



Anatomy-Guided Weakly-Supervised Abnormality Localization in Chest X-rays

Ke Yu¹(✉), Shantanu Ghosh¹, Zhexiong Liu¹, Christopher Deible²,
and Kayhan Batmanghelich¹

¹ University of Pittsburgh, Pittsburgh, PA, USA
yu.ke@pitt.edu

² University of Pittsburgh Medical Center, Pittsburgh, PA, USA

Abstract. Creating a large-scale dataset of abnormality annotation on medical images is a labor-intensive and costly task. Leveraging *weak supervision* from readily available data such as radiology reports can compensate lack of large-scale data for anomaly detection methods. However, most of the current methods only use image-level pathological observations, failing to utilize the relevant *anatomy mentions* in reports. Furthermore, Natural Language Processing (NLP)-mined weak labels are noisy due to label sparsity and linguistic ambiguity. We propose an Anatomy-Guided chest X-ray Network (AGXNet) to address these issues of weak annotation. Our framework consists of a cascade of two networks, one responsible for identifying anatomical abnormalities and the second responsible for pathological observations. The critical component in our framework is an anatomy-guided attention module that aids the downstream observation network in focusing on the relevant anatomical regions generated by the anatomy network. We use Positive Unlabeled (PU) learning to account for the fact that lack of mention does not necessarily mean a negative label. Our quantitative and qualitative results on the MIMIC-CXR dataset demonstrate the effectiveness of AGXNet in disease and anatomical abnormality localization. Experiments on the NIH Chest X-ray dataset show that the learned feature representations are transferable and can achieve the state-of-the-art performances in disease classification and competitive disease localization results. Our code is available at <https://github.com/batmanlab/AGXNet>.

Keywords: Weakly-supervised learning · PU learning · Disease detection · Class activation map · Residual attention

1 Introduction

There is considerable interest in developing automated abnormality detection systems for chest X-rays (CXR) in order to improve radiologists' workflow effi-

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-16443-9_63.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2022
L. Wang et al. (Eds.): MICCAI 2022, LNCS 13435, pp. 658–668, 2022.
https://doi.org/10.1007/978-3-031-16443-9_63




Fig. 1. Schematic diagram of the proposed AGXNet. Our architecture comprises of two classification networks, g_a for anatomical abnormalities and g_o for pathological observations. Relations $(a_i, o_k) = P$, $(a_j, o_k) = P$ are parsed from the report and represent that the observation o_k is annotated as *Present* in two anatomical landmarks a_i, a_j , respectively. We obtain the anatomy-guided attention map H^k by aggregating CAMs of a_i, a_j and incorporate the H^k as a residual attention into g_o . The symbols f_a, f_o denote the intermediate anatomy and observation features, respectively.

ciency and reduce observational oversights [3, 6, 13, 19]. Typically, training a high-precision detection model using deep learning requires high-quality annotations. However, collecting large-scale annotations by clinical experts is time-consuming and prohibitively expensive. This motivates weakly-supervised learning (WSL) methods [2, 14, 16, 24] that leverage weak supervision from paired CXR reports that are readily available on a large scale. There are, however, two unaddressed challenges: (1) The entire report is often summarized to a small set of disease labels, which misses the opportunity of incorporating anatomical context, and (2) Weak labels derived from radiology reports using Natural Language Processing (NLP) are noisy due to the sparsity of the labels and the ambiguity of the language. Improper handling of labeling noise can result in underperforming models. We propose a novel WSL framework to bridge these two gaps.

A variety of automated text labelers [9, 11, 18, 24] have been developed to extract image-level observations from CXR reports. However, they do not consider *anatomy mentions* that provide important contexts for associated observations. Aiming to fill this gap, the recently developed pipelines, Chest Imagenome [26] and RadGraph [10] extract fine-grained “*observation-located at-anatomy*” relations, (e.g., “*opacity in the right lower lung*”), from CXR reports. Yet none of them has been incorporated into a weakly-supervised disease detection method. We build our WSL framework based on the RadGraph dataset. Radiologists typically employ a systematic approach [21] when reading CXR images to ensure that no significant abnormality is missed. This approach is essentially documented in CXR reports which typically have imaging observa-

tions paired with the related anatomical locations. We design an architecture that reflects this process.

