

Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/tciv20

Cross modality medical image synthesis for improving liver segmentation

Muhammad Rafiq , Hazrat Ali , Ghulam Mujtaba , Zubair Shah & Shoaib Azmat

To cite this article: Muhammad Rafiq , Hazrat Ali , Ghulam Mujtaba , Zubair Shah & Shoaib Azmat (2025) Cross modality medical image synthesis for improving liver segmentation, Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization, 13:1, 2476702, DOI: [10.1080/21681163.2025.2476702](https://doi.org/10.1080/21681163.2025.2476702)

To link to this article: <https://doi.org/10.1080/21681163.2025.2476702>



© 2025 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 11 Mar 2025.



Submit your article to this journal [↗](#)



Article views: 80



View related articles [↗](#)



View Crossmark data [↗](#)

Cross modality medical image synthesis for improving liver segmentation

Muhammad Rafiq^a, Hazrat Ali^b, Ghulam Mujtaba^a, Zubair Shah^c and Shoaib Azmat^a

^aDepartment of Electrical and Computer Engineering, COMSATS University Islamabad, Abbottabad, Pakistan; ^bComputing Science and Mathematics, University of Stirling, Stirling, UK; ^cCollege of Science and Engineering, Hamad Bin Khalifa University, Doha, Qatar

ABSTRACT

Deep learning-based computer-aided diagnosis (CAD) of medical images requires large datasets. However, the lack of large publicly available labelled datasets limits the development of deep learning-based CAD systems. Generative Adversarial Networks (GANs), in particular, CycleGAN, can be used to generate new cross-domain images without paired training data. However, most CycleGAN-based synthesis methods lack the potential to overcome alignment and asymmetry between the input and generated data. We propose a two-stage technique for the synthesis of abdominal MRI using cross-modality translation of abdominal CT. We show that the synthetic data can help improve the performance of the liver segmentation network. We increase the number of abdominal MRI images through cross-modality image transformation of unpaired CT images using a CycleGAN inspired deformation invariant network called EssNet. Subsequently, we combine the synthetic MRI images with the original MRI images and use them to improve the accuracy of the U-Net on a liver segmentation task. We train the U-Net on real MRI images and then on real and synthetic MRI images. Consequently, by comparing both scenarios, we achieve an improvement in the performance of U-Net. In summary, the improvement achieved in the Intersection over Union (IoU) is 1.17%. The results show the potential to address the data scarcity challenge in medical imaging.

ARTICLE HISTORY

Received 22 November 2024
Accepted 1 March 2025

KEYWORDS

Computer aided diagnosis;
CycleGAN; medical imaging;
MRI; segmentation

1. Introduction

Deep learning models have advanced the current state of artificial intelligence (AI) in medical imaging. In general, we require a significant number of labelled images to train a deep learning model. However, acquiring medical image data with ground-truth labels is costly and time-consuming, as one must rely on manual annotation by trained radiologists. In addition, privacy concerns also limit the public sharing of medical data. On the other hand, training a deep learning model on a small number of images limits the generalisation of the model. For small medical datasets, synthesising newer images using the available cross-modality data may help to improve the model's training. Furthermore, certain features of human tissues may be better visualised using only certain modalities. For example, compared to computed tomography (CT) scans, magnetic resonance imaging (MRI) can provide better insights into soft tissues (Stallard 2019). However, an MRI scanner is a complex device that is costly to operate and only available at specialised centres (BetterHealth 2012). On the other hand, CT is comparatively less expensive than MRI and is, therefore, more accessible (Eurostat 2021). One possible approach to increasing the number of MRI images is to transform CT into MRI. This domain-to-domain translation can aid in gaining additional MRI data and, hence, reveal features in human tissues that would otherwise be not available. Consequently, this increase in data can improve the performance of a deep learning model. It is important to note that

increasing images by this domain-to-domain image synthesis is different from data augmentation; in data augmentation, the changes are in the geometric and photometric transformations of the same image, while in image synthesis, new images are generated from original images to increase the number of images in a dataset.

Recently, Generative adversarial networks (GANs) have demonstrated good potential to generate synthetic image data (e.g (Yi et al. 2019; Gui et al. 2021; Ali et al. 2022; Ali and Shah 2022)). GAN consists of a generative network G and a discriminative network D . The generative network G takes the noise data and produces synthetic data, and the discriminative network D takes input both from the generator G and from the training dataset and outputs the probability of whether the input has come from the real data (Goodfellow et al. 2020). In each iteration of the training, the discriminator is updated to make it better at discriminating between synthetic and real samples, and accordingly, the generator is updated based on its performance in fooling the discriminator. In other words, the two models are trained together in a two-player minimax game. The training of GAN (Goodfellow et al. 2020) is prone to instability. DCGAN (Radford et al. 2015) introduced a set of constraints to address the instability problem. The GAN also has the problem of generating images unconditionally, i.e. with no control over the output. The CGAN model provided control over the data synthesis process by conditioning the GAN based on supplementary information (Mirza and

Osindero 2014). This conditional supplementary information may be a class label or data from another modality. ACGAN (Odena et al. 2017) takes both the class label and the noise as input and generates output images. The above conditional GANs are application-specific. To generalise CGAN, image-to-image translation GAN (pix2pix GAN) maps an input image from one domain to an output image in a different domain. The pix2pix GAN has demonstrated effectiveness on a wide variety of problems (Isola et al. 2017). The pix2pix-GAN was recently used by (Yang et al. 2018) to improve the registration and segmentation of MRI images. Pix2pix-GAN can be utilised to convert a T1-weighted MRI image into a T2-weighted MRI image, as demonstrated by (Dar et al. 2019). Similarly, Liu et al. (2020) synthesised brain MRI image types, including T1, T2, DWI and Flair, using paired images based on a pix2pix-GAN-inspired architecture. Other architectures inspired by pix2pix-GAN include MI-GAN (Iqbal and Ali 2018), which generates retinal vessel images using paired training data to improve the performance of the vessel segmentation, while CovidGAN (Waheed et al. 2020) generates chest X-ray (CXR) images using paired data to improve the CNN performance for COVID-19 detection. Furthermore, in (Lei et al. 2020) and (Fu et al. 2020), paired CT and MRI images are used to improve the performance of multiorgan pelvic segmentation. However, pix2pix-GAN requires paired samples of data from the source and target domains, which are not commonly available in the medical imaging field.

The CycleGAN model extracts features from the input domain image and learns to translate these features to the target domain image without any paired training samples (Zhu et al. 2017). Shen et al. (2020, 2020), used CycleGAN as the baseline framework for MRI image-to-image translation (T1 to T2 and vice versa). Zhang et al. (2022), used a large set of publicly available paediatric structural brain MRI images to construct a switchable CycleGAN model for image synthesis between multi-contrast brain MRI (T1w and T2w) images. SLA-StyleGAN (Wang et al. 2018) used unpaired data to synthesise skin lesion images of high quality, which in turn improved the classification accuracy of skin lesions. Park et al. (2020) proposed an improved variant of CycleGAN based on contrastive learning. Guo et al. (2025) used the model from Park et al. (2020), for data augmentation to improve the accuracy of kidney segmentation. They generated more kidney ultrasound images using unpaired CT images to improve segmentation accuracy on U-Net using both real and synthesised ultrasound images. Hong et al. (2022) developed a specialised domain adaptation framework based on adversarial learning between unpaired CT and MRI. They synthesised MRI images from CT and segmented the liver in abdominal MRI using only CT labels, without any MRI labels.

In short, for synthesising cross-domain medical images (e.g. CT to MRI, MRI to CT, CT to PET, etc.), CycleGAN and its variants mentioned above have proven to be effective for domain-to-domain translation tasks on unpaired data. However, most CycleGAN-based synthesis methods lack the potential to overcome alignment and asymmetry between the input and generated data. It has been noted that CycleGAN may replicate the data's 'domain-specific deformations' (Chartsias et al. 2017; Zhao et al. 2021), which means that there are geometric

changes (scaling, rotation, translation and shear) between the source and target domains. DicyclicGAN (Chartsias et al. 2017) addressed the problem of domain-specific deformation. It used deformable convolutional (DC) layers along with new cycle-consistency losses. This model has the capacity to synthesise output data that is aligned with input data. In this model, 'deformation' and 'image translation' parameters were combined into a single network. However, the DC layers of DicyclicGAN only learn relatively unchanged and specific deformations. The improved version of DicyclicGAN was named TPSDicyc (Wang et al. 2019), formed on a thin-plate-spline (TPS). To learn the relative distortion between the input and output data, TPSDicyc used a segregated spatial transformation network (STN). The authors of the TPSDicyc technique used publicly available multi-sequence brain MR data and multi-modality abdominal data to evaluate their method, which gave good results for generating synthetic images that are aligned with the source data. The alignment issue was also tackled by (Zhao et al. 2021) using the alignment step, converting the source data in such a way that the position, scale and viewing angle are identical to the output data before synthesising MR images from CT images. An improved end-to-end network, which is also based on CycleGAN, named EssNet, was introduced in (Huo et al. 2018). EssNet synthesised unpaired MRI to CT images and CT splenomegaly segmentation concurrently in the absence of ground truth labels in CT. EssNet had a segmentation network part in addition to the CycleGAN network part. In EssNet, the alignment and domain-specific deformation issues were resolved through simultaneous training of the CycleGAN and segmentation networks. The model was trained on unpaired CT and MRI data and used ground truth only from MRI data. It generated synthetic CT images as well as the corresponding labels by leveraging the ground truth labelling of MRI images. Zhang et al. (2018) proposed a variant of EssNet where they used two segmentors instead of one, utilising both source and target labels for cross-modality domain adaptation between heart MRI and CT. Similarly, Liu et al. (2024) also proposed a variant of EssNet where they used two segmentors instead of one, utilising source labels only for cross-modality domain adaptation between brain and abdomen MRI and CT. More related efforts on GANs for the synthesis of medical image data for different organs are reviewed in (Ali et al. 2022; Ali and Shah 2022).

