

# SCALP - Supervised Contrastive Learning for Cardiopulmonary Disease Classification and Localization in Chest X-rays using Patient Metadata

Ajay Jaiswal Tianhao Li Cyprian Zander Yan Han Justin F. Rousseau Yifan Peng Ying Ding  
UT Austin UT Austin MIS, Germany UT Austin UT Austin Weill Cornell Medicine UT Austin

**Abstract**—Computer-aided diagnosis plays a salient role in more accessible and accurate cardiopulmonary diseases classification and localization on chest radiography. Millions of people get affected and die due to these diseases without an accurate and timely diagnosis. Recently proposed contrastive learning heavily relies on data augmentation, especially positive data augmentation. However, generating clinically-accurate data augmentations for medical images is extremely difficult because the common data augmentation methods in computer vision, such as sharp, blur, and crop operations, can severely alter the clinical settings of medical images. In this paper, we proposed a novel and simple data augmentation method based on patient metadata and supervised knowledge to create clinically accurate positive and negative augmentations for chest X-rays. We introduce an end-to-end framework, SCALP, which extends the self-supervised contrastive approach to a supervised setting. Specifically, SCALP pulls together chest X-rays from the same patient (positive keys) and pushes apart chest X-rays from different patients (negative keys). In addition, it uses ResNet-50 along with the triplet-attention mechanism to identify cardiopulmonary diseases, and Grad-CAM++ to highlight the abnormal regions. Our extensive experiments demonstrate that SCALP outperforms existing baselines with significant margins in both classification and localization tasks. Specifically, the average classification AUCs improve from 82.8% (SOTA using DenseNet-121) to 83.9% (SCALP using ResNet-50), while the localization results improve on average by 3.7% over different IoU thresholds.

**Index Terms**—Thoracic Disorder, Contrastive Learning, Chest-Xray, Classification, Bounding Box

## I. INTRODUCTION

Chest X-rays (CXR) are one of the most common imaging tools used to examine cardiopulmonary diseases. Currently, CXRs diagnosis primarily relies on professional knowledge and meticulous observations of expert radiologists. Automated systems for medical image classification face several challenges. First, it heavily depends on manually-annotated training data which requires highly specialized radiologists to do manual annotation. Radiologists are already overloaded with their diagnosis duties and their hourly charge is costly. Second, common data augmentation methods in computer vision [1], [2] such as crop, mask, blur, and color jitter can significantly alter medical images and generate inaccurate clinical images. Third, unlike images in the general domain, there is subtle variability across medical images. In addition, a significant amount of the variance is localized in small regions.

Thus, there is *an unmet need* for deep learning models to capture the subtle differences across diseases by attending to discriminative features present in these localized regions.

Recently, contrastive learning frameworks which heavily rely on data augmentation techniques [1]–[3] have become promising due to their ability to capture fine-grained discriminative features in the latent space. This paper proposes a novel and simple data augmentation method based on patient metadata and supervised knowledge such as disease labels to create clinically accurate positive samples for chest X-rays. The supervised classification loss helps SCALP create decision boundaries across different diseases while the patient-based contrastive loss helps SCALP learn discriminative features across different patients. Compared to other baselines [4], [5], SCALP uses simpler ResNet-50 architecture and performs significantly better. SCALP uses GradCAM++ [6] to generate activation maps that indicate the spatial location of the cardiopulmonary diseases. The highlights of our contribution are:

- Our augmentation **technique** for contrastive learning utilizes both patient metadata and supervised disease labels to generate clinically accurate positive and negative keys. *Positive keys* are generated by taking two chest radiographs of the same patient  $P$  while *negative keys* are generated using radiographs from patients other than  $P$  and having the same disease as  $P$ .
- A novel unified **framework** to simultaneously improve cardiopulmonary diseases classification and localization. We go beyond the conventional two-staged training (pre-training and fine-tuning) involved in contrastive learning. We demonstrate that single-staged end-to-end supervised contrastive learning can improve existing baselines significantly.
- We propose an innovative rectangular *Bounding Box* generation **algorithm** using pixel-thresholding and dynamic programming.