The second challenge is that weak labels extracted from CXR reports are inherently noisy. Typically, only a subset of weak labels is mentioned in a report, and unlabeled data is handled using a basic zero-imputation strategy. However, lack of mention does not necessarily mean a negative label. For example, a CXR report that contains mentions of strong visual clues for pneumonia, such as *opacity* or *consolidation*, may not specifically establish *pneumonia* as a diagnosis due to human variability in the reporting process. We also consider uncertainty mentions (e.g., *may*, *possible*, *can't exclude*) as unlabeled, where the noise originates from the intrinsic ambiguities in reports. We formulate this problem as a Positive Unlabeled (PU) learning [1] where the learner has access to a small set of examples with high-confidence labels (either positive or negative) and a large amount of unlabeled data mixed by positive and negative examples with an unknown mixture ratio.

In this paper, we present Anatomy-Guided chest X-ray Networks (AGXNet), a novel WSL framework that leverages information from pathological observations and their associated anatomical landmarks mentioned in CXR reports. Our architecture consists of a cascade of two networks, with the upstream anatomy network tasked with identifying anatomical abnormalities in CXR images and the downstream observation network tasked with identifying pathological observations. The inclusion of an anatomy-guided attention (AGA) module aims to aid the observation network paying attention to abnormalities in the context of associated anatomical landmarks. During training, the AGA module also provides a feedback loop to the upstream anatomy network, thus mutually improving the quality of both types of features. In addition, we adopt a PU learning technique to estimate the fraction of positive samples in the unlabeled data and use a self-training approach to iteratively reduce noise in labels. Our model is trained end-to-end. We evaluate the proposed framework on the MIMIC-CXR [12] dataset. Results show that the proposed AGXNet model outperforms both supervised and weakly-supervised baselines in disease localization. In-depth ablation studies demonstrate that both AGA module and PU learning can help to improve localization accuracy. We further evaluate the pre-trained encoders on the NIH Chest X-ray dataset [24] and show that the transferred models achieve the state-of-the-art (SOTA) results in disease classification and competitive performance in disease localization.

2 Methods

We parse the report into an adjacency matrix that encodes the relations between the observations and anatomical landmarks. The AGA module connects two networks that are used to predict the presence of observations and anatomical abnormalities in CXR images. We use PU learning to explicitly address noise in unlabeled observations. The proposed framework is illustrated in Fig. 1.

Report Representation. We use the RadGraph [10] to parse reports to anatomical landmark tokens $\{a_j\}$, observation tokens $\{o_k\}$, and encode their relations using an adjacency matrix A . An entry $A(j, k)$ can take one of the three variants P, N , or U , representing whether the observation o_k is *Present*, *Absent* or *Unlabeled* in the anatomical landmark a_j . Note, for o_k that is mentioned without a specified location (e.g., “*No evidence of pneumothorax*”), we link it to a special anatomical landmark token named *unspecified*. Figure A.1 shows an example adjacent matrix. We summarize the image-level labels for observation and anatomical landmark using the adjacency matrix. An observation o_k is assigned a positive label ($y(o_k) = 1$) if column $A(\cdot, k)$ has at least one P entry, a negative label ($y(o_k) = 0$) if column $A(\cdot, k)$ has no P entry and at least one N entry, and assigned as unlabeled ($y(o_k) = u$) if column $A(\cdot, k)$ only has U entries. For anatomy, an anatomical landmark is labeled as abnormal $y(a_j) = 1$ if row $A(j, \cdot)$ has at least one P entry, otherwise it is labeled as normal $y(a_j) = 0$.

Anatomy Network. The anatomy network g_a is responsible for identifying abnormalities from anatomical landmarks. The weighted binary cross-entropy (BCE) loss for a_j is given by:

$$\mathcal{L}_{a_j}(\theta_a, w_a^j) = -b_j^+ \sum_{y(a_j)=1} \ln(g_a(x)) - b_j^- \sum_{y(a_j)=0} \ln(1 - g_a(x)), \quad (1)$$

where θ_a is the parameter of anatomy encoder network, w_a^j is the classification weight corresponding to a_j , x is an image, $g_a(x)$ is the predicted probability of being abnormal in a_j , b_j^+ and b_j^- are balancing factors introduced in [20].

