

# Iterative Attention Mining for Weakly Supervised Thoracic Disease Pattern Localization in Chest X-Rays

Jinzheng Cai<sup>1</sup>, Le Lu<sup>2</sup>, Adam P. Harrison<sup>2</sup>, Xiaoshuang Shi<sup>1</sup>,  
Pingjun Chen<sup>1</sup>, and Lin Yang<sup>1</sup>

<sup>1</sup> University of Florida, Gainesville, FL, 32611, USA

<sup>2</sup> AI-Infra, NVIDIA Corp, Bethesda, MD, 20814, USA

**Abstract.** Given image labels as the only supervisory signal, we focus on harvesting/mining, thoracic disease localizations from chest X-ray images. Harvesting such localizations from existing datasets allows for the creation of improved data sources for computer-aided diagnosis and retrospective analyses. We train a convolutional neural network (CNN) for image classification and propose an attention mining (AM) strategy to improve the model’s sensitivity or saliency to disease patterns. The intuition of AM is that once the most salient disease area is blocked or hidden from the CNN model, it will pay attention to alternative image regions, while still attempting to make correct predictions. However, the model requires to be properly constrained during AM, otherwise, it may overfit to uncorrelated image parts and forget the valuable knowledge that it has learned from the original image classification task. To alleviate such side effects, we then design a knowledge preservation (KP) loss, which minimizes the discrepancy between responses for X-ray images from the original and the updated networks. Furthermore, we modify the CNN model to include multi-scale aggregation (MSA), improving its localization ability on small-scale disease findings, *e.g.*, lung nodules. We validate our method on the publicly-available ChestX-ray14 dataset, outperforming a class activation map (CAM)-based approach, and demonstrating the value of our novel framework for mining disease locations.

## 1 Introduction

Automatic analysis of chest X-rays is critical for diagnosis and treatment planning of thoracic diseases. Recently, several methods applying deep learning for automatic chest X-ray analysis [8,5,11,14,7] have been proposed. In particular, much work has focused on the ChestX-ray14 dataset [11], which is an unprecedentedly large-scale and rich dataset but only provides image-level labels for the far majority of the samples. On the other hand, harvesting abnormality locations in this dataset is an important goal, as that provides an even richer source of data for training computer-aided diagnosis system and/or performing retrospective data analyses. Harvesting disease locations can be conducted through a weakly supervised image classification approach [11]; or, in our case we reformulate it as a label supervised *pattern-mining problem*, to gain higher localization accuracy. Toward this end, we propose an integrated and novel framework that combines attention mining, knowledge preservation, and multi-scale aggregation that improves upon current efforts to accurately localize disease patterns.

Recent work on chest X-rays have focused on both classification and localization. Along with the ChestX-ray14 dataset, Wang *et al.* [11] also propose

a class activation map (CAM)-based [16] approach using convolutional neural network (CNNs) to perform weakly supervised disease localization. To improve image classification accuracy, Rajpurkar *et al.* [8] introduce an ultra-deep CNN architecture while Yao *et al.* [14] design a new learning objective that exploits dependencies among image labels. Other work investigate methods to automatically generate X-ray reports [4,10]. The proposed framework is a complementary or orthogonal development from the above advances [8,14] since we mine “free” disease locations in the form of bounding boxes given image-level labels. It also can further benefit downstream applications like [4] and [10].

In terms of related work, our attention mining (AM) approach is closely related to an adversarial erasing scheme proposed in [12] that forces the network to discover other salient image regions by erasing the most representative area of the object class in question. In a similar spirit, we propose AM to locate multiple suspicious disease regions inside a chest X-ray. However, different from [12], AM drops out corresponding pixels in the activation maps so as to leave the original X-ray images unchanged. More importantly, AM is designed to seamlessly couple with multi-label classification, where activation maps are required to be blocked in a class-wise manner. Next, to alleviate the side effects caused by dropping out activation maps, we exploit methods to prevent the network from forgetting its originally learned knowledge on recognizing and localizing disease patterns. Distilling a network’s knowledge is proposed in [2] to transfer the learned parameters from multiple models to a new, typically smaller sized, model. A similar technique is used in [9] to regularize the CNN model for incremental learning with new image categories, keeping the network’s output of old image categories mostly unchanged. In our method, we minimize the  $\ell_2$ -distance of output logits between the original and updated networks to achieve knowledge preservation (KP). Distinct from [2] and [9], we use the logits not only from the last output layer but also the intermediate network layers, in order to introduce stronger regularizations. Finally, we propose a multi-scale aggregation (MSA) because we notice that the localization accuracy of lung nodules in [11] is not as good as the other disease findings, which we believe results from the coarse resolution of the attention maps, *i.e.*, CAMs. Inspired by recent work [6,15] we modify the CNN to generate attention maps with doubled resolution, improving the detection performance of small-scale targets.