The recent state-of-the-art techniques discussed above use unpaired image-to-image domain adaptation to achieve image synthesis, where synthetic images from one modality with sufficient labelled data aid segmentation tasks for another modality. Our aim in this work is to use cross-modality synthesis for MRI data augmentation. MRI image synthesis through the domain-to-domain translation of CT scans can increase the size of the dataset and thus improve the generalisation of the AI models. In this work, we use a CycleGAN-based EssNet architecture for cross-modality translation of abdominal CT to abdominal MRI that can help improve the performance of a segmentation network. We train the EssNet using unpaired CT and MRI images and use manual labels for the CT images only. The simultaneous training of the CycleGAN and segmentation network parts of EssNet solves the alignment and domain-specific

deformation issues. To evaluate the improvements resulting from cross-modality translation, we use the U-Net biomedical segmentation model for liver segmentation using the original and the generated MRI images. Through segmentation performance in terms of dice coefficient and Intersection over Union (IoU), we demonstrate the improvements that occur as a result of using the generated data. To the best of our knowledge, this is the first method that performs cross-modality abdominal (CT to MRI) image synthesis to improve the segmentation accuracy of a model. The main contributions of this work are:

- We present a new two-stage EssNet+U-Net architecture that can synthesise abdominal MRI using image transformation of unpaired CT images and perform liver segmentation.
- We perform a cross-modality image transformation to generate synthetic MRI images that are free from alignment and domain-specific deformation issues.
- Through empirical evaluation, we demonstrate the potential of synthetic MRI images and their effectiveness in improving the performance of U-Net for liver segmentation.

The remainder of the paper is organised as follows: [Section 2](#) presents the proposed methodology in detail, followed by a discussion on datasets in [Section 3](#). [Section 4](#) presents the results obtained with the proposed method and also presents the results for ablation studies on the proposed method. Finally, [Section 4](#) concludes the paper.

2. Methods

Huo et al. (2018) proposed an end-to-end synthesis and segmentation network, which performs unpaired MRI to CT image synthesis and does CT splenomegaly segmentation concurrently in the absence of ground truth labels in CT. Inspired by EssNet, we propose a two-stage (EssNet+U-Net) architecture, as shown in [Figure 1](#). In the proposed approach, we use EssNet to transform unpaired abdominal CT images into MRI images and perform MRI liver segmentation concurrently in the absence of ground truth labels for the MRI. The purpose of MRI liver segmentation is to help the generator network generate high-quality MRI images without alignment and domain-specific deformation issues. The images generated by EssNet, along with the real MRI images, are fed into the U-Net model to improve the segmentation accuracy. Unlike EssNet, which performed unpaired MRI to CT image synthesis and segmented CT splenomegaly, our aim is to use EssNet to achieve unpaired abdominal CT to MRI aligned image synthesis only and then use these newly generated images to improve the performance of U-Net segmentation. In the following text, we first describe the EssNet network details, followed by the U-Net.

2.1. EssNet

[Figure 2](#) shows the EssNet architecture. The EssNet model consists of a CycleGAN part, which has two generators and two discriminators and a segmentation network part. Part 1 of [Figure 2](#) is basically CycleGAN which generates synthetic images, while Part 2 of [Figure 2](#) segments the synthetic images, which also helps to overcome the alignment problem of

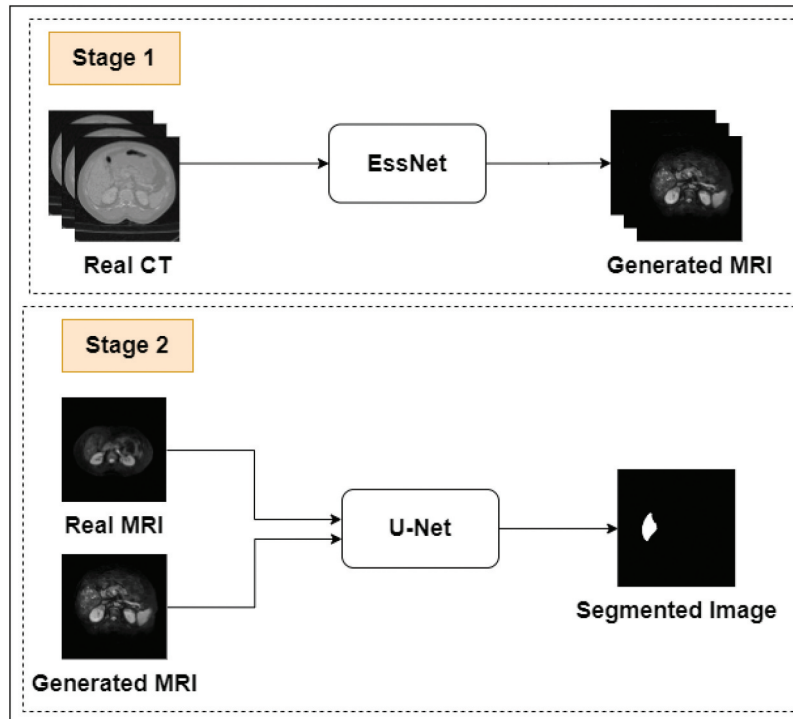


Figure 1. Our proposed two-stage (EssNet+U-Net) approach. Stage 1: the approach uses CycleGAN-based EssNet architecture to generate abdominal MRI. Stage 2: the generated MRI is combined with real MRI to improve the U-Net segmentation.

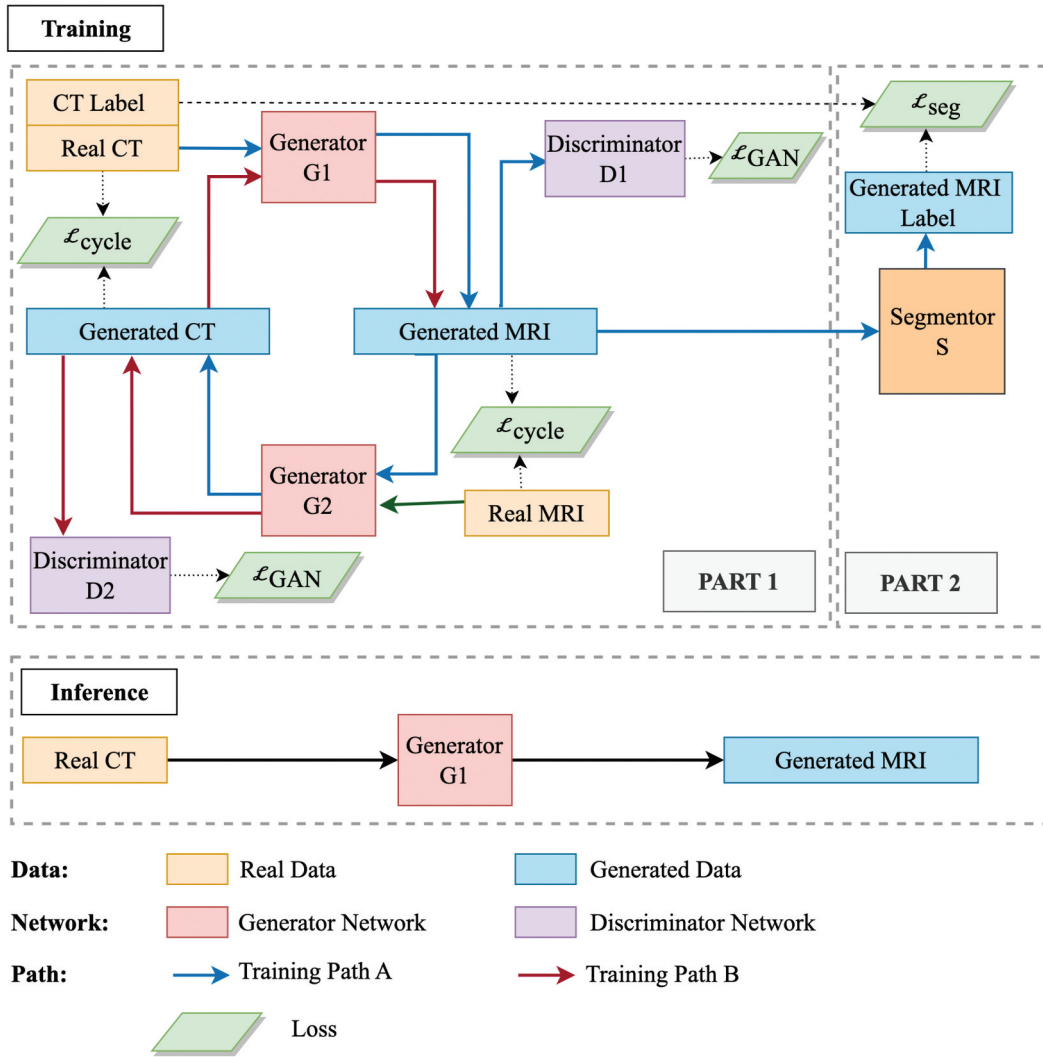


Figure 2. Overall workflow of the EssNet architecture. Part 1 of the training module is the synthesis part, which is basically CycleGAN. G_1 and G_2 are the two generators, while D_1 and D_2 are the two discriminators. The segmentation part for end-to-end training is in part 2. In the testing module, real CT is fed to generate additional MRI images. (adapted from Huo et al. (2018)).

generated images by CycleGAN. Two generators (G_1 and G_2) are used to generate CT to MRI images and vice versa. The two adversarial discriminators (D_1 and D_2) are used for checking the probability of the input image being fake (generated) or real. D_1 assesses whether the MRI image is real or fake, whereas D_2 determines the status of the CT image. Two training routes (Path A and Path B) are presented in forward cycles when this framework is applied to unpaired CT and MRI. In the testing phase, we feed real CT images to the trained network to generate additional MRI images. The complete architecture of EssNet is discussed below, where we first discuss the generator/segmentor, followed by the discriminator and then the loss functions.