## II. RELATED WORK

### A. Medical Image Diagnosis

In the past decade, machine learning and deep learning have played a vital role in analyzing medical data, primarily medical imaging data. Recognition of anomalies and their localization has been a prevalent task for image analysis. Recent surveys

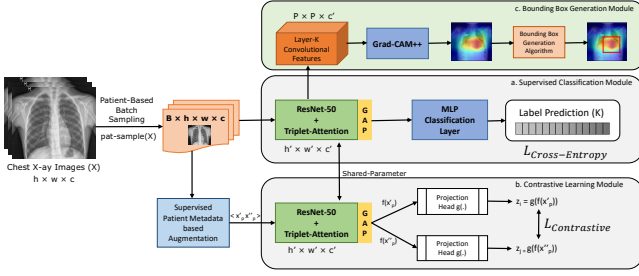


Fig. 1. Model Overview. The input images are sampled in batches with a constraint that no two images in a batch are from the same patient. Learning is performed using a shared encoder (Resnet-50) with triplet attention and joint loss from the supervised classification and contrastive learning module.

[7], [8] have illustrated the success of CNNs for classification and localization of several diseases in numerous medical imaging datasets varying from chest X-rays, MRIs, and CT scans. With the availability of large public chest X-rays datasets such as [9], [10], many researchers have explored the task of thoracic disease classification [4], [11]–[13]. CheXNet [4] uses 121-layer CNN trained on ChestXray14 [9] for pneumonia detection. [12] and [14] tried to improve the localization results using manually annotated localization data.

### B. Contrastive Learning for Chest X-rays

In the medical domain, prior work has found the performance improvement on applying contrastive learning to the chest X-rays [15], [16]. [15] presented an adaptation of MoCo [2] for chest X-ray dataset by pre-training it on [17]. The closest recent work to our knowledge [18] proposed to use patient metadata did not utilize supervised label information associated with chest X-rays. Recently, [19] has proposed supervised contrastive learning which extends the self-supervised batch contrastive approach to a fully-supervised setting, allowing us to effectively leverage label information. Inspired by [19] supervised contrastive loss, we propose to generate data augmentation by exploiting patient data and class labels together.

## III. THE PROPOSED APPROACH

Given chest X-rays with the cardiopulmonary disease labels, we aim to design a unified model that simultaneously classifies and localizes cardiopulmonary disease. We formulate both tasks in the same prediction framework and train them using a joint contrastive and supervised binary cross-entropy loss. More specifically, each image in our training data is labeled with an 8-dim vector  $y = [y_1, \dots, y_k, \dots, y_K]$ ,  $y_k \in \{0, 1\}$ ,  $K = 8$  for each image, and  $y_k$  indicates the presence or absence of 8 cardiopulmonary diseases. Our model produces a probability distribution of over 8 diseases for each image in the test set along with a heatmap with the localization information. The heatmap is passed to the BB-generation algorithm to generate rectangular bounding boxes indicating the presence of pathology.

### A. Image Model

As shown in Figure 1, we use the residual neural network (ResNet-50) architecture considering its manageable training with limited GPU resources and popularity in numerous image classification and object detection challenges. Inspired by Triplet attention [20], we incorporated the lightweight attention module in our ResNet-50 architecture to use cross-dimension interaction between the spatial dimensions and channel dimension for better localization. Image input with shape  $h \times w \times c$  produces a feature tensor with shape  $h' \times w' \times c'$ , where  $h, w$ , and  $c$  are the height, width, and the number of channels of the input image, respectively while  $h' = h/32, w' = w/32$ , and  $c' = 2048$ . Our framework is composed of two parallel modules: the supervised classification module and the contrastive learning module. Both modules share the same ResNet-50 encoder, i.e., the same set of parameters for encoding the chest X-ray image inputs.

1) *Supervised Classification Module*:: This module is responsible for learning the high-dimensional decision boundaries across different cardiopulmonary diseases. The encoded input chest X-ray images pass through a global average pooling layer to generate a 2048-dimensional feature vector. The feature vector is fed to a non-linear MLP layer to generate the probability distribution over 8 cardiopulmonary diseases. We calculate  $L_{Cross-Entropy}$  by summing loss from each class.

