MR-to-CT Synthesis using Cycle-Consistent Generative Adversarial Networks

Jelmer M. Wolterink*
Image Sciences Institute
UMC Utrecht
Utrecht, The Netherlands

Anna M. Dinkla
Department of Radiotherapy
UMC Utrecht
Utrecht, The Netherlands

Mark H.F. Savenije
Department of Radiotherapy
UMC Utrecht
Utrecht, The Netherlands

Peter R. Seevinck
Image Sciences Institute
UMC Utrecht
Utrecht, The Netherlands

Cornelis A.T. van den Berg
Department of Radiotherapy
UMC Utrecht
Utrecht, The Netherlands

Ivana Išgum
Image Sciences Institute
UMC Utrecht
Utrecht, The Netherlands

1 Introduction

Radiotherapy treatment planning requires a magnetic resonance (MR) volume for segmentation of tumors and organs at risk, and a spatially corresponding computed tomography (CT) volume for dose planning. Separate acquisition of these volumes is time-consuming, costly and a burden to the patient. Alternatively, a synthetic CT image derived from an MR image would enable MR-only radiotherapy treatment planning, thereby omitting the additional acquisition of a CT scan [1]. Previously proposed convolutional neural networks (CNNs) for CT synthesis required pairs of aligned MR and CT images for training based on a voxel-wise loss [2, 4, 5]. However, the availability of MR-CT training pairs may be limited and misalignment within training pairs may cause blurring in synthesized CT images.

We leverage recent advances in generative adversarial network (GAN) research and propose to use a CycleGAN model for CT synthesis [6], which can be trained without the need for paired training data and voxel-wise correspondence between MR and CT. Instead, the synthesis CNN is trained based on the overall quality of the synthesized image as determined by an adversarial discriminator CNN. To prevent synthesis of images that look real but bear little similarity to the input image, an additional synthesis CNN translates the synthesized image back into an image that is as similar as possible to the original image.

2 Materials and Methods

This study included brain MR and CT images of 24 patients who underwent radiotherapy treatment planning for brain tumors. For evaluation purposes, MR and CT images of the same patient were aligned using rigid registration based on mutual information following a clinical procedure and resampled to a $183 \times 288 \times 288$ voxel volume with a voxel size of $1.00 \times 0.87 \times 0.87$ mm^3. A head mask excluding surrounding air was obtained in the CT image and propagated to the MR image.

We used a CycleGAN model with a forward and a backward cycle (Fig. 1) [6], each containing three CNNs. In the forward cycle, CNN Syn_{CT} translates an input MR image I_{MR} into a CT image, Syn_{MR} translates a synthesized CT image $Syn_{CT}(I_{MR})$ back into an MR image, and Dis_{CT} discriminates between synthesized $Syn_{CT}(I_{MR})$ and real CT images I_{CT}. The backward cycle is trained jointly with the forward cycle, translating CT images into MR images and back into CT images. The backward cycle uses the same CNNs Syn_{CT} and Syn_{MR}, and an additional discriminator network Dis_{MR} that aims to distinguish synthesized MR images from real MR images.