2.1.1. Architecture of generator

The two generators (G_1 and G_2) are constructed with the 9-block ResNet (defined in (Johnson et al. 2016; Zhu et al. 2017)), and have an encoder-decoder architecture, as shown in Figure 3. G_1 transforms CT images to MRI, whereas G_2 transforms MRI images to CT. The segmentation network S is

concatenated with G_1 as an additional forward branch. The 9-block ResNet (Johnson et al. 2016; Zhu et al. 2017) is employed as S , and its network structure is identical to that of G_1 and G_2 . The segmentation is then performed on the generated MRI. ResNet employs the idea of residual blocks, which incorporate shortcut skip connections. By redefining the layers as residual functions that learn about the inputs to the layer, the authors of ResNet (He et al. 2016) demonstrate that residual networks are simpler to optimise and can gain more accuracy at a noticeably higher depth.

2.1.2. Architecture of discriminator

The architecture of the two adversarial discriminators (D_1 and D_2) is patch-based, as introduced in (Isola et al. 2017) and shown in Figure 4. D_1 assesses whether the MRI image is real or fake, whereas D_2 determines the nature of the CT image. In the patch discriminator, after feeding one input image to the network, it provides us with the probabilities of two outcomes: real or fake. However, it does not utilize a scalar output; instead, it uses an $N \times N$ output matrix. The output of the discriminator is

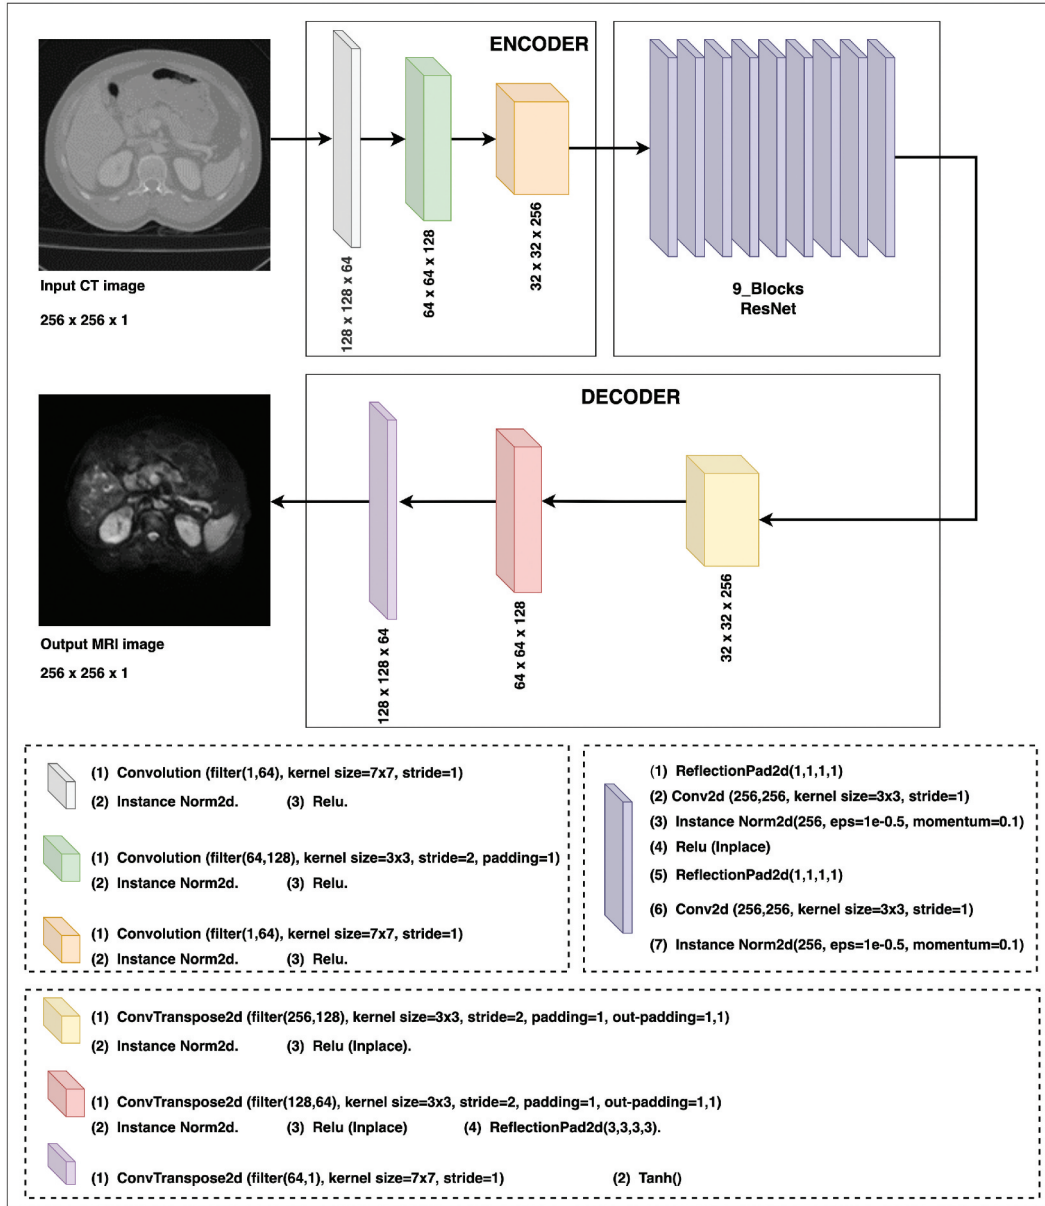


Figure 3. The generator architecture includes an encoder block, a 9-block ResNet and a decoder block.

an $N \times N$ matrix corresponding to the $N \times N$ patches of the input image. Each patch is 70×70 pixels, and this size is fixed. Here, $N \times N$ can vary based on the dimension of the input image.

2.2. Loss function

A combination of five loss functions is used to accomplish the training. The losses include two adversarial loss functions, two cycle-consistency loss functions and one segmentation loss function, as discussed below.

2.2.1. Adversarial loss function

The two adversarial loss functions, referred to as $\mathcal{L}_{GAN}(\cdot)$, are defined as follows:

$$\mathcal{L}_{GAN}(G_1, D_1, A, B) = E_{y \sim B}[\log D_1(y)] + E_{x \sim A}[\log(1 - D_1(G_1(x)))] \quad (1)$$

$$\mathcal{L}_{GAN}(G_2, D_2, B, A) = E_{x \sim A}[\log D_2(x)] + E_{y \sim B}[\log(1 - D_2(G_2(y)))] \quad (2)$$

Equations 1 and 2 are standard adversarial cross-entropy loss GAN equations. In Equation 1, G_1 the input image x of modality A and attempts to generate images $G_1(x)$ that resembles images y of modality B, while D_1 attempts to differentiate between generated samples $G_1(x)$ and real images y from modality B. Equation 2 does the same thing with the roles of modality A and B reversed. The probabilities that the data instances y of modality B and x of modality A, in our case an MRI image and a CT image, are real, are estimated by the two

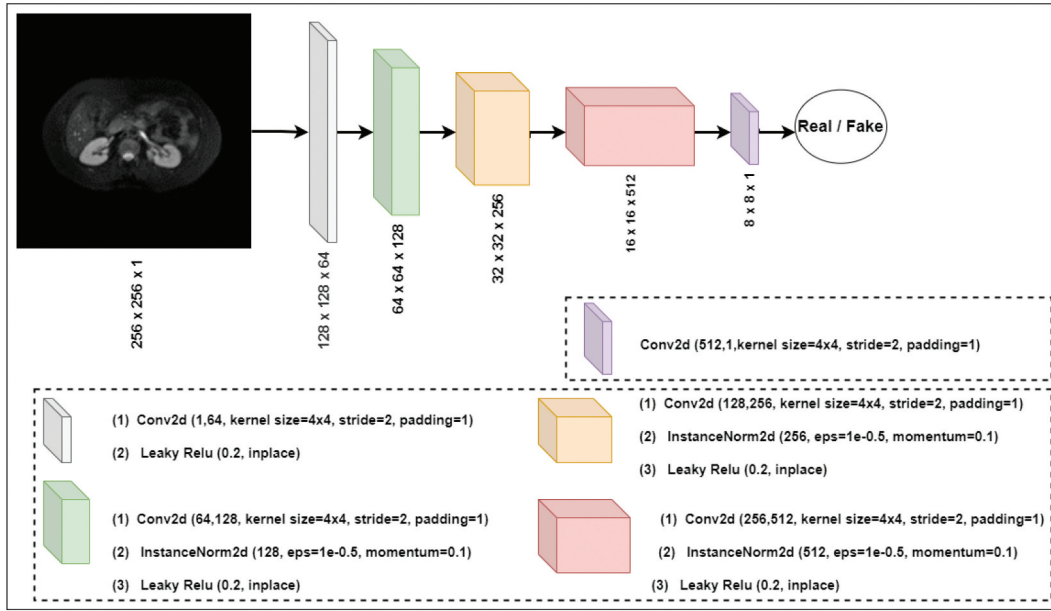


Figure 4. Architecture of a patch-based discriminator. The discriminator has five convolution blocks.

discriminators $D_1(y)$ and $D_2(x)$. For input instances x and y , the two generators' outputs are $G_1(x)$ and $G_2(y)$. $1 - D_1(G_1(x))$ and $1 - D_2(G_2(y))$ estimate the probability that the input instances are synthetic. Both $E_{y \sim B}$ and $E_{x \sim A}$ represent the values that should be expected across all real data instances. The generator tries to reduce the given loss function, whereas the discriminator tries to maximise it, similar to a min-max game.