AGA. We introduce the AGA module to guide the downstream observation network g_o focusing on the relevant anatomical regions mentioned in the radiology report. In detail, we construct the AGA map H^k for observation o_k by aggregating the class activation maps (CAMs) [27] for locations $\{a_j\}$ where o_k was positively observed, i.e., $A(a_j, o_k) = P$. Formally,

$$H^k(\theta_a, w_a^j) = \sum_j \mathbb{1}\{A(a_j, o_k) = P\} \underbrace{\sum_c w_a^j f_a(:, :, c; \theta_a)}_{\text{CAM}(a_j)}, \quad (2)$$

where $f_a(:, :, c)$ represents the activation of channel c in anatomy feature map f_a , and $\mathbb{1}\{A(a_j, o_k) = P\}$ is an indicator function. We then transform the values of H^k to $[0, 1]$ using the min-max normalization.

Observation Network. The observation network g_o is responsible for identifying the presence of pathological observations in CXR images. We incorporate the H^k into g_o as a residual attention map and modify the observation feature map f_o as follows, $f_{o_k}' = (1 + \beta \cdot H^k) \odot f_o$, where $\odot$ indicates element-wise multiplication for spatial positions, and β is a scaling hyperparameter. The weighted

BCE loss for o_k is given by:

$$\mathcal{L}_{o_k}(\Theta) = -b_k^+ \sum_{y(o_k)=1} \ln(g_o(x, H^k)) - b_k^- \sum_{y(o_k) \in \{0, u\}} \ln(1 - g_o(x, H^k)), \quad (3)$$

where $\Theta = (\theta_o, w_o^k, \theta_a, w_a^j)$, θ_o is the parameter of observation encoder network, w_o^k is the classification weight corresponding to o_k , $g_o(x, H^k)$ is the predicted probability of o_k being present in the image, b_k^+ and b_k^- are balancing factors. It is important to note that mapping all unlabeled examples to negative ($|y(o_k) \in \{0, u\}| = |N| + |U|$) would overlook the noise in the unlabeled data, which can degrade model performance. We explicitly handle the randomness present in unlabeled data and formulate this problem using a PU learning approach.

PU Learning. The distribution of unlabeled data P_u can be decomposed as, $P_u = \alpha P_p + (1 - \alpha)P_n$, where α denotes the mixture proportion of positive examples in the unlabeled data, and P_p, P_n denotes the class-conditional distribution for positive and negative class, respectively. We adopt the method *Best Bin Estimation* (BBE) proposed by Garg *et al.* [5] to estimate α . In short, let X_p, X_n, X_u denote the positive, negative and unlabeled samples for o_k in the validation data set, $\hat{F}_p(z), \hat{F}_u(z)$ denote the empirical cumulative distributions of the predicted probabilities of the observation network, named Z_p, Z_u . The mixture proportion α is estimated by minimizing the upper confidence bound of the ratio, namely $(1 - \hat{F}_u(z))/(1 - \hat{F}_p(z))$. We integrate the BBE with an iterative self-training approach summarized as follows: (1) warm-start training with treating all unlabeled samples as negative, (2) estimating α using BBE, (3) removing α fraction of unlabeled training samples scored as most positive and relabeling the rest $1 - \alpha$ unlabeled samples as negative, (4) updating the model using positive samples ($|P|$) and provisional negative samples ($|N| + (1 - \alpha)|U|$). We repeat steps 2 to 4 until the classifiers reach the best validation performance.

Optimization. The final loss is given by: $\mathcal{L} = \frac{1}{N_a} \sum_j^{N_a} \mathcal{L}_{a_j} + \frac{1}{N_o} \sum_k^{N_o} \mathcal{L}_{o_k}$, where N_a, N_o are the number of anatomy and observation labels used in training, respectively. The network is trained end-to-end. Importantly, during training, the gradients $\partial \mathcal{L}_{o_k} / \partial \theta$, $\partial \mathcal{L}_{o_k} / \partial w_a^j$ provide a feedback loop to the anatomy network through the AGA module, mutually reinforcing the anatomy and observation features. During inference time, we do not require any text input and set $H^k = 0$.

3 Experiments

Experiments are carried out to evaluate the performance of AGXNet in abnormality classification and localization. We conduct ablation studies to validate the efficiency of the AGA module and PU learning. We also test the robustness and transferability of the learned anatomy and observation features.