*j.m.wolterink@umcutrecht.nl

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

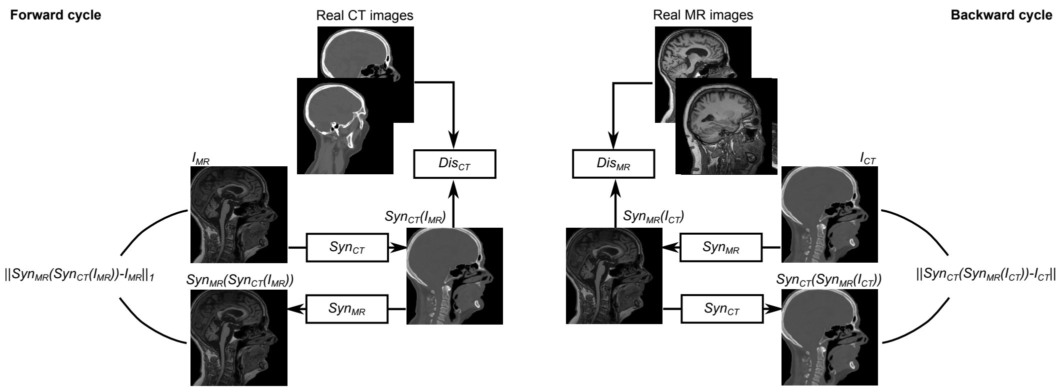

Figure 1: The CycleGAN model consists of a forward cycle and a backward cycle. In the forward cycle, an input MR image I_{MR} is transformed into a CT image, compared with reference CT images and transformed back into an MR image. In the backward cycle, an input CT image is transformed into an MR image, compared with reference MR images and transformed back into a CT image.

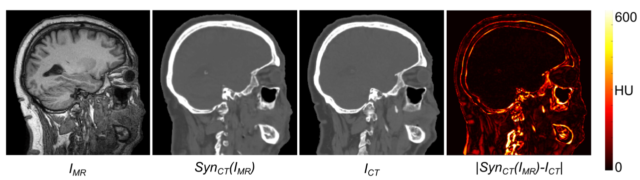

Figure 2: Input MR, synthesized and reference CT images, and absolute error between the CT images.

The synthesis and discriminator networks are optimized using adversarial objectives. While Dis_{CT} and Dis_{MR} aim to distinguish synthesized CT or MR images from real images, networks Syn_{CT} and Syn_{MR} try to synthesize images that cannot be distinguished from real images. In addition, Syn_{CT} and Syn_{MR} aim to minimize the difference of images that are transformed back to their original domain with respect to the original image.

3 Experiments and Results

The model was trained using images of 18 patients and tested using images of 6 patients. From each volume, 183 sagittal 2D images were extracted. The PyTorch implementation provided by the authors of [6] was used in all experiments[2]. Parameter optimization using Adam [3] took 52 hours on a single NVIDIA Titan X GPU. MR to CT synthesis with a trained model took around 10 seconds.

Figure 2 shows an example MR input image, the output of CT synthesis, and the corresponding reference image. The model learned to differentiate between structures with similar intensity values in MR but not in CT, such as bone, cerebrospinal fluid and air. The difference image shows that differences may be partly caused by the reduced image quality in the neck area and misalignment between the MR image and the reference CT image. The average mean absolute error (MAE) between synthesized and reference CT images was 73.7 $\pm$ 2.3 HU. In comparison, errors were substantially higher for a model trained with paired training data and a voxel-wise loss term, which resulted in an average MAE of 89.4 $\pm$ 6.8 HU. In addition, we found that MR images reconstructed from synthesized CT images were remarkably similar to the original MR image.

4 Discussion and Conclusion

We have shown that a CNN can be trained to synthesize a CT image from an MR image using unpaired and unaligned MR and CT training images. In contrast to previous work, the model learns to synthesize realistically-looking images guided only by the performance of an adversarial discriminator network and the similarity of back-transformed output images to the original input image.

[2] https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix

References

[1] Jens M. Edmund and Tufve Nyholm. A review of substitute CT generation for MRI-only radiation therapy. *Radiat. Oncol.*, 12(1):28, 2017.

[2] Xiao Han. MR-based synthetic CT generation using a deep convolutional neural network method. *Med. Phys.*, 44(4):1408–1419, 2017.

[3] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

[4] Dong Nie, Xiaohuan Cao, Yaozong Gao, Li Wang, and Dinggang Shen. Estimating CT image from MRI data using 3D fully convolutional networks. In *MICCAI DLMIA Workshop*, volume 10008 of *LNCS*, pages 170–178, Cham, 2016. Springer International Publishing.

[5] Dong Nie, Roger Trullo, Jun Lian, Caroline Petitjean, Su Ruan, Qian Wang, and Dinggang Shen. Medical image synthesis with context-aware generative adversarial networks. In *MICCAI*, volume 10435 of *LNCS*, pages 417–425, Cham, 2017. Springer International Publishing.

[6] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.