



OPEN ACCESS

EDITED BY

Henning Müller,
University of Applied Sciences Western
Switzerland (HES-SO Valais), Switzerland

REVIEWED BY

Yashin Dicente Cid,
Roche Diagnostics S.L., Spain,
Alexandros Karargyris,
Hôpitaux Universitaires de Strasbourg, France

*CORRESPONDENCE

Ricardo Bigolin Lanfredi
✉ ricbl@sci.utah.edu

RECEIVED 03 November 2022

ACCEPTED 05 June 2023

PUBLISHED 22 June 2023

CITATION

Bigolin Lanfredi R, Schroeder JD and Tasdizen T
(2023) Localization supervision of chest x-ray
classifiers using label-specific eye-tracking
annotation.

Front. Radiol. 3:1088068.

doi: 10.3389/fradi.2023.1088068

COPYRIGHT

© 2023 Bigolin Lanfredi, Schroeder and
Tasdizen. This is an open-access article
distributed under the terms of the [Creative
Commons Attribution License \(CC BY\)](#). The use,
distribution or reproduction in other forums is
permitted, provided the original author(s) and
the copyright owner(s) are credited and that the
original publication in this journal is cited, in
accordance with accepted academic practice.
No use, distribution or reproduction is
permitted which does not comply with these
terms.

Localization supervision of chest x-ray classifiers using label-specific eye-tracking annotation

Ricardo Bigolin Lanfredi^{1*}, Joyce D. Schroeder²
and Tolga Tasdizen¹

¹Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT, United States,

²Department of Radiology and Imaging Sciences, University of Utah, Salt Lake City, UT, United States

Convolutional neural networks (CNNs) have been successfully applied to chest x-ray (CXR) images. Moreover, annotated bounding boxes have been shown to improve the interpretability of a CNN in terms of localizing abnormalities. However, only a few relatively small CXR datasets containing bounding boxes are available, and collecting them is very costly. Opportunely, eye-tracking (ET) data can be collected during the clinical workflow of a radiologist. We use ET data recorded from radiologists while dictating CXR reports to train CNNs. We extract snippets from the ET data by associating them with the dictation of keywords and use them to supervise the localization of specific abnormalities. We show that this method can improve a model's interpretability without impacting its image-level classification.

KEYWORDS

eye tracking, chest x-ray (CXR), interpretability, annotation, localization, gaze

1. Introduction

Along with the success of deep learning methods in medical image analysis, interpretability methods have been used to validate that models are working as expected (1). The interpretability of deep learning models has also been listed as a critical research priority for artificial intelligence in medical imaging (2). The employment of bounding box annotations during training has been shown to improve a model's ability to highlight abnormalities and, consequently, their interpretability (3). However, bounding boxes for medical images are costly to acquire since they require expert annotation, whereas image-level labels can readily be extracted from radiology reports. This fact is exemplified by the relatively small size of bounding box datasets for chest x-rays (CXRs) (4, 5) when compared to the size of CXR datasets with image-level labels. Eye-tracking (ET) data, on the other hand, contain implicit information about the location of labels, and its collection may be easier to scale up than bounding boxes if the acquisition of gaze from radiologists is implemented in clinical practice.

CXRs are the most common medical imaging exam in the United States (6). This type of imaging has also had much attention from deep learning practitioners, with successful technical results (1, 7). Despite their universality, reading a CXR is considered one of the hardest interpretations performed by radiologists (8), with high inter-rater variability in reported abnormalities (9, 10). Moreover, abnormalities in CXRs can appear in a vast diversity of locations, including the lungs, mediastinum, pleural space, vessels, airways, and ribs (11). They can also be described as hundreds of different findings (11). The diverse aspect of the report has been simplified for use in deep learning applications,

where, in several cases, a simplified subset of the most common labels has been automatically extracted from reports for use in a multi-label formulation (1, 12, 13). Since radiologists pay close attention to several areas when dictating a CXR report, scanning almost the whole image for signs of several abnormalities, the ET data accumulated during the full report dictation might highlight several areas with no evidence of abnormality. Therefore, the use of the temporal aspect of the report, by processing the ET data with the dictation-transcription timestamps, may achieve a more precise localization for specific abnormalities.

We propose to use ET data with timestamped dictations of radiology reports to identify when the presence of specific abnormalities was dictated, identify the times when radiologists would have visually attended to such abnormality, and extract the associated gaze locations. The extracted information can be used as label-specific annotation for supervising models to highlight abnormalities spatially. The localization supervision is performed using a combination of a multiple instance learning loss over the last spatial layer of a convolutional neural network (CNN) (3) and a multi-task learning loss (14), adding an output representing label-specific ET maps. To complement the annotations of an ET dataset, we employ a large dataset of CXRs with image-level labels in a weak supervision formulation. We evaluate the classification performance and the ability to localize abnormalities of a model trained with data annotated by ET. This model is compared against baselines using no annotated data and hand-annotated data. We show that using ET data during training improves the localization performance of generated interpretable heatmaps without compromising area under the receiver operating characteristic curve (AUC) classification scores and that this type of data might have value in replacing hand-labeling, depending on the costs and benefits of each type of data collection. In addition, to the best of our knowledge, this study offers the first estimation of how the value of ET data compares to the value of hand-annotated localization data when a very large dataset with image-level labels is available for weak supervision. The main contributions of this paper are:

- developing knowledge of the complexities of the use of eye-tracking data for annotation in radiology, including the

proposed method of considering the lag between a radiologist's gaze and dictation for accumulating ET data specific for each abnormality label; and

- informing about expected relative value between ET data and manual annotations for the decision of starting future more comprehensive eye-tracking studies.

2. Materials and methods

2.1. Extracting localization information from eye-tracking data

We designed a pipeline to extract disease locations from ET data. This pipeline has two main parts: extracting label mentions in reports and generating an ET heatmap for a given detected label. A representation of the pipeline is shown in Figure 1. The pipeline requires the ET dataset to contain timestamps, transcriptions of report dictations, and fixations, i.e., locations in the image where radiologists stabilized their gaze for some time.