Experimental Details. We first evaluate the proposed AGXNet on the MIMIC-CXR dataset [12]. The RadGraph’s inference dataset [10] provides annotations automatically generated by DYGIE++ [23] model for 220,763 MIMIC-CXR reports. We select the corresponding 220,763 frontal images from the

MIMIC-CXR dataset and obtain their adjacency matrix representations from the RadGraph annotations. The adjacency matrix of each sample has 48 rows (46 anatomical landmarks, 1 *unspecified*, 1 *other anatomies* representing anatomy tokens in the tail distribution) and 64 columns (63 mostly mentioned observations and 1 *other observations* representing observation tokens in the tail distribution). Appx. B provides additional details of anatomical labels. For evaluation of disease localization, we use a held out set [22] of MIMIC-CXR images with 390 bounding boxes (BBox) for pneumonia (196/390) and pneumothorax (194/390) annotated by board-certified radiologists. For evaluation of anatomical abnormality localization, we utilize the anatomy BBox from the Chest ImaGenome dataset [26], which were extracted by an atlas-based detection pipeline. We use a 80%-10%-10% train-validation-test split with no patient shared across splits. For both anatomy and observation encoders in AGXNet, we use DenseNet-121 [8] with pre-trained weights from ImageNet [4] as the backbone. We train our framework to predict presence of abnormality in the 46 anatomical landmarks and the presence of two diseases (i.e., pneumonia and pneumothorax) that have ground truth BBox in [22]. The β is set to be 0.1 based on validation results. We optimize the networks using SGD with momentum = 0.9, weight decay = $10e^{-4}$ and stop training once the validation error reaches minimum. The learning rate is set to be 0.01 and divided by 10 every 6 epochs. We resize all images to 512×512 without any data augmentation and set the batch size as 16. The model is implemented in PyTorch and trained on a single NVIDIA GPU with 32G of memory.

Evaluation Metric. We produce disease-specific CAMs for pneumonia and pneumothorax, and anatomy-specific CAMs for the anatomical landmarks. We apply a thresholding-based bounding box generation method and extract isolated regions in which pixel values are greater than the 95% quantile of the CAM’s pixel value distribution. We evaluate the generated boxes against the ground truth BBox using intersection over union ratio (IoU). A generated box is considered as a true positive when $\text{IoU} > T(\text{IoU})$, where $T(*)$ is a threshold.

Table 1. Disease localization on MIMIC-CXR. Results are reported as the average over 5 independent runs with standard deviations. The highest values are highlighted in bold, and the best results without BBox annotation are underlined.

Disease	Model	Supervised	IoU @ 0.1		IoU @ 0.25		IoU @ 0.5	
			Recall	Precision	Recall	Precision	Recall	Precision
Pneumothorax	RetinaNet [15]	✓	0.52 ± 0.04	0.19 ± 0.00	0.41 ± 0.04	0.15 ± 0.00	0.26 ± 0.00	0.09 ± 0.00
	CheXpert [9]	✗	0.68 ± 0.01	0.40 ± 0.01	0.47 ± 0.01	0.28 ± 0.00	0.14 ± 0.01	0.08 ± 0.01
	AGXNet w/o AGA	✗	0.67 ± 0.04	0.41 ± 0.03	0.54 ± 0.04	0.33 ± 0.03	<u>0.16 ± 0.02</u>	<u>0.10 ± 0.01</u>
	AGXNet w/ AGA	✗	0.75 ± 0.01	0.44 ± 0.01	0.60 ± 0.02	0.35 ± 0.01	<u>0.16 ± 0.00</u>	0.09 ± 0.00
	AGXNet w/ AGA + PU	✗	0.74 ± 0.03	0.43 ± 0.02	0.58 ± 0.03	0.34 ± 0.02	0.13 ± 0.01	0.07 ± 0.01
Pneumonia	RetinaNet [15]	✓	0.66 ± 0.04	0.27 ± 0.04	0.62 ± 0.04	0.26 ± 0.09	0.42 ± 0.04	0.17 ± 0.03
	CheXpert [9]	✗	0.73 ± 0.01	0.40 ± 0.01	0.51 ± 0.01	0.28 ± 0.01	0.19 ± 0.01	0.10 ± 0.01
	AGXNet w/o AGA	✗	0.69 ± 0.01	0.44 ± 0.01	0.59 ± 0.02	0.37 ± 0.02	0.18 ± 0.02	0.11 ± 0.01
	AGXNet w/ AGA	✗	0.67 ± 0.02	0.46 ± 0.01	0.58 ± 0.02	0.39 ± 0.02	0.18 ± 0.03	0.12 ± 0.02
	AGXNet w/ AGA + PU	✗	0.73 ± 0.01	0.47 ± 0.01	0.62 ± 0.01	0.40 ± 0.01	0.20 ± 0.01	<u>0.13 ± 0.01</u>

