Published in final edited form as:

Med Image Comput Comput Assist Interv. 2017 September ; 10435: 568–576. doi: 10.1007/978-3-319-66179-7_65.

Discriminative Localization in CNNs for Weakly-Supervised Segmentation of Pulmonary Nodules

Xinyang Feng¹, Jie Yang¹, Andrew F. Laine¹, and Elsa D. Angelini^{1,2}

¹Department of Biomedical Engineering, Columbia University, New York, NY, USA

²ITMAT Data Science Group, NIHR Imperial BRC, Imperial College, London, UK

Abstract

Automated detection and segmentation of pulmonary nodules on lung computed tomography (CT) scans can facilitate early lung cancer diagnosis. Existing supervised approaches for automated nodule segmentation on CT scans require voxel-based annotations for training, which are labor- and time-consuming to obtain. In this work, we propose a weakly-supervised method that generates accurate voxel-level nodule segmentation trained with image-level labels only. By adapting a convolutional neural network (CNN) trained for image classification, our proposed method learns discriminative regions from the activation maps of convolution units at different scales, and identifies the true nodule location with a novel candidate-screening framework. Experimental results on the public LIDC-IDRI dataset demonstrate that, our weakly-supervised nodule segmentation framework achieves competitive performance compared to a fully-supervised CNN-based segmentation method.

1 Introduction

Lung cancer is a major cause of cancer-related deaths worldwide. Pulmonary nodules refer to a range of lung abnormalities that are visible on lung computed tomography (CT) scans as roughly round opacities, and have been regarded as crucial indicators of primary lung cancers [1]. The detection and segmentation of pulmonary nodules in lung CT scans can facilitate early lung cancer diagnosis, timely surgical intervention and thus increase survival rate [2].

Automated detection systems that locate and segment nodules of various sizes can assist radiologists in cancer malignancy diagnosis. Existing supervised approaches for automated nodule segmentation require voxel-level annotations for training, which are labor-intensive and time-consuming to obtain. Alternatively, image-level labels, such as a binary label indicating the presence of nodules, can be obtained more efficiently. Recent work [3,4] studied nodule segmentation using weakly labeled data without dense voxel-level annotations. Their methods, however, still rely on user inputs for additional information such as exact nodule location and estimated nodule size during the segmentation.

Convolutional neural networks (CNNs) have been widely used for supervised image classification and segmentation tasks. It was very recently discovered in a study [5] on natural images that CNNs trained on semantic labels for image classification task (“what”), have remarkable capability in identifying the discriminative regions (“where”) when combined with a global average pooling (GAP) operation. This method utilizes the up-sampled weighted activation maps from the last convolutional layer in a CNN. It demonstrated the localization capability of CNNs for detecting relatively large-sized targets within image, which is not the general scenario in medical imaging domain where pathological changes are more various in size and rather subtle to capture. However, this work sheds light on weakly-supervised disease detection.

In this work, we exploit CNN for accurate and fully-automated segmentation of nodules in a weakly-supervised manner with binary slice-level labels only. Specifically, we adapt a classic image classification CNN model to detect slices with nodule, and simultaneously learn the discriminative regions from the activation maps of convolution units at different scales for coarse segmentation. We then introduce a candidate-screening framework utilizing the same network to generate accurate localization and segmentation. Experimental results on the public LIDC-IDRI dataset [6,7] demonstrate that, despite the largely reduced amount of annotations required for training, our weakly-supervised nodule segmentation framework achieves competitive performance compared to a CNN-based fully-supervised segmentation method.

2 Method

The framework is overviewed in Fig. 1. There are two stages: training stage and segmentation stage. In the first stage, we train a CNN model to classify CT slices as with or without nodule. The CNN is composed of a fully convolutional component, a convolutional layer + global average pooling layer (Conv + GAP) structure, and a final fully-connected (FC) layer. Besides providing a binary classification, the CNN generates a nodule activation map (NAM) showing potential nodule localizations, using a weighted average of the activation maps with the weights learnt in the FC layer. In the second stage, coarse segmentation of nodule candidates is generated within a spatial scope defined by the NAM. For fine segmentation, each nodule candidate is masked out from the image alternately. By feeding the masked image into the same network, a residual NAM (called R-NAM) is generated and used to select the true nodule. Shallower layers in the CNN can be concatenated into the classification task through skip architecture and Conv + GAP structure, extending the one-GAP CNN model to multi-GAP CNN that is able to generate NAMs with higher resolution.

2.1 Nodule Activation Map

In a classification-oriented CNN, while the shallower layers represent general appearance information, the deep layers encode discriminative information that is specific to the classification task. Benefiting from the convolutional structure, spatial information can be retained in the activations of convolutional units. Activation maps of deep convolutional layers, therefore, enable discriminative spatial localization of the class of interest. In our

case, we locate nodules with a specially generated weighted activation map called nodule activation map.

