



CT-Realistic Lung Nodule Simulation from 3D Conditional Generative Adversarial Networks for Robust Lung Segmentation

Dakai Jin, Ziyue Xu^(✉), Youbao Tang, Adam P. Harrison,
and Daniel J. Mollura

National Institutes of Health, Bethesda, MD, USA
ziyue.xu@nih.gov

Abstract. Data availability plays a critical role for the performance of deep learning systems. This challenge is especially acute within the medical image domain, particularly when pathologies are involved, due to two factors: (1) limited number of cases, and (2) large variations in location, scale, and appearance. In this work, we investigate whether augmenting a dataset with artificially generated lung nodules can improve the robustness of the progressive holistically nested network (P-HNN) model for pathological lung segmentation of CT scans. To achieve this goal, we develop a 3D generative adversarial network (GAN) that effectively learns lung nodule property distributions in 3D space. In order to embed the nodules within their background context, we condition the GAN based on a volume of interest whose central part containing the nodule has been erased. To further improve realism and blending with the background, we propose a novel multi-mask reconstruction loss. We train our method on over 1000 nodules from the LIDC dataset. Qualitative results demonstrate the effectiveness of our method compared to the state-of-art. We then use our GAN to generate simulated training images where nodules lie on the lung border, which are cases where the published P-HNN model struggles. Qualitative and quantitative results demonstrate that armed with these simulated images, the P-HNN model learns to better segment lung regions under these challenging situations. As a result, our system provides a promising means to help overcome the data paucity that commonly afflicts medical imaging.

Keywords: Lung nodule · CT · GAN · Dataset bottleneck
Lung segmentation

Z. Xu—This work is supported by NIH Intramural Research. We also thank NVidia for the donation of a Tesla K40 GPU.

1 Introduction

Deep learning has achieved significant recent successes. However, large amounts of training samples, which sufficiently cover the population diversity, are often necessary to produce high quality results. Unfortunately, data availability in the medical image domain, especially when pathologies are involved, is quite limited due to several reasons: significant image acquisition costs, protections on sensitive patient information, limited numbers of disease cases, difficulties in data labeling, and large variations in locations, scales, and appearances. Although efforts have been made towards constructing large medical image datasets, options are limited beyond using simple automatic methods [8], huge amounts of radiologist labor [1], or mining from radiologist reports [14]. Thus, it is still an open question on how to generate effective and sufficient medical data samples with limited or no expert-intervention.

One enticing alternative is to generate synthetic training data. However, historically synthetic data is less desirable due to shortcomings in realistically simulating true cases. Yet, the advent of generative adversarial networks (GANs) [4] has made game-changing strides in simulating real images and data. This ability has been further expanded with developments on fully convolutional [13] and conditional [10] GANs. In particular, Isola *et al.* extend the conditional GAN (CGAN) concept to predict pixels from known pixels [6]. Within medical imaging, Nie *et al.* use a GAN to simulate CT slices from MRI data [11], whereas Wolterink *et al.* introduce a bi-directional CT/MRI generator [15]. For lung nodules, Chuquicusma *et al.* train a simple GAN to generate simulated images from random noise vectors, but do not condition based on surrounding context [2].




Fig. 1. Lung nodule simulation using the 3D CGAN. (a) A VOI centered at a lung nodule; (b) 2D axial view of (a); (c) same as (b), but with central sphere region erased; (d-e) simulated lung nodule using a plain $L1$ reconstruction loss and the 3D CGAN with multi-mask $L1$ loss coupled with adversarial loss, respectively.

