

DeepEM: Deep 3D ConvNets With EM For Weakly Supervised Pulmonary Nodule Detection

Wentao Zhu*, Yeeleng S. Vang*, Yufang Huang[†], and Xiaohui Xie*

*University of California, Irvine [†]Lenovo AI Lab
{wentaoz1,ysvang,xhx}@uci.edu, yufang@lenovo.com

Abstract. Recently deep learning has been witnessing widespread adoption in various medical image applications. However, training complex deep neural nets requires large-scale datasets labeled with ground truth, which are often unavailable in many medical image domains. For instance, to train a deep neural net to detect pulmonary nodules in lung computed tomography (CT) images, current practice is to manually label nodule locations and sizes in many CT images to construct a sufficiently large training dataset, which is costly and difficult to scale. On the other hand, electronic medical records (EMR) contain plenty of partial information on the content of each medical image. In this work, we explore how to tap this vast, but currently unexplored data source to improve pulmonary nodule detection. We propose DeepEM, a novel deep 3D ConvNet framework augmented with expectation-maximization (EM), to mine weakly supervised labels in EMRs for pulmonary nodule detection. Experimental results show that DeepEM can lead to 1.5% and 3.9% average improvement in free-response receiver operating characteristic (FROC) scores on LUNA16 and Tianchi datasets, respectively, demonstrating the utility of incomplete information in EMRs for improving deep learning algorithms.¹

Keywords: Deep 3D convolutional nets, weakly supervised detection, DeepEM (deep 3D ConvNets with EM), pulmonary nodule detection

1 Introduction

Lung cancer is the most common cause of cancer-related death in men. Low-dose lung computed tomography (CT) screening provides an effective way for early diagnosis and can sharply reduce the lung cancer mortality rate. Advanced computer-aided diagnosis (CAD) systems are expected to have high sensitivities while maintaining low false positive rates to be truly useful. Recent advance in deep learning provides new opportunities to design more effective CAD systems to help facilitate doctors in their effort to catch lung cancer in their early stages.

The emergence of large-scale datasets such as the LUNA16 [16] has helped to accelerate research in nodule detection. Typically, nodule detection consists

¹ <https://github.com/uci-cbcl/DeepEM-for-Weakly-Supervised-Detection.git>




Fig. 1. Illustration of DeepEM framework. Faster R-CNN is employed for nodule proposal generation. Half-Gaussian model and logistic regression are employed for central slice and lobe location respectively. In the E-step, we utilize all the observations, CT slices, and weak label to infer the latent variable, nodule proposals, by maximum a posteriori (MAP) or sampling. In the M-step, we employ the estimated proposals to update parameters in the Faster R-CNN and logistic regression.

of two stages: nodule proposal generation and false positive reduction. Traditional approaches generally require hand-designed features such as morphological features, voxel clustering and pixel thresholding [11, 7, 10]. More recently, deep convolutional architectures were employed to generate the candidate bounding boxes. Setio et al. proposed multi-view convolutional network for false positive reduction [15]. Several work employed 3D convolutional networks to handle the challenge due to the 3D nature of CT scans. The 3D fully convolutional network (FCN) was proposed to generate region candidates and deep convolutional network with weighted sampling was used in the false positive reduction stage [19, 3, 9, 17]. CASED proposed curriculum adaptive sampling for 3D U-net training in nodule detection [8, 14]. Ding et al. used Faster R-CNN to generate candidate nodules, followed by 3D convolutional networks to remove false positive nodules [2]. Due to the effective performance of Faster R-CNN [13], Faster R-CNN with a U-net-like encoder-decoder scheme was proposed for nodule detection [19].

A prerequisite to utilization of deep learning models is the existence of an abundance of labeled data. However, labels are especially difficult to obtain in the medical image analysis domain. There are multiple contributing factors: a) labeling medical data typically requires specially trained doctors; b) marking lesion boundaries can be hard even for experts because of low signal-to-noise ratio in many medical images; and c) for CT and magnetic resonance imaging (MRI) images, the annotators need to label the entire 3D volumetric data, which can be costly and time-consuming. Due to these limitations, CT medical image datasets are usually small, which can lead to over-fitting on the training set and, by extension, poor generalization performance on test sets [21].