2.1.1. Label mention extraction

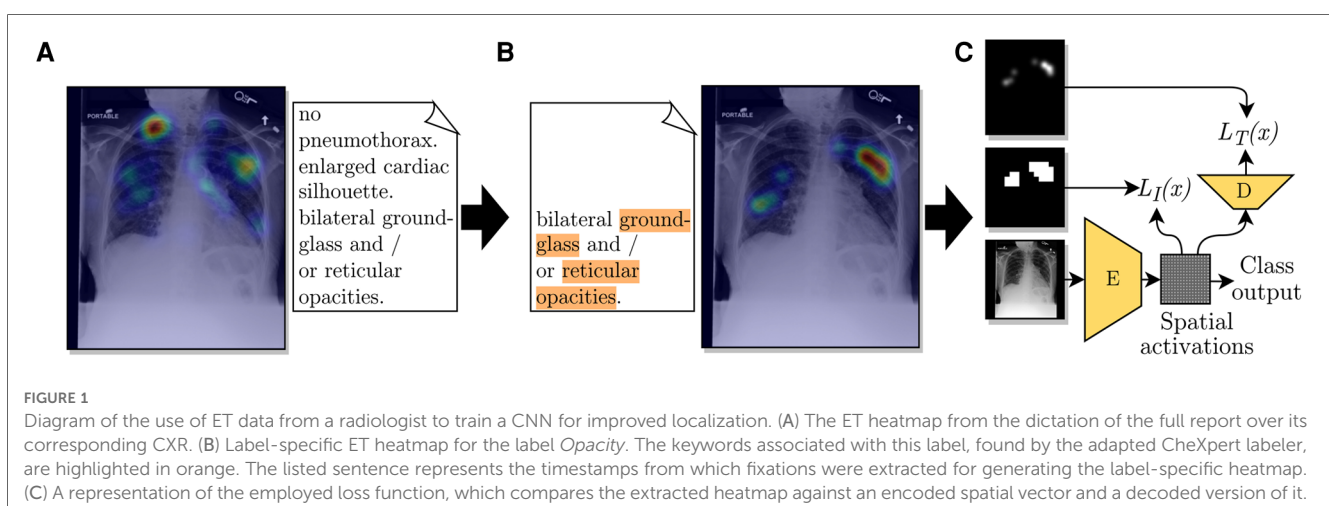
To extract labels from reports, we adopted a modified version of the CheXpert labeler (12), which uses a set of hand-crafted rules to detect label mentions and negation structures.

2.1.2. Fixation association with each mention

From observing the gaze of radiologists on a few examples, we noticed patterns that seemed to be in common for all radiologists:

- for the first moments after being shown a CXR, radiologists looked all over the image without dictating anything;
- when dictating, radiologists usually looked at regions corresponding to the content of the current sentence or the following sentence (when near the end of the dictation of the current sentence).

From these two observations, we decided to generate ET heatmaps for detected labels from the fixations of the sentences where the label was mentioned, the previous sentence, and the pause between sentences. We accumulated fixations within a limit of



1.5 s previous to the start of the mentioning sentence and up to the last mention in the mentioning sentence. An illustration of the method for choosing which fixations were included in the heatmaps is given in Figure 2. An example with a case from the ET data we used are given in Figures 1A,B. From applying this extraction method, abnormality mentions within the same report sentence are associated with the same localization heatmap.

2.1.3. Heatmap generation

Heatmaps were generated by placing Gaussians over each fixation location with a standard deviation of one degree of visual angle, following Le Meur et al. (15). Fixations had the amplitude of their Gaussians weighted by their duration. The heatmap for each detected mention of a label was normalized to have a maximum value of 1. Multiple mentions of a label for the same CXR were aggregated with a maximum function.

2.2. Multiple instance learning

We used the multiple instance learning loss term from Li et al. (3) to train an encoder with the extracted ET heatmaps. The encoder E , as represented in Figure 1C, was built to output a grid of cells, where each cell represented a multi-label classifier for label presence in the homologous region of the image. The image-level prediction $C_k(x)$ for label k was formulated as

$$C_k(x) = 1 - \prod_{j \in \Gamma_x} (1 - \sigma(\gamma_{jkx})), \tag{1}$$

where γ_{jkx} is the logit output for class k and grid cell j for image x , Γ_x is the set of all grid cells for image x , and $\sigma(\cdot)$ is the sigmoid function. Equation (1) is a soft version of the Boolean OR function, assigning a positive image-level label when at least one of the grid cells was found to contain that class.

During training, the loss function depended on the presence of a localization annotation. For images annotated with localization (A), grid cells were trained to match a resized version of the annotation, as shown in Figure 1C. We used the loss

$$L_{kA}(x) = -\log \left(\prod_{j \in B_{kx}} (\sigma(\gamma_{jkx})) \prod_{j \in \Gamma_x - B_{kx}} (1 - \sigma(\gamma_{jkx})) \right), \tag{2}$$

where B_{kx} is the set of grid cells labeled as containing evidence of disease k for image x and $L_{kA}(x)$ is the output of the loss function for annotated images of class k .

For images that did not contain localization annotations (U), the loss depended on the image-level label. For positive images, at least one grid cell should be positive. We used the loss

$$L_{kU^+}(x) = -\log(C_k(x)), \tag{3}$$

where $L_{kU^+}(x)$ is the output of the loss function for unannotated images labeled as positive for class k . For negative images, all grid cells should be negative. We used

$$L_{kU^-}(x) = -\log \left(\prod_{j \in \Gamma_x} (1 - \sigma(\gamma_{jkx})) \right) \tag{4}$$

where $L_{kU^-}(x)$ is the output of the loss function for unannotated images labeled as negative for class k . The multiple instance learning loss term $L_I(x)$ was then formulated as

$$L_I(x) = \mathbb{E}_{x \in X, k \in K} [\lambda_A \mathbb{1}_{kA}(x) L_{kA}(x) + \mathbb{1}_{kU^+}(x) L_{kU^+}(x) + \mathbb{1}_{kU^-}(x) L_{kU^-}(x)], \tag{5}$$

where λ_A is a hyperparameter controlling the relative importance of annotated images during training, X is the set of all images, annotated and unannotated, K is the set of all classes, and $\mathbb{1}_{kA}(x)$, $\mathbb{1}_{kU^+}(x)$, and $\mathbb{1}_{kU^-}(x)$ are the output of indicator functions that were 1 when the image x was annotated, positive for class k (unannotated) and negative for class k (unannotated), respectively.