2.2.2. Cycle consistency loss function

Adversarial loss alone is insufficient for producing high-quality images, but the cycle consistency loss solves this problem. Cycle consistency loss is used for unpaired image-to-image translation in generative adversarial networks. Two cycle-consistency loss functions, referred to as $\mathcal{L}_{cycle}(\cdot)$, are utilised to compare the real images with the reconstructed images.

$$\mathcal{L}_{cycle}(G_1, G_2, A) = E_{x \sim A} [\| G_2(G_1(x) - x) \|_1] \quad (3)$$

$$\mathcal{L}_{cycle}(G_2, G_1, B) = E_{y \sim B} [\| G_1(G_2(y) - y) \|_1] \quad (4)$$

From Equations (3) and (4): we wish to learn a mapping between $G_1(x)$ and $G_2(y)$ for two domains x and y . We wish to reinforce the idea that these translations should be inverses of each other and bijections. Therefore, the inclusion of both the cycle consistency losses for $G_2(G_1(x)) \approx x$ and $G_1(G_2(y)) \approx y$ are encouraged. In other words, if we pass synthesised image $G_1(x)$ (input modality A, output modality B) to G_2 (input modality B, output modality A), the original image x should be reconstructed and vice versa. By requiring forwards and backward consistency, the loss narrows the range of potential mapping functions.

2.2.3. Segmentation loss function

A segmentation loss function, referred to as $\mathcal{L}_{seg}(\cdot)$, is defined as:

$$\mathcal{L}_{seg}(S, G_1, A) = \sum_i N_i \cdot \log(S(G_1(x_i))) \quad (5)$$

The equation is the cross-entropy segmentation loss, where $S(G_1(x_i))$ shows the segmented output of segmentor S , which segments the modality B image generated by generator G_1 , N is the pixel-wise ground truth of modality A for image x , and i is the pixel index. The total loss function is subsequently defined below where the lambdas are hyper parameters set empirically:

$$\mathcal{L}_{total} = \lambda_1 \cdot \mathcal{L}_{GAN}(G_1, D_1, A, B) + \lambda_2 \cdot \mathcal{L}_{GAN}(G_2, D_2, B, A) + \lambda_3 \cdot \mathcal{L}_{cycle}(G_1, G_2, A) + \lambda_4 \cdot \mathcal{L}_{cycle}(G_2, G_1, B) + \lambda_5 \cdot \mathcal{L}_{seg}(S, G_1, A) \quad (6)$$

2.3. U-Net architecture

We used the U-Net (Ronneberger et al. 2015) architecture to test improvements in segmentation accuracy using artificially generated images. U-Net consists of an expanding path and a contracting path. The contracting route follows the convolutional network architectural conventions. On the contracting path repeatedly, two 3×3 convolutions followed by a rectified linear unit (ReLU) are used for feature extraction, and a 2×2 max pooling operation with stride 2 is used for downsampling. With each downsampling step, we increase the number of feature channels by a factor of two. At each stage of the expanding path, the process begins with upsampling the feature map, followed by a 2×2 up-convolution. This is then concatenated with the feature map from the contracting path, which has been cropped to match in size. Afterwards, two 3×3 convolutional layers are applied sequentially, each followed by a ReLU activation function. Throughout the expanding path, the number of feature channels is progressively halved at each step, which contrasts with the contracting path, where the number of channels is increased. Finally, in the last layer, each of the 64-component feature vectors is mapped to the desired number of

output classes using a 1×1 convolution. The overall architecture of the network consists of 23 convolutional layers in total.

3. Datasets

This study uses two datasets: abdominal MRI (T1 and T2 weighted) and abdominal CT. The two datasets are included in the Combined (CT-MR) Healthy Abdominal Organ Segmentation (CHAOS) dataset (Kavur et al. 2021). Both datasets include the MRI/CT of multiple healthy persons. The datasets were retrospectively and arbitrarily gathered from Dokuz Eylul University (DEU) Hospital's picture archiving and communication system (PACS). The CT and MR datasets are unrelated to one another since they were gathered from unregistered individuals who are distinct from one another (Kavur et al. 2021).

3.1. Abdominal CT dataset

There are 40 different persons' CT images in this dataset. The images in the datasets are acquired from healthy individuals, so the liver is free of lesions or other disorders. The CT images were taken from the person's upper abdomen. Each person's 3D CT DICOM file includes a collection of 2D DICOM images representing different CT cross-sections. Each DICOM image is 16-bit with a resolution of 512×512 . A few samples from this dataset are shown in Figure 5. The data of 20 persons is kept for the training set, and the data of the remaining 20 persons is kept for the test set. In our work, we used only the training set for both training and testing because of the availability of ground truth CT images in this training set. The total number of images in this training set is 2874. We used those 2050 images in which the liver can be seen, as we want to do liver segmentation.

3.2. Abdominal MR dataset

In this dataset, 120 3D DICOM files from two distinct MRI sequences are included. The MRI sequences are T1-DUAL in phase (40 3D DICOM files), out phase (40 3D DICOM files) and T2-SPIR (40 3D DICOM files). Each person's 3D DICOM file includes a collection of 2D DICOM images representing different MRI cross-sections. These MRI images were also taken from the person's abdomen. Each image in this dataset is 12-bit with a resolution of 256×56 . A few samples from this dataset are shown in Figure 6. From this dataset, we use T2-SPIR, which has 20 files for training and 20 files for testing. However, the ground truth is available for the training set images only, so we use samples from the training set only. The total number of images in this training set is 623. Of these, the number of images suitable for our tasks is 408, as the liver is visible in these images.

4. Results

The experiments were performed on an NVIDIA GeForce RTX 2060 with 6GB GDDR6 memory. For EssNet, the authors' online implementation based on PyTorch was adapted, where the loss function lambdas were set to $\lambda_1 = 1, \lambda_2 = 1, \lambda_3 = 10, \lambda_4 = 10$ based on the CycleGAN implementation whereas λ_5 was simply set to 1. The optimiser used was Adam with a learning rate of 0.0001 for G_1, G_2 and S , and 0.0002 for D_1 and D_2 . The number of input and output channels for all networks was one, except for the Segmentor S . For S , we used two output channels: one for the background and one for the liver. The model was trained for 100 epochs with a batch size of 1 to remain within the GPU's memory limits. For the U-Net, the Keras-TensorFlow implementation was employed. The Adam optimiser was used with a learning rate of 0.0001. The model was trained for 300

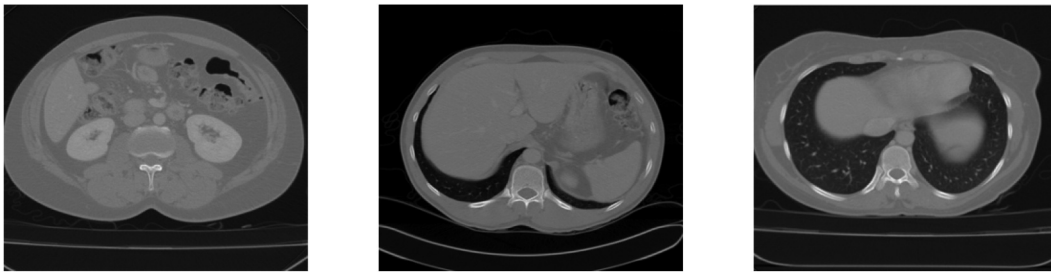


Figure 5. Samples of abdominal CT images. These are the original CT images from the abdominal CT dataset and show different organs of the abdomen.

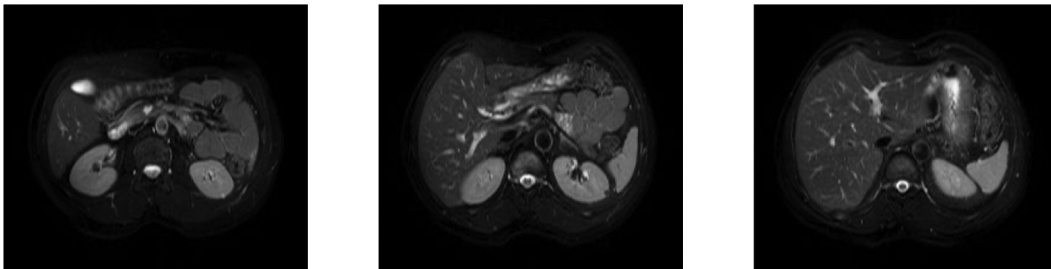


Figure 6. Samples of abdominal MRI images from the T2-SPIR sequence. These are the original MRI images from the abdominal MR dataset and show different organs of the abdomen.

epochs with a batch size of 2, and the steps per epoch were equal to the number of images divided by the batch size.

