

Attention U-Net Based Adversarial Architectures for Chest X-ray Lung Segmentation

Gusztáv Gaál¹ and Balázs Maga² and András Lukács³

Abstract. X-ray is by far the most common among medical imaging modalities, being faster, more accessible, and more cost-effective compared to other radiographic methods. Chest X-ray (CXR) is the most commonly requested test due to its contribution to the early detection of lung cancer. The most important biomarker in detecting cancer of the lung are nodules, and in finding those, lung segmentation of chest X-rays is essential. Another goal is interpretability, helping radiologists integrate computer-aided detection methods into their diagnostic pipeline, greatly reducing their workload. For this reason, a robust algorithm to perform this otherwise arduous segmentation task is much desired in the field of medical imaging. In this work, we present a novel deep learning approach that uses state-of-the-art fully convolutional neural networks in conjunction with an adversarial critic model. Our network generalized well to CXR images of unseen datasets with different patient profiles, achieving a final DSC of 97.5% on the *JSRT CXR* dataset.

1 INTRODUCTION

X-ray is the most commonly performed radiographic examination, being significantly easier to access, cheaper and faster to carry out than computed tomography (CT), diagnostic ultrasound and magnetic resonance imaging (MRI), as well as having lower dose of radiation compared to a CT scan. According to the publicly available, official data of the National Health Service ([2]), in the period from February 2017 to February 2018, the count of imaging activity was about 41 million in England, out of which almost 22 million was plain X-ray. Many of these imaging tests might contribute to early diagnosis of cancer, amongst which chest X-ray is the most commonly requested one by general practitioners. In order to identify lung nodules, lung segmentation of chest X-rays is essential, and this step is vital in other diagnostic pipelines as well, such as calculating the cardiothoracic ratio, which is the primary indicator of cardiomegaly. For this reason, a robust algorithm to perform this otherwise arduous segmentation task is much desired in the field of medical imaging.

Semantic segmentation aims to solve the challenging problem of assigning a pre-defined class to each pixel of the image. This task requires a high level of visual understanding, in which state-of-the-art performance is attained by methods utilizing Fully Convolutional Networks (FCN) [7]. In [8], adversarial training is used to enhance segmentation of colored images. This idea was incorporated to [13] in order to segment chest X-rays with a fully convolutional, residual neural network. Recently, Mask R-CNN [4] is utilized to realize

instance segmentation on chest X-rays and obtained state-of-the-art results [12, 5].

2 DEEP LEARNING APPROACH

2.1 Network Architecture

Our goal is to produce accurate organ segmentation masks on chest X-rays, meaning for input images we want pixel-wise dense predictions regarding if the given pixel is either part of the left lung, the right lung, the heart, or none of the above.

For this purpose Fully Convolutional Networks (FCNs) are known to significantly outperform other widely used registration-based methods. Specifically we applied a U-Net architecture, thus enabling us to efficiently compute the segmentation mask in the same resolution as the input images. The fully convolutional architecture also enables the use images of different resolutions, since unlike standard convolutional networks, FCNs don't contain input-size dependent layers. In [9] it has been shown that for medical image analysis tasks the integration of the proposed Attention Gates (AGs) improved the accuracy of the segmentation models, while preserving computational efficiency. The architecture of the proposed Attention U-Net is described by Figure 1. Without the use of AGs, it's common practice to use cascade CNNs, selecting a Region Of Interest (ROI) with another CNN where the target organ is likely contained. With the use of AGs we eliminate the need for such a preselecting network, instead the Attention U-Net learns to focus on most important local features, and dulls down the less relevant ones. We note that the dulling of less relevant local features also result in decreased false positive rates.