2) *Contrastive Learning Module*:: In addition to the inter-class variance of abnormalities in chest X-rays (i.e., feature differences between different diseases which are captured by classification loss), chest X-rays also have a high intra-class variance (i.e., differences in the X-rays of different patients having the same disease). To capture these intra-class variances, we introduce a supervised contrastive learning module to learn discriminative intra-class features. Our Supervised Patient Metadata based Augmentation module (Section III-C) generates two augmented views  $\langle x'_p, x''_p \rangle$  for each image in the batch. After being encoded by a shared encoder  $f(\cdot)$ , both views are fed to the global average pooling layer to generate feature embedding  $f(x'_p)$  and  $f(x''_p)$ . These feature embeddings are then transformed through the non-linear projection head similar to [1] to generate  $g(f(x'_p))$  and  $g(f(x''_p))$ .  $L_{Contrastive}$  loss is calculated by maximizing the agreement between  $g(f(x'_p))$  and  $g(f(x''_p))$  [3].

### B. Triplet Attention

To augment the quality of the localization by exploiting attention from the cross-dimension interaction in feature tensors, we integrate the Triplet Attention [20] into our architecture. Cross-Dimension Interaction involves computing attention weights for each dimension in tensor against every other dimension to capture the spatial and channel attention. In simple terms, spatial attention tells *where the channel to focus on*, while the channel attention tells *what channel to focus on*. With a minimal overload of few learnable parameters, the triplet attention mechanism successfully captures the interaction between the spatial and channel dimension of the

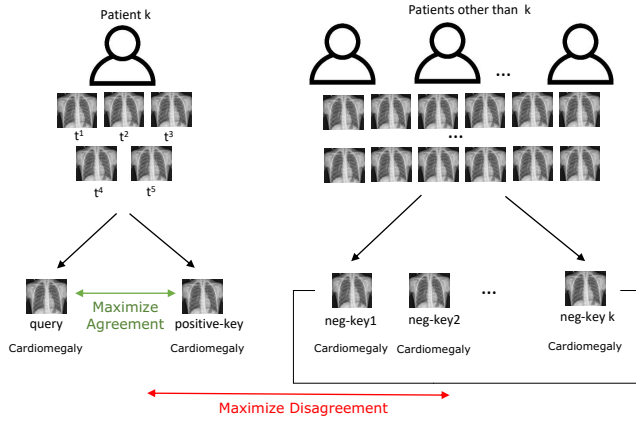


Fig. 2. Overview of our Supervised Patient-Metadata Based Contrastive Augmentation. Positive and Negative Keys are constructed using Patient Metadata (association of patient id with chest X-rays) and Supervised disease labels of chest X-rays.

input tensor. Following [20], the input tensor with dimension  $H \times W \times C$  in SCALP uses a branching mechanism to capture dependencies between  $(C, H)$ ,  $(C, W)$ , and  $(H, W)$ .

### C. Supervised Patient-Metadata Based Augmentation

Recently proposed Supervised Contrastive Learning [19] provides an innovative way to leverage label information for generating positive and negative keys effectively. In our case, we have acutely used patient metadata and label information simultaneously to generate positive and negative keys for the SCALP contrastive learning module. The motivation behind introducing contrastive learning in SCALP is to capture intra-class discriminative features between patients diagnosed with the same disease. Figure 2 provides an overview of our novel technique.

1) *Positive Sampling*: Data augmentation in medical imaging is sensitive to augmentation techniques such as random crop, Gaussian blur, and color jitter proposed in prior self-supervised learning [1]–[3]. For medical images, such operations may either change the disease label or are not meaningful for grayscale X-ray images. Since our goal is to incorporate discriminative features between two different patients to SCALP, we randomly select two chest X-ray studies having the same disease label of a patient  $P$  using patient metadata. The first study is called *query*, and the second is called a *positive key*. We use the contrastive loss to maximize agreement between them in latent space.

2) *Negative Sampling*: As shown in Figure 2, we select  $k$  negative keys for each patient  $P$  in our input batch. *Negative keys* are selected randomly from the pool of chest X-ray studies from patients except  $P$ , having the same diagnosis as *query*. This helps SCALP to distinguish the subtle differences between patients who are diagnosed with the same disease. Contrastive loss tries to push these *negative keys* away from *query* during training.

### D. Loss Function

SCALP is trained using the linear combination of supervised classification and contrastive loss. For the supervised classification, we use binary cross-entropy loss. For contrastive learning, we use the extended version of NT-Xent loss [1].