Evaluating Disease Localization. We trained a RetinaNet [15] using the annotated BBox from [22] as the supervised baseline and a DenseNet-121 using CheXpert [9] disease labels for CAM-based localization as the WSL baseline. We investigated three variants of AGXNet to understand the effects of AGA module and PU learning. Table 1 shows that AGXNet trained with PU learning achieved the best localization results for pneumonia at $T(\text{IoU}) = 0.1$ and 0.25 , while AGXNet trained without PU learning performed best in pneumothorax localization at the same IoU thresholds. Note that the label noise in pneumothorax is intrinsically low, thus adding PU learning may not significantly improve its localization accuracy. RetinaNet achieved the best localization results at $T(\text{IoU}) = 0.5$, probably due to its direct predictions for the coordinates of BBox.


Table 2. Ablation studies on AGA and PU Learning. Results are averaged over 5 runs. α : mixture proportion of positive samples. AA: anatomic abnormality. LAL: left apical lung. LL: left lung. RAL: right apical lung. RLL: right lower lung. RL: right lung.

Model	Classification				AA localization				
	Pneumonia		Pneumothorax		Accuracy @ IoU = 0.25				
	α	AUPRC	α	AUPRC	LAL	LL	RAL	RLL	RL
AGXNet w/o AGA	–	0.57	–	0.55	0.02	0.10	0.08	0.43	0.12
AGXNet w/ AGA	–	0.57	–	0.57	0.28	0.22	0.49	0.63	0.22
AGXNet w/ AGA + PU	0.14	0.62	0.03	0.57	0.38	0.23	0.50	0.69	0.22

Ablation Studies. Table 2 shows additional results of ablation studies for AGXNet, including (1) disease classification performance on the positive and negative samples in the test set using the area under the precision-recall curve (AUPRC), and (2) anatomical abnormality localization accuracy at $T(\text{IoU}) = 0.25$ for $\{a_j\}$, where $A(a_j, o_k) = P$ and $o_k \in \{\text{pneumonia, pneumothorax}\}$. **Effect of AGA:** Results in Table 1 and Table 2 show that the AGA module significantly improved the results of pneumothorax detection and anatomical abnormality detection, suggesting that the attention module mutually enhanced both types of features. Figure 2 shows the qualitative comparison of two models with (M2) and without (M1) the AGA module. M2 correctly detects pneumothorax and the relevant abnormal anatomical landmarks in both examples, while M1 fails to do so. Furthermore, in the second example, a pigtail catheter is applied to treat pneumothorax and acts as the only discriminative feature used by M1, while M2 is more robust against this shortcut and detects disease in the correct anatomical location. **Effect of PU Learning:** Table 2 shows that the estimated mixture proportion of positive examples (α) is significantly higher in pneumonia (14%) than in pneumothorax (3%), reflecting the fact that there is considerable variability in pneumonia diagnoses [17]. Accordingly, results in Table 1 and Table 2 show that using the PU learning technique improves the performance of pneumonia classification and localization, while it is less effective for pneumothorax whose label noise is intrinsically low.

Table 3. The test AUROCs for the NIH Chest X-ray disease classification task.

Method	Atelectasis	Cardiomegaly	Effusion	Infiltration	Mass	Nodule	Pneumonia	Pneumothorax	Mean
Scratch	0.79	0.91	0.88	0.69	0.81	0.70	0.70	0.82	0.79
Wang et al. [24]	0.72	0.81	0.78	0.61	0.71	0.67	0.63	0.81	0.72
Wang et al. [25]	0.73	0.84	0.79	0.67	0.73	0.69	0.72	0.85	0.75
Liu et al. [16]	0.79	0.91	0.88	0.69	0.81	0.73	0.75	0.89	0.80
Rajpurkar et al. [20]	0.82	0.91	0.88	0.72	0.86	0.78	0.76	0.89	0.83
Han et al. [7]	0.84	0.93	0.88	0.72	0.87	0.79	0.77	0.90	0.84
AGXNet - Ana.	0.84	0.91	0.90	0.71	0.86	0.80	0.74	0.86	0.83
AGXNet - Obs.	0.85	0.92	0.90	0.72	0.86	0.80	0.76	0.87	0.84
AGXNet - Both	0.85	0.92	0.90	0.72	0.87	0.80	0.76	0.88	0.84

**Fig. 2.** AGXNet with (M2) or without (M1) the AGA module for pneumothorax detection. Heatmaps indicate the CAMs. The blue and green boxes stand for ground truth anatomy and disease annotations. The red boxes are generated from the CAMs. (Color figure online)**Table 4.** The localization accuracy for the NIH Chest X-ray disease localization task. The highest values are highlighted in bold, and the best results without BBox annotation are underlined.