By contrast, medical institutions have large amount of weakly labeled medical images. In these databases, each medical image is typically associated with an electronic medical report (EMR). Although these reports may not contain explicit information on detection bounding box or segmentation ground truth, it often includes the results of diagnosis, rough locations and summary descrip-

tions of lesions if they exist. We hypothesize that these extra sources of weakly labeled data may be used to enhance the performance of existing detector and improve its generalization capability.

There are previous attempts to utilize weakly supervised labels to help train machine learning models. Deep multi-instance learning was proposed for lesion localization and whole mammogram classification [20]. The two-stream spatio-temporal ConvNet was proposed to recognize heart frames and localize the heart using only weak labels for whole ultrasound image of fetal heartbeat [5]. Different pooling strategies were proposed for weakly supervised localization and segmentation respectively [18, 4, 1]. Papandreou et al. proposed an iterative approach to infer pixel-wise label using image classification label for segmentation [12]. Self-transfer learning co-optimized both classification and localization networks for weakly supervised lesion localization [6]. Different from these works, we consider nodule proposal as latent variable and propose DeepEM, a new deep 3D convolutional nets with Expectation-Maximization optimization, to mine the big data source of weakly supervised label in EMR as illustrated in Fig. 1. Specifically, we infer the posterior probabilities of the proposed nodules being true nodules, and utilize the posterior probabilities to train nodule detection models.

2 DeepEM for Weakly Supervised Detection

Notation We denote by $\mathbf{I} \in \mathbb{R}^{h \times w \times s}$ the CT image, where h , w , and s are image height, width, and number of slices respectively. The nodule bounding boxes for \mathbf{I} are denoted as $\mathbf{H} = \{\mathbf{H}_1, \mathbf{H}_2, \dots, \mathbf{H}_M\}$, where $\mathbf{H}_m = \{x_m, y_m, z_m, d_m\}$, the (x_m, y_m, z_m) represents the center of nodule proposal, d_m is the diameter of the nodule proposal, and M is the number of nodules in the image \mathbf{I} . In the weakly supervised scenario, the nodule proposal \mathbf{H} is a latent variable, and each image \mathbf{I} is associated with weak label $\mathbf{X} = \{\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_M\}$, where $\mathbf{X}_m = \{loc_m, z_m\}$, $loc_m \in \{1, 2, 3, 4, 5, 6\}$ is the location (right upper lobe, right middle lobe, right lower lobe, left upper lobe, lingula, left lower lobe) of nodule \mathbf{H}_m in the lung, and z_m is the central slice of the nodule.

For fully supervised detection, the objective function is to maximize the log-likelihood function for observed nodule ground truth \mathbf{H} given image \mathbf{I} as

$$\mathcal{L}(\boldsymbol{\theta}) = \log P(\mathbf{H} \cup \bar{\mathbf{H}} | \mathbf{I}; \boldsymbol{\theta}) = \frac{1}{M} \sum_{m=1}^M \log P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}) + \frac{1}{N} \sum_{n=1}^N \log P(\bar{\mathbf{H}}_n | \mathbf{I}; \boldsymbol{\theta}), \quad (1)$$

where $\bar{\mathbf{H}} = \{\bar{\mathbf{H}}_1, \bar{\mathbf{H}}_2, \dots, \bar{\mathbf{H}}_N\}$ are hard negative nodule proposals [13], $\boldsymbol{\theta}$ is the weights of deep 3D ConvNet. We employ Faster R-CNN with 3D Res18 for the fully supervised detection because of its superior performance.

For weakly supervised detection, nodule proposal \mathbf{H} can be considered as a latent variable. Using this framework, image \mathbf{I} and weak label $\mathbf{X} = \{(loc_1, z_1), (loc_2,$

$z_2), \dots, (loc_M, z_M)\}$ can be considered as observations. The joint distribution is

$$\begin{aligned} P(\mathbf{I}, \mathbf{H}, \mathbf{X}; \boldsymbol{\theta}) &= P(\mathbf{I}) \prod_{m=1}^M (P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}) P(\mathbf{X}_m | \mathbf{H}_m; \boldsymbol{\theta})) \\ &= P(\mathbf{I}) \prod_{m=1}^M (P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}) P(loc_m | \mathbf{H}_m; \boldsymbol{\theta}) P(z_m | \mathbf{H}_m; \boldsymbol{\theta})). \end{aligned} \quad (2)$$