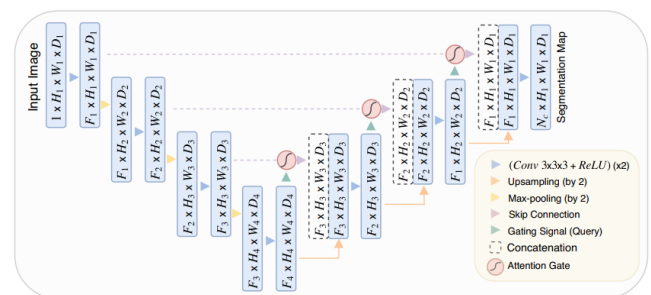


Figure 1. Schematic architecture of the Attention U-Net [9]

In order to enhance the performance of Attention U-Net, we further experimented with adversarial techniques, motivated by [13]. In

¹ Eötvös Loránd University, Hungary, email: guzzzti@gmail.com

² Eötvös Loránd University, Hungary, email: mbalazs0701@gmail.com

³ Eötvös Loránd University, Hungary, email: lukacs@cs.elte.hu

that work, the authors first designed a Fully Convolutional Network (FCN) for the lung segmentation task, and noted that in certain cases the network tends to segment abnormal and incorrect organ shapes. For example, the apex of the ribcage might be mistaken as an internal rib bone, resulting in the mask “bleeding out” to the background, which has similar intensity as the lung field. To address this issue, they developed an adversarial scheme, leading to a model which they call Structure Correcting Adversarial Network (SCAN). This archi-

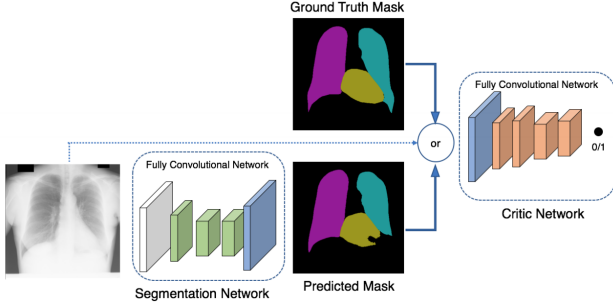


Figure 2. Schematic architecture of the Structure Correcting Adversarial Networks [13]

ture is based on the idea of the General Adversarial Networks [3]. They use the pretrained Fully Convolutional Network as a generator of a General Adversarial Network, and they also train a critic network which is fed the ground truth mask, the predicted mask and optionally the original image. The critic network has roughly the same architecture, resulting in similar capacity. This approach forces the generator to segment more realistic masks, eventually removing obviously wrong shapes.

In our work, besides the standard Attention U-Net, we also created a network of analogous structure, in which the FCN used in [13] is replaced by the Attention U-Net. We did not introduce any modification in the critic model design, such experiments are left to future work.

2.2 Tversky Loss

In the field of medical imaging, Dice Score Coefficient (DSC) is probably the most widespread and simple way to measure the overlap ratio of the masks and the ground truth, and hence to compare and evaluate segmentations. Given two sets of pixels X, Y , their DSC is

$$DSC(X, Y) = \frac{2|X \cap Y|}{|X| + |Y|}.$$

If Y is in fact the result of a test about which pixels are in X , we can rewrite it with the usual notation true/false positive (TP/FP), false negative (FN) to be

$$DSC(X, Y) = \frac{2TP}{2TP + FN + FP}.$$

We would like to use this concept in our setup. The class c we would like to segment corresponds to a set, but it is more appropriate to consider its indicator function g , that is $g_{i,c} \in \{0, 1\}$ equals 1 if and only if the i th pixel belongs to the object. On the other hand, our prediction is a probability for each pixel denoted by $p_{i,c} \in [0, 1]$. Then the Dice Score of the prediction in the spirit of the above description

is defined to be

$$DSC = \frac{\sum_{i=1}^N p_{i,c} g_{i,c} + \varepsilon}{\sum_{i=1}^N (p_{i,c} + g_{i,c}) + \varepsilon},$$

where N is the total number of pixels, and ε is introduced for the sake of numerical stability and to avoid division by 0. The linear Dice Loss (DL) of the multiclass prediction is then

$$DL = \sum_c (1 - DSC_c).$$

A deficiency of Dice Loss is that it penalizes false negative and false positive predictions equally, which results in high precision but low recall. For example practice shows that if the region of interests (ROI) are small, false negative pixels need to have a higher weight than false positive ones. Mathematically this obstacle is easily overcome by introducing weights α, β as tuneable parameters, resulting in the definition of Tversky similarity index [11]:

$$TI_c = \frac{\sum_{i=1}^N p_{i,c} g_{i,c} + \varepsilon}{\sum_{i=1}^N p_{i,c} g_{i,c} + \alpha \sum_{i=1}^N p_{i,\bar{c}} g_{i,c} + \beta \sum_{i=1}^N p_{i,c} g_{i,\bar{c}} + \varepsilon},$$

where $p_{i,\bar{c}} = 1 - p_{i,c}$ and $g_{i,\bar{c}} = 1 - g_{i,c}$, that is the overline simply stands for describing the complement of the class.