In the first stage, we trained the EssNet network on the original 350 MRI and 350 CT images. We generated 2050 synthetic MRI images from the original CT images. The EssNet network was used only for generating new images from the original images. The arrangement of training and testing images for EssNet is shown in Table 1. Figure 7 shows the original input CT images along with their ground truth and the generated output MRI images of EssNet. It can be seen that there are no alignment or domain-specific deformation issues in the generation. However, as Figure 8 shows, there are some minor defects in the synthesised images. In each row, the first image shows the original CT image, the second

Table 1. Number of training and test images as input to EssNet, and the number of output images by EssNet.

Train/Test	MRI	CT	Generated MRI Images
Train	350	350	–
Test	–	2050	2050

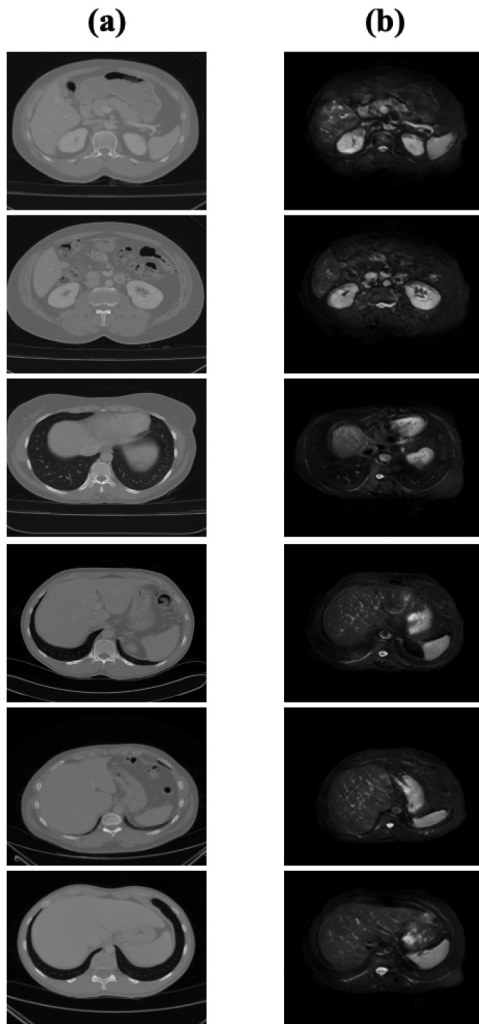


Figure 7. Sample results for CT to MRI translation using EssNet on CHAOS dataset. Column (a) has the real CT images, whereas column (b) has the generated MRI images.

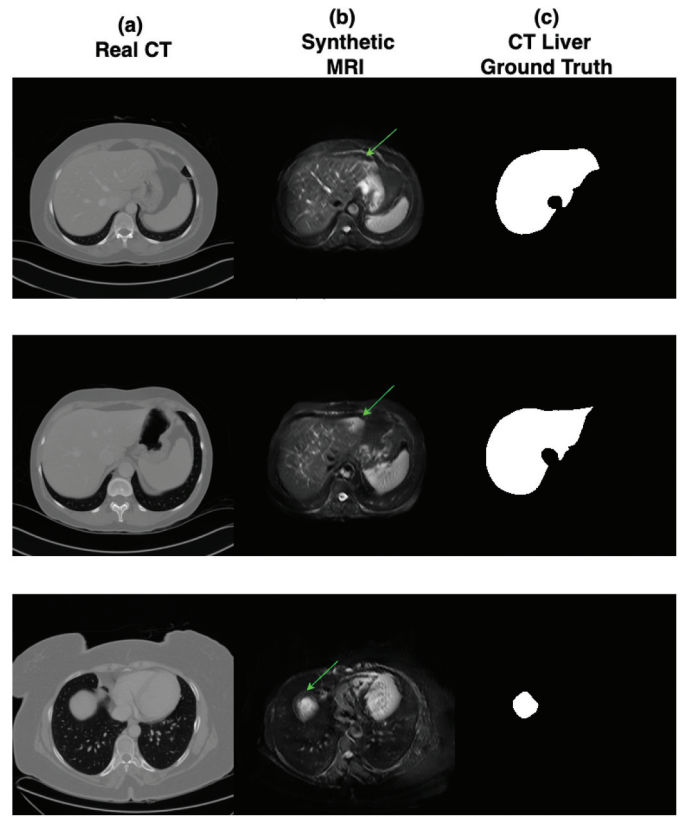


Figure 8. Defects in CT to MRI translation as pointed to by arrows. Column (a) has the real CT images, column (b) has the generated MRI images, whereas column (c) has the CT liver ground truth segmentation mask.

image shows the synthesised MRI image and the third shows the segmentation ground truth for the liver in the CT image. As indicated in the image by the arrow, a minor part of the liver in the synthesised image is missing, which can affect segmentation accuracy when these images are used. Therefore, in the future, more accurate synthesis of images is required, which will improve the accuracy of deep learning tasks on medical images, especially in cases where data is scarce. In the second stage, we used the U-Net architecture for liver segmentation. First, we trained the U-Net on 350 real MRI images. Then, we tested the U-Net on 58 real MRI images and calculated the dice coefficient and IoU. After this, we trained the U-Net again on real and synthetic MRI images while gradually increasing the number of synthetic images in the training data, as shown in Table 2. Finally, we tested the segmentation performance of the U-Net on 58 real MRI images. For performance evaluation, we computed the dice coefficient and IoU. The results reported in Table 3 show that the dice coefficient and IoU are improving

Table 2. Arrangements for training and testing images for U-Net.

Train/Test	Real MRI Images	Generated MRI Images	Combined MRI Images (Real + Generated)
6*Training	350	–	350
	350	354	704
	350	714	1064
	350	1034	1384
	350	1487	1837
	350	2050	2400
Testing	58	–	58

Table 3. Results of segmentation with U-Net using varying sizes of training sets of MRI data. The size of the test set remains the same.

Number of training images (MRI)	Dice	IoU
Real only (350)	0.9459	0.8974
Combined (704)	0.9467	0.8989
Combined (1064)	0.9524	0.9091
Combined (1384)	0.9485	0.9020
Combined (1837)	0.9475	0.9002
Combined (2400)	0.9505	0.9058

until the combined 1064 scenario. After this, the dice coefficient and IoU decrease. Therefore, it can be concluded that performance is saturated for U-Net when we used combined 1064 images. From Figure 9, we can see that when U-Net

segmentation models are trained using both real and synthetic data, the segmentation results show a close resemblance to the ground truth, as compared to when the U-Net is only trained on real images. The AUC-ROC curve is presented in Figure 10, which shows the separability between the two classes. From Figure 10, we can see that the AUC is very high, which means that the prediction capability of the model is very good.

The number of parameters for the EssNet model in our setting is 28,256,644 (28 million). The training time for 100 epochs is 6 h and 23 min, i.e. 230 seconds per epoch. The test time for generating 714 images, as synthesised 714 plus real 350 gives us 1064 training images, is 352 s, i.e. 2.02 frames per second (FPS). The number of parameters for the U-Net