To model $P(z_m | \mathbf{H}_m; \boldsymbol{\theta})$, we propose using a half-Gaussian distribution based on nodule size distribution because z_m is correct if it is within the nodule area (center slice of \mathbf{H}_m as $z_{\mathbf{H}_m}$, and nodule size σ can be empirically estimated based on existing data) for nodule detection in Fig. 2(a). For lung lobe prediction $P(loc_m | \mathbf{H}_m; \boldsymbol{\theta})$, a logistic regression model is used based on relative value of nodule center $(x_{\mathbf{H}_m}, y_{\mathbf{H}_m}, z_{\mathbf{H}_m})$ after lung segmentation. That is

$$P(z_m, loc_m | \mathbf{H}_m; \boldsymbol{\theta}) = \frac{2}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{|z_m - z_{\mathbf{H}_m}|^2}{2\sigma^2}\right) \frac{\exp(\mathbf{f}(\mathbf{H}_m)\boldsymbol{\theta}_{loc_m})}{\sum_{loc_m=1}^6 \exp(\mathbf{f}(\mathbf{H}_m)\boldsymbol{\theta}_{loc_m})}, \quad (3)$$

where $\boldsymbol{\theta}_{loc_m}$ is the associated weights with lobe location loc_m for logistic regression, feature $\mathbf{f}(\mathbf{H}_m) = (\frac{x_{\mathbf{H}_m}}{x_{\mathbf{I}}}, \frac{y_{\mathbf{H}_m}}{y_{\mathbf{I}}}, \frac{z_{\mathbf{H}_m}}{z_{\mathbf{I}}})$, and $(x_{\mathbf{I}}, y_{\mathbf{I}}, z_{\mathbf{I}})$ is the total size of image \mathbf{I} after lung segmentation. In the experiments, we found the logistic regression converges quickly and is stable.

The expectation-maximization (EM) is a commonly used approach to optimize the maximum log-likelihood function when there are latent variables in the model. We employ the EM algorithm to optimize deep weakly supervised detection model in equation 2. The expected complete-data log-likelihood function given previous estimated parameter $\boldsymbol{\theta}'$ in deep 3D Faster R-CNN is

$$\begin{aligned} Q(\boldsymbol{\theta}; \boldsymbol{\theta}') &= \frac{1}{M} \sum_{m=1}^M \mathbb{E}_{P(\mathbf{H}_m | \mathbf{I}, z_m, loc_m; \boldsymbol{\theta}')} [\log P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}) \\ &\quad + \log P(z_m, loc_m | \mathbf{H}_m; \boldsymbol{\theta})] + \mathbb{E}_{Q(\bar{\mathbf{H}}_n | \mathbf{z})} [\log P(\bar{\mathbf{H}}_n | \mathbf{I}; \boldsymbol{\theta})], \end{aligned} \quad (4)$$

where $\mathbf{z} = \{z_1, z_2, \dots, z_m\}$. In the implementation, we only keep hard negative proposals far away from weak annotation \mathbf{z} to simplify $Q(\bar{\mathbf{H}}_n | \mathbf{z})$. The posterior distribution of latent variable \mathbf{H}_m can be calculated by

$$P(\mathbf{H}_m | \mathbf{I}, z_m, loc_m; \boldsymbol{\theta}') \propto P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}') P(z_m, loc_m | \mathbf{H}_m; \boldsymbol{\theta}'). \quad (5)$$

Because Faster R-CNN yields a large number of proposals, we first use hard threshold (-3 before sigmoid function) to remove proposals of small confident probability, then employ non-maximum suppression (NMS) with intersection over union (IoU) as 0.1. We then employ two schemes to approximately infer the latent variable \mathbf{H}_m : maximum a posteriori (MAP) or sampling.