1) *Supervised classification Loss*: Our disease classification is a multi-class classification problem where we have 8 disease types. Multiple diseases can often be identified in one chest X-ray image and diseases are not mutually exclusive. We, therefore, define 8 binary classifiers for each class/disease type. Since all images in our dataset have 8 labels, the loss function for class  $k$  can be expressed as minimizing the binary cross-entropy as:

$$L_k = -y_k \cdot \log p(k|I) - (1 - y_k) \cdot \log(1 - p(k|I)) \quad (1)$$

where  $y_k$  is the ground-truth label of the  $k$ -th class, and  $I$  is the input image. To enable end-to-end training across all the classes, we sum up the class-wise loss to calculate total supervised loss as:

$$L_{Cross-Entropy} = \sum_{n=1}^{batch-size} \sum_{k=1}^8 L_k^n \quad (2)$$

2) *Contrastive Loss*: Our contrastive loss extends the normalized temperature scaled cross-entropy loss (NT-Xent). Using our patient-based sampling, we randomly select a batch of  $N$  chest X-ray images belonging to  $N$  patients. We derive the contrastive loss on the pairs of augmented examples from the batch. Let  $I'_{PD}$  be an image in batch belonging to patient  $P$  with disease  $D$ ,  $sim(x, y)$  denotes similarity between  $x$  and  $y$ , and  $f(\cdot)$  denotes ResNet-50 encoder,  $g(\cdot)$  denotes the projection head. The loss function  $l(I'_{PD})$  for a positive pair of example  $\langle I'_{PD}, I''_{PD} \rangle$  is defined as:

$$l(I'_{PD}) = -\log \frac{e^{(sim(g(f(I'_{PD})), g(f(I''_{PD}))))/\tau)}}{\sum_k \mathbb{I}_{[p \neq P, d=D]} e^{(sim(g(f(I'_{PD})), g(f(I''_{pd}))))/\tau)}} \quad (3)$$

where  $\mathbb{I}_{[p \neq P, d=D]} \in \{0, 1\}$  is an indicator function evaluating to 1 iff  $p \neq P$  and  $d = D$ ,  $\tau$  is the temperature parameter. The final contrastive loss is calculated as the sum over all instances in batch:

$$L_{Contrastive} = \sum_{k=1}^{batch-size} l(I_{PD}^k) \quad (4)$$

Eventually, we treat SCALP learning as the optimization of both contrastive and supervised cross-entropy loss together. Total loss for SCALP is defined as:

$$L_{Total} = \lambda \times L_{Cross-Entropy} + (1 - \lambda) \times L_{Contrastive} \quad (5)$$

### E. Bounding Box Generation Algorithm

We propose an innovative and time-efficient approach  $\mathcal{O}(n^2)$  to generate regular-shaped rectangular *bounding boxes* on chest X-rays indicating the approximate spatial location of the predicted cardiopulmonary disease. As shown in Figure 3, we feed the  $k$ -th layer of our image encoder (ResNet-50) to

Method	Atelectasis	Cardiomegaly	Effusion	Infiltration	Mass	Nodule	Pneumonia	Pneumothorax	Mean
Wang et. al. [9]	0.72	0.81	0.78	0.61	0.71	0.67	0.63	0.81	0.718
Wang et. al. [21]	0.73	0.84	0.79	0.67	0.73	0.69	0.72	0.85	0.753
Yao et. al. [22]	0.77	0.90	0.86	0.70	0.79	0.72	0.71	0.84	0.786
Raj. et. al. [4]	0.82	0.91	0.88	0.72	0.86	0.78	0.76	0.89	0.828
Kum. et. al. [23]	0.76	0.91	0.86	0.69	0.75	0.67	0.72	0.86	0.778
Liu et. al. [13]	0.79	0.87	0.88	0.69	0.81	0.73	0.75	0.89	0.801
Seyed et. al. [11]	0.81	0.92	0.87	0.72	0.83	0.78	0.76	0.88	0.821
Our model (std)	0.79 $\pm 0.01$	0.92 $\pm 0.00$	0.79 $\pm 0.01$	0.89 $\pm 0.01$	0.88 $\pm 0.02$	0.87 $\pm 0.00$	0.77 $\pm 0.01$	0.81 $\pm 0.02$	0.839

TABLE I

COMPARISON WITH THE BASELINE MODELS FOR AUC OF EACH CLASS AND AVERAGE AUC. FOR EACH COLUMN, RED VALUES DENOTE THE BEST RESULTS. NOTE THAT THE BEST BASELINE WITH MEAN AUC 0.828 USES THE DENSENET-121 ARCHITECTURE, WHILE OUR MODEL IS TRAINED USING A COMPARATIVELY SIMPLE AND LIGHT-WEIGHT RESNET-50 ARCHITECTURE.