2.2.1. Avoiding numerical underflow and balanced range normalization

To avoid numerical underflow and have a more uniform range for the output of models, Li et al. (3) suggested normalizing the factors of the products in Equations (1) to (5), i.e., $\sigma(\gamma_{jk})$ and $1 - \sigma(\gamma_{jk})$, to the range [0.98, 1]. After running tests, we achieved better results by balancing this normalization, changing the range of each product factor to $[0.0056738^{(1/n_t)}, 1]$, where n_t is the number of factors being multiplied. This range allows all

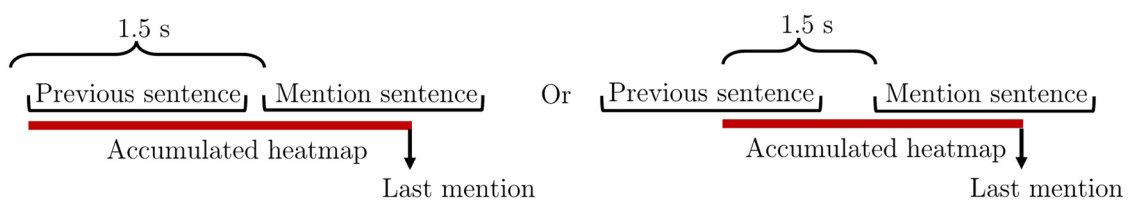


FIGURE 2 Multiple instance learning technique. Method of accumulation of fixations for generating heatmaps for each sentence that mentions at least one of the abnormality labels. The chosen fixations could be between the start of the previous sentence and the last mention of the current sentence or between 1.5 s before the start of the current sentence and the last mention of the current sentence, whichever has the shortest duration.

products to have a similar expected range and keeps the same [0.98,1] range when $n_t = 256$.

2.3. Multi-task learning

Inspired by a work by Karargyris et al. (14), we added another loss term to our method. With the intuition of giving more supervision to the representations calculated by encoder E and, consequently, improving its representations, we added the task of predicting a high-resolution ET map for each label, performed with the help of decoder D , as shown in Figure 1C. From testing the network’s performance, we modified the method in that class outputs were calculated according to Equation (1) instead of adding fully connected layers as suggested by Karargyris et al. (14). Another difference in our method was that our decoder output had one channel for each of the ten labels in our classification task. The output of decoder D could then provide estimations for the localization of abnormalities. In other words, the output of D is an interpretability output: an alternative to the spatial activations or other interpretability methods, such as GradCAM (16). Decoder D had an architecture with three blocks, each composed of a sequence of a bilinear upsampling layer, a convolution layer, and a batch normalization layer. The loss $L_T(x)$ added for this task is the pixel-level cross-entropy between the output of decoder D and the label-specific ET map in the same resolution, 256×256 , formulated as

$$L_T(x) = -\mathbb{E}_{x \in X, k \in K} [\mathbb{1}_{kA^+}(x) \mathbb{E}_{j \in \Gamma_x} [\mathbb{1}_{j \in B_{kx}} \log(\sigma(\gamma_{jkx})) + \mathbb{1}_{j \notin B_{kx}} \log(1 - \sigma(\gamma_{jkx}))]]], \tag{6}$$

where $\mathbb{1}_{kA^+}(x)$ is the output of an indicator function that was 1 when the image x was annotated and positive for class k . This loss was used to train both decoder D and encoder E and, as shown in Equation (6), was applied only for channels corresponding to positive ground-truth labels.

2.4. Multi-resolution architecture

As shown in Figure 3, for our encoder E we adapted the Resnet-50 (17) architecture by replacing its average pooling and

last linear layer with two convolutional layers separated by batch normalization (18) and ReLU activation (CNN Block 5 from Figure 3). To improve the results for labels with small findings in the original image, we modified the network such that spatial maps with 32×32 resolution were used as inputs to CNN Block 5.

2.5. Loss function

The final loss function $L(\cdot)$, to be minimized while training, is given by

$$L(x) = L_I(x) + \lambda_T L_T(x), \tag{7}$$

where λ_T is a hyperparameter controlling the relative importance of $L_T(x)$. The classification output of our model only influences the $L_I(x)$ term.

2.6. Datasets

We used two datasets in our study. The REFLACX dataset (19–21) provides ET data and reports from five radiologists for CXRs from the MIMIC-CXR-JPG dataset (21–23). Additionally, the REFLACX dataset contains image-level labels and radiologist-drawn abnormality ellipses, which can be used to validate the locations highlighted by our tested models. Except for the experiments described in Sections 2.7 and 2.8, we used examples from Phase 3 from the REFLACX dataset. The MIMIC-CXR-JPG dataset, which contains patients who visited the emergency department of the Beth Israel Deaconess Medical Center between 2011 and 2016, was also utilized for its unannotated CXRs and image-level labels. Images from the MIMIC-CXR-JPG dataset were filtered using the same criteria as the REFLACX dataset: only labeled frontal images from studies with a single frontal image were considered. The test sets for both datasets were kept the same. A few subjects from the training set of the REFLACX dataset were assigned to its validation set so that around 10% of the REFLACX dataset was part of the validation set. The same subjects were also assigned to the validation set of the MIMIC-CXR-JPG dataset. The train, validation, and test splits had, respectively, 1,724, 277, and 506 images for the annotated set