DeepEM with MAP We only use the proposal of maximal posterior probability to calculate the expectation.

$$\hat{\mathbf{H}}_m = \arg \max_{\mathbf{H}_m} P(\mathbf{H}_m | \mathbf{I}; \boldsymbol{\theta}') P(z_m, loc_m | \mathbf{H}_m; \boldsymbol{\theta}') \quad (6)$$

Algorithm 1 DeepEM for Weakly Supervised Detection

Input: Fully supervised dataset $D_F = \{(\mathbf{I}, \mathbf{H})_i\}_{i=1}^{N_F}$, weakly supervised dataset $D_W = \{(\mathbf{I}, \mathbf{X})_i\}_{i=1}^{N_W}$, 3D Faster R-CNN and logistic regression parameters $\boldsymbol{\theta}$.

- 1: *Initialization*: Update weights $\boldsymbol{\theta}$ by maximizing equation 1 using data from D_F .
- 2: *for epoch = 1 to #TotalEpochs*:
 - ▷▷▷ **Weakly supervised training**
 - 3: Use Faster R-CNN model $\boldsymbol{\theta}'$ to obtain proposal probability $P(\mathbf{H}_m|\mathbf{I}; \boldsymbol{\theta}')$ for weakly supervised data sampled from D_W .
 - 4: Remove proposals with small probabilities and NMS.
 - 5: *for m = 1 to M*: ▷▷▷ Each weak label
 - 6: Calculate $P(z_m, loc_m|\mathbf{H}_m; \boldsymbol{\theta})$ for each proposal by equation 3.
 - 7: Estimate posterior distribution $P(\mathbf{H}_m|\mathbf{I}, z_m, loc_m; \boldsymbol{\theta}')$ by equation 5 with normalization.
 - 8: Employ MAP by equation 6 or Sampling to obtain the inference of \mathbf{H}_m .
 - 9: Obtain the expect log-likelihood function by equation 4 using the estimated proposal (MAP) or by equation 7 (Sampling).
 - 10: Update parameter by equation 8.
 - ▷▷▷ **Fully supervised training**
 - 11: Update weights $\boldsymbol{\theta}$ by maximizing equation 1 using fully supervised data D_F .

DeepEM with Sampling We approximate the distribution by sampling \hat{M} proposals $\hat{\mathbf{H}}_m$ according to normalized equation 5. The expected log-likelihood function in equation 4 becomes

$$Q(\boldsymbol{\theta}; \boldsymbol{\theta}') = \frac{1}{M\hat{M}} \sum_{m=1}^M \sum_{\hat{\mathbf{H}}_m}^{\hat{M}} (\log P(\hat{\mathbf{H}}_m|\mathbf{I}; \boldsymbol{\theta}) + \log P(z_m, loc_m|\hat{\mathbf{H}}_m; \boldsymbol{\theta})) \quad (7)$$

$$+ \mathbb{E}_{Q(\hat{\mathbf{H}}_n|\mathbf{z})} [\log P(\hat{\mathbf{H}}_n|\mathbf{I}; \boldsymbol{\theta})].$$

After obtaining the expectation of complete-data log-likelihood function in equation 4, we can update the parameters $\boldsymbol{\theta}$ by

$$\hat{\boldsymbol{\theta}} = \arg \max Q(\boldsymbol{\theta}; \boldsymbol{\theta}'). \quad (8)$$

The M-step in equation 8 can be conducted by stochastic gradient descent commonly used in deep network optimization for equation 1. Our entire algorithm is outlined in algorithm 1.

3 Experiments

We used 3 datasets, LUNA16 dataset [16] as fully supervised nodule detection, NCI NLST² dataset as weakly supervised detection, Tianchi Lung Nodule Detection³ dataset as holdout dataset for test only. LUNA16 dataset is the largest publicly available dataset for pulmonary nodules detection [16]. LUNA16 dataset removes CTs with slice thickness greater than 3mm, slice spacing inconsistent or

² <https://biometry.nci.nih.gov/cdas/datasets/nlst/>

³ <https://tianchi.aliyun.com/>




Fig. 2. (a) Empirical estimation of half-Gaussian model for $P(z_m | H_m; \theta)$ on LUNA16. (b) FROC (%) comparison among Faster R-CNN, DeepEM with MAP, DeepEM with Sampling on LUNA16.