Gradient-weighted Class Activation Mapping (GradCAM++) [6] to extract the attention maps/heatmaps. Due to the simplicity of intensity distributions in these heatmaps, we first scale heatmaps to the range [0, 255] and apply an ad-hoc threshold to convert these heatmaps into a binary matrix. Pixel values are converted to 1 if its intensity is greater than the threshold and 0 otherwise. Many previous works [9], [24] use only intensity threshold to generate bounding boxes which lead to many false positives. We use dynamic programming to generate a set of  $k$  candidate rectangles for bounding boxes and eliminate false positives by selecting the candidate which has the highest average intensity per pixel.

#### IV. EXPERIMENTS

##### A. Dataset and Preprocessing

NIH Chest X-ray dataset [9] consists of 112,120 chest X-rays collected from 30,805 patients, and each image is labeled with 8 cardiopulmonary disease labels. The NIH dataset also includes high-quality bounding box annotations for 880 images by radiologists. We separate these 880 images from our entire dataset, and they are used only to evaluate disease localization. Our method does not require any training data related to bounding boxes which is a significant difference compared to other existing baseline methods [12], [13] which use some percentage of these images for training. We follow the same protocol as [12], to shuffle our dataset (excluding images with BB annotations) into three subsets: 70% for training, 10% for validation, and 20% for testing. In order to prevent data leakage across patients, we make sure that there is no patient overlap between our subsets.

##### B. Implementation Details

We used the ResNet-50 model with triplet attention and initialized it with pre-trained weights provided by [20]. Our MLP layer is a two-layered fully-connected network with RELU non-linearity, while the projection head is defined similarly to [1]. Both our projection head and MLP layer are randomly initialized. We have used 0.01 learning rate and weight decay of  $10^{-6}$  and  $10^{-4}$  for contrastive and classification loss, respectively. SCALP uses an SGD optimizer and learning rate scheduler with step size and gamma value of 10 and 0.1, respectively.

##### Algorithm 1: Bounding Box Generation Algorithm

- 1 **Input:**  $k$ -th layer attention map/heatmap from ResNet-50
- 2 **Output:** coordinates (x1, y1, x2, y2) of the bounding box
- 3 Scale heatmap intensities to [0, 255] and create a mask matrix with the same dimension as the heatmap.
- 4 **if**  $pixel > 180$  **then**
- 5    $mask[pixel] = 1$
- 6 **else**
- 7    $mask[pixel] = 0$
- 8 Using dynamic programming [25], generate  $k$  maximum area rectangles as candidate BB.
- 9 Expand candidate rectangles uniformly across the edge till newly added *ratio (0s count, 1s count)*  $> 1$
- 10 Select the rectangle with the maximum average pixel intensity mapped in the heatmap and return its coordinates.

##### C. Disease Identification

SCALP classification is a multi-label classification problem. It assigns one or more labels among 8 cardiopulmonary diseases. We conduct a 3-fold cross-validation (Table I). We compare SCALP with reference models, which have published state-of-the-art performance of disease classification on the NIH dataset. We have used Area under the Receiver Operating Characteristics (AUC) to estimate the performance of our model in Table I. Our results also present the 3-fold cross-validation to show the robustness of our model. Compared to other baselines, SCALP achieves a mean AUROC score of 0.839 using ResNet-50 across the 8 different classes, which is 0.011 higher than the SOTA (uses DenseNet-121) on disease classification.