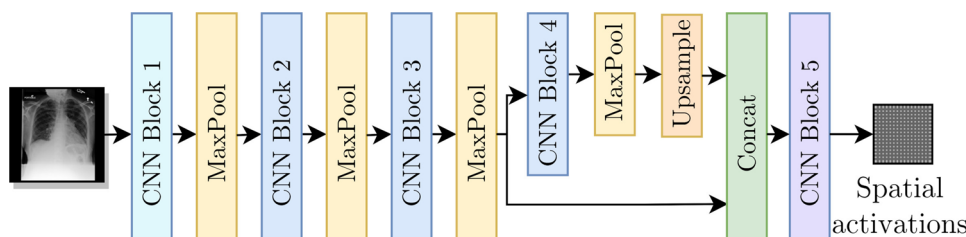


FIGURE 3 Encoder E as a modified Resnet-50 architecture to include the multi-resolution branches.

and 187,519, 4,275, and 2,701 images for the unannotated set. The use of both datasets did not require ethics approval because they are publicly available de-identified datasets. All ellipses and ET data we used originated from the REFLACX dataset and the split sizes of that dataset reflect the number of such annotations that we had available.

The sets of labels from the annotated and unannotated datasets were different. We decided to use the ten labels listed in **Table 1**. We provide, in **Table 2**, a list of the labels from each dataset that were considered equivalent to each of the ten labels used in this study and, in **Table 3**, the number of examples of each label present in the datasets.

2.7. Labeler

The set of labels from the REFLACX dataset is slightly different from the ones provided by the CheXpert labeler. With the help of a cardiothoracic subspecialty-trained radiologist, we modified the labeler to output a new set of labels. Modifications were also made to improve the identification of the already present labels after observing common mistakes on a separate validation set composed of 20% of Phase 1 and Phase 2 from the REFLACX dataset. We adjusted rules for negation finding and added/adapted expressions to match and unmatch labels.¹

2.8. Location extraction

We tested several methods for extracting label-specific localization of abnormalities from the eye-tracking data. All methods involved the accumulation of fixations into heatmaps, with different starting and ending accumulation times, after extracting the label's mention time from the dictation. For a first stage of validation, the starting times we considered were:

- MAX(Start of mention sentence - TIME, Start of the previous sentence),
- MAX(First mention in the sentence - TIME, Start of the previous sentence),
- MAX(End of mention sentence - TIME, Start of the previous sentence),
- start of first report sentence,
- start of the previous sentence,
- end of the previous sentence,
- start of mention sentence,
- start of the recording of data for that CXR,

where TIME is a time delay assuming the values of 2.5 s, 5.0 s, and 7.5 s. The end times we considered were:

TABLE 1 Per-label AUC metric on the test set for the two baselines and our method.

Label	<i>Unannotated</i>	<i>Ellipse</i>	<i>ET model (ours)</i>
AMC	0.658 [0.654, 0.662]	0.665 [0.660, 0.670]	0.660 [0.655, 0.665]
Atelectasis	0.749 [0.744, 0.754]	0.751 [0.748, 0.753]	0.748 [0.746, 0.751]
ECS	0.768 [0.764, 0.772]	0.768 [0.765, 0.771]	0.765 [0.760, 0.771]
Consolidation	0.704 [0.697, 0.711]	0.709 [0.703, 0.716]	0.710 [0.706, 0.713]
Edema	0.839 [0.837, 0.841]	0.835 [0.833, 0.837]	0.838 [0.836, 0.840]
Fracture	0.714 [0.690, 0.737]	0.710 [0.696, 0.724]	0.714 [0.701, 0.727]
Lung Lesion	0.760 [0.747, 0.774]	0.751 [0.740, 0.763]	0.745 [0.738, 0.752]
Opacity	0.784 [0.782, 0.786]	0.785 [0.783, 0.788]	0.783 [0.779, 0.787]
Pleural abnormality	0.868 [0.865, 0.870]	0.869 [0.867, 0.871]	0.868 [0.865, 0.871]
Pneumothorax	0.825 [0.811, 0.838]	0.810 [0.799, 0.820]	0.819 [0.810, 0.827]

TABLE 2 List of labels that were grouped to form the labels from the analysis presented in this paper.

Labels used in our models	Labels from public datasets (REFLACX and MIMIC-CXR)
Abnormal Mediastinal Contour (AMC)	abnormal mediastinal contour (AMC), enlarged cardiomeastinum
Atelectasis	atelectasis
Enlarged Cardiac Silhouette (ECS)	enlarged cardiac silhouette (ECS), cardiomegaly
Consolidation	consolidation
Edema	pulmonary edema, edema
Fracture	fracture, acute fracture
Lung Lesion	lung nodule or mass, lung lesion
Opacity	pulmonary edema, edema, lung nodule or mass, atelectasis, consolidation, groundglass opacity, interstitial lung disease, pneumonia, lung opacity
Pleural Abnormality	pleural abnormality, pleural other, pleural effusion
Pneumothorax	pneumothorax

- start of mention sentence,
- end of mention sentence,
- end of the first mention,
- end of the last mention.

We tested all combinations between starting times and end times with a duration of 0 s or more. We compared the extracted heatmaps with the validation hand-annotated ellipses using the IoU metric with a validated threshold. After this validation, we finetuned, as a second stage of validation, the time delay by testing more times (0.5 s, 0.75 s, 1 s, 1.25 s, 1.5 s, 1.75 s, 2 s, 2.5 s, 3 s, 3.5 s, 4 s, 4.5 s, 5 s).

2.9. Validation and evaluation

For our experiments,² we used PyTorch 1.10.2 (24). Hyperparameters commonly used for CXR classifiers were employed during training and were not tuned for any tested

¹The final set of rules can be found in our code repository at <https://github.com/ricbl/eye-tracking-localization>.

²The code for our experiments can be found at <https://github.com/ricbl/eye-tracking-localization>.