missing slices, and consist of 888 low-dose lung CTs which have explicit patient-level 10-fold cross validation split. NLST dataset consists of hundreds of thousands of lung CT images associated with electronic medical records (EMR). In this work, we focus on nodule detection based on image modality and only use the central slice and nodule location as weak supervision from the EMR. As part of data cleansing, we remove negative CTs, CTs with slice thickness greater than 3mm and nodule diameter less than 3mm. After data cleaning, we have 17,602 CTs left with 30,951 weak annotations. In each epoch, we randomly sample $\frac{1}{16}$ CT images for weakly supervised training because of the large numbers of weakly supervised CTs. Tianchi dataset contains 600 training low-dose lung CTs and 200 validation low-dose lung CTs for nodule detection. The annotations are location centroids and diameters of the pulmonary nodules, and do not have less than 3mm diameter nodule, which are the same with those on LUNA16 dataset.

Parameter estimation in $P(z_m | H_m; \theta)$ If the current z_m is within the nodule, it is a true positive proposal. We can model $|z_m - z_{H_m}|$ using a half-Gaussian distribution shown as the red dash line in Fig. 2(a). The parameters of the half-Gaussian is estimated from the LUNA16 data empirically. Because LUNA16 removes nodules of diameter less than 3mm, we use the truncated half-Gaussian to model the central slice z_m as $\max(|z_m - z_{H_m}| - \mu, 0)$, where μ is the mean of related Gaussian as the minimal nodule radius with 1.63.

Performance comparisons on LUNA16 We conduct 10-fold cross validation on LUNA16 to validate the effectiveness of DeepEM. The baseline method is Faster R-CNN with 3D Res18 network denoted as **Faster R-CNN** [13, 19]. Then we employ it to model $P(H_m | I; \theta')$ for weakly supervised detection scenario. Two inference scheme for H_m are used in DeepEM denoted as **DeepEM (MAP)** and **DeepEM (Sampling)**. In the proposal inference of DeepEM with Sampling, we sample two proposals for each weak label because the average number of nodules each CT is 1.78 on LUNA16. The evaluation metric, Free receiver operating characteristic (FROC), is the average recall rate at the average num-

Table 1. FROC (%) comparisons among Faster R-CNN with 3D ResNet18, DeepEM with MAP, DeepEM with Sampling on Tianchi.

Fold	0	1	2	3	4	5	6	7	8	9	Average
Faster R-CNN	72.8	70.8	69.8	71.9	76.4	73.0	71.3	74.7	72.9	71.3	72.5
DeepEM (MAP)	77.2	75.8	75.8	74.9	77.0	75.5	77.2	75.8	76.0	74.7	76.0
DeepEM (Sampling)	77.4	75.8	75.9	75.0	77.3	75.0	77.3	76.8	77.7	75.8	76.4




Fig. 3. Detection visual comparison among Faster R-CNN, DeepEM with MAP and DeepEM with Sampling on nodules randomly sampled from Tianchi. DeepEM provides more accurate detection (central slice, center and diameter) than Faster R-CNN.

ber of false positives at 0.125, 0.25, 0.5, 1, 2, 4, 8 per scan, which is the official evaluation metric for LUNA16 and Tianchi [16].

From Fig. 2(b), DeepEM with MAP improves about 1.3% FROC over Faster R-CNN and DeepEM with Sampling improves about 1.5% FROC over Faster R-CNN on average on LUNA16 when incorporating weakly labeled data from NLST. We hypothesize the greater improvement of DeepEM with Sampling over DeepEM with MAP is that MAP inference is greedy and can get stuck at a local minimum while the nature of sampling may allow DeepEM with Sampling to escape these local minimums during optimization.

Performance comparisons on holdout test set from Tianchi We employed a holdout test set from Tianchi to validate each model from 10-fold cross validation on LUNA16. The results are summarized in Table 1. We can see DeepEM utilizing weakly supervised data improves 3.9% FROC on average over Faster R-CNN. The improvement on holdout test data validates DeepEM as an effective model to exploit potentially large amount of weak data from electronic medical records (EMR) which would not require further costly annotation by expert doctors and can be easily obtained from hospital associations.