To understand the importance of contrastive module for disease classification, we trained SCALP with and without the contrastive loss and evaluated performance on the test set of NIH data. Table II presents the significant drop of  $> 9\%$  AUC when we exclude the contrastive module from the SCALP pipeline. This demonstrates the importance of our innovative association of contrastive modules with the



T(IoU)	Model	Atelectasis	Cardiomegaly	Effusion	Infiltration	Mass	Nodule	Pneumonia	Pneumothorax	Mean
0.1	Baseline [9]	0.69	0.94	0.66	0.71	0.40	0.14	0.63	0.38	0.569
	Our Model	0.62	0.97	0.64	0.81	0.51	0.12	0.8	0.29	<b>0.595</b>
0.2	Baseline [9]	0.47	0.68	0.45	0.48	0.26	0.05	0.35	0.23	0.371
	Our Model	0.42	0.92	0.42	0.6	0.25	0.04	0.56	0.18	<b>0.434</b>
0.3	Baseline [9]	0.24	0.46	0.30	0.28	0.15	0.04	0.17	0.13	0.221
	Our Model	0.29	0.78	0.23	0.37	0.13	0.01	0.4	0.05	<b>0.283</b>
0.4	Baseline [9]	0.09	0.28	0.20	0.12	0.07	0.01	0.08	0.07	0.115
	Our Model	0.18	0.55	0.12	0.19	0.09	0.01	0.25	0.02	<b>0.176</b>
0.5	Baseline [9]	0.05	0.18	0.11	0.07	0.01	0.01	0.03	0.03	0.061
	Our Model	0.07	0.33	0.04	0.10	0.04	0.0	0.14	0.10	<b>0.102</b>
0.6	Baseline [9]	0.02	0.08	0.05	0.02	0.00	0.01	0.02	0.03	0.029
	Our Model	0.02	0.14	0.02	0.04	0.03	0.00	0.07	0.00	<b>0.040</b>
0.7	Baseline [9]	0.01	0.03	0.02	0.00	0.00	0.00	0.01	0.02	0.011
	Our Model	0.01	0.04	0.01	0.03	0.01	0.00	0.02	0.00	<b>0.015</b>

TABLE II

DISEASE LOCALIZATION UNDER VARYING IOU ON THE NIH CHEST X-RAY DATASET. NOTE THAT SINCE OUR MODEL DOESN'T USE ANY GROUND TRUTH BOUNDING BOX INFORMATION, TO FAIRLY EVALUATE THE PERFORMANCE OF OUR MODEL, WE ONLY CONSIDER THE PREVIOUS METHODS' RESULTS WITH THE SAME SETTINGS AS SCALP.

	SCALP w/o Contrastive	SCALP
Atelectasis	0.751	0.79
Cardiomegaly	0.850	0.92
Effusion	0.833	0.79
Infiltration	0.670	0.89
Mass	0.694	0.88
Nodule	0.640	0.87
Pneumonia	0.700	0.77
Pneumothorax	0.792	0.81
<b>Mean</b>	0.7413	0.839

TABLE III

AUC COMPARISON OF SCALP WITH AND WITHOUT CONTRASTIVE LEARNING MODULE.

$\lambda$	0.99	0.90	0.85	0.80	0.75	0.70
<b>AUC</b>	0.766	0.794	0.822	<b>0.839</b>	0.818	0.785

TABLE IV

AUC COMPARISON OF SCALP FOR VARYING  $\lambda$  IN EQUATION 7. A HIGHER VALUE OF  $\lambda$  IMPLIES LOWER WEIGHT TO THE CONTRASTIVE LOSS. SCALP ACHIEVES THE BEST PERFORMANCE WHEN 80% WEIGHT IS GIVEN TO CLASSIFICATION LOSS AND 20% WEIGHT IS GIVEN TO CONTRASTIVE LOSS.

classification pipeline. Our experiments in Table IV prove our hypothesis that both contrastive and cross-entropy loss is important. In a calculated ratio, they help SCALP learn both disease-level and patient-level discriminative visual features.

#### D. Disease Localization

The NIH dataset has 880 images labeled by radiologists with the bounding box information. We have used this dataset to evaluate the performance of SCALP for disease localization. Many prior works [12], [13] have used a fraction of ground truth (GT) bounding boxes for training and evaluated their system on the remaining. To ensure a robust evaluation, we do not use any GT for training, and Table II presents our evaluation results on all 880 images. For localization, we evaluated our detected regular rectangular regions against the

annotated ground truth (GT) bounding boxes, using intersection over union ratio (IoU). The localization is defined as correct only if  $\text{IoU} > \text{T(IoU)}$ . We evaluate SCALP for different thresholds ranging from  $\{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7\}$  as shown in Table II. A higher IoU threshold is preferred for disease localization because clinical usage requires high accuracy. Note that SCALP mean performance for 8 diseases is significantly better than the baseline under all IoU thresholds. When the IoU is set to 0.1, SCALP outperforms the baseline in terms of Cardiomegaly, Infiltration, Mass, and Pneumonia.

