the UNet one ($91.18\% \pm 0.68$ vs $90.76\% \pm 0.55$ of mean Dice), in the test set, the latter seems to be slightly better likely due to the fact that transformers usually need more data for finer generalization [36].

Differently from some other related works [26, 27] our framework do not just exploit assistant modality to improve the segmentation of the target one, but it keeps the ability of producing high-quality segmentation for the assistant modality by switching the input domain. To confirm that, In Table 3, WHS and Mean Dice of MRI segmentation is reported, where the baseline is considered to be the model trained only with assistant modality and used to fine-tune the target modality. Also the assistant modality segmentation accuracy benefits of our framework, as well as of joint-training, improving with respect to the baseline model. Moreover, Figure 6 also shows some qualitative improvements of the MRI segmentation, especially from baseline to C-ViT. As previously mentioned, transformers are "data-hungry" [54], in the sense that they suffer from significant performance drops on small-size datasets. We have shown that, with our conditional framework applied to ViT, it is possible to relax the need of a huge dataset for training by gathering several small datasets of different modalities.

5. Conclusions

In this work, we proposed a simple framework that aims to segment different types of medical images, while reducing the model overhead and the need of registered data for training, using a single cross-modality conditional model trained with interleaved mixed data. We introduced a general definition of a conditional model, based on CIN, that can be applied to all state-of-the-art medical image segmentation architectures. Furthermore, we developed a new C-ViT encoder that can be used to create conditional ViT-based models. We validated our framework using a public multi-modal dataset for heart substructures segmentation, which collects both MRI and CT data. Our framework achieves new state-of-the-art performance for the cross-modality medical image segmentation with assistant and target modality. We have shown that not only the first one helps the learning on the latter when few annotated data are available, but also the segmentation of the assistant modality benefits of our framework. We believe that our work can be an important contribution for developing robust cross-modality medical image segmentation methods, to tackle cases in which the acquisition of a specific kind of image is limited by environmental or patient-specific restrictions. Furthermore, in the future, the proposed framework may be extended in an Unsupervised Domain Adaptation (UDA) fashion, to make possible the training on a single labeled modality with adaptation to unlabeled domains.

6. Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 945304-Cofund AI4theSciences hosted by PSL University.

This work was granted access to the HPC/AI resources of IDRIS under the allocation 2022-AD011013902 made by GENCI.