## 2 Methods

Our framework is visually depicted in Fig. 1.

### 2.1 Disease Pattern Localization with Attention Mining

Starting from the output of CNN’s last convolutional layer, we denote the feature map as  $X \in R^{N \times W \times H \times D}$ , where  $N$ ,  $W$ ,  $H$ , and  $D$  are the mini-batch size, width, height, and feature map channel, respectively. We then split the classification layer of the CNN into  $C$  branches because feature map erasure is required to be class specific. For now, we assume a binary erasure mask is available, which is defined as  $M^c \in R^{N \times W \times H \times 1}$ , where  $c \in C = \{1, \dots, C\}$  is the index of a specific disease type (see Section 3.2 for details to generate  $M^c$ ). Zeroed regions in  $M^c$  mark spatial regions to drop out of  $X$ . For the  $c^{th}$  disease,  $M^c$  first replicated

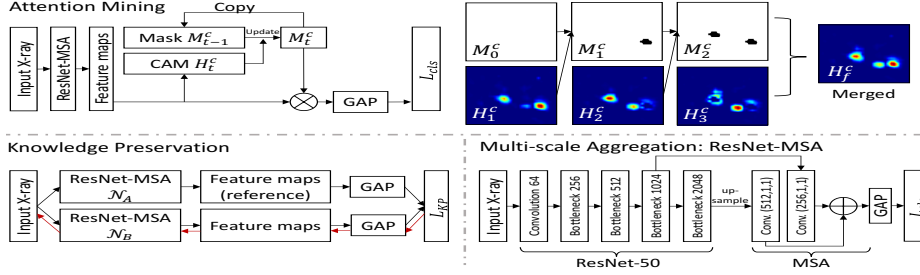


Fig. 1: Architectures of the proposed attention mining (AM), knowledge preservation (KP), and multi-scale aggregation (MSA). Red arrows in the KP module indicate the path of back-propagation. The convolution parameters for MSA are shown as (number of filters, kernel size, stride). See Sec. 2 for details.

across its  $4^{th}$  dimension  $D$  times as  $\hat{M}^c$ , and then the erased feature map is,

$$\hat{X}^c = X \odot \hat{M}^c, \quad (1)$$

where  $\odot$  is element-wise multiplication. The new feature map  $\hat{X}^c$  is then fed into the  $c^{th}$  network branch for binary classification, with the loss defined as,

$$L^c = \frac{1}{N} \left[ h \left( \sigma \left( (w^c)^T g(\hat{X}^c) \right), y^c \right) \right], \quad (2)$$

where  $g(\cdot)$  is global average pooling (GAP) [16] over the  $W$  and  $H$  dimensions,  $w^c \in R^{D \times 1}$  is the network parameter of the  $c^{th}$  branch,  $\sigma(\cdot)$  is the sigmoid activation function,  $y^c \in \{0, 1\}^N$  are the labels of class  $c$  in a mini-batch, and  $h(\cdot)$  is the cross entropy loss function. Thus, the total classification loss is defined as,

$$L_{cls} = \frac{1}{C} \sum_{c \in \mathcal{C}} L^c. \quad (3)$$

While AM can help localize pathologies, the CNN model may overfit to spurious regions after erasure, causing the model to classify an X-ray by remembering its specific image part rather than actual disease patterns. We address this with a knowledge preservation (KP) method described below.

## 2.2 Incremental Learning with Knowledge Preservation

We explore two methods of KP. Given a mini-batch of  $N$  images, a straightforward way to preserve the learned knowledge is to use only the first  $n$  images for AM and leave the later  $N - n$  untouched. If the ratio  $n/N$  is set to be small enough (*e.g.*, 0.125 in our implementation), the CNN's updates can possibly be alleviated from overfitting to uncorrelated image parts. We refer to this vanilla implementation of knowledge preservation as KP-Vanilla.