Tversky Loss is obtained from Tversky index as Dice Loss was obtained from Dice Score Coefficient:

$$TL = \sum_c (1 - TI_c).$$

Another issue with the Dice Loss is that it struggles to segment small ROIs as they do not contribute to the loss significantly. This difficulty was addressed in [1], where the authors introduced the quantity Focal Tversky Loss in order to improve the performance of their lesion segmentation model:

$$FTL = \sum_c (1 - TI_c)^\gamma,$$

where $\gamma \in [1, 3]$. In practice, if a pixel with is misclassified with a high Tversky index, the Focal Tversky Loss is unaffected. However, if the Tversky index is small and the pixel is misclassified, the Focal Tversky Loss will decrease significantly.

2.3 Training

The explanation of the training of our structure correcting network is a bit longer to explain, we directly follow the footsteps of [13]. Let S, D be the segmentation network and the critic network, respectively. The data consist of the input images \mathbf{x}_i and the associated mask labels \mathbf{y}_i , where \mathbf{x}_i is of shape $[H, W, 1]$ for a single-channel gray-scale image with height H and width W , and \mathbf{y}_i is of shape $[H, W, C]$ where C is the number of classes including the background. Note that for each pixel location (j, k) , $y_i^{jkc} = 1$ for the labeled class channel c while the rest of the channels are zero ($y_i^{jkc'} = 0$ for $c' \neq c$). We use $S(\mathbf{x}) \in [0, 1]^{[H, W, C]}$ to denote the class probabilities predicted by S at each pixel location such that the

class probabilities normalize to 1 at each pixel. Let $D(\mathbf{x}_i, \mathbf{y})$ be the scalar probability estimate of \mathbf{y} coming from the training data. They defined the optimization problem as

$$\min_S \max_D \left\{ J(S, D) := \sum_{i=1}^N J_s(S(\mathbf{x}_i), \mathbf{y}_i) - \lambda \left[J_d(D(\mathbf{x}_i, \mathbf{y}_i), 1) + J_d(D(\mathbf{x}_i, S(\mathbf{x}_i)), 0) \right] \right\}, \quad (1)$$

where

$$J_s(\hat{\mathbf{y}}, \mathbf{y}) := \frac{1}{HW} \sum_{j,k} \sum_{c=1}^C -y^{jkc} \ln y^{jkc}$$

is the multiclass cross-entropy loss for predicted mask $\hat{\mathbf{y}}$ averaged over all pixels.

$$J_d(\hat{t}, t) := -t \ln \hat{t} + (1 - t) \ln(1 - \hat{t})$$

is the binary logistic loss for the critic’s prediction. λ is a tuning parameter balancing pixel-wise loss and the adversarial loss. We can solve equation (1) by alternate between optimizing S and optimizing D using their respective loss functions. This is a point where we introduced a modification: instead of using the multiclass cross-entropy loss $J_s(\hat{\mathbf{y}}, \mathbf{y})$ in the first term, we applied the Focal Tversky Loss $FTL(\hat{\mathbf{y}}, \mathbf{y})$.

Now since the first term in equation (1) does not depend on D , we can train our critic network by *minimizing* the following objective with respect to D for a fixed S :

$$\sum_{i=1}^N J_d(D(\mathbf{x}_i, \mathbf{y}_i), 1) + J_d(D(\mathbf{x}_i, S(\mathbf{x}_i)), 0).$$

Moreover, given a fixed D , we train the segmentation network by minimizing the following objective with respect to S :

$$\sum_{i=1}^N FTL(S(\mathbf{x}_i), \mathbf{y}_i) + \lambda J_d(D(\mathbf{x}_i, S(\mathbf{x}_i)), 1).$$

Following the recommendation in [3], we use $J_d(D(\mathbf{x}_i, S(\mathbf{x}_i)), 1)$ in place of $-J_d(D(\mathbf{x}_i, S(\mathbf{x}_i)), 0)$, as it leads to stronger gradient signals. After tests on the value of λ we decided to use $\lambda = 0.1$.