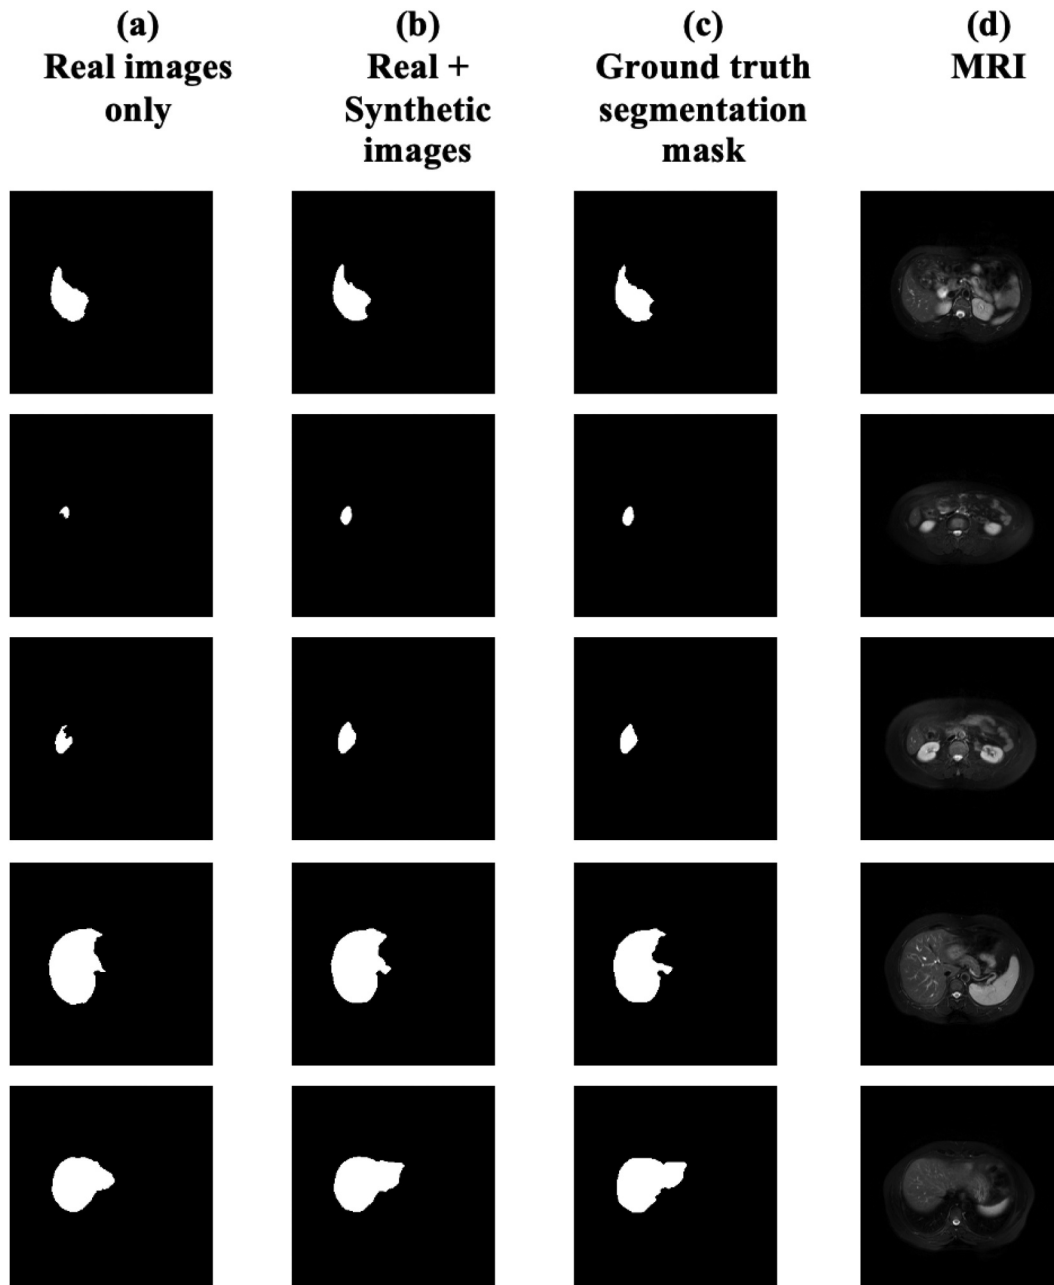


Figure 9. Liver segmentation results using real and synthetic MRI. Column (a) has the segmentation results of U-Net trained only on real images; column (b) has the segmentation results of U-Net trained on both real and synthetic images; column (c) has the original ground truth, whereas column (d) has the input MRI images. A close resemblance with the ground truth can be seen for segmentation when the model is trained using real as well as synthetic data.

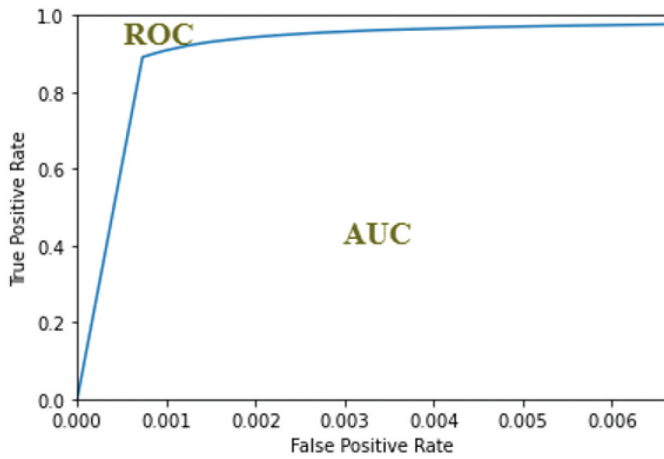


Figure 10. The AUC-ROC curve. The high AUC shows that the separability between the two classes is very high.

Table 4. Computational cost of the models.

	Trainable Parameters	Train Time (seconds per epoch)	Inference (frames per second)
EssNet	28 M	230	2.02
U-Net	31 M	80	21

model in our implementation is 31,031,685 (31 million). With 1064 training images, the training time for 300 epochs is 6 h and 39 min, i.e. 80 seconds per epoch. The test time for segmenting 58 images is 2.76 s i.e. 21 FPS, so the end result is a model that can segment images relatively quickly at test time. The parameters and times are summarised in Table 4. It should be noted that these training times are for an older GeForce RTX 2060 GPU, and the models are expected to train much faster on modern GPUs.

4.1. Ablation studies

To illustrate the viability of the proposed model, we performed an ablation study on the synthesised MRI images from CT images. We made changes to the EssNet model by removing the segmentation network (part 2 in Figure 2), i.e. we used the CycleGAN only. We generated MRI images from CT images through CycleGAN and kept the training and testing images the same as EssNet, as discussed in the results section. Figure 11 shows the results of CycleGAN for CT to MR synthesis. There are some alignment and domain-specific deformations in the generated images. To see the effect of this deformation, we trained U-Net using both real and synthetic MRI images and calculated the dice coefficient and IoU. The results reported in Table 5 show that using real and synthetic MRI images did not bring any improvements in dice coefficient and IoU compared to the results for real images only. Figure 12 also shows that there is little or no difference in the output segmentation

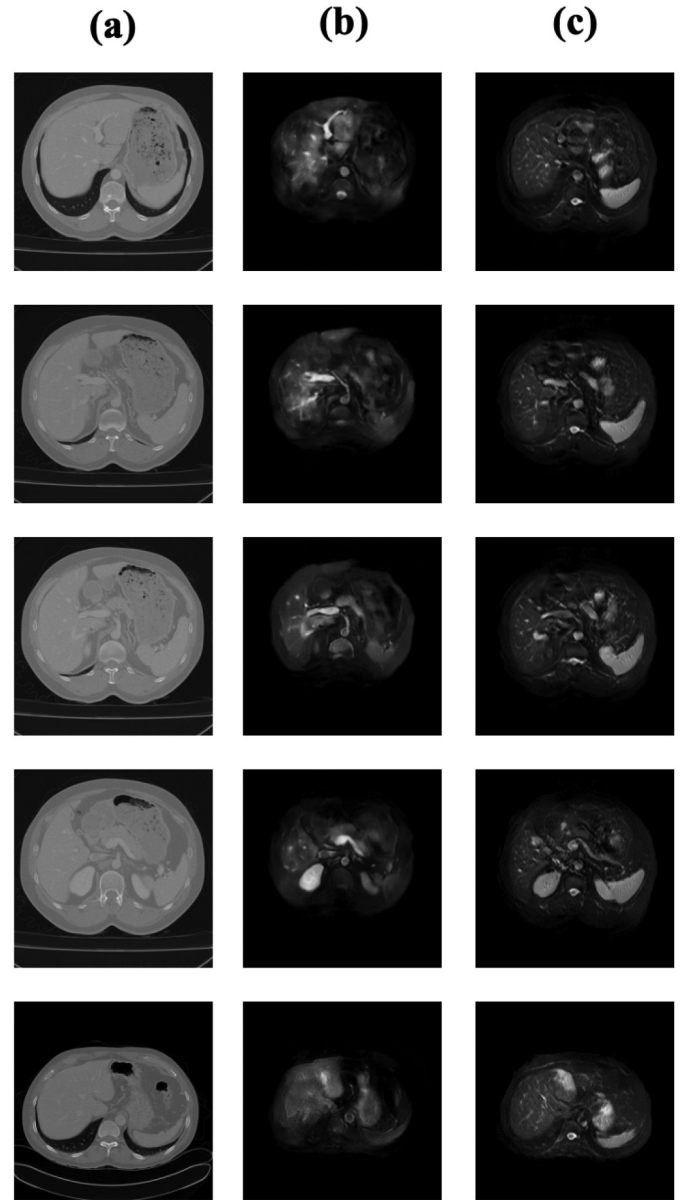


Figure 11. Synthesis results for ablation studies after removing the segmentation part of the EssNet. Column (a) has the real CT images. Column (b) has the MRI images generated by CycleGAN without the segmentation part. Column (c) has the synthetic MRI images generated by EssNet architecture.

results of U-Net, trained on only real images and trained on both real and synthetic images, and that the segmentation results do not show a close resemblance to the ground truth. From this ablation study, we can say that EssNet based method is better at synthesising cross-modality images, as these images have no alignment or deformation issues.

4.2. Limitations

The results presented in this study are derived from a dataset originating from a single hospital, and their generalisability to other datasets/hospitals remains unverified. While the framework demonstrates promising performance within this dataset, caution is advised when extrapolating these findings. The segmentation performance presented in this work is limited to the

Table 5. The output segmentation results of U-Net, trained on only real images and trained on both real and generated images by CycleGAN.