We investigate a stronger regularizer for KP by constraining the outputs of intermediate network layers. Our main idea is to make the CNN's activation to the later  $N - n$  images be mostly unchanged. Formally, we denote the original network before AM updates as  $\mathcal{N}_A$  and the updated model as  $\mathcal{N}_B$ . Initially,  $\mathcal{N}_A$  and  $\mathcal{N}_B$  are identical to each other, but  $\mathcal{N}_B$  is gradually altered as it learns to classify the blocked feature maps during AM. Considering outputs from the  $k^{th}$  layer of  $\mathcal{N}_A$  and  $\mathcal{N}_B$  as  $X_{A(k)}, X_{B(k)} \in R^{(N-n) \times W \times H \times C}$  for the later  $N - n$  images, we define the distance between  $X_{A(k)}$  and  $X_{B(k)}$  as the  $\ell_2$ -distance between their

GAP features as,

$$L^k = \frac{1}{N-n} \|g(X_{A(k)}) - g(X_{B(k)})\|_2. \quad (4)$$

When multiple network layers are chosen, the total loss from KP is,

$$L_{KP} = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} L^k, \quad (5)$$

where  $\mathcal{K}$  is the indices set of the selected layers, and  $|\mathcal{K}|$  is its cardinality. Finally, the objective for  $\mathcal{N}_B$  training is a weighted combination of  $L_{cls}$  and  $L_{KP}$ ,

$$L = L_{cls} + \lambda L_{KP}, \quad (6)$$

where  $\lambda$  balances the classification and KP loss. Empirically we find the model updates properly when the value of  $\lambda L_{KP}$  is roughly a half of  $L_{cls}$ , *i.e.*,  $\lambda = 0.5$ .

### 2.3 Multi-Scale Aggregation

Our final contribution uses multi-scale aggregation (MSA) to improve the performance of locating small-scale objects, *e.g.*, lung nodules. Taking ResNet-50 [3] as the backbone network, we implement MSA using the outputs of the last two bottlenecks, and refer to the modified network as ResNet-MSA. Given the output of the last bottleneck, denoted as  $X_k \in R^{N \times W/2 \times H/2 \times 2048}$ , we feed it into a  $1 \times 1$  convolutional layer to reduce its channel dimension to 512 and also upsample its width and height by 2 using bilinear interpolation. The resulting feature map is denoted as  $\bar{X}_k \in R^{N \times W \times H \times 512}$ . Similarly, the output of the penultimate bottleneck,  $X_{k-1} \in R^{N \times W \times H \times 1024}$ , is fed into another  $1 \times 1$  convolutional layer to lower its channel dimension to 256, producing  $\bar{X}_{k-1} \in R^{N \times W \times H \times 256}$ . Finally, we concatenate them to produce an aggregated feature map  $X = [\bar{X}_k, \bar{X}_{k-1}]$  for AM. However, MSA is not restricted to bilinear upsampling, as deconvolution [6] can also be used for upsampling, where we use  $3 \times 3$  convolutions. However, as our experiments will demonstrate, the improvements are marginal, leaving bilinear as an efficient option. On the other hand, the channel dimensions of  $X_k$  and  $X_{k-1}$  are largely reduced in order to fit the models into limited GPU memory.


## 3 Experimental Results and Analysis

The proposed method is evaluated on the ChestX-ray14 dataset [11], which contains 51,709 and 60,412 X-ray images of subjects with thoracic and no diseases, respectively. 880 images are marked with bounding boxes (bboxes) corresponding to 984 disease patterns of 8 types, *i.e.*, atelectasis (AT), cardiomegaly (CM), pleural effusion (PE), infiltration (Infiltrat.), mass, nodule, pneumonia (PNA), and pneumothorax (PTx). We first use the same data split as [11] to train base models, *i.e.*, ResNet-50, and ResNet-MSA. Later during AM, the 880 bbox images are then incorporated into the training set to further fine-tune models. We notice that the AM strategy is originally designed to mine disease locations in training images. However, for the purpose of conducting quantitative analysis, we use the bbox images during AM, but only using image labels for training, while leaving the bboxes aside to evaluate localization results.

For ease of comparison, we use the same evaluation metrics as [11]. Given the ground truth and the localized bboxes of a disease, its localization accuracy (Acc.)