Concerning training schedule, we found that following pretraining the generator for 50 epochs, we can train the adversarial network for 50 epochs, in which we perform 1 optimization step on the critic network after each 5 optimization step on the generator. This choice of balance is also borrowed from [13], however, we note that the training of our network is much faster.

3 DATASETS

For training- and validation data, we used the Japanese Society of Radiological Technology (JSRT) dataset [10], as well as the Montgomery- and Shenzhen dataset [6], all of which are public datasets of chest X-rays with available organ segmentation masks reviewed by expert radiologists.

The JSRT dataset contains a total of 247 images, of which 154 contains lung nodules. The X-rays are all in 2048×2048 resolution, and have 12-bit grayscale levels. Both lung and heart segmentation masks are available for this dataset.

The Montgomery dataset contains 138 chest X-rays, of which 80 X-rays are from healthy patients, and 58 are from patients with tuberculosis. The X-rays have either a resolution of 4020×4892 or 4892×4020 , and have 12-bit grayscale levels as well. In the case of

this dataset, only lung segmentation masks are publicly available. The Shenzhen dataset contains a total of 662 chest X-rays, of which 326 are of healthy patients, and in a similar fashion, 336 are of patients with tuberculosis. The images vary in sizes, but all are of high resolution, with 8-bit grayscale levels. Only lung segmentation masks are publicly available for the dataset.

3.1 Preprocessing Data

X-rays are grayscale images with typically low contrast, which makes their analysis a difficult task. This obstacle might be overcome by using some sort of histogram equalization technique. The idea of standard histogram equalization is spreading out the the most frequent intensity values to a higher range of the intensity domain $[0, 255]$ by modifying the intensities so that their cumulative distribution function (CDF) on the complete modified image is as close to the CDF of the uniform distribution as possible. Improvements might be made by using adaptive histogram equalization, in which the above method is not utilized globally, but separately on pieces of the image, in order to enhance local contrasts. However, this technique might overamplify noise in near-constant regions, hence our choice was to use Contrast Limited Adaptive Histogram Equalization (CLAHE), which counteracts this effect by clipping the histogram at a predefined value before calculating the CDF, and redistribute this part of the image equally among all the histogram bins.

Applying CLAHE to an X-ray image has visually appealing results, as displayed in Figure 3. As our experiments displayed, it does not merely help human vision, but also neural networks.




Figure 3. Example of chest X-ray images before and after CLAHE

The images were then resized to 512×512 resolution and mapped to $[-1, 1]$ before being fed to our network.

4 EXPERIMENTS AND RESULTS

The aforementioned Attention U-Net architecture was implemented using Keras-TensorFlow Python libraries, to which we have fed our dataset and trained for 40 epochs with 8 X-ray scans in each batch. Our optimizer of choice was *stochastic gradient descent*, having found that *Adam* failed to converge in many cases. As loss function, we applied *Focal Tversky Loss*.

We have found that applying various data augmentation techniques such as flipping, rotating, shearing the image as well as increasing or decreasing the brightness of the image were of no help and just resulted in slower convergence.

Using the Attention U-Net infrastructure, we managed to reach a dice score of 0.9628 for the lungs. Unlike in [13], where no major preprocessing was done, with our preprocessing method, the

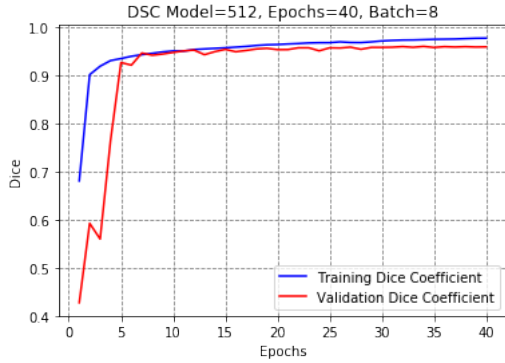


Figure 4. Epoch-wise dice score coefficient

Table 1. Dice scores of different architectures over different datasets.