TABLE 3 Number of positive examples for each of the splits for both employed datasets: REFLACX (R) and MIMIC-CXR-JPG (M).

Label	Train R	Val R	Test R	Train M	Val M	Test M
Abnormal Mediastinal Contour	59	8	27	13,551	351	289
Atelectasis	503	69	189	46,981	1,196	754
Enlarged Cardiac Silhouette	340	55	173	41,681	1,138	821
Consolidation	543	95	196	12,311	352	253
Edema	222	39	124	33,383	971	841
Fracture	34	5	27	4,028	62	71
Lung Lesion	77	13	41	6,015	114	106
Opacity	865	131	337	97,640	2,560	1,861
Pleural Abnormality	486	69	205	50,914	1,426	1,036
Pneumothorax	48	6	13	9,653	258	96

method. Models were trained for 60 epochs with the AMSGrad Adam optimizer (25) using a learning rate equal to 0.001 and weight decay of 0.00001. A batch size of 20 images was chosen for the use of GPUs with 16GB of memory or more. Images were resized such that their longest dimension had 512 pixels, whereas the other dimension was padded with black pixels to reach a length of 512 pixels. For training, images were augmented with rotation up to 45 degrees, translation up to 15%, and scaling up to 15%. The grid supervised by loss $L_I(x)$ had 1,024 cells (32×32). We used the max-pooling operation to convert the ET heatmap annotations to the same dimension. We thresholded the ET heatmaps at 0.15. This number was chosen after visual analysis of their histograms of intensities. We used $\lambda_A = 3$ and $\lambda_T = 300$ after validation of AUC and IoU values for our proposed method considering the following values: 0.3, 1, 3, 10, 30, and 100 for λ_A and 0.1, 0.3, 1, 3, 10, 30, 100, 300, 1,000, 3,000 for λ_T . We trained models with five different seeds and report their average results and 95% confidence intervals. Experiments were run in internal servers containing Nvidia GPUs (TITAN RTX 24GB, RTX A6000 48GB, Tesla V100-SXM2 16 GB). Each training run took approximately two to three days in one GPU.

As baselines, we evaluated a model trained without the annotated data (*Unannotated*) and a model trained with data annotated by the drawn ground truth (GT) truth ellipses (*Ellipse*). The ellipses were represented by binary heatmaps and were processed in the same way as the ET heatmaps. The loss function, CNN architecture, and training hyperparameters were the same for all methods. We did not include a cross-

entropy classification loss baseline because it achieved lower scores than the presented methods. The best epoch for each method was chosen using the average AUC on the validation set. The best model heatmap threshold for each method and label was calculated using the average validation intersection over union (IoU) over the five seeds, considering the full range of thresholds.

We evaluated our model (*ET model*) and the two baselines by calculating the test AUC for image-level labels of the MIMIC-CXR-JPG dataset and the test IoU for localization of abnormalities, compared against the drawn GT ellipses. IoU was calculated individually per positive label for all images with a positive label. We calculated three heatmaps for each label and input image: the output of decoder D , the spatial activations, and the output of the GradCAM method (16). We tested which heatmap had the best IoU validation results for each of the three reported training methods. We report results for the *Ellipse* and *ET model* using the output of decoder D and for the *Unannotated* model using the spatial activations. The heatmaps for each method were upsampled to the resolution of the GT ellipses using nearest-neighbor interpolation.

3. Results

3.1. Labeler

Labeler quality estimations, after modifications, are shown in Table 4 and were calculated with the rest of the data from Phase 2, representing 80% of the cases. Results were variable depending on the label, and misdetections should be expected when using this labeler.

3.2. Location extraction

For the first validation, the highest-scoring accumulated heatmaps used a starting time of MAX(Start of mention sentence -2.5 s, Start of the previous sentence) and an ending at the end of the last mention present in the sentence. For the second stage of the location extraction validation, when we tested delay times in a higher resolution, the time with the best IoU was 1.5 s with an IoU of 0.233, justifying our approach as presented in Section 2.1.

TABLE 4 Results of the label detection with a modified version of the CheXpert labeler (12).

Label	Recall	Precision	Label	Recall	Precision
AMC	0.67	0.73	Interstitial Lung Disease	0.75	0.27
Acute Fracture	1.00	1.00	Lung Nodule or Mass	0.50	0.50
Atelectasis	0.87	0.64	Pleural Abnormality	0.97	0.98
Consolidation	0.96	0.77	Pneumothorax	0.89	1.00
ECS	0.91	0.92	Pulmonary Edema	0.89	0.86
Groundglass Opacity	0.79	0.75			

AMC stands for *Abnormal Mediastinal Contour* and ECS for *Enlarged Cardiac Silhouette*.

TABLE 5 Results on the test set comparing the *ET model* with the two baselines.

Method	<i>Unannotated</i>	<i>Ellipse</i>	<i>ET model (ours)</i>
AUC	0.767 [0.763, 0.771]	0.765 [0.763, 0.768]	0.765 [0.763, 0.767]
IoU	0.201 [0.198, 0.204]	0.335 [0.330, 0.339]	0.256 [0.253, 0.260]

AUC and IoU were averaged over the scores of all labels.

3.3. Comparison with baselines

Results, averaged over all labels, are presented in [Table 5](#). The average AUC for the *ET model* was not significantly different from the baselines. Regarding localization, the IoU values showed that training with the ET data was significantly better than training without annotated data and worse than training with the hand-labeled localization ellipses. Results for AUC and IoU of individual labels are presented in [Tables 1, 6](#). AUC was stable among all methods for almost all labels. Successful and unsuccessful heatmaps generated by our trained models are shown in [Figure 4](#).

When training the *Ellipse* model with only 15% of the annotated dataset, we achieved an IoU of .257 [.248,.266]. Therefore, given the IoU provided for the *ET model* in [Table 5](#), the best estimation for the value of ET data is 15%, i.e., around one-seventh, of the value of the hand-annotated data.