In this work, we explore using CGAN to augment training data for specific tasks. For this work, we focus on pathological lung segmentation, where the recent progressive holistically nested network (P-HNN) has demonstrated state of the art results [5]. However, P-HNN can struggle when there are relatively large (*e.g.*, >5 mm) peripheral nodules touching the lung boundary. This is mainly because these types of nodule are not common in Harrison *et al.*'s [5] training set. To improve P-HNN's robustness, we generate synthetic 3D lung

nodules of different sizes and appearances, at multiple locations, that naturally blend with surrounding tissues (see Fig. 1 for an illustration). We develop a 3D CGAN model that learns nodule shape and appearance distributions directly in 3D space. For the generator, we use a U-Net-like [3] structure, where the input to our CGAN is a volume of interest (VOI) cropped from the original CT image with the central part, containing the nodule, erased (Fig. 1(c)). We note that filling in this region with a realistic nodule faces challenges different than generating a random 2D nodule image from scratch [2]. Our CGAN must generate realistic and natural 3D nodules conditioned upon and consistent with the surrounding tissue information. To produce high quality nodule images and ensure their natural blending with surrounding lung tissues, we propose a specific multi-mask reconstruction loss that complements the adversarial loss.

The main contributions of this work are: (1) we formulate lung nodule generation using a 3D GAN conditioned on surrounding lung tissues; (2) we design a new multi-mask reconstruction loss to generate high quality realistic nodules alleviating boundary discontinuity artifacts; (3) we provide a feasible way to help overcome difficulties in obtaining data for “edge cases” in medical images; and (4) we demonstrate that GAN-synthesized data can improve training of a discriminative model, in this case for segmenting pathological lungs using P-HNN [5].

2 Methods

Figure 2 depicts an overview of our method. Below, we outline the CGAN formulation, architecture, and training strategy used to generate realistic lung nodules.

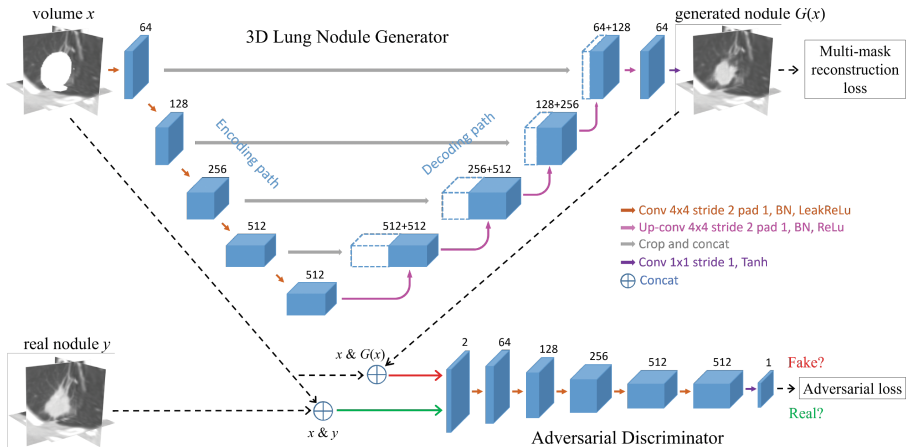


Fig. 2. 3D CGAN architecture for lung nodule generation. The input is the original CT VOI, y , containing a real nodule and the same VOI, x , with the central region erased. Channel numbers are placed next to each feature map.

2.1 CGAN Formulation

In their original formulation, GANs [4] are generative models that learn a mapping from a random noise vector z to an output image y . The generator, G , tries to produce outputs that fool a binary classifier discriminator D , which aims to distinguish real data from generated “fake” outputs. In our work, the goal is to generate synthetic 3D lung nodules of different sizes, with various appearances, at multiple locations, and have them naturally blend with surrounding lung tissues. For this purpose, we use a CGAN conditioned on the image x , which is a 3D CT VOI cropped from a specific lung location. Importantly, as shown in Fig. 1(c), we erase the central region containing the nodule. The advantage of this conditional setting is that the generator not only learns the distribution of nodule properties from its surrounding context, but it also forces the generated nodules to naturally fuse with the background context. While it is possible to also condition on the random vector z , we found it hampered performance. Instead, like Isola *et al.* [6], we use dropout to inject randomness into the generator.