Dataset	SCAN	ATTN U-Net	Ours (Adv. ATTN)
JSRT	97.3 \pm 0.8%	96.3 \pm 0.7%	97.6 \pm 0.5%
All	-	95.8 \pm 0.4%	96.2 \pm 0.4%
All / JSRT	-	96.6 \pm 0.6	97.8 \pm 0.6%

network performed very well even if the test- and the validation sets were of different datasets. This is extremely important for real world applications, as X-ray images of different machines are significantly different, largely dependent on the specific calibration of each machine, thus it is no trivial task to have X-rays accurately evaluated that are from machines from which no images were in the training set.

We note that even though introducing the adversarial scheme in our setting increased the dice scores, the improvement was not as drastic as in the case of the FCN and SCAN. By checking the masks generated by the vanilla Attention U-Net, we found that this phenomenon can be attributed to the fact that while the FCN occasionally produces abnormally shaped masks, due to our preprocessing steps the Attention U-Net does not commit this mistake. Consequently, the adversarial scheme is responsible for subtle shape improvements only, which is indicated by the Dice Score less spectacularly.

5 FUTURE WORK

So far we have not experimented with the architecture of the critic network, we found the performance of the architecture in [13] completely satisfying. However, it would be desirable to carry out further tests in this direction in order to achieve better understanding of the role of adversarial scheme.

REFERENCES

[1] Nabila Abraham and Naimul Mefraz Khan, ‘A novel focal Tversky loss function with improved attention U-Net for lesion segmentation’, in *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pp. 683–687. IEEE, (2019).

[2] NHS England and NHS Improvement, ‘Diagnostic imaging dataset statistical release’, (2019).

[3] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, ‘Generative adversarial nets’, in *Advances in neural information processing systems*, pp. 2672–2680, (2014).

[4] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick, ‘Mask R-CNN’, in *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969, (2017).

[5] Qinhuo Hu, Luís Fabrício de F Souza, Gabriel Bandeira Holanda, Shara SA Alves, Francisco Hércules dos S Silva, Tao Han, and Pedro P Rebouças Filho, ‘An effective approach for CT lung segmentation using mask region-based convolutional neural networks’, *Artificial Intelligence in Medicine*, 101792, (2020).

[6] Stefan Jaeger, Sema Candemir, Sameer Antani, Yi-Xiang J Wang, Pu-Xuan Lu, and George Thoma, ‘Two public chest X-ray datasets for computer-aided screening of pulmonary diseases’, *Quantitative imaging in medicine and surgery*, 4(6), 475, (2014).

[7] Jonathan Long, Evan Shelhamer, and Trevor Darrell, ‘Fully convolutional networks for semantic segmentation’, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, (2015).

[8] Pauline Luc, Camille Couprie, Soumith Chintala, and Jakob Verbeek, ‘Semantic segmentation using adversarial networks’, *arXiv preprint arXiv:1611.08408*, (2016).

[9] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al., ‘Attention U-Net: Learning where to look for the pancreas’, *arXiv preprint arXiv:1804.03999*, (2018).

[10] Junji Shiraishi, Shigehiko Katsuragawa, Junpei Ikezoe, Tsuneo Matsumoto, Takeshi Kobayashi, Ken-ichi Komatsu, Mitate Matsui, Hiroshi Fujita, Yoshie Kodera, and Kunio Doi, ‘Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists’ detection of pulmonary nodules’, *American Journal of Roentgenology*, 174(1), 71–74, (2000).

[11] Amos Tversky, ‘Features of similarity.’, *Psychological review*, 84(4), 327, (1977).

[12] Jie Wang, Zhigang Li, Rui Jiang, and Zhen Xie, ‘Instance segmentation of anatomical structures in chest radiographs’, in *2019 IEEE 32nd International Symposium on Computer-Based Medical Systems (CBMS)*, pp. 441–446. IEEE, (2019).

[13] Nanqing Dong, Wei Dai, B. Zeya Wang, Xiaodan Liang, Hao Zhang, and Eric P Xing, ‘SCAN: Structure correcting adversarial network for organ segmentation in chest X-rays’, in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings*, volume 11045, p. 263. Springer, (2018).