Number of training images (MRI)	Dice	IoU
Real images only (350)	0.9459	0.8974
Combined (1064) Real+Synthetic	0.9460	0.8976

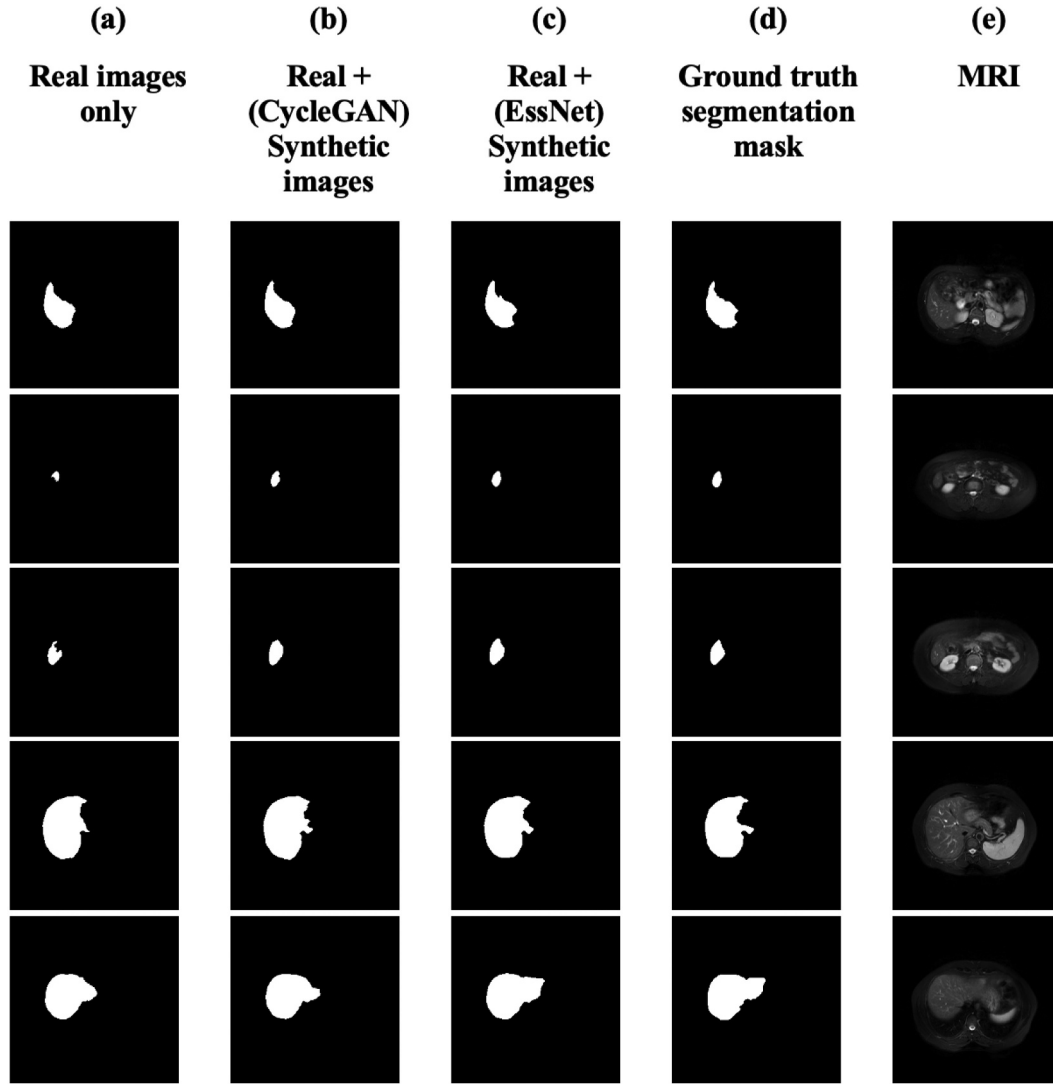


Figure 12. Segmentation results for ablation studies on MRI dataset. Column (a) has the output segmentation results of U-Net only trained on real images; column (b) has the output segmentation results of U-Net trained on both real and generated images (CycleGAN), column (c) has the output segmentation results of U-Net trained on both real and generated images (EssNet), column (d) has the original ground truth, whereas column (e) has the original input MRI images. Results in (b) do not show a close resemblance to the ground truth for segmentation when the model is trained using real as well as synthesis data (CycleGAN).

use of the U-Net model only and does not explore other advanced methods for segmentation. We believe it is justified as the focus of this work has been to show the effectiveness of the synthetic data, and the use of more advanced segmentation techniques has been left out for future experiments.

5. Conclusion

Medical image synthesis using GANs is an exciting and rapidly evolving field that holds great potential for improving medical image analysis and computer-aided diagnosis. In this work, we proposed a two-stage (EssNet+U-Net) architecture for improving the accuracy of liver segmentation through cross-modality medical image synthesis. We use the EssNet model to transform unpaired abdominal CT images to MRI images in such a way that there are no alignment issues or domain-specific deformation issues. The images generated by EssNet are fed into the U-Net model along with the real images to improve segmentation. We evaluated the performance of U-Net in both cases, i.e.

only trained on real images and trained on both real and synthetic images. Experimental results showed the effectiveness of our proposed method, where the segmentation accuracy results improved when we trained U-Net on real plus synthetic images. In summary, the improvement achieved is 1.17% in IoU and 0.65% in the dice coefficient. We believe that these results can serve as motivation for future research into generating cross-modality images in order to increase the number of images in a dataset, particularly, in clinical applications, without the need to acquire additional data where training data is scarce and acquiring additional data is costly. The improvements of the order 1.17% and 0.65% in the IoU and the dice coefficient are relatively smaller; however, it is not uncommon to report such marginal improvements for segmentation tasks. Since the focus of the work has been to demonstrate the effectiveness of synthetic image data, only U-Net was used for the segmentation task. It is expected that more advanced segmentation methods would bring further improvements; however, their use is left out for future work.

Author contributions

CRedit: **Muhammad Rafiq**: Formal analysis, Investigation, Methodology, Writing – original draft; **Hazrat Ali**: Conceptualization, Funding acquisition, Resources, Supervision, Writing – original draft, Writing – review & editing; **Ghulam Mujtaba**: Supervision, Validation, Writing – review & editing; **Zubair Shah**: Validation, Writing – review & editing; **Shoaib Azmat**: Conceptualization, Resources, Supervision, Writing – review & editing.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Notes on contributors

Muhammad Rafiq completed his MSc in Electrical Engineering from COMSATS University Islamabad, Pakistan. He is a research assistant in AI for healthcare. His research interests include Deep Learning, Generative AI, and Medical Imaging.

Dr. Hazrat Ali is a Lecturer in Artificial Intelligence and Data Science. He is an experienced educator and researcher with over a decade of experience in the field. His extensive research portfolio spans various areas of Generative AI, Medical AI, healthcare, and computer vision. He has authored more than 60 peer-reviewed articles published in top-tier international journals and conferences. He is the lead editor of the book *Advances in Deep Generative Models for Medical AI*, published by Springer Nature. He is a senior member of IEEE.

Dr. Ghulam Mujtaba received BSc degree in Computer Systems Engineering from GIKIEST, Pakistan, in 2003. He did PhD in Electrical Engineering from Loughborough University, United Kingdom, in 2011. Currently, he is an Assistant Professor in the Department of Electrical and Computer Engineering, COMSATS University Islamabad, Abbottabad Campus. His research interests include artificial intelligence and machine learning.

Dr. Zubair Shah is an Assistant Professor in the College of Science and Engineering specializing in artificial intelligence and big data analytics with a focus on health informatics. He holds an MS degree in Computer System Engineering from Politecnico di Milano, Italy, and a PhD from the University of New South Wales, Australia. Before joining HBKU, Dr. Shah was a Research Fellow at the Australian Institute of Health Innovation, Macquarie University, from 2017 to 2019. His research focuses on applying AI to healthcare, particularly in developing AI-driven applications that enhance disease diagnosis and prediction. He has secured critical funding from HBKU Innovation Center and QNRF to advance these efforts, including work on diabetes management and multimodal AI systems combining electronic health records (EHR), medical imaging, and multi-omics data. His team has developed innovative deep-learning techniques, including methods for identifying IDH mutations from histopathology slides and exploring the use of generative models like GANs and neural diffusion for medical data augmentation.

Dr. Shoaib Azmat received his Ph.D. and M.S. degrees in Electrical and Computer Engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2014 and 2011, respectively. Prior to this, he received his B.Sc. degree in Computer Systems Engineering from the University of Engineering & Technology, Peshawar, Pakistan. Currently, he is an Associate Professor of Computer Engineering at COMSATS University, Abbottabad, Pakistan. His research interests include computer vision, machine learning, and digital image processing, with research publications in reputed international journals.

Authors' contributions

H. A. and S. A. designed the research. M. R. conducted the experiments. M. R., H. A. and S. A. did the analysis and interpretation of the results. G. M. and S. A. supervised the work. M. R. and S. A. wrote the

manuscript. H. A., G. M. and Z. S. revised the manuscript. All authors reviewed the manuscript.

Data availability statement

All data analysed during this study is publicly available from <https://chaos.grand-challenge.org/>.