The adversarial loss for CGANs can then be expressed as

$$L_{CGAN}(G, D) = E_{x,y}[\log D(x, y)] + E_x[\log(1 - D(x, G(x)))], \quad (1)$$

where y is the original VOI and G tries to minimize this objective against an adversarial discriminator, D , that tries to maximize it. Like others [6, 12], we also observe that an additional reconstruction loss is beneficial, as it provides a means to learn the latent representation from surrounding context to recover the missing region. However, reconstruction losses tend to produce blurred results because it tends to average together multiple modes in the data distribution [6]. Therefore, we combine the reconstruction and adversarial loss together, making the former responsible for capturing the overall structure of the missing region while the latter learns to pick specific data modes based on the context. We use the $L1$ loss, since the $L2$ loss performed poorly in our experiments.

Since the generator is meant to learn the distribution of nodule appearances in the erased region, it is intuitive to apply $L1$ loss only to this region. However, completely ignoring surrounding regions during generator’s training can produce discontinuities between generated nodules and the background. Thus, to increase coherence we use a new multi-mask $L1$ loss. Formally, let M be the binary mask where the erased region is filled with 1’s. Let N be a dilated version of M . Then, we assign a higher $L1$ loss weight to voxels where $N - M$ is equal to one:

$$L_{L1}(G) = E_{x,y}[M \odot \|y - G(x)\|_1] + \alpha E_{x,y}[(N - M) \odot \|y - G(x)\|_1], \quad (2)$$

where \odot is the element-wise multiplication operation and $\alpha \geq 1$ is a weight factor. We find that a dilation of 3 to 6 voxels generally works well. By adding the specific multi-mask $L1$ loss, our final CGAN objective is

$$G^* = \arg \min_G \max_D L_{CGAN}(G, D) + \lambda L_{L1}(G), \quad (3)$$

where α and λ are determined experimentally. We empirically find $\alpha = 5$ and $\lambda = 100$ works well in our experiments.

2.2 3D CGAN Architecture

Figure 2 depicts our architecture, which builds off of Isola *et al.*'s 2D work [6], but extends it to 3D images. More specifically, the generator consists of an encoding path with 5 convolutional layers and a decoding path with another 5 deconvolutional layers where short-cut connections are added in a similar fashion to U-net [3]. The encoding path takes an input VOI x with missing regions and produces a latent feature representation, and the decoding path takes this feature representation and produces the erased nodule content. We find that without shortcut connections, our CGAN models do not converge, suggesting that they are important for information flow across the network and for handling fine-scale 3D structures, confirmed by others [7]. To inject randomness, we apply dropout on the first two convolutional layers in the decoding path.

The discriminator also contains an encoding path with 5 convolutional layers. We also follow the design principles of Radford *et al.* [13] to increase training stability, which includes strided convolutions instead of pooling operations, LeakyReLU's in the encoding path of G and D , and a Tanh activation for the last output layer of G .

2.3 CGAN Optimization

We train the CGAN model end-to-end. To optimize our networks, we use the standard GAN training approach [4], which alternates between optimizing G and D , as we found this to be the most stable training regimen. As suggested by Goodfellow *et al.* [4], we train G to maximize $\log D(x, G(x))$ rather than minimize $\log(1 - D(x, G(x)))$. Training employs the Adam optimizer [9] with a learning rate 0.0001 and momentum parameters $\beta_1 = 0.5$ and $\beta_2 = 0.999$ for both the generator and discriminator.

3 Experiments and Results

We first validate our CGAN using the LIDC dataset [1]. Then, using artificially generated nodules, we test if they can help fine-tune the state-of-the-art P-HNN pathological lung segmentation method [5].