Table 1: To compare different MSA setups, each table cell show localization  $Acc.$  given  $T(IoU)=0.3$  using all bboxes in  $\mathcal{B}$ . (See Sec. 3.1 and Sec. 3.2 for details.)

Method	AT	CM	PE	Infiltrat.	Mass	Nodule	PNA	PTx
512-baseline	0.21	<b>0.81</b>	<b>0.37</b>	0.37	0.21	<u>0.04</u>	<b>0.38</b>	<b>0.35</b>
512-bilinear	0.21	0.62	0.34	<b>0.54</b>	<b>0.35</b>	<b>0.24</b>	0.37	0.29
512-deconv.	<b>0.28</b>	0.55	0.35	0.50	0.32	0.20	0.35	0.27
1024-baseline	0.21	0.19	0.33	0.37	0.35	0.09	0.23	0.09
1024-bilinear	0.11	0.10	0.30	0.30	0.15	0.22	0.05	0.08
1024-deconv.	0.07	0.10	0.23	0.40	0.13	0.01	0.05	0.03



(a) object mined at  $t = 2$  (b) object mined at  $t = 3$  (c) the failure case  
 Fig. 2: Visualization of heatmaps generated during attention mining. The ground truth and the automatic bboxes are colored in red and green, respectively.

and average false positive (AFP) is calculated by comparing the intersection over union (IoU) ratios with a predefined threshold, *i.e.*,  $T(IoU)$ . Finally, all of our deep learning implementations are built upon Tensorflow [1] and Tensorpack [13].

### 3.1 Multiple Scale Aggregation

We first test the impact of MSA prior to the application of AM and KP, implementing the bilinear interpolation and deconvolution variants. We also test two different input image resolutions:  $1024 \times 1024$  and  $512 \times 512$ , where the latter is downsampled from the original images using bilinear interpolation. Before applying MSA, we fine-tune the base network ResNet-50 with a learning rate of 0.1 for 50 epochs. Mini-batch sizes for 1024 and 512 inputs are 32 and 64, respectively. Then, to initialize MSA, we fix the network parameters below MSA and tune the other layers for 10 epochs. Finally, we have the whole ResNet-MSA updated end-to-end until the validation loss plot plateaus. Since we mainly focus on investigating AM and KP, no further modification has been taken for the network architecture, and thus the ResNet-MSA achieves similar classification performance as reported in [11] (see supplementary materials for details).

The results of different MSA setups are reported in Table 1, where the “baseline” refers to the original ResNet-50, the “bilinear” and “deconv.” refer to ResNet-MSA with bilinear upsampling, and deconvolution operation, respectively. Prefixes denote the input resolution. As can be seen, the 512 variants perform better than their 1024 counterparts. This is likely because the receptive field size of the MSA layers with “1024-” input is too small to capture sufficient contextual information. Note that for the “512-” input, the two MSA configurations outperform the baseline by a large margin for the infiltration, mass, and nodule categories. This is supporting our design intuition that MSA can help locate small-scale disease patterns more accurately. Because of the efficiency of bilinear upsampling, we select it to be MSA variant of choice, which will be further fine-tuned with AM and KP using Equation (6).

Table 2: Comparison of localization results, which are reported in the form of “Acc.-AFP”. The “ref.” and “base.” are presenting the referred method proposed by Wang *et al.* [11] and the baseline method presented in Sec.3.2 in the manuscript, respectively. The proposed method is presented as “ours”. The best performance, which has the highest Acc. and lowest AFP is presented in bold.