One-GAP CNN—For a given image I , we represent the activation of unit k at spatial location (x, y) in the last convolutional layer as $a_k(x, y)$. The activation of each unit k is summarized through a spatially global average pooling operation as $A_k = \sum_{(x,y)} a_k(x, y)$. The feature vector constituted of A_k is followed by a FC layer, which generates the nodule classification score (i.e. input to the softmax function for nodule class) as:

$$S_{\text{nodule}} = \sum_k w_{k,\text{nodule}} A_k = \sum_k w_{k,\text{nodule}} \sum_{(x,y)} a_k(x, y) \quad (1)$$

where the weights $w_{k,\text{nodule}}$ learnt in the FC layer essentially measure the importance of unit k in the classification task. As spatial information is retained in the activation maps through $a_k(x, y)$, a weighted average of the activation maps results in a robust nodule activation map:

$$\text{NAM}(x, y) = \sum_k w_{k,\text{nodule}} a_k(x, y) \quad (2)$$

The nodule classification score can be directly linked with the NAM by:

$$S_{\text{nodule}} = \sum_{(x,y)} \sum_k w_{k,\text{nodule}} a_k(x, y) = \sum_{(x,y)} \text{NAM}(x, y) \quad (3)$$

By simply up-sampling the NAM to the size of the input image I , we can identify the discriminative image region that is most relevant to nodule.

Multi-GAP CNN—Although activation maps of the last convolutional layer carry most discriminative information, they are usually greatly down-sampled from the original image resolution due to pooling operations. We hereby introduce a multi-GAP CNN model that takes advantage of shallower layers with higher spatial resolution. Similar to the idea of the skip architecture proposed in fully-convolutional network (FCN) [8], shallower layers can be directed to the final classification task skipping the following layers. We also add a Conv + GAP structure following the shallow layers. The concatenation of feature vectors generated by each GAP layer is fed into the final FC layer. The NAM generated from the multi-GAP CNN model (multi-GAP NAM) is a weighted activation map involving activations at multiple scales.

2.2 Segmentation

Coarse Segmentation—For slices classified as “nodule slice”, nodule candidates are screened within a spatial scope C defined by the most prominent blob in the NAM processed via watershed. They are then coarsely segmented using an iterated conditional mode (ICM) based multi-phase segmentation method [9], with the phase number equal to four as determined by global intensity distribution.

Fine Segmentation—The NAM indicates a potential but not exact nodule location. To identify the true nodule from the coarse segmentation results, i.e. which nodule candidate triggered the activation, we generate residual NAMs (R-NAMs) by masking each nodule candidate R_j alternately and feeding the masked image I/R_j into the same network. Significant change of activations within C indicates the exclusion of a true nodule. Formally, we generate the fine segmentation by selecting the nodule candidate R_k following:

$$R_k = \operatorname{argmax}_{R_j} \sum_{(x,y) \in C} \left[\text{NAM}_I(x,y) - \text{NAM}_{I/R_j}(x,y) \right]^2 \quad (4)$$

where NAM_I is the original NAM, and NAM_{I/R_j} is the R-NAM generated by masking nodule candidate R_j . Our current implementation targets the segmentation of one nodule per NAM. Incidence of slices with two nodules is ~1% within slices with nodules. No slices contain more than two nodules in our dataset.

Multi-GAP Segmentation—For the multi-GAP CNN model, we observed a slight drop in classification accuracy compared with the one-GAP CNN model (see Sect. 3.2), which is expected since features from shallower layers are more general and less discriminative. In light of this, we further propose a multi-GAP segmentation method by training both a one-GAP CNN model and a multi-GAP CNN model to combine the discriminative capability of the one-GAP system and finer localization of the multi-GAP system.

Specifically, segmentation is performed on slices classified as “nodule slice” by the one-GAP CNN model for its higher classification accuracy. To define the screening scope for coarse segmentation, we first use the one-GAP NAM to generate a baseline scope C_1 . If there is a prominent blob C_{multi} within C_1 in the multi-GAP NAM, we define the final scope C as C_{multi} to eliminate redundant nodule candidates with more localized spatial constraints. When the multi-GAP NAM fails to identify any discriminative regions within C_1 , the final screening scope C remains C_1 . The R-NAM of the masked image is generated by the one-GAP CNN model and compared with one-GAP NAM within C_1 . Figure 2 illustrates 1-/2-/3-GAP NAMs, the corresponding screening scopes C and coarse segmentation results on a sample slice. While multi-GAP NAM enables finer localization, one-GAP NAM has better discriminative power.