References

- Ali H, Biswas MR, Mohsen F, Shah U, Alamgir A, Mousa O, Shah Z. 2022. The role of generative adversarial networks in brain MRI: a scoping review. *Insights Imag*. 13(1):98. doi: [10.1186/s13244-022-01237-0](https://doi.org/10.1186/s13244-022-01237-0).
- Ali H, Shah Z. 2022. Combating COVID-19 using generative adversarial networks and artificial intelligence for medical images: scoping review. *JMIR Med Inf*. 10(6):e37365. doi: [10.2196/37365](https://doi.org/10.2196/37365).
- BetterHealth. 2012. MRI scan. Vic.gov.au. [accessed 2023 Apr 30]. <https://www.betterhealth.vic.gov.au/health/conditionsandtreatments/mri-scan>.
- Chartsias A, Joyce T, Dharmakumar R, Tsafaris SA. 2017. Adversarial image synthesis for unpaired multi-modal cardiac data. *Simulation and Synthesis in Medical Imaging: Second International Workshop, SASHIMI 2017, Held in Conjunction with MICCAI 2017, Proceedings 2*; Sep 10; Québec City, QC, Canada: Springer International Publishing. p. 3–13.
- Dar SU, Yurt M, Karacan L, Erdem A, Erdem E, Cukur T. 2019. Image synthesis in multi-contrast MRI with conditional generative adversarial networks. *IEEE Trans Med Imag*. 38(10):2375–2388. doi: [10.1109/TMI.2019.2901750](https://doi.org/10.1109/TMI.2019.2901750).
- Eurostat. 2021 July 2. Availability of CT and MRI units in hospitals. [accessed 2023 Apr 30]. <https://ec.europa.eu/eurostat/web/products-eurostat-news/-/ddn-20210702-2>.
- Fu Y, Lei Y, Wang T, Tian S, Patel P, Jani AB, Curran WJ, Liu T, Yang X. 2020. Pelvic multi-organ segmentation on cone-beam CT for prostate adaptive radiotherapy. *Med Phys*. 47(8):3415–3422. doi: [10.1002/mp.14196](https://doi.org/10.1002/mp.14196).
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Bengio Y. 2020. Generative adversarial networks. *Commun ACM*. 63(11):139–144. doi: [10.1145/3422622](https://doi.org/10.1145/3422622).
- Gui J, Sun Z, Wen Y, Tao D, Ye J. 2021. “A Review on Generative Adversarial Networks: Algorithms, Theory, and Applications”. *IEEE transactions on knowledge and data engineering*. 35(4):3313–3332. doi:[10.1109/TKDE.2021.3130191](https://doi.org/10.1109/TKDE.2021.3130191).
- Guo S, Sheng X, Chen H, Zhang J, Peng Q, Wu M, Fischer K, Tasian GE, Fan Y, Yin S. 2025. A novel cross-modal data augmentation method based on contrastive unpaired translation network for kidney segmentation in ultrasound imaging. *Med Phys*. 1–11. doi: [10.1002/mp.17663](https://doi.org/10.1002/mp.17663).
- He K, Zhang X, Ren S, Sun J. 2016. Deep residual learning for image recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV. p. 770–778. doi: [10.1109/CVPR.2016.90](https://doi.org/10.1109/CVPR.2016.90).
- Hong J, Yu SC-H, Chen W. 2022. Unsupervised domain adaptation for cross-modality liver segmentation via joint adversarial learning and self-learning. *Appl Soft Comput*. 121:108729. doi: [10.1016/j.asoc.2022.108729](https://doi.org/10.1016/j.asoc.2022.108729).
- Huo Y, Xu Z, Bao S, Assad A, Abramson RG, Landman BA. 2018 Apr. Adversarial synthesis learning enables segmentation without target modality ground truth. 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018) Washington, DC; IEEE. p. 1217–1220. doi: [10.1109/ISBI.2018.8363790](https://doi.org/10.1109/ISBI.2018.8363790).
- Iqbal T, Ali H. 2018. Generative adversarial network for medical images (MI-GAN). *J Med Syst*. 42(11):1–11. doi: [10.1007/s10916-018-1072-9](https://doi.org/10.1007/s10916-018-1072-9).
- Isola P, Zhu J-Y, Zhou T, Efros AA. 2017. “Image-to-Image Translation with conditional adversarial networks,”. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR); Honolulu, HI, USA; 5967–5976. doi: [10.1109/CVPR.2017.632](https://doi.org/10.1109/CVPR.2017.632).
- Johnson J, Alahi A, Fei-Fei L. 2016. Perceptual losses for real-time style transfer and super-resolution. *Computer Vision– ECCV 2016: 14th European Conference, Proceedings, Part II 14*; Oct 11–14; Amsterdam, The Netherlands: Springer International Publishing. p. 694–711.
- Kavur AE, Gezer NS, BarŞ M, Aslan S, Conze PH, Groza V, Pham DD, Chatterjee S, Ernst P, Özkan S, et al. 2021. CHAOS challenge-combined

- (CT-MR) healthy abdominal organ segmentation. *Med Image Anal.* 69:101950. doi: [10.1016/j.media.2020.101950](https://doi.org/10.1016/j.media.2020.101950).
- Lei Y, Wang T, Tian S, Dong X, Jani AB, Schuster D, Curran WJ, Patel P, Liu T, Yang X. 2020 Feb 4. Male pelvic multi-organ segmentation aided by CBCT-based synthetic MRI. *Phys Med Biol.* 65(3):035013. doi: [10.1088/1361-6560/ab63bb](https://doi.org/10.1088/1361-6560/ab63bb).
- Liu X, Pan J, Li X, Wei X, Liu Z, Pan Z, Tang J. 2024. Feb. Attention based cross-domain synthesis and segmentation from unpaired medical images. *IEEE Trans Emerging Top Comput Intel.* 8(1):917–929. doi: [10.1109/TETCI.2023.3296499](https://doi.org/10.1109/TETCI.2023.3296499).
- Liu X, Yu A, Wei X, Pan Z, Tang J. 2020. Oct. Multimodal MR image synthesis using gradient prior and adversarial learning. *IEEE J Sel Top Signal Process.* 14(6):1176–1188. doi: [10.1109/JSTSP.2020.3013418](https://doi.org/10.1109/JSTSP.2020.3013418).
- Mirza M, Osindero S. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
- Odena A, Olah C, Shlens J. 2017 July. Conditional image synthesis with auxiliary classifier GANs. *International conference on machine learning*, Sydney, Australia; PMLR. p. 2642–2651.
- Park T, Efros AA, Zhang R, Zhu J-Y. 2020. Contrastive learning for unpaired image-to-image translation. *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, Proceedings, Part IX; Aug 23– 28; Berlin, Heidelberg: Springer-Verlag.* p. 319–345.
- Radford A, Metz L, Chintala S. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- Ronneberger O, Fischer P, Brox T. 2015. U-net: convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Proceedings, Part III 18, Oct 5–9; Munich, Germany: Springer International Publishing.* p. 234–241.
- Shen Z, Chen Y, Huang TS, Zhou SK, Georgescu B, Liu X. 2020. "One-to-one mapping for unpaired image-to-image translation,". 2020 IEEE Winter Conference on Applications of Computer Vision (WACV); Snowmass, CO, USA; 1159–1168, doi: [10.1109/WACV45572.2020.9093622](https://doi.org/10.1109/WACV45572.2020.9093622).
- Shen Z, Zhou SK, Chen Y, Georgescu B, Liu X, Huang T. 2020. One-to-one mapping for unpaired image-to-image translation. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision.* p. 1170–1179.
- Stallard J. 2019 May 10. CT Vs MRI: what's the difference? And how do doctors choose which imaging method to use? *Memorial Sloan Kettering Cancer Center*; [accessed 2023 Mar 31]. <https://www.mskcc.org/news/ct-vs-mri-what-s-difference-and-how-do-doctors-choose-which-imaging-method-use>.
- Waheed A, Goyal M, Gupta D, Khanna A, Al-Turjman F, Pinheiro PR. 2020. Covidgan: data augmentation using auxiliary classifier gan for improved COVID-19 detection. *IEEE Access.* 8:91916–91923. doi: [10.1109/ACCESS.2020.2994762](https://doi.org/10.1109/ACCESS.2020.2994762).
- Wang C, Macnaught G, Papanastasiou G, MacGillivray T, Newby D. 2018. Unsupervised learning for cross-domain medical image synthesis using deformation invariant cycle consistency networks. *Simulation and Synthesis in Medical Imaging: Third International Workshop, SASHIMI 2018, Held in Conjunction with MICCAI 2018, Proceedings 3; Sep 16; Granada, Spain: Springer International Publishing.* p. 52–60.
- Wang C, Papanastasiou G, Tsaftaris S, Yang G, Gray C, Newby D, MacGillivray T. 2019. Tpsdicyc: improved deformation invariant cross-domain medical image synthesis. *Machine Learning for Medical Image Reconstruction: Second International Workshop, MLMIR 2019, Held in Conjunction with MICCAI 2019, Proceedings 2; Oct 17; Shenzhen, China: Springer International Publishing.* p. 245–254.
- Yang Q, Li N, Zhao Z, Fan X, Chang EC, Xu Y. 2018. MRI image-to-image translation for cross-modality image registration and segmentation. *arXiv preprint arXiv:1801.06940*.
- Yi X, Walia E, Babyn P. 2019. Generative adversarial network in medical imaging: a review. *Med Image Anal.* 58:101552. doi: [10.1016/j.media.2019.101552](https://doi.org/10.1016/j.media.2019.101552).
- Zhang H, Li H, Dillman JR, Parikh NA, He L. 2022. Multi-contrast MRI image synthesis using switchable cycle-consistent generative adversarial networks. *Diagnostics.* 12(4):816. doi: [10.3390/diagnostics12040816](https://doi.org/10.3390/diagnostics12040816).
- Zhang Z, Yang L, Zheng Y. 2018. Translating and segmenting multimodal medical volumes with cycle- and shape-consistency generative adversarial network. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition; Salt Lake City, (UT), USA. p. 9242–9251.
- Zhao C, Shuai R, Ma L, Liu W, Hu D, Wu M. 2021. Dermoscopy image classification based on StyleGAN and DenseNet201. *IEEE Access.* 9:8659–8679. doi: [10.1109/ACCESS.2021.3049600](https://doi.org/10.1109/ACCESS.2021.3049600).
- Zhu J-Y, Park T, Isola P, Efros AA. 2017. "Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks,". 2017 IEEE International Conference on Computer Vision (ICCV); Venice, Italy; 2242–2251, doi: [10.1109/ICCV.2017.244](https://doi.org/10.1109/ICCV.2017.244).