3.1 3D CGAN Performance

The LIDC dataset contains 1018 chest CT scans of patients with observed lung nodules, totaling roughly 2000 nodules. Out of these, we set aside 22 patients and their 34 accompanying nodules as a test set. For each nodule, there can be multiple radiologist readers, and we use the union of the masks for such cases. True nodule images, y , are generated by cropping cubic VOIs centered at each nodule with 3 random scales between 2 and 2.5 times larger than the maximum dimension of the nodule mask. All VOIs are then resampled to a fixed size of $64 \times 64 \times 64$. Conditional images, x , are derived by erasing the pixels within a

sphere of diameter 32 centered at the VOI. We exclude nodules whose diameter is less than 5 mm, since small nodules provide very limited contextual information after resampling and our goal is to generate relatively large nodules. This results in roughly 4300 training sample pairs. We train the cGAN for 12 epochs.

We tested against three variants of our method: (1) only using an all-image $L1$ loss; (2) using both the adversarial and all-image $L1$ loss. This is identical to Isola *et al.*'s approach [6], except extended to 3D; (3) using the same combined objective in (3), but not using the multi-mask version, *i.e.*, only using the first term of equation (2). As reconstruction quality hinges on subjective assessment [6], we visually examine nodule generation on our test set. Selected examples are shown in Fig. 3. As can be seen, our proposed CGAN produces realistic nodules of high quality with various shapes and appearances that naturally blend with surrounding tissues, such as vessels, soft tissue, and parenchyma. In contrast, when only using the reconstruction $L1$ loss, results are considerably blurred with very limited variations in shape and appearance. Results from Isola *et al.*'s method [6] improve upon the $L1$ only loss; however, it has obvious inconsistencies/misalignments with the surrounding tissues and undesired sampling artifacts that appear inside the nodules. It is possible that by forcing the generator to reconstruct the entire image, it distracts the generator from learning the nodule appearance distribution. Finally, when only performing the $L1$ loss on the erased region, the artifacts seen in Isola *et al.*'s are not exhibited; however, there are stronger border artifacts between the M region and the rest of the VOI. In contrast, by incorporating a multi-mask loss, our method can produce nodules with realistic interiors and without such border artifacts.




Fig. 3. Examples results: (a) original images; (b) input after central region erased; (c) only $L1$ loss, applied to the entire image; (d) Isola *et al.*'s method [6]; (e) CGAN with $L1$ loss, applied only to the erased region; (f) our CGAN with multi-mask $L1$ loss.

3.2 Improving Pathological Lung Segmentation

With the CGAN trained, we test whether our CGAN benefits pathological lung segmentation. In particular, the P-HNN model shared by Harrison *et al.* [5] can struggle when peripheral nodules touch the lung boundary, as these were not well represented in their training set. Prior to any performed experiments, we selected 34 images from the LIDC dataset exhibiting such peripheral nodules. Then, we randomly chose 42 LIDC subjects from relatively healthy subjects with no large nodules. For each of these, we pick 30 random VOI locations, centering within (8,20)mm to the lung boundary with random size ranging (32, 80)mm. VOIs are resampled to $64 \times 64 \times 64$ voxels and simulated lung nodules are generated in each VOI, using the same process as in Sect. 3.1, except the trained CGAN is only used for inference. The resulting VOIs are resampled back to their original resolution and pasted back to the original LIDC images, and then the axial slices containing the simulated nodules are used as training data (~ 10000 slices) to fine-tune the P-HNN model for 4-5 epochs. For comparison, we also fine-tune P-HNN using images generated by the $L1$ -only loss and also Isola *et al.*'s CGAN.




Fig. 4. Lung segmentation results on LIDC patients with peripheral lung nodules. All metrics are measured on a size 64 pixel VOI centered on the nodule.




Fig. 5. Examples P-HNN lung segmentations. (a) ground truth; (b) original model; (c-e) fine-tuned with $L1$ loss only, Isola *et al.* [6], and proposed CGAN, respectively.