3.4. Ablation study

We present in [Table 7](#) the results of an ablation study with each modification added to the original method from Li et al. (3), which is presented in Section 2.2, including the label specific heatmaps from Section 2.1, the balanced range normalization from Sections 2.2.1, the multi-resolution architecture from Section 2.4, and the multi-task learning from Section 2.3. [Table 7](#) presents the elements added to the model in chronological order of addition to our project in its first five rows. We show an advantage from each modification for all methods. We also added a row with the removal of only the label-specific heatmaps to show that they had a big impact on

TABLE 6 Per-label IoU metric on the test set for the two baselines and our method.

Label	<i>Unannotated</i>	<i>Ellipse</i>	<i>ET model (ours)</i>
AMC	0.111 [0.100, 0.122]	0.297 [0.264, 0.331]	0.214 [0.187, 0.242]
Atelectasis	0.245 [0.241, 0.250]	0.385 [0.378, 0.392]	0.335 [0.320, 0.350]
ECS	0.386 [0.346, 0.427]	0.747 [0.743, 0.751]	0.379 [0.357, 0.401]
Consolidation	0.242 [0.235, 0.249]	0.380 [0.374, 0.385]	0.324 [0.314, 0.334]
Edema	0.314 [0.299, 0.330]	0.466 [0.460, 0.472]	0.401 [0.396, 0.406]
Fracture	0.012 [0.006, 0.018]	0.004 [0.004, 0.004]	0.007 [0.006, 0.008]
Lung Lesion	0.113 [0.104, 0.123]	0.203 [0.193, 0.214]	0.213 [0.194, 0.231]
Opacity	0.260 [0.255, 0.264]	0.387 [0.382, 0.391]	0.341 [0.336, 0.347]
Pleural abnormality	0.210 [0.202, 0.218]	0.297 [0.283, 0.311]	0.246 [0.240, 0.251]
Pneumothorax	0.114 [0.103, 0.125]	0.180 [0.173, 0.187]	0.103 [0.089, 0.118]

the final IoU. Its removal caused a decrease of around 0.062 (24.2%) on the IoU, reaching an IoU similar to the *Unannotated* model.

Furthermore, we evaluated the impact of the quality of the label mention extraction, performed using the modified version of the CheXpert labeler. To check how IoU scores change with respect to the quality of the labeler, we intentionally lowered the quality of Pleural Abnormality label extraction, for which both recall and precision were high in [Table 4](#). We randomly removed the mention of approximately 45% of the cases with mentions of Pleural Abnormality and randomly assigned a mention of another label as a Pleural Abnormality mention in approximately 5% of the cases where there was no mention of Pleural Abnormality. With these changes, we estimate that the recall for mentions of Pleural Abnormality is 0.51, and the precision is 0.86. We trained a model with these modified mentions for Pleural Abnormality and the original mentions for the other nine abnormality labels. We achieved an IoU of 0.227 [0.205,0.249], likely to lie between the *Unannotated* (0.210 [0.202, 0.218]) and *ET model* (0.246 [0.240, 0.251]), shown in [Table 6](#). There was no impact on the Pleural Abnormality AUC (0.869 [0.868, 0.870]).

4. Discussion

Other studies have applied ET data for localizing abnormalities and improving the localization of models. Sember et al. (26) showed that radiologists looked at the location of a label when indicating the presence of tumors in MRIs. However, we use a less restrictive and more challenging scenario, with freeform reports and multiple types of abnormalities reported. Saab et al. (27) showed that, when gaze data are aggregated in hand-crafted features and used in a multi-task setup, the GradCAM heatmap of a model overlaps more often with the location of pneumothoraces. Li et al. (28) developed an attention-guided network for glaucoma diagnosis split into three sequential stages: a prediction of an attention map supervised by an ET heatmap, an intermediary classification network trained to refine the attention map through guided backpropagation, and a final classification network. Wang et al. (29) performed osteoarthritis grading by enforcing the class activation map (CAM) heatmap to be similar to the ET map, allowing for uncertainty in the ET map. We tested adding this method to our loss, but no improvement was seen. All of these methods used ET datasets where radiologists focused on a single task/abnormality, making their ET data intrinsically label-specific and their setup distant from clinical practice. The method we propose uses dictations to identify moments when the radiologist looks at evidence of multiple abnormalities, allowing for the use of ET data collected during clinical report dictation. We used a multi-task formulation as part of our loss following Karargyris et al. (14). However, even though they used a dataset where radiologists looked for several types of abnormalities, they showed only the impact of a single ET heatmap for all labels. Our study focuses