T(IoU)	Method	Supervised	Atelectasis	Cardiomegaly	Effusion	Infiltration	Mass	Nodule	Pneumonia	Pneumothorax	Mean
0.1	Scratch	✗	0.46	0.85	0.61	0.40	0.48	0.11	0.56	0.25	0.47
	Han et al. [7]	✓	0.72	0.96	0.88	0.93	0.74	0.45	0.65	0.64	0.75
	Wang et al. [24]	✗	0.69	0.94	0.66	<u>0.71</u>	0.40	0.14	0.63	0.38	0.57
	AGXNet - Ana.	✗	0.64	0.99	0.66	0.69	<u>0.69</u>	<u>0.29</u>	0.73	0.29	0.62
	AGXNet - Obs.	✗	0.61	0.99	<u>0.70</u>	0.70	<u>0.69</u>	0.23	0.72	0.55	0.65
	AGXNet - Both	✗	<u>0.71</u>	1.00	0.69	0.68	0.68	0.22	0.71	<u>0.62</u>	<u>0.66</u>
0.3	Scratch	✗	0.10	0.45	0.34	0.15	0.14	0.00	0.31	0.06	0.19
	Han et al. [7]	✓	0.39	0.85	0.60	0.67	0.43	0.21	0.40	0.45	0.50
	Wang et al. [24]	✗	0.24	0.46	0.30	0.28	0.15	<u>0.04</u>	0.17	0.13	0.22
	AGXNet - Ana.	✗	0.20	0.48	0.40	0.43	<u>0.28</u>	0.00	0.46	0.13	0.30
	AGXNet - Obs.	✗	0.22	<u>0.62</u>	0.45	<u>0.46</u>	0.27	0.01	0.54	0.31	<u>0.36</u>
	AGXNet - Both	✗	<u>0.27</u>	0.56	<u>0.47</u>	0.45	0.27	0.01	0.50	<u>0.36</u>	<u>0.36</u>

Transfer Learning on NIH Chest X-ray. We pre-trained an AGXNet + PU model on the MIMIC-CXR dataset using the 46 anatomy labels and 8 observation labels listed in C.1. We then fine-tuned the encoder(s) and re-trained a classifier using the NIH Chest X-ray dataset [24]. We investigated three variants

of fine-tuning regimes: (1) only using anatomy encoder, (2) only using observation encoder, and (3) using both encoders and their concatenated embeddings. We compare our transferred models with a baseline model trained from scratch using the NIH Chest X-ray dataset and a series of relevant baselines. We evaluate the classification performance using the area under the ROC curve (AUROC) score and the localization accuracy at $T(\text{IoU}) = 0.1, 0.3$ across the 8 diseases. Table 3 shows that all variants of fine-tuned AGXNet models achieve disease classification performances comparable to the SOTA method [7], demonstrating that both learned anatomy and observation features are robust and transferable. Table 4 shows that all variants of fine-tuned AGXNet models outperform learning from scratch and the existing CAM-based baseline method [24] in disease localization. Note that the model proposed by Han et al. [7] achieved higher accuracy by utilizing the ground truth BBox during training, therefore it is not directly comparable and should be viewed as an upper bound method.

4 Conclusion

In this work, we propose a novel WSL framework to incorporate anatomical contexts mentioned in radiology reports to facilitate disease detection on corresponding CXR images. In addition, we use a PU learning approach to explicitly handle noise in unlabeled data. Experimental evaluations on the MIMIC-CXR dataset show that the addition of anatomic knowledge and the use of PU learning improve abnormality localization. Experiments on the NIH Chest X-ray datasets demonstrate that the learned anatomical and pathological features are transferable and encode robust classification and localization information.

Acknowledgements. This work was partially supported by NIH Award Number 1R01HL141813-01, NSF 1839332 Tripod+X, and SAP SE. We are grateful for the computational resources provided by Pittsburgh SuperComputing grant number TG-ASC170024.