References

- [1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A Next-generation Hyperparameter Optimization Framework, July 2019. arXiv:1907.10902 [cs, stat].
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer Normalization, July 2016. arXiv:1607.06450 [cs, stat].
- [3] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS’11*, pages 2546–2554, Red Hook, NY, USA, Dec. 2011. Curran Associates Inc.
- [4] M. Jorge Cardoso, Wenqi Li, Richard Brown, Nic Ma, Eric Kerfoot, Yiheng Wang, Benjamin Murrey, Andriy Myronenko, Can Zhao, Dong Yang, Vishwesh Nath, Yufan He, Ziyue Xu, Ali Hatamizadeh, Andriy Myronenko, Wentao Zhu, Yun Liu, Mingxin Zheng, Yucheng Tang, Isaac Yang, Michael Zephyr, Behrooz Hashemian, Sachidanand Alle, Mohammad Zalbani Darestani, Charlie Budd, Marc Modat, Tom Vercauteren, Guotai Wang, Yiwen Li, Yipeng Hu, Yunguan Fu, Benjamin Gorman, Hans Johnson, Brad Genereaux, Barbaros S. Erdal, Vikash Gupta, Andres Diaz-Pinto, Andre Dourson, Lena Maier-Hein, Paul F. Jaeger, Michael Baumgartner, Jayashree Kalpathy-Cramer, Mona Flores, Justin Kirby, Lee A. D. Cooper, Holger R. Roth, Daguang Xu, David Bericat, Ralf Floca, S. Kevin Zhou, Haris Shuaib, Keyvan Farahani, Klaus H. Maier-Hein, Stephen Aylward, Prerna Dogra, Sebastien Ourselin, and Andrew Feng. MONAI: An open-source framework for deep learning in healthcare, Nov. 2022. arXiv:2211.02701 [cs].
- [5] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L. Yuille, and Yuyin Zhou. TransUNet: Transformers Make Strong Encoders for Medical Image Segmentation, Feb. 2021. arXiv:2102.04306 [cs].
- [6] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. CrDoCo: Pixel-level Domain Transfer with Cross-Domain Consistency, Jan. 2020. arXiv:2001.03182 [cs].
- [7] Onat Dalmaç, Mahmut Yurt, and Tolga Çukur. ResViT: Residual vision transformers for multi-modal medical image synthesis. *IEEE Transactions on Medical Imaging*, 41(10):2598–2614, Oct. 2022. arXiv:2106.16031 [cs, eess].
- [8] Salman UH. Dar, Mahmut Yurt, Levent Karacan, Aykut Erdem, Erkut Erdem, and Tolga Cukur. Image Synthesis in Multi-Contrast MRI With Conditional Generative Adversarial Networks. *IEEE Transactions on Medical Imaging*, 38(10):2375–2388, Oct. 2019.
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, June 2021. arXiv:2010.11929 [cs].
- [10] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A Learned Representation For Artistic Style, Feb. 2017. arXiv:1610.07629 [cs] version: 5.
- [11] Xinjie Fan, Qifei Wang, Junjie Ke, Feng Yang, Boqing Gong, and Mingyuan Zhou. Adversarially Adaptive Normalization for Single Domain Generalization, June 2021. arXiv:2106.01899 [cs].
- [12] Lucas Fidon, Wenqi Li, Luis C. Garcia-Peraza-Herrera, Jinnendra Ekanayake, Neil Kitchen, Sebastien Ourselin, and Tom Vercauteren. Scalable Multimodal Convolutional Networks for Brain Tumour Segmentation. In Maxime Descoteaux, Lena Maier-Hein, Alfred Franz, Pierre Jannin, D. Louis Collins, and Simon Duchesne, editors, *Medical Image Computing and Computer Assisted Intervention - MICCAI 2017*, Lecture Notes in Computer Science, pages 285–293, Cham, 2017. Springer International Publishing.
- [13] Varun Godbole, George E. Dahl, Justin Gilmer, Christopher J. Shallue, and Zachary Nado. Deep Learning Tuning Playbook, Feb. 2023. original-date: 2023-01-18T23:32:32Z.
- [14] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks, June 2014. arXiv:1406.2661 [cs, stat].
- [15] Zhe Guo, Xiang Li, Heng Huang, Ning Guo, and Quanzheng Li. Medical Image Segmentation Based on Multi-Modal Convolutional Neural Network: Study on Image Fusion Schemes. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 903–907, Apr. 2018. arXiv:1711.00049 [cs].
- [16] Ali Hatamizadeh, Vishwesh Nath, Yucheng Tang, Dong Yang, Holger Roth, and Daguang Xu. Swin UNETR: Swin Transformers for Semantic Segmentation of Brain Tumors in MRI Images, Jan. 2022. arXiv:2201.01266 [cs, eess].
- [17] Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, Bennett Landman, Holger Roth, and Daguang Xu. UNETR: Transformers for 3D Medical Image Segmentation, Oct. 2021. arXiv:2103.10504 [cs, eess].
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity Mappings in Deep Residual Networks, July 2016. arXiv:1603.05027 [cs].
- [19] Derek L G Hill, Philipp G Batchelor, Mark Holden, and David J Hawkes. Medical image registration. *Medical image registration*, June 2000.
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the Knowledge in a Neural Network, Mar. 2015. arXiv:1503.02531 [cs, stat].
- [21] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Dar-