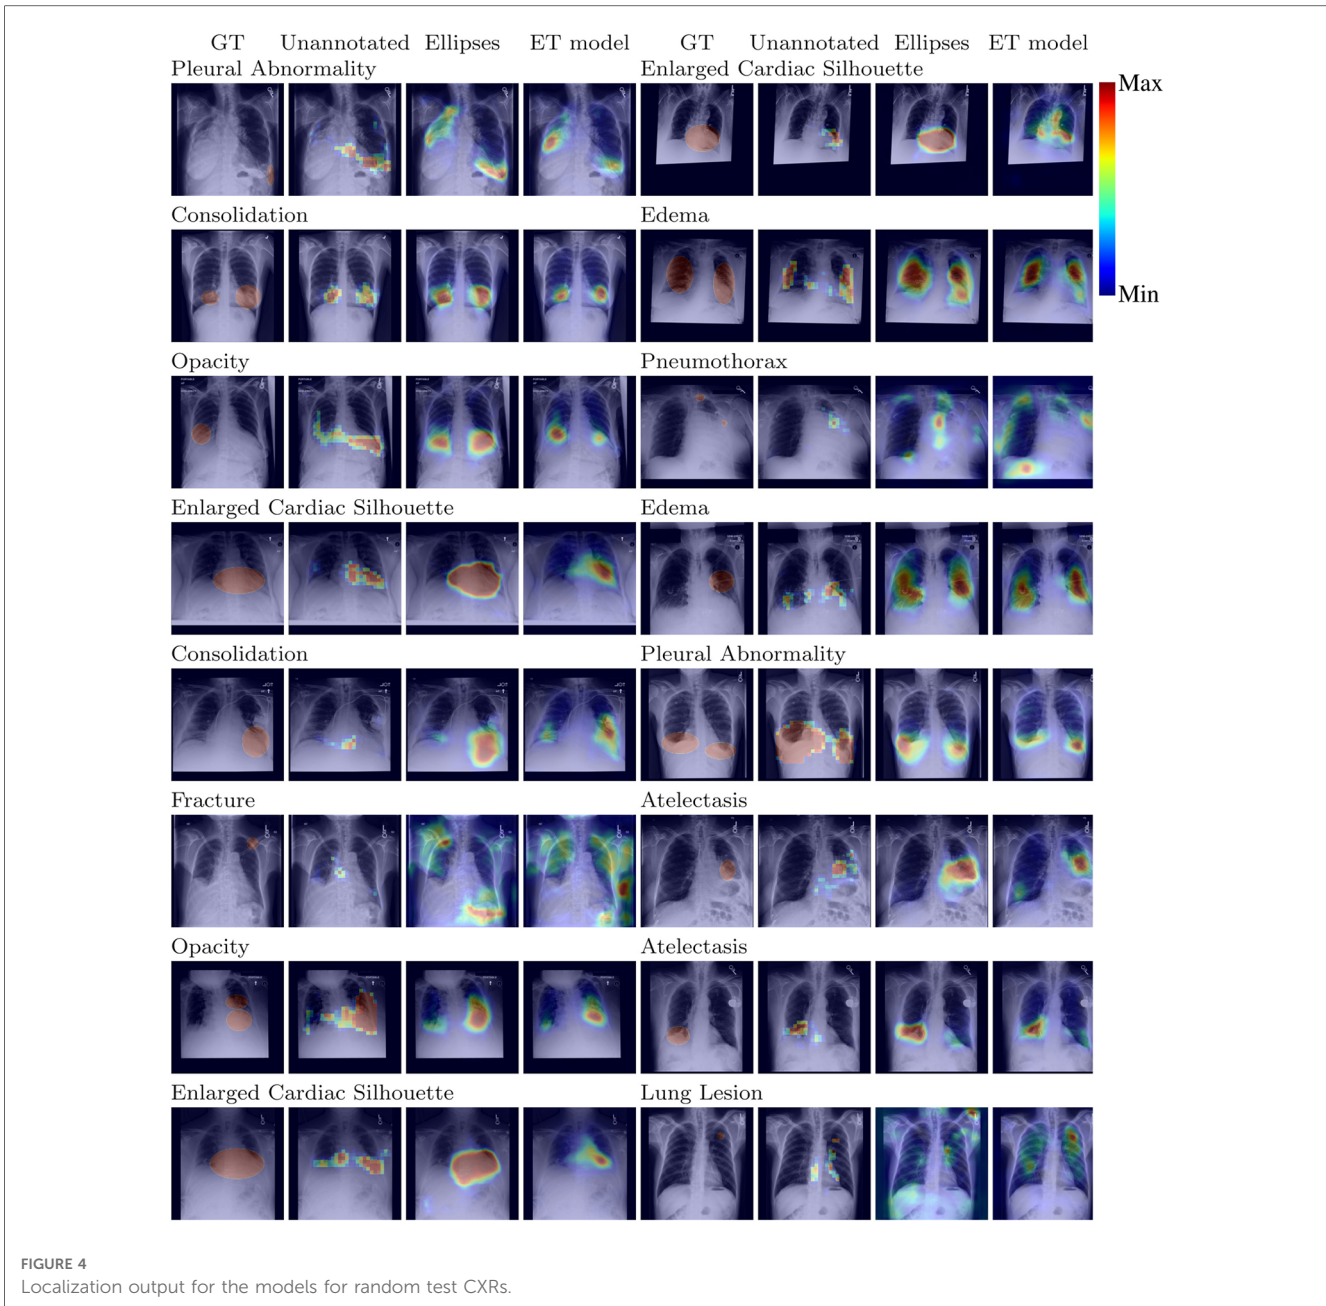


FIGURE 4 Localization output for the models for random test CXRs.

TABLE 7 IoU metric on the test set indicating the advantage of using each of the modifications to the multiple instance learning (MIL) method proposed by Li et al. (3).

MIL	LSH	BRN	MRA	MTL	IoU for ET model
✓	×	×	×	×	.165 [.162, .168]
✓	✓	×	×	×	.180 [.176, .185]
✓	✓	✓	×	×	.200 [.197, .203]
✓	✓	✓	✓	×	.218 [.216, .221]
✓	×	✓	✓	✓	.194 [.192, .197]
✓	✓	✓	✓	✓	.256 [.252, .260]

We tested different combinations of the following methods: label-specific heatmaps (LSH) (our contribution), balanced range normalization (BRN) (our contribution), multi-resolution architecture (MRA) (our contribution), and multi-task learning (MTL) (14).

on more complex uses of the ET data, with the generation of label-specific heatmaps.

Contrary to other works, such as Karargyris et al. (14) and Li et al. (3), we achieved no improvements in classification performance in our setup when applying a variety of localization losses to our model. One of the reasons we might achieve different levels of improvement from Karargyris et al. (14) is that we use a much larger dataset for training the model, weakly including most of the MIMIC-CXR-JPG dataset. The use of abundant unannotated data might reduce the impact of the annotated data on the final model. However, improvements in the ability to localize the abnormalities were still shown in our experiments.