T(IoU)	Method	AT	CM	Effusion	Infiltra.	Mass	Nodule	PNA	PTx
0.1	ref.	0.69-0.89	0.94-0.59	0.66-0.83	0.71-0.62	0.40-0.67	0.14-0.61	0.63-1.02	0.38-0.49
	base.	0.55-0.88	<b>0.99-0.08</b>	0.57-0.82	0.47-0.61	0.36-0.65	0.25-0.59	<b>0.73-1.01</b>	0.42-0.48
	ours	0.68-0.88	0.97-0.18	0.65-0.82	0.52-0.61	<b>0.56-0.65</b>	<b>0.46-0.59</b>	0.65-1.01	<b>0.43-0.48</b>
0.2	ref.	0.47-0.98	0.68-0.72	0.45-0.91	0.48-0.68	0.26-0.69	0.05-0.62	0.35-1.08	0.23-0.52
	base.	0.36-0.97	<b>0.99-0.08</b>	0.33-0.90	0.22-0.67	0.26-0.68	0.09-0.61	0.48-1.07	<b>0.36-0.50</b>
	ours	<b>0.51-0.97</b>	0.90-0.28	<b>0.52-0.90</b>	0.44-0.67	<b>0.47-0.68</b>	<b>0.27-0.61</b>	<b>0.54-1.07</b>	0.24-0.50
0.3	ref.	0.24-1.40	0.46-0.78	0.30-0.95	0.28-0.72	0.15-0.71	0.04-0.62	0.17-1.11	0.13-0.53
	base.	0.22-1.38	<b>0.96-0.12</b>	0.18-0.93	0.13-0.71	0.19-0.69	0.04-0.61	0.31-1.09	<b>0.21-0.52</b>
	ours	<b>0.33-1.38</b>	0.85-0.35	<b>0.34-0.93</b>	<b>0.28-0.71</b>	<b>0.33-0.69</b>	<b>0.11-0.61</b>	<b>0.39-1.09</b>	0.16-0.52
0.4	ref.	0.09-1.08	0.28-0.81	0.20-0.97	0.12-0.75	0.07-0.72	0.01-0.62	0.07-1.12	0.07-0.54
	base.	0.12-1.06	<b>0.92-0.15</b>	0.05-0.95	0.02-0.73	0.12-0.71	<b>0.03-0.61</b>	0.17-1.11	<b>0.12-0.53</b>
	ours	<b>0.23-1.06</b>	0.73-0.47	0.18-0.95	<b>0.20-0.73</b>	<b>0.18-0.71</b>	<b>0.03-0.61</b>	<b>0.23-1.11</b>	0.11-0.53
0.5	ref.	0.05-1.09	0.18-0.84	0.11-0.99	0.07-0.76	0.01-0.72	0.01-0.62	0.03-1.13	0.03-0.55
	base.	0.06-1.07	<b>0.68-0.39</b>	0.02-0.97	0.02-0.74	0.06-0.71	0.01-0.61	0.06-1.12	<b>0.08-0.53</b>
	ours	<b>0.11-1.07</b>	0.60-0.60	0.10-0.97	<b>0.12-0.74</b>	<b>0.07-0.71</b>	<b>0.03-0.61</b>	<b>0.17-1.12</b>	<b>0.08-0.53</b>
0.6	ref.	0.02-1.09	0.08-0.85	0.05-1.00	0.02-0.76	0.00-0.72	0.01-0.62	0.02-1.13	0.03-0.55
	base.	0.01-1.08	<b>0.48-0.60</b>	0.00-0.99	0.02-0.75	0.04-0.71	<b>0.01-0.61</b>	0.03-1.12	0.06-0.53
	ours	<b>0.03-1.08</b>	0.44-0.76	<b>0.05-0.99</b>	<b>0.06-0.75</b>	<b>0.05-0.71</b>	<b>0.01-0.61</b>	<b>0.05-1.12</b>	<b>0.07-0.53</b>
0.7	ref.	0.01-1.10	0.03-0.86	0.02-1.01	0.00-0.77	0.00-0.72	0.00-0.62	0.01-1.13	0.02-0.55
	base.	0.00-1.08	<b>0.18-0.84</b>	0.00-0.99	0.01-0.75	<b>0.01-0.71</b>	<b>0.00-0.61</b>	0.01-1.12	0.01-0.53
	ours	<b>0.01-1.08</b>	0.17-0.84	0.01-0.99	<b>0.02-0.75</b>	<b>0.01-0.71</b>	<b>0.00-0.61</b>	<b>0.02-1.12</b>	<b>0.02-0.53</b>

Table 3: Ablation study of attention mining (AM) and knowledge preservation (KP). Each table cell shows the Acc. by using all of the bboxes in  $\mathcal{B}$  (in Sec. 3.2).