References

1. Bekker, J., Davis, J.: Learning from positive and unlabeled data: a survey. *Mach. Learn.* **109**(4), 719–760 (2020). <https://doi.org/10.1007/s10994-020-05877-5>
2. Bhalodia, R., et al.: Improving pneumonia localization via cross-attention on medical images and reports. In: de Bruijne, M., et al. (eds.) *MICCAI 2021*. LNCS, vol. 12902, pp. 571–581. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87196-3_53
3. Castellino, R.A.: Computer aided detection (CAD): an overview. *Cancer Imaging* **5**(1), 17 (2005)
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition, pp. 248–255. IEEE (2009)
5. Garg, S., Wu, Y., Smola, A.J., Balakrishnan, S., Lipton, Z.: Mixture proportion estimation and PU learning: a modern approach. In: *Advances in Neural Information Processing Systems* 34 (2021)

6. Gromet, M.: Comparison of computer-aided detection to double reading of screening mammograms: review of 231,221 mammograms. *Am. J. Roentgenol.* **190**(4), 854–859 (2008)
7. Han, Y., et al.: Knowledge-augmented contrastive learning for abnormality classification and localization in chest X-rays with radiomics using a feedback loop. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 2465–2474 (2022)
8. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4700–4708 (2017)
9. Irvin, J., et al.: CheXpert: a large chest radiograph dataset with uncertainty labels and expert comparison. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 590–597 (2019)
10. Jain, S., et al.: RadGraph: extracting clinical entities and relations from radiology reports (2021)
11. Jain, S., et al.: VisualCheXbert: addressing the discrepancy between radiology report labels and image labels. In: *Proceedings of the Conference on Health, Inference, and Learning*, pp. 105–115 (2021)
12. Johnson, A.E., et al.: MIMIC-CXR-JPG - chest radiographs with structured labels (version 2.0.0). *PhysioNet* (2019). <https://doi.org/10.13026/8360-t248>
13. Lakhani, P., Sundaram, B.: Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks. *Radiology* **284**(2), 574–582 (2017)
14. Li, Z., et al.: Thoracic disease identification and localization with limited supervision. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8290–8299 (2018)
15. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2980–2988 (2017)
16. Liu, J., Zhao, G., Fei, Y., Zhang, M., Wang, Y., Yu, Y.: Align, attend and locate: chest X-ray diagnosis via contrast induced attention network with limited supervision. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10632–10641 (2019)
17. Neuman, M.I., et al.: Variability in the interpretation of chest radiographs for the diagnosis of pneumonia in children. *J. Hosp. Med.* **7**(4), 294–298 (2012)
18. Peng, Y., Wang, X., Lu, L., Bagheri, M., Summers, R., Lu, Z.: NegBio: a high-performance tool for negation and uncertainty detection in radiology reports. *AMIA Summits Transl. Sci. Proc.* **2018**, 188 (2018)
19. Qin, C., Yao, D., Shi, Y., Song, Z.: Computer-aided detection in chest radiography based on artificial intelligence: a survey. *Biomed. Eng. Online* **17**(1), 1–23 (2018)
20. Rajpurkar, P., et al.: CheXNet: radiologist-level pneumonia detection on chest X-rays with deep learning. *arXiv preprint arXiv:1711.05225* (2017)
21. Sait, S., Tombs, M.: Teaching medical students how to interpret chest X-rays: the design and development of an e-learning resource. *Adv. Med. Educ. Pract.* **12**, 123 (2021)
22. Tam, L.K., Wang, X., Turkbey, E., Lu, K., Wen, Y., Xu, D.: Weakly supervised one-stage vision and language disease detection using large scale pneumonia and pneumothorax studies. In: Martel, A.L., et al. (eds.) *MICCAI 2020. LNCS*, vol. 12264, pp. 45–55. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59719-1_5

23. Wadden, D., Wennberg, U., Luan, Y., Hajishirzi, H.: Entity, relation, and event extraction with contextualized span representations. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 5784–5789, November 2019. <https://doi.org/10.18653/v1/D19-1585>. <https://aclanthology.org/D19-1585>
24. 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: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2097–2106 (2017)
25. Wang, X., Peng, Y., Lu, L., Lu, Z., Summers, R.M.: TieNet: text-image embedding network for common thorax disease classification and reporting in chest X-rays. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 9049–9058 (2018)
26. Wu, J.T., et al.: Chest ImaGenome dataset for clinical reasoning. In: Thirty-Fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2) (2021)
27. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2921–2929 (2016)