3 Experimental Results

3.1 Data and Experimental Setup

Data used in this study contains 1,010 thoracic CT scans from the public LIDC-IDRI database. Details about this database, such as acquisition protocols and quality evaluations, can be found in [6]. Lungs were segmented and each axial slice was cropped to 384×384 pixels centering on the lung mask. Nodules were delineated by up to four experts. Voxel-level annotations are used to generate slice-level labels, and are used as ground truth for segmentation evaluation. Nodules with diameter < 3 mm are excluded [10]. Given the high false positive rate of nodule detection, we select slices with nodule if there were overlapped annotations by at least two experts, and select slices without nodule if no expert reported a

nodule in the slice. Annotations from different experts were merged using the STAPLE algorithm [11]. A total of $N_{\text{slice}} = 8,345$ slices with nodule are selected, and an equal number of slices without nodule are randomly extracted. The total number of voxels belonging to nodule is $N_{\text{voxel}} = 1,658,981$. Segmentation evaluation is focused on slices with one nodule. Rare cases of slices with two nodules are discussed in the end of Sect. 3.2. Training, validation and test sets are generated by distributing the full set of subjects in a ratio of 4:1:1 through stratified sampling so that they have non-overlapping subjects and similar distribution of nodule occurrence.

3.2 Segmentation Performance

We compare our framework with a fully-supervised CNN-based method (see below). True positive rate (TPR) of nodule detection, false positive rate (FPR) of “nodule” detected on slices *without* nodule, false positive rate ($\text{FPR}_{\text{nodule}}$) of “nodule” detected on slices *with* nodule, Dice overlap of nodule segmentation over all slices with nodule (Dice), Dice over truly detected nodules (TP Dice) and absolute difference of segmented areas over truly detected nodules (TP DOA) are reported in Table 1. Furthermore, TP Dice and TP DOA versus nodule size are reported in Fig. 3.

Weakly-Supervised Segmentation Based on NAM—Our network is based on VGG16Net architecture [12], implemented in TensorFlow. The last pooling layer `pool5` and the FC layers `fc6`, `fc7`, `fc8` are removed [5]. The weights of remaining VGG16Net layers are initialized based on the model pre-trained on ImageNet. The Conv + GAP structure is added after `conv5_3` layer for 1-GAP CNN, added after `conv5_3` and `conv4_3` layers for 2-GAP CNN, and added after `conv5_3`, `conv4_3`, and `conv3_3` layers for 3-GAP CNN. The learning rate of the newly added layers is 10 times the learning rate of the remaining VGG16Net layers. We trained using stochastic gradient descent with momentum. The initial learning rate (10^{-2} for 1-GAP, 2×10^{-3} for 2-GAP, 10^{-3} for 3-GAP), learning decay (0.99), batch size (30) were set by grid search based on classification accuracy on the validation set. The best accuracy values are 88.4% for 1-GAP CNN, 86.6% for 2-GAP model, and 84.4% for 3-GAP model on the test set.

Comparison with Fully-Supervised Segmentation—An adapted model based on U-net architecture [13] is used as a fully-supervised CNN-based model for comparison. The cost function is the negative mean Dice coefficient across mini-batch. The algorithm was optimized with Adam method. The initial learning rate (2×10^{-4}), learning decay (0.999), and batch size (20) were determined with grid search based on average Dice on the validation set.

Two-Nodule Detection—For slices with two nodules, our framework can detect nodules by segmenting the top two activation blobs in the NAM. We tested the detection on a total of 108 slices with two nodules. The 2-GAP model achieves the best detection performance, where both nodules are correctly detected in 50 slices, and one of the two nodules is correctly detected in another 42 slices. With adequate training data, our framework can extend to multi-class classification to automatically determine the number of nodules to segment in the slice.

4 Discussions and Conclusions

In this work we have proposed an original design for lung nodule segmentation, extending a classification-trained CNN model with GAP operations, to learn discriminative regions at different resolution scales utilizing only weakly labeled training data (present or not of a lung nodule). Coarse-to-fine segmentation extracts nodule candidates using an ICM deformable model, and determines the true nodule exploiting a novel candidate-screening framework. Compared with voxel-based labels, the number of labeling required for our method is reduced by $N_{\text{voxel}}/N_{\text{slice}} \sim 100$ times. Detection performance of our weakly-supervised framework compares very favorably with a fully-supervised CNN model (higher TPR and much lower FPR). Our average segmentation accuracy on detected nodules is also very high and gets very close to the benchmark method for larger nodules. Fully-supervised CNN achieves, on average, more accurate segmentation when correctly detecting the nodule, which is expected since voxel-level annotation utilized during training provides more power to deal with various intensity patterns, especially at edges. On the other hand, standard deviations are smaller with the proposed method, hence indicates fewer large mistakes.