T(IoU)	Method	AT	CM	PE	Infiltrat.	Mass	Nodule	PNA	PTx
AM	t=1	0.57	0.96	<b>0.84</b>	0.78	0.58	<b>0.65</b>	0.66	0.68
	t=2	0.65	<b>0.97</b>	0.82	0.77	0.60	0.61	0.72	<b>0.72</b>
	t=3	<b>0.68</b>	<b>0.97</b>	0.83	<b>0.79</b>	<b>0.62</b>	0.57	<b>0.73</b>	<b>0.72</b>
	t=1	<b>0.34</b>	0.73	0.47	0.40	<b>0.36</b>	<b>0.33</b>	0.33	0.30
	t=2	0.32	0.82	0.47	0.40	0.34	0.25	0.42	0.36
	t=3	0.33	<b>0.85</b>	<b>0.48</b>	<b>0.42</b>	0.34	0.20	<b>0.44</b>	<b>0.38</b>
KP	w/o KP	0.67	<b>1.00</b>	0.48	0.58	0.54	0.47	0.70	0.71
	KP-Vanilla	0.65	0.97	0.78	<b>0.87</b>	0.61	0.48	<b>0.73</b>	<b>0.72</b>
	KP	<b>0.68</b>	0.97	<b>0.83</b>	0.79	<b>0.62</b>	<b>0.57</b>	<b>0.73</b>	<b>0.72</b>
	w/o KP	0.22	<b>0.99</b>	0.11	0.20	0.20	0.06	0.29	0.34
	KP-Vanilla	0.26	0.73	0.41	<b>0.46</b>	0.32	0.10	0.42	0.36
	KP	<b>0.33</b>	0.85	<b>0.48</b>	0.42	<b>0.34</b>	<b>0.20</b>	<b>0.44</b>	<b>0.38</b>

### 3.2 Disease Pattern Localization with AM and KP

In our implementation, we develop the attention mining (AM) basing on the class activation map (CAM) approach [16] that obtains class-specific heat maps. Specifically, the binary erasure mask,  $M^c$  is initialized to be all 1, denoted as  $M_0^c$ . The AM procedure is then iteratively performed  $T$  times, and at time step  $t$ , the intermediate CAMs are generated as,

$$H_t^c = (X \odot \hat{M}_{t-1}^c)w^c, \quad (7)$$

where the inner product is executed across the channel dimension. These CAMs are then normalized to  $[0, 1]$  and binarized with a threshold of 0.5.  $M_t^c$  is then updated from  $M_{t-1}^c$ , except that pixel locations of the connected component that contains the global maximum of the binarized CAM are now set to 0.

There are different options to generate a final heatmap aggregated from all  $T$  CAMs. We choose to have them averaged. However, when  $t > 1$  regions have been erased, as per Equation (7). Thus, we fill in these regions from the corresponding un-erased regions from prior heatmaps. If we define the complement of the masks as  $\bar{M}_t^c = (1 - \hat{M}_t^c)$ , then the final heatmap  $H_f^c$  is calculated using

$$H_f^c = \frac{1}{T} \sum_{t=1}^T \left[ H_t^c + \sum_{t'=1}^{t-1} (H_{t'}^c \odot \bar{M}_{t'}^c) \right]. \quad (8)$$

Empirically, we find  $T = 3$  works best in our implementation.

**Bbox Generation:** To convert CAMs into bboxes, we have 3 bboxes generated from each  $H_f^c$  by adjusting the intensity threshold. For image  $i$ , the bboxes are then ranked as  $\{bbox_i^c(1), bbox_i^c(2), bbox_i^c(3)\}$  in descending order based on the mean  $H_f^c$  intensity values inside the bbox areas. These are then arranged into an aggregated list across all test images from the  $c^{th}$  category:

$\mathcal{B} = \{bbox_1^c(1), \dots, bbox_N^c(1), bbox_1^c(2), \dots, bbox_N^c(2), bbox_1^c(3), \dots, bbox_N^c(3)\}$ . Thereafter, these bboxes are sequentially selected from  $\mathcal{B}$  to calculate  $Acc$ . until the AFP reaches its upper bound, which is the corresponding AFP value reported in [11]. Here, we choose to generate 3 bboxes from each image as it is large enough to cover the corrected locations, while an even larger alternate will greatly increase the AFP value. However, in some cases,  $H_f^c$  would just allow to generate fewer than 3 bboxes, for instance, see Fig. 2(a).