- rell. CyCADA: Cycle-Consistent Adversarial Domain Adaptation, Dec. 2017. arXiv:1711.03213 [cs].
- [22] Xun Huang and Serge Belongie. Arbitrary Style Transfer in Real-time with Adaptive Instance Normalization, July 2017. arXiv:1703.06868 [cs].
- [23] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift, Mar. 2015. arXiv:1502.03167 [cs].
- [24] Peter Izsak, Moshe Berchansky, and Omer Levy. How to Train BERT with an Academic Budget, Sept. 2021. arXiv:2104.07705 [cs].
- [25] Shruti Jadon. A survey of loss functions for semantic segmentation. In *2020 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB)*, pages 1–7, Oct. 2020. arXiv:2006.14822 [cs, eess].
- [26] Jue Jiang, Yu-Chi Hu, Neelam Tyagi, Pengpeng Zhang, Andreas Rimner, Gig S. Mageras, Joseph O. Deasy, and Harini Veeraraghavan. Tumor-aware, Adversarial Domain Adaptation from CT to MRI for Lung Cancer Segmentation. *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, 11071:777–785, Sept. 2018.
- [27] Kang Li, Lequan Yu, Shujun Wang, and Pheng-Ann Heng. Towards Cross-modality Medical Image Segmentation with Online Mutual Knowledge Distillation, Oct. 2020. arXiv:2010.01532 [cs, eess].
- [28] Liam Li, Kevin Jamieson, Afshin Rostamizadeh, Ekaterina Gonina, Moritz Hardt, Benjamin Recht, and Amey Talwalkar. A System for Massively Parallel Hyperparameter Tuning, Mar. 2020. arXiv:1810.05934 [cs, stat].
- [29] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal Loss for Dense Object Detection, Feb. 2018. arXiv:1708.02002 [cs].
- [30] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows, Aug. 2021. arXiv:2103.14030 [cs].
- [31] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization, Jan. 2019. arXiv:1711.05101 [cs, math].
- [32] Priyanka Malhotra, Sheifali Gupta, Deepika Koundal, Atef Zaguia, and Wegayehu Enbeyle. Deep Neural Networks for Medical Image Segmentation. *Journal of Healthcare Engineering*, 2022:e9580991, Mar. 2022. Publisher: Hindawi.
- [33] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation, June 2016. arXiv:1606.04797 [cs].
- [34] Yoshihiko Ozaki, Yuki Tanigaki, Shuhei Watanabe, Masahiro Nomura, and Masaki Onishi. Multiobjective Tree-Structured Parzen Estimator. *Journal of Artificial Intelligence Research*, 73:1209–1250, Apr. 2022.
- [35] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at Once: Enhancing Learning and Generalization Capacities via IBN-Net, Mar. 2020. arXiv:1807.09441 [cs].
- [36] Namuk Park and Songkuk Kim. How Do Vision Transformers Work?, June 2022. arXiv:2202.06709 [cs].
- [37] Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding Transfer Learning for Medical Imaging. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation, May 2015. arXiv:1505.04597 [cs].
- [39] Mattia Segu, Alessio Tonioni, and Federico Tombari. Batch Normalization Embeddings for Deep Domain Generalization, May 2021. arXiv:2011.12672 [cs].
- [40] Sandeep Singh Sengara, Christopher Meulengrachtb, Mikael Ploug Boesemb, Anders Føhrby Overgaardb, Henrik Gudbergseb, Janus Damm Nybingb, and Erik Bjørnager Dam. UNet Architectures in Multiplanar Volumetric Segmentation – Validated on Three Knee MRI Cohorts, Mar. 2022. arXiv:2203.08194 [cs, eess].
- [41] Seonguk Seo, Yumin Suh, Dongwan Kim, Geeho Kim, Jongwoo Han, and Bohyung Han. Learning to Optimize Domain Specific Normalization for Domain Generalization, July 2020. arXiv:1907.04275 [cs, stat].
- [42] Nahian Siddique, Sidike Paheding, Colin P. Elkin, and Vijay Devabhaktuni. U-Net and Its Variants for Medical Image Segmentation: A Review of Theory and Applications. *IEEE Access*, 9:82031–82057, 2021. Conference Name: IEEE Access.
- [43] Leslie N. Smith. Cyclical Learning Rates for Training Neural Networks, Apr. 2017. arXiv:1506.01186 [cs].
- [44] Paul Suetens. *Fundamentals of Medical Imaging*. Cambridge University Press, Cambridge, 3 edition, 2017.
- [45] S. Suganyadevi, V. Seethalakshmi, and K. Balasamy. A review on deep learning in medical image analysis. *International Journal of Multimedia Information Retrieval*, 11(1):19–38, Mar. 2022.
- [46] Hongfei Sun, Qianyi Xi, Rongbo Fan, Jiawei Sun, Kai Xie, Xinye Ni, and Jianhua Yang. Synthesis of pseudo-CT images from pelvic MRI images based on an MD-CycleGAN model for radiotherapy. *Physics in Medicine & Biology*, 67(3):035006, Jan. 2022. Publisher: IOP Publishing.
- [47] Nima Tajbakhsh, Jae Y. Shin, Suryakanth R. Gurudu, R. Todd Hurst, Christopher B. Kendall, Michael B. Gotway, and Jianming Liang. Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning? *IEEE Transactions on Medical Imaging*, 35(5):1299–1312, May 2016. Conference Name: IEEE Transactions on Medical Imaging.
- [48] Leonardo Tanzi, Pietro Piazzolla, Francesco Porpiglia, and Enrico Vezzetti. Real-time deep learning semantic segmentation during intra-operative surgery for 3D augmented reality assistance. *International Journal of Computer Assisted Radiology and Surgery*, 16(9):1435–1445, Sept. 2021.
- [49] Nicholas J. Tustison, Brian B. Avants, Philip A. Cook, Yuanjie Zheng, Alexander Egan, Paul A. Yushkevich, and James C. Gee. N4ITK: Improved N3 Bias Correction. *IEEE transactions on medical imaging*, 29(6):1310–1320, June 2010.