Visualizations We compare Faster R-CNN with the proposed DeepEM visually in Fig. 2(b). We randomly choose nodules from Tianchi. From Fig. 2(b), DeepEM yields better detection for nodule center and tighter nodule dia-

ter which demonstrates DeepEM improves the existing detector by exploiting weakly supervised data.

4 Conclusion

In this paper, we have focused on the problem of detecting pulmonary nodules from lung CT images, which previously has been formulated as a supervised learning problem and requires a large amount of training data with the locations and sizes of nodules precisely labeled. Here we propose a new framework, called DeepEM, for pulmonary nodule detection by taking advantage of abundantly available weakly labeled data extracted from EMRs. We treat each nodule proposal as a latent variable, and infer the posterior probabilities of proposal nodules being true ones conditioned on images and weak labels. The posterior probabilities are further fed to the nodule detection module for training. We use an EM algorithm to train the entire model end-to-end. Two schemes, maximum a posteriori (MAP) and sampling, are used for the inference of proposals. Extensive experimental results demonstrate the effectiveness of DeepEM for improving current state of the art nodule detection systems by utilizing readily available weakly supervised detection data. Although our method is built upon the specific application of pulmonary nodule detection, the framework itself is fairly general and can be readily applied to other medical image deep learning applications to take advantage of weakly labeled data.

Acknowledgement

We gratefully acknowledge the sharing of pulmonary CT imaging data from National Lung Screening Trial (NLST). We also thank NVIDIA for supporting this research.

References

1. Bilen, H., et al.: Weakly supervised deep detection networks. In: CVPR (2016)
2. Ding, J., et al.: Accurate pulmonary nodule detection in computed tomography images using deep convolutional neural networks. In: MICCAI (2017)
3. Dou, Q., et al.: Automated pulmonary nodule detection via 3d convnets with online sample filtering and hybrid-loss residual learning. In: MICCAI (2017)
4. Feng, X., et al.: Discriminative localization in cnns for weakly-supervised segmentation of pulmonary nodules. In: MICCAI (2017)
5. Gao, Y., Noble, J.A.: Detection and characterization of the fetal heartbeat in free-hand ultrasound sweeps with weakly-supervised two-streams convolutional networks. In: MICCAI. pp. 305–313. Springer (2017)
6. Hwang, S., Kim, H.E.: Self-transfer learning for weakly supervised lesion localization. In: MICCAI. pp. 239–246. Springer (2016)
7. Jacobs, C., et al.: Automatic detection of subsolid pulmonary nodules in thoracic computed tomography images. Medical image analysis (2014)

8. Jesson, A., et al.: Cased: Curriculum adaptive sampling for extreme data imbalance. In: MICCAI (2017)
9. Liao, F., et al.: Evaluate the malignancy of pulmonary nodules using the 3d deep leaky noisy-or network. arXiv preprint (2017)
10. Lopez Torres, E., et al.: Large scale validation of the m5l lung cad on heterogeneous ct datasets. Medical physics (2015)
11. Murphy, K., et al.: A large-scale evaluation of automatic pulmonary nodule detection in chest ct using local image features and k-nearest-neighbour classification. Medical image analysis (2009)
12. Papandreou, G., et al.: Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In: ICCV (2015)
13. Ren, S., et al.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NIPS (2015)
14. Ronneberger, O., et al.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI (2015)
15. Setio, A.A.A., et al.: Pulmonary nodule detection in ct images: false positive reduction using multi-view convolutional networks. IEEE TMI (2016)
16. Setio, A.A.A., et al.: Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the luna16 challenge. Medical image analysis 42, 1–13 (2017)
17. Tang, H., Kim, D., Xie, X.: Automated pulmonary nodule detection using 3d deep convolutional neural networks. ISBI (2018)
18. Wang, X., et al.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: CVPR (2017)
19. Zhu, W., Liu, C., Fan, W., Xie, X.: Deeplung: Deep 3d dual path nets for automated pulmonary nodule detection and classification. In: IEEE WACV (2018)
20. Zhu, W., Lou, Q., Vang, Y.S., Xie, X.: Deep multi-instance networks with sparse label assignment for whole mammogram classification. In: MICCAI (2017)
21. Zhu, W., et al.: Adversarial deep structured nets for mass segmentation from mammograms. In: IEEE ISBI (2018)