Note that our innovative bounding box generation algorithm successfully eliminates dispersed attention, and identifies regions where maximum attention is concentrated. For example, in "Effusion", the generated heatmap has dispersed attention on both sides of the lungs. However, attention intensity is concentrated in the left side of the lung. Our algorithm is able to generate a bounding box on the left side of the lung and have a high overlap with the ground-truth. Similarly, for "Infiltration" and "Nodule", many undesirable patches of attention have been eliminated which is helpful in improving the IoU evaluation of SCALP. The attention maps generated by SCALP are sharp and focused compared to our reference baseline [9]. Overall, our results show that the predicted disease localizations have significant alignment with the ground truth and can serve as interpretable cues for the disease classification.

#### V. CONCLUSION

In this work, we propose a simple and effective end-to-end framework SCALP using supervised contrastive learning to identify cardiopulmonary diseases in chest X-ray. We go beyond two-stage training (pre-training and fine-tuning), and demonstrate that an end-to-end supervised contrastive training using two images from the same patient as a positive pair, can significantly outperform SOTA on disease classification. SCALP can jointly model disease identification and

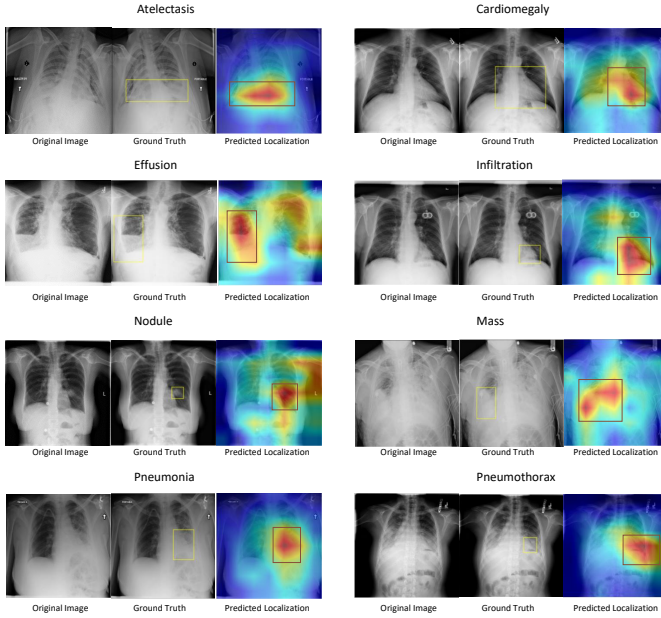


Fig. 3. Examples of visualization of localization on the test images. We plot the results of diseases near the thoracic. The attention maps are generated from the fourth layer of SCALP’s encoder and overlapped with its corresponding original radiology image. The ground-truth and the predicted bounding boxes are shown in yellow and red color respectively.




Fig. 4. Effect of Batch Size on SCALP performance for Disease Classification. Prior works [1], [2] have verified that contrastive learning benefits from a larger batch size. SCALP shows a similar trend with increasing batch size.

localization using the linear combination of contrastive and classification loss. We also propose a time-efficient Bounding Box generation algorithm that generates bounding boxes from the attention map of SCALP. Our extensive qualitative and quantitative results demonstrate the effectiveness of SCALP and its state-of-the-art performance.

## VI. ACKNOWLEDGMENT

This work is supported by Amazon Machine Learning Research Award 2020. It also was supported by the National Library of Medicine under Award No. 4R00LM013001.

## REFERENCES

[1] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.

[2] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, “Momentum contrast for unsupervised visual representation learning,” in *CVPR*, 2020, pp. 9729–9738.

[3] X. Chen, H. Fan, R. Girshick, and K. He, “Improved baselines with momentum contrastive learning,” *arXiv preprint arXiv:2003.04297*, 2020.

[4] P. Rajpurkar, J. Irvin, K. Zhu, B. Yang, H. Mehta, T. Duan, D. Ding, A. Bagul, C. Langlotz, K. Shpanskaya, M. P. Lungren, and A. Y. Ng, “Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning,” 2017.