- [50] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance Normalization: The Missing Ingredient for Fast Stylization, Nov. 2017. arXiv:1607.08022 [cs].
- [51] Vanya V. Valindria, Nick Pawlowski, Martin Rajchl, Ioannis Lavdas, Eric O. Aboagye, Andrea G. Rockall, Daniel Rueckert, and Ben Glocker. Multi-modal Learning from Unpaired Images: Application to Multi-organ Segmentation in CT and MRI. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 547–556, Mar. 2018.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need, Dec. 2017. arXiv:1706.03762 [cs].
- [53] Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F. Wong, and Lidia S. Chao. Learning Deep Transformer Models for Machine Translation, June 2019. arXiv:1906.01787 [cs].
- [54] Wen Wang, Jing Zhang, Yang Cao, Yongliang Shen, and Dacheng Tao. Towards Data-Efficient Detection Transformers, Aug. 2022. arXiv:2203.09507 [cs].
- [55] Xue Wang, Zhanshan Li, Yongping Huang, and Yingying Jiao. Multimodal medical image segmentation using multi-scale context-aware network. *Neurocomputing*, 486:135–146, May 2022.
- [56] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. HuggingFace’s Transformers: State-of-the-art Natural Language Processing, July 2020. arXiv:1910.03771 [cs].
- [57] Yuxin Wu and Kaiming He. Group Normalization, June 2018. arXiv:1803.08494 [cs].
- [58] Huixian Zhang, Hailong Li, Jonathan R. Dillman, Nehal A. Parikh, and Lili He. Multi-Contrast MRI Image Synthesis Using Switchable Cycle-Consistent Generative Adversarial Networks. *Diagnostics*, 12(4):816, Apr. 2022. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.
- [59] Yiwen Zhang, Liming Zhong, Hai Shu, Zhenhui Dai, Kaiyi Zheng, Zefeiyun Chen, Qianjin Feng, Xuetao Wang, and Wei Yang. Cross-Task Feedback Fusion GAN for Joint MR-CT Synthesis and Segmentation of Target and Organs-At-Risk. *IEEE Transactions on Artificial Intelligence*, pages 1–12, 2022. Conference Name: IEEE Transactions on Artificial Intelligence.
- [60] Zizhao Zhang, Lin Yang, and Yefeng Zheng. Translating and Segmenting Multimodal Medical Volumes with Cycle- and Shape-Consistency Generative Adversarial Network, Mar. 2019. arXiv:1802.09655 [cs].
- [61] Yefeng Zheng. Cross-modality medical image detection and segmentation by transfer learning of shapel priors. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pages 424–427, Apr. 2015. ISSN: 1945-8452.
- [62] Tongxue Zhou, Su Ruan, and Stéphane Canu. A review: Deep learning for medical image segmentation using multi-modality fusion. *Array*, 3-4:100004, Sept. 2019. arXiv:2004.10664 [cs, eess, stat].
- [63] Ziqi Zhou, Lei Qi, Xin Yang, Dong Ni, and Yinghuan Shi. Generalizable Cross-modality Medical Image Segmentation via Style Augmentation and Dual Normalization. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20824–20833, New Orleans, LA, USA, June 2022. IEEE.
- [64] Xiahai Zhuang, Lei Li, Christian Payer, Darko Stern, Martin Urschler, Mattias P. Heinrich, Julien Oster, Chunliang Wang, Orjan Smedby, Cheng Bian, Xin Yang, Pheng-Ann Heng, Aliasghar Mortazi, Ulas Bagci, Guanyu Yang, Chenchen Sun, Gaetan Galisot, Jean-Yves Ramel, Thierry Brouard, Qianqian Tong, Weixin Si, Xiangyun Liao, Guodong Zeng, Zenglin Shi, Guoyan Zheng, Chengjia Wang, Tom MacGillivray, David Newby, Kawal Rhode, Sebastien Ourselin, Raad Mohiaddin, Jennifer Keegan, David Firmin, and Guang Yang. Evaluation of Algorithms for Multi-Modality Whole Heart Segmentation: An Open-Access Grand Challenge, Feb. 2019. arXiv:1902.07880 [cs].