NAM can act as an efficient screening framework that can be incorporated with patch-level labels for false positive reduction [10], or with a small amount of voxel-level labels to learn fine segmentation contour. Future work will also extend NAM to 3D CNN to take advantage of the 3D contextual information.

A machine learning model requiring only weakly-labeled data is key for a sustainable development of CAD systems, as expert time is scarce and expensive and as scanners continue to evolve significantly. Our work used transfer learning from a CNN trained on natural images; with more annotated data, it will be possible to train a fully dedicated network that is likely to be more effective.

Acknowledgments

Thanks NIH R01-HL121270 for funding.

References

1. MacMahon H, Austin JH, et al. Guidelines for management of small pulmonary nodules detected on CT scans: a statement from the Fleischner society 1. *Radiology*. 2005; 237(2):395–400. [PubMed: 16244247]
2. Henschke CI, McCauley DI, et al. Early lung cancer action project: overall design and findings from baseline screening. *Lancet*. 1999; 354(9173):99–105. [PubMed: 10408484]
3. Messay T, Hardie RC, et al. Segmentation of pulmonary nodules in computed tomography using a regression neural network approach and its application to the Lung Image Database Consortium and Image Database Resource Initiative dataset. *Med Image Anal*. 2015; 22(1):48–62. [PubMed: 25791434]
4. Anirudh R, Thiagarajan JJ, et al. Lung nodule detection using 3D convolutional neural networks trained on weakly labeled data. *SPIE Medical Imaging*. 2016:978532.
5. Zhou, B., Khosla, A., et al. *IEEE CVPR*. 2016. Learning deep features for discriminative localization.
6. Armato SG, McLennan G, et al. The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans. *Med Phys*. 2011; 38(2):915–931. [PubMed: 21452728]

7. Clark K, et al. The Cancer Imaging Archive (TCIA): maintaining and operating a public information repository. *J Digit Imaging*. 2013; 26(6):1045–1057. [PubMed: 23884657]
8. Shelhamer E, Long J, et al. Fully convolutional networks for semantic segmentation. *IEEE Trans Pattern Anal Mach Intell*. 2016 PP(99), 1.
9. Israel-Jost V, Breton E, et al. Vectorial multi-phase mouse brain tumor segmentation in T1–T2 MRI. *IEEE ISBI*. 2008:5–8.
10. Setio AAA, Traverso A, et al. Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the LUNA16 challenge. 2016 arXiv preprint arXiv:1612.08012.
11. Warfield SK, Zou KH, et al. Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE Trans Med Imaging*. 2004; 23(7):903–921. [PubMed: 15250643]
12. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. 2014 arXiv preprint arXiv:1409.1556.
13. Ronneberger, O., Fischer, P., Brox, T. U-net: convolutional networks for biomedical image segmentation. In: Navab, N.Hornegger, J.Wells, WM., Frangi, AF., editors. *MICCAI 2015 LNCS*. Vol. 9351. Springer; Cham: 2015. p. 234–241.

Fig. 1.

(A) Training: a CNN model is trained to classify CT slices and generate nodule activation maps (NAMs); (B) Segmentation: for test slices classified as "nodule slice", nodule candidates are screened using a spatial scope defined by the NAM for coarse segmentation. Residual NAMs (R-NAMs) are generated from images with masked nodule candidates for fine segmentation.

Fig. 2. Illustration of 1-/2-/3-GAP NAMs, the screening scopes C and coarse segmentation results on a sample slice.

Fig. 3.

TP Dice and TP DOA (mean and standard deviation) versus nodule size.

Table 1

Comparison of segmentation performance

Method	TPR	FPR	FPR _{nodule}	Dice		TP Dice		TP DOA (mm ²)	
				mean ± SD		mean ± SD		mean ±SD	
1-GAP Coarse	0.77 [*]	0.11 ^{*,†}	0.08 [*]	0.46 (±0.31)		0.61 (±0.20)		57.6 (±71.1)	
2-GAP Coarse	0.76	–	0.09	0.50 (±0.34)		0.66 (±0.18)		41.6 (±53.6)	
3-GAP Coarse	0.75	–	0.11	0.50 (±0.32)		0.67 (±0.18)		40.1 (±50.9)	
1-GAP Fine	0.75	–	0.11	0.54 (±0.34)		0.73 (±0.15)		30.7 (±52.8)	
2-GAP Fine	0.75	–	0.11	0.55 [*] (±0.33)		0.74 [*] (±0.14)		29.2 [*] (±46.8)	
3-GAP Fine	0.74	–	0.12	0.54 (±0.34)		0.74 (±0.14)		29.3 (±46.4)	
U-net	0.74	0.29	0.26	0.56 (±0.38)		0.76 (±0.19)		28.3 (±44.8)	

^{*} best performance within our framework; **boldfaced** = overall best performance.

[†] 1-GAP model is used for nodule slice-level detection within our framework.