

Capturing Variabilities from Computed Tomography Images with Generative Adversarial Networks

Umair Javaid^{1,2} and John A. Lee^{1,2} *

1- Université Catholique de Louvain - ICTEAM
Place du Levant 3 L5.03.02, 1348 Louvain-la-Neuve - Belgium

2- Université Catholique de Louvain - IREC/MIRO
Avenue Hippocrate 55 B1.54.07, 1200 Brussels - Belgium

Abstract. With the advent of Deep Learning (DL) techniques, especially Generative Adversarial Networks (GANs), data augmentation and generation are quickly evolving domains that have raised much interest recently. However, the DL techniques are data demanding and since, medical data is not easily accessible, they suffer from data insufficiency. To deal with this limitation, different data augmentation techniques are used. Here, we propose a novel unsupervised data-driven approach for data augmentation that can generate 2D Computed Tomography (CT) images using a simple GAN. The generated CT images have good global and local features of a real CT image and can be used to augment the training datasets for effective learning. In this proof-of-concept study, we show that our proposed solution using GANs is able to capture some of the global and local CT variabilities. Our network is able to generate visually realistic CT images and we aim to further enhance its output by scaling it to a higher resolution and potentially from 2D to 3D.

1 Introduction

Deep learning algorithms are data demanding; meaning more data leads to more effectiveness. Medical image classification and segmentation tasks using deep learning often have limited data due to scarcity of annotated data. This is particularly true in the health care industry where access to medical data is strictly protected due to privacy concerns. To bridge this limitation, data augmentation techniques can be used.

A common practice for image data augmentation consists of geometric augmentations such as reflection, rotation, scaling and translation [1] of the input images. Other approaches are the use of medical phantoms [2], [3] in which medical data is simulated with a certain degree of realism. These approaches however, are implemented manually so as to extend the training datasets. Therefore, there is a need of an automated data augmentation technique, which is generative in nature and offers an acceptable degree of realism.

Recently, GANs [4] have raised much interest in the domain of data generation and are being used extensively for data augmentation. A standard GAN

*Umair Javaid is funded by the F.N.R.S. Télévie grant no. 7.4625.16. John A. Lee is a Senior Research Associate with the F.N.R.S.

consists of two convolutional neural networks namely, Discriminator D and a Generator G , which are trained to learn simultaneously. Following the minimax strategy, one network (D) learns how to distinguish real samples from the ones generated by the generator where G produces samples that can trick D into believing they are real.

In this proof-of-concept study, we address the challenges of data augmentation by proposing an automatic pipeline that is generative and offers a certain degree of realism with the medical data. We explore the potential of GANs to capture the variabilities in CT images and generate identically looking CT images. The generated images can be used to augment the training datasets which automate the augmentation processing and potentially leverage the scarcity of medical data.

The dataset examined in this study is the Lung CT images dataset publicly available at the Lung Image Database Consortium image collection (LIDC) [5]. The dataset consists of 267 Lung CT scans of dimensions 512×512 .

Our contribution in this study is two-fold. First, we propose an automated unsupervised data-driven approach using a simple GAN that can generate realistic CT images, which can be used to augment the training datasets for effective learning. To our knowledge, this is the first implementation for generating full CT image from a lower-resolution (in our case, 40×40 resolution). Second, using a simple GAN, we propose a robust framework that can be scaled to a higher resolution without changing any parameters of the model.

The remainder of this paper is organized as follows. Section 2 reviews the related work and outlines recent advancements in the domain of data generation using GANs. We present the proposed Methodology in Section 3. Experimental results are shown in Section 4, followed by Conclusion and Future Perspectives in Sections 5.

2 Related Work

Goodfellow et al. [4] proposed an adversarial approach to learn deep generative models namely, Generative Adversarial Networks (GANs). A number of GAN architectures have been proposed to enable stable training of generative models for producing realistic images. Recent progresses like Radford et al. [6] have shown improved performances on synthesizing realistic images by using multiple strided and fractionally strided convolutional layers for both discriminator and generator networks. D. Jiwoong et al. [7] make use of recurrent neural networks to train a generative model in order to produce high-quality images.