Since Wang *et al.* [11] were not tackling the disease localization in the way as data-mining, direct comparison to their results is not appropriate as they incorporated the bbox images in their test set. Instead, we use our method prior to the application of AM steps as the baseline, which is, for all intents and purposes, Wang *et al.*'s approach [11] applied to the *data-mining problem*. It is presented as the “baseline” method in Table 2. More specifically, it is set up as the 1<sup>st</sup> time step of AM with KP-Vanilla and then fine-tuned until it is converged on the bbox images. As shown in Table 2, our method reports systematic and consistent quantitative performance improvements over the “baseline” method, except slightly degrades on the category of CM, demonstrating the impact of our AM and KP enhancements. Meanwhile, comparing with the results in [11], our method achieves significant improvements by using no extra manual annotations. More Importantly, the results in Table 2 indicate our method would also be effective when implemented to mine disease locations in the training images.

Figure 2 depicts three atelectasis cases visualizing the AM process. As can be seen, AM improves upon the baseline results,  $H_1^c$ , by discovering new regions after erasing that correlate with the disease patterns. More qualitative results can be found in the supplementary materials.

**Ablation Study:** We further investigate the impact of AM and KP. First, we compare the three steps in AM. Since the erasure map  $M_0^c$  is initialized with all 1s, then the time step  $t = 1$  is treated as the baseline. Table 3 shows the localization results at two IoU thresholds, 0.1 and 0.3. As can be seen, significant improvements of AM are observed in the AT, CM, mass, PNA, and PTx disease patterns. Next, we compare the KP, KP-Vanilla and an implementation without

KP, where the ResNet-MSA is tuned in AM using only the bbox images. In particular, we set  $\mathcal{K}$  by using the outputs of the 3<sup>rd</sup> and 4<sup>th</sup> bottleneck, the MSA, and the classification layers. As Table 3 presents, KP performs better than KP-Vanilla in the AT, PE, mass, nodule, PNA and PTx categories.

## 4 Conclusion

We present a novel localization data-mining framework, combining AM, KP, and MSA. We introduce a powerful means to harvest disease locations from chest X-ray datasets. By showing improvements over a standard CAM-based approach, our method can mine localization knowledge in existing large-scale datasets, potentially allowing for the training of improved computer-aided diagnosis tools or more powerful retrospective analyses. Future work includes improving the MSA, possibly by using the atrous convolution [15]. Additionally, we find that when the activation map fails to localize disease in none of the AM steps, our method will not locate the correct image region as demonstrated in Figure 2(c). To address this issue, we may consider semi-supervised learning, like the use of bboxes in [5], as a complementary means to discover those difficult cases.

## References

1. Abadi, M., Agarwal, A., Barham, P., et al.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015), <https://www.tensorflow.org/>
2. Geoffrey, H., et al.: Distilling the knowledge in a neural network. In: NIPS (2014)
3. He, K., et al.: Deep residual learning for image recognition. In: IEEE CVPR (2016)
4. Jing, B., Xie, P., Xing, E.: On the automatic generation of medical imaging reports. arXiv:1711.08195 (2017)
5. Li, Z., Wang, C., Han, M., Xue, Y., et al.: Thoracic disease identification and localization with limited supervision. IEEE CVPR (2018)
6. Noh, H., Hong, S., Han, B.: Learning deconvolution network for semantic segmentation. In: IEEE ICCV. pp. 1520–1528 (2015)
7. Pesce, E., Ypsilantis, P.P., et al.: Learning to detect chest radiographs containing lung nodules using visual attention networks. arXiv:1712.00996 (2017)
8. Rajpurkar, P., Irvin, J., Zhu, K., et al.: Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. arXiv:1711.05225 (2017)
9. Shmelkov, K., Cordelia Schmid, K.A.: Incremental learning of object detectors without catastrophic forgetting. In: IEEE ICCV. pp. 3420–3429 (2017)
10. Wang, X., Peng, Y., et al.: Tienet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays. In: IEEE CVPR (2018)
11. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: IEEE CVPR (2017)
12. Wei, Y., et al.: Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In: IEEE CVPR. pp. 6488–6496 (2017)
13. Wu, Y., et al.: Tensorpack. <https://github.com/tensorpack/> (2016)
14. Yao, L., Poblens, E., Dagunts, D., et al.: Learning to diagnose from scratch by exploiting dependencies among labels. arXiv:1710.10501 (2017)
15. Yu, F., Koltun, V.: Multi-Scale Context Aggregation by Dilated Convolutions. In: ICLR (2016)
16. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: IEEE CVPR. pp. 2921–2929 (2016)