[5] C. Ma, H. Wang, and S. C. H. Hoi, “Multi-label thoracic disease image classification with cross-attention networks,” 2020.

[6] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra, “Grad-cam: Why did you say that?” *arXiv preprint arXiv:1611.07450*, 2016.

[7] S. M. Anwar, M. Majid, A. Qayyum, M. Awais, M. Alnowami, and M. K. Khan, “Medical image analysis using convolutional neural networks: A review,” *Journal of Medical Systems*, vol. 42, no. 11, 2018.

[8] M. Rezaei, H. Yang, and C. Meinel, “Deep learning for medical image analysis,” *ArXiv*, vol. abs/1708.08987, 2017.

[9] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, “Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases,” in *CVPR*, 2017, pp. 2097–2106.

[10] A. E. W. Johnson, T. J. Pollard, N. R. Greenbaum, M. P. Lungren, C. ying Deng, Y. Peng, Z. Lu, R. G. Mark, S. J. Berkowitz, and S. Horng, “Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs,” 2019.

[11] L. Seyyed-Kalantari, G. Liu, M. McDermott, and M. Ghassemi, “Chexclusion: Fairness gaps in deep chest x-ray classifiers,” *arXiv preprint arXiv:2003.00827*, 2020.

[12] Z. Li, C. Wang, M. Han, Y. Xue, W. Wei, L.-J. Li, and L. Fei-Fei, “Thoracic disease identification and localization with limited supervision,” in *CVPR*, 2018, pp. 8290–8299.

[13] J. Liu, G. Zhao, Y. Fei, M. Zhang, Y. Wang, and Y. Yu, “Align, attend and locate: Chest x-ray diagnosis via contrast induced attention network with limited supervision,” in *ICCV*, October 2019.

[14] K. Chaitanya, E. Erdil, N. Karani, and E. Konukoglu, “Contrastive learning of global and local features for medical image segmentation with limited annotations,” *arXiv preprint arXiv:2006.10511*, 2020.

[15] H. Sowrirajan, J. Yang, A. Y. Ng, and P. Rajpurkar, “Moco-cxr: Moco pretraining improves representation and transferability of chest x-ray models,” 2021.

[16] A. Sriram, M. Muckley, K. Sinha, F. Shamout, J. Pineau, K. J. Geras, L. Azour, Y. Aphinyanaphongs, N. Yakubova, and W. Moore, “Covid-19 prognosis via self-supervised representation learning and multi-image prediction,” 2021.

[17] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, D. B. Larson, C. P. Langlotz, B. N. Patel, M. P. Lungren, and A. Y. Ng, “Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison,” 2019.

[18] Y. N. T. Vu, R. Wang, N. Balachandar, C. Liu, A. Y. Ng, and P. Rajpurkar, “Medaug: Contrastive learning leveraging patient metadata improves representations for chest x-ray interpretation,” 2021.

[19] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, “Supervised contrastive learning,” 2021.

[20] D. Misra, T. Nalamada, A. U. Arasanipalai, and Q. Hou, “Rotate to attend: Convolutional triplet attention module,” 2020.

[21] X. Wang, Y. Peng, L. Lu, Z. Lu, and R. M. Summers, “Tienet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays,” in *CVPR*, 2018, pp. 9049–9058.

[22] L. Yao, E. Poblentz, D. Dagunts, B. Covington, D. Bernard, and K. Lyman, “Learning to diagnose from scratch by exploiting dependencies among labels,” *arXiv preprint arXiv:1710.10501*, 2017.

[23] P. Kumar, M. Grewal, and M. M. Srivastava, “Boosted cascaded convnets for multilabel classification of thoracic diseases in chest radiographs,” in *International Conference Image Analysis and Recognition*. Springer, 2018, pp. 546–552.

[24] Y. Han, C. Chen, L. Tang, M. Lin, A. Jaiswal, Y. Ding, and Y. Peng, “Using radiomics as prior knowledge for abnormality classification and localization in chest x-rays,” 11 2020.

[25] xxx, “Maximum size Rectangle in Binary Matrix,” <https://www.geeksforgeeks.org/maximum-size-rectangle-binary-sub-matrix-1s/>.