Figure 4 depicts quantitative results. First, as the chart demonstrates, fine-tuning using *all* CGAN variants improves P-HNN’s performance on peripheral lung nodules. This confirms the value in using simulated data to augment training datasets. Moreover, the quality of nodules is also important, since the results using nodules generated by only an all-image $L1$ loss have the least improvement. Importantly, out of all alternatives, our proposed CGAN produces the greatest improvements in Dice scores, Hausdorff distances and average surface distances. For instance, our proposed CGAN allows P-HNN’s mean Dice scores to improve from 0.964 to 0.989, and reduces the Hausdorff and average surface distance by 2.4 mm and 1.2 mm, respectively. In particular, worse case performance is also much better for our proposed system, showing it can help P-HNN deal with edge cases. In terms of visual quality, Fig. 5 depicts two examples. As these examples demonstrate, our proposed CGAN allows P-HNN to produce considerable improvements in segmentation mask quality at peripheral nodules, allowing it to overcome an important limitation.

4 Conclusion

We use a 3D CGAN, coupled with a novel multi-mask loss, to effectively generate CT-realistic high-quality lung nodules conditioned on a VOI with an erased central region. Our new multi-mask $L1$ loss ensures a natural blending of the generated nodules with the surrounding lung tissues. Tests demonstrate the superiority of our approach over three competitor CGANs on the LIDC dataset, including Isola *et al.*’s state-of-the-art method [6]. We further use our proposed CGAN to generate a fine-tuning dataset for the published P-HNN model [5], which can struggle when encountering lung nodules adjoining the lung boundary. Armed with our CGAN images, P-HNN is much better able to capture the true lung boundaries compared to both its original state and when it is fine-tuned using the other CGAN variants. As such, our CGAN approach can provide an effective and generic means to help overcome the dataset bottleneck commonly encountered within medical imaging.

References

1. Armato, S.G., McLennan, G., Bidaut, L.: The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on ct scans. *Med. Phys.* **38**(2), 915–931 (2011)
2. Chuquicuma, M.J., Hussein, S., Burt, J., Bagci, U.: How to fool radiologists with generative adversarial networks? a visual turing test for lung cancer diagnosis. In: *Proceedings of IEEE ISBI*, pp. 240–244 (2017)
3. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *MICCAI 2016*. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_49

4. Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems*, pp. 2672–2680 (2014)
5. Harrison, A.P., Xu, Z., George, K., Lu, L., Summers, R.M., Mollura, D.J.: Progressive and multi-path holistically nested neural networks for pathological lung segmentation from CT images. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10435, pp. 621–629. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_71
6. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of IEEE CVPR*, pp. 1125–1134 (2017)
7. Jin, D., Xu, Z., Harrison, A.P., George, K., Mollura, D.J.: 3D convolutional neural networks with graph refinement for airway segmentation using incomplete data labels. In: Wang, Q., Shi, Y., Suk, H.-I., Suzuki, K. (eds.) *MLMI 2017*. LNCS, vol. 10541, pp. 141–149. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67389-9_17
8. Karwoski, R.A., Bartholmai, B., Zavaletta, V.A., et al.: Processing of ct images for analysis of diffuse lung disease in the lung tissue research consortium. In: *Proceedings of SPIE*, vol. 6916 (2008)
9. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
10. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint [arXiv:1411.1784](https://arxiv.org/abs/1411.1784) (2014)
11. Nie, D., et al.: Medical image synthesis with context-aware generative adversarial networks. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10435, pp. 417–425. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_48
12. Pathak, D., Krahenbuhl, P., Donahue, J., et al.: Context encoders: Feature learning by inpainting. In: *Proceedings of IEEE CVPR*, pp. 2536–2544 (2016)
13. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint [arXiv:1511.06434](https://arxiv.org/abs/1511.06434) (2015)
14. Wang, X., Peng, Y., Lu, L., et al.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: *Proceedings of IEEE CVPR*, pp. 3462–3471 (2017)
15. Wolterink, J.M., Dinkla, A.M., Savenije, M.H.F., Seevinck, P.R., van den Berg, C.A.T., Išgum, I.: Deep MR to CT synthesis using unpaired data. In: Tsafaris, S.A., Gooya, A., Frangi, A.F., Prince, J.L. (eds.) *SASHIMI 2017*. LNCS, vol. 10557, pp. 14–23. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68127-6_2