Other than designing different GAN architectures, several authors have tried to rectify GAN output by improving the training cycle with different criteria. For example, enhanced loss functions that aid in better convergence and learning of GANs. Moreover, Zhao et al. [8] presented an energy-based GAN that offers optimal discriminator learning by minimizing an energy function where the energy is computed by an auto-encoder structured discriminator. Other researchers like Nowozin et al. [9] and Chen et al. [10] analyze GANs from an

information theory perspective; learning representations that capture different latent factors to generate data samples.

GJ. Qi [11] has recently proposed a solution using an advanced loss function. Rather than having a fixed loss function, the paper shows that a loss function can be learnt and allows better separation between generated and real examples. Arjovsky et al. [12] proposed WGAN in which they use Wasserstein distance (a.k.a. Earth-Mover distance). It offers stable training of GANs but increases the training time. Bertholet et al. [13] recently proposed Boundary Equilibrium GAN by proposing a new equilibrium enforcing method that is paired with a loss derived from the Wasserstein distance to train GAN.

There are several studies that use GANs to produce medical data. For example, Nie et al. [14] use GANs to estimate CT images from MR images. Guibas et al. [15] make use of a dual-stage GAN to generate retinal fundi images. Chuquicusma et al. [16] use GANs to generate realistic lung nodule samples. However, no study has been done for generating a full CT image and to our knowledge, this is the first attempt to generate full CT images from low-resolution input images.

In this study, we show that our simple GAN with a standard loss function is able to learn the CT image features and can generate almost realistic CT images from low-resolution input images.

3 Methodology

GANs are generative models that can learn mapping from a random noise vector z to an output image. Generator G takes as input a random noise vector z (uniformly distributed) and outputs image $X_{\text{fake}} = G(Z)$. Depending upon the feedback of D , G improves its output images until the D cannot distinguish them from real images. Simultaneously, D receives as input either the real (training image) or fake sample (generated image) and classifies the data produced by G either as real or fake. The generator is initialized with a random noise vector while D is trained with a small set of real images (training data). The objective of a GAN can be expressed as $\min_G \max_D V(D, G)$ with

$$V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))], \quad (1)$$

where G tries to minimize this objective against an adversarial D that tries to maximize it.

Our proposed GAN is a simple architecture with few number of layers in both the D and G networks that use the standard Cross Entropy loss function.

D consists of 4 convolutional layers and two dense layers whereas G is a fully convolutional network - composed of 6 convolutional layers. Following [6], both networks use LeakyReLU as activation function except the last layer of D that uses Sigmoid activation whereas, Linear activation is used at G output. Dropout is used in D and batch normalization is used in G .

The detailed architecture of our network is as follows:

Discriminator

1. Conv with 256 filters and 3×3 kernel. LeakyReLU.
 - Strided convolution with stride of 2.
 - Dropout.
2. Conv with 128 filters and 3×3 kernel. LeakyReLU.
 - Dropout.
3. Conv with 64 filters and 3×3 kernel. LeakyReLU.
 - Dropout.
4. Conv with 32 filters and 3×3 kernel. LeakyReLU.
 - Dropout.
5. Fully connected layer with output dimension 128. LeakyReLU.
6. Fully connected layer with output dimension 1. Sigmoid activation.

Generator

1. Transposed Conv with 256 filters and 3×3 kernel. Stride = 2.
 - Batch Normalization. LeakyReLU.
2. Transposed Conv with 128 filters and 3×3 kernel. Stride = 2.
 - Batch Normalization. LeakyReLU.
3. Transposed Conv with 64 filters and 3×3 kernel.
 - Batch Normalization. LeakyReLU.
4. Transposed Conv with 32 filters and 3×3 kernel.
 - Batch Normalization. LeakyReLU.
5. Transposed Conv with 16 filters and 3×3 kernel.
 - Batch Normalization. LeakyReLU.
6. Transposed Conv with 3 filters and 3×3 kernel. Linear Activation.

4 Experiments and Results

The input images were downsampled from 512×512 to 40×40 . As a baseline analysis, only 32 CT images were used as input. Afterwards, to test the learning capability of our network, it was trained on 76 and 129 images. No pre-processing was applied to the training images besides intensity scaling to the range $[0, 2]$.

All networks were trained from the scratch using mini-batch stochastic gradient descent (SGD) with a mini-batch size of 16. All weights were initialized from a zero-centered Gaussian distribution with standard deviation 0.02. G uses a noise vector to represent all the variations. The noise vector used was 100 dimensional, uniformly distributed in the range $[-1, 1]$.

