3D Medical Image Synthesis using Generative Adversarial Networks

Irina Sánchez
Department of Signal Theory and Communications,
Technical University of Catalonia (UPC)
Barcelona, Spain 08034
irina.sanchez@alu-etsetb.upc.edu

Verónica Vilaplana
Department of Signal Theory and Communications,
Technical University of Catalonia (UPC)
Barcelona, Spain 08034
veronica.vilaplana@upc.edu

ABSTRACT

In this work we propose an architecture for 3D medical image synthesis based on Generative Adversarial Networks.

ACM Reference format:
Irina Sánchez and Verónica Vilaplana. 2017. 3D Medical Image Synthesis using Generative Adversarial Networks. In *Proceedings of womENcourage 2017, Barcelona, Spain, September 2017,* 1 pages.

1 INTRODUCTION

Medical image synthesis is useful for many different tasks. One application is super-resolution (SR) imaging, where the goal is to create a high-resolution (HR) image from a low-resolution one (LR). Due to economical, technological or physical limitations, it may not be easy to obtain images at the desired resolution . A solution is to use SR techniques, since HR images can significantly improve the diagnosis ability for treatment and also automatic detection and segmentation results. Another application of image synthesis could be the estimation of Computed Tomography (CT) from MRI.

While different approaches have been proposed to solve these problems, solutions based on deep learning models are currently being investigated motivated by their success in many computer vision tasks. In particular, Generative Adversarial Networks (GAN) [4] are a very promising approach for image generation, and have been used for super-resolution [2] and CT synthesis [3].

Generative Adversarial Networks propose a scenario with two different players, a Discriminant (D) and a Generator (G). G generates images from a uniform or Gaussian distribution G(**z**), while D tries to recognize if an image is from the input data or it is generated by G. G will have to learn how to cheat D, making the images perceptually closer to the input data, while D will have to learn features of the data in order to recognize efficiently the real samples from the generated ones.

Recently, different architectures and loss functions that try to improve the quality of the images generated using GANs have been presented [1][2] [5]. This motivates us to use some of these ideas in order to address a super-resolution problem by using a 3D SRGAN architecture. In future works, this same network could be used for CT-MRI synthesis.

This work has been partially supported by the project MALEGRA, TEC2016-75976-R, financed by the Spanish Ministerio de Economía, Industria y Competitividad and the European Regional Development Fund (ERDF)
womENcourage 2017, September 2017, Barcelona, Spain

2 METHOD

The architecture proposed is a 3D CNN highly influenced by the SRGAN network [2]. The discriminator network is composed by 8 convolutional layers with increasing number of 3x3x3 filter kernels, increasing by a factor of 2 from 64 to 512 filters. Strided convolutions are used to reduce the image resolution each time the number of features is doubled. After the convolutional layers, we can find two dense layers and a final sigmoid activation, that will output a probability saying if the input image is real or fake. In the case of the generator, we use 6 residual blocks composed by a convolution of 32 filters of size 3x3x3, batch normalization, a LeakyReLU activation and another convolution with the same parameters and batch normalization. For the upsampling phase we perform a convolution followed by a LeakyReLU activation and then a tile replication followed by another convolution and activation.

A very critical point is the definition of the cost function; LS-GAN [5] states that the sigmoid cross entropy loss function causes problems of vanishing gradients. Using a loss based on Least Square Error, will make the generator to generate samples that are closer to the real ones: the images that are far away from the decision boundary, even being well classified, will be penalized because they are perceptually far from the real images. We take advantage of that work to define the losses of both agents, D and G:

Discriminative loss:

$$l_D = \frac{1}{2}[D(x) - 1]^2 + \frac{1}{2}[D(G(z)]^2 \tag{1}$$

Generative loss: We use two different terms for the loss function of G, a content loss that is responsible of generating images that are similar to the real ones and an adversarial term that is trying to cheat the discriminator:

$$l_G = w_{content} * l_{content} + w_{adv} * [D(G(z) - 1]^2 \tag{2}$$

Currently, our system is producing promising results and we are working on the optimization of its hyperparameters. SR images generated by our network will be shown in the poster.

REFERENCES

[1] L. Metz A. Radford and S. Chintala. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *Interna- tional Conference on Learning Representations (ICLR)* (2015).

[2] F. Huszar J. Caballero A. Cunningham A. Acosta A. Aitken A. Tejani J. Totz Z. Wang C. Ledig, L. Theis and W. Shi. 2016. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. (2016).

[3] C. Petitjean S. Ruan D. Nie, R. Trullo and D.Shen. 2016. Medical Image Synthesis with Context-Aware Generative Adversarial Networks. (2016).

[4] M. Mirza B. Xu D. Warde-Farley S. Ozair A. Courville I. Goodfellow, J. Pouget-Abadie and Y. Bengio. 2014. Generative adversarial nets. *Advances in Neural Information Processing Systems (NIPS)* (2014), 2672–2680.

[5] H. Xie RYK.Lau X.Mao, Q. Li and Z. Wang. 2017. Least Squares Generative Adversarial Networks. (2017).