In the LeakyReLU, the slope of the leak was set to 0.2 in both networks. RMSProp was used as an optimizer where the learning rate for discriminator was set to 0.0001 and for generator it was set to 0.0002. A small value of L2 norm (i.e. 10^{-5}) was also used for both networks to avoid over-fitting. The value of batch normalization was set to 0.9 whereas a dropout of 0.6 was used in the discriminator network.

Results from some of our experiments are given as follows. A comparison between the different convolutional kernel size is shown in Fig. 1 followed by the effect of using RMSProp and ADAM as optimizers in Fig. 2. Fig. 3 shows the scalability of our network to a higher resolution i.e., from 40×40 to 64×64 .

We implemented our GAN in TensorFlow and for all the experiments, it was trained on a single NVIDIA GTX 1070 GPU. Training our model is not expensive, all the results shown were generated after 1 hour of training (30,000 epochs).



Fig. 1: Effect of kernel size on the generated images. Leftmost is the ground truth, middle is generated using 3×3 kernel size and the rightmost image is generated using 5×5 kernel size.



Fig. 2: Choice of optimizer influences the generated images. Leftmost is the ground truth, middle is generated using RMSProp and the rightmost image is generated using ADAM.



Fig. 3: Our proposed network is able to scale to a higher resolution without changing any parameter(s). Leftmost is the ground truth, middle is 40×40 and rightmost image is the generated 64×64 image.

5 Conclusion and Future Perspectives

In this study, we proposed an automated augmentation framework using GANs. We showed that our approach is able to generate visually realistic CT images from low-resolution input images.

As a proof-of-concept, we used a very simple GAN architecture that is able to capture CT variabilities. We showed that our simple GAN offers good scalability to higher resolution images without changing any parameters.

In future, we want to scale our network to higher resolutions, do a comparative study between a number of recent GAN architectures and an image quality analysis. As CT images are volumetric, we will also explore image generation in the 3D space which will retain the information lost in 2D space. Once higher resolution images are generated, we aim to use them in our training datasets for medical image segmentation tasks using Deep Learning.

References

- [1] R.C Gonzalez et al. Morphological image processing. *Digital image processing*, 3:627–688, 2008.
- [2] L. E Court et al. Use of a realistic breathing lung phantom to evaluate dose delivery errors. *Medical physics*, 37(11):5850–5857, 2010.
- [3] J Chang et al. Development of a deformable lung phantom for the evaluation of deformable registration. *Journal of applied clinical medical physics*, 11(1):281–286, 2010.
- [4] I Goodfellow et al. Generative Adversarial Nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [5] S G Armato III et al. Lung image database consortium: developing a resource for the medical imaging research community. *Radiology*, 232(3):739–748, 2004.
- [6] A Radford et al. Unsupervised representation learning with deep convolutional Generative Adversarial Networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [7] D J Im et al. Generating images with recurrent adversarial networks. *arXiv preprint arXiv:1602.05110*, 2016.
- [8] J Zhao et al. Energy-based Generative Adversarial Network. *arXiv preprint arXiv:1609.03126*, 2016.
- [9] S Nowozin et al. f-GAN: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, pages 271–279, 2016.
- [10] X Chen et al. InfoGAN: Interpretable representation learning by information maximizing Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, pages 2172–2180, 2016.
- [11] GJ Qi. Loss-sensitive Generative Adversarial Networks on Lipschitz densities. *arXiv preprint arXiv:1701.06264*, 2017.
- [12] M Arjovsky et al. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.
- [13] D Berthelot et al. BEGAN: Boundary Equilibrium Generative Adversarial Networks. *arXiv preprint arXiv:1703.10717*, 2017.
- [14] D Nie et al. Medical image synthesis with context-aware Generative Adversarial Networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 417–425. Springer, 2017.
- [15] J T Guibas et al. Synthetic medical images from dual Generative Adversarial Networks. *arXiv preprint arXiv:1709.01872*, 2017.
- [16] M JM Chuquicuma et al. How to fool radiologists with Generative Adversarial Networks? a visual turing test for lung cancer diagnosis. *arXiv preprint arXiv:1710.09762*, 2017.