We tried to apply the same method to train localization with the 1,064 images from the dataset shared by Karargyris et al. (14). However, the performance was similar to the *Unannotated* model. There are several reasons why the method might not be generalizable to the other dataset, including:

- It is possible that the method choices were only well-adapted to some of the five radiologists in the REFLACX dataset and not generic for all radiologists, including the single one in the other dataset.
- Some of the method choices might have only been well-adapted to the characteristics of the ET data of the REFLACX dataset. The other dataset, for example, has data collected at 60 Hz instead of 1,000 Hz and has two cases with the first fixations happening after the first mention of an abnormality.
- The size of the other dataset might be too small, or the distribution of cases might be different from the validation set of the REFLACX dataset, which we used to evaluate performance. For example, the dataset shared by Karargyris et al. (14) includes only posterior anterior (PA) CXRs and had only one case where the labeler identified a mention of Pneumothorax.

The method we proposed for producing label-specific localization annotations from ET data and for training models to produce heatmaps that match the annotation improved the interpretability of deep learning models for CXRs, as measured by comparing produced heatmaps against hand-annotated localization of abnormalities. From the IoU achieved by training the model with 15% of the available bounding boxes, we showed that, in our setup, around seven CXRs with ET data provide the same level of efficacy in localization supervision as one CXR with expert-annotated ellipses, showing the possible value of using this type of data for scaling up annotations. This relative value might be used in calculations involving the costs and benefits of each data collection method when deciding on how to get annotations. As shown by our ablation study, the use of label-specific annotations was essential to the added value of using the ET data. We also showed in the ablation study that the performance of the label extraction algorithm has a corresponding impact on the improvements over IoU. These results show that there is an opportunity for improvement of the IoU results for labels that had a low recall and/or precision on **Table 4**. Recent advances in large language models (30) suggest that these models could significantly improve the accuracy of label extraction in the future.

The ET data were relatively noisy, and the achieved IoU for our label-specific training heatmaps was relatively low, limiting the achieved IoU for our proposed method of using ET data for training. In future work, we will investigate other methods of extracting the localization information to reduce the noise in the data, including methods of unsupervised alignment.

Data availability statement

Publicly available datasets were analyzed in this study. This data can be found here: <https://www.physionet.org/content/reflectx-xray-localization/1.0.0/>

<https://www.physionet.org/content/mimic-cxr-jpg/2.0.0/>.

Ethics statement

Ethical review and approval was not required for the study on human participants in accordance with the local legislation and institutional requirements. Written informed consent for participation was not required for this study in accordance with the national legislation and the institutional requirements.

Author contributions

RBL wrote the manuscript, coded and conducted the experiments, and ran the analyses. JDS participated in the study design and provided feedback for the manuscript. TT is the project PI, coordinating the study design, leading discussions about the project, and editing the manuscript. All authors reviewed the manuscript. All authors contributed to the article and approved the submitted version.

Funding

This research was funded by the National Institute of Biomedical Imaging and Bioengineering of the National Institutes of Health under Award Number R21EB028367.

Acknowledgments

Yichu Zhou participated in the modification and evaluation of the CheXpert labeler. This article has appeared in a preprint (31).

Conflict of interest

The author TT declared that they were an editorial board member of *Frontiers*, at the time of submission. This had no impact on the peer review process and the final decision.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Rajpurkar P, Irvin J, Zhu K, Yang B, Mehta H, Duan T, et al. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning (2017).
- Langlotz CP, Allen B, Erickson BJ, Kalpathy-Cramer J, Bigelow K, Cook TS, et al. A road map for foundational research on artificial intelligence in medical imaging: from the 2018 NIH/RSNA/ACR/the academy workshop. *Radiology*. (2019) 291:781–91. doi: 10.1148/radiol.2019190613
- Li Z, Wang C, Han M, Xue Y, Wei W, Li L, et al. Thoracic disease identification, localization with limited supervision. In: *2018 IEEE Conference on Computer Vision, Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018*. IEEE Computer Society (2018). p. 8290–8299. doi: 10.1109/CVPR.2018.00865
- Nguyen HQ, Lam K, Le LT, Pham HH, Tran DQ, Nguyen DB, et al. Vindr-cxr: Anopen dataset of chest x-rays with radiologist's annotations (2021).
- Wang X, Peng Y, Lu L, Lu Z, Bagheri M, Summers RM. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21–26, 2017*. IEEE Computer Society (2017). p. 3462–3471. doi: 10.1109/CVPR.2017.369
- Mettler FA, Bhargavan M, Faulkner K, Gilley DB, Gray JE, Ibbott GS, et al. Radiologic and nuclear medicine studies in the united states and worldwide: frequency, radiation dose, and comparison with other radiation sources—1950–2007. *Radiology*. (2009) 253:520–31. doi: 10.1148/radiol.2532082010. PMID: 19789227
- Lakhani P, Sundaram B. Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks. *Radiology*. (2017) 284:574–82. doi: 10.1148/radiol.2017162326
- Nyboe J. Evaluation of efficiency in interpretation of chest x-ray films. *Bull World Health Organ*. (1966) 35:535–45. PMID: 5297553.
- Balabanova Y, Coker R, Fedorin I, Zakharova S, Plavinskij S, Krukov N, et al. Variability in interpretation of chest radiographs among Russian clinicians, implications for screening programmes: observational study. *BMJ*. (2005) 331:379–82. doi: 10.1136/bmj.331.7513.379
- Quekel LG, Kessels AG, Goei R, van Engelshoven JM. Detection of lung cancer on the chest radiograph: a study on observer performance. *Eur J Radiol*. (2001) 39:111–6. doi: 10.1016/S0720-048X(01)00301-1
- Bustos A, Pertusa A, Salinas JM. Padchest: A large chest x-ray image dataset with multi-label annotated reports. *Med Image Anal*. (2020) 66:101797. doi: 10.1016/j.media.2020.101797
- Irvin J, Rajpurkar P, Ko M, Yu Y, Ciurea-Ilcus S, Chute C, et al. CheXpert: A large chest radiograph dataset with uncertainty labels, expert comparison. In: *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27–February 1, 2019*. AAAI Press (2019). p. 590–597. doi: 10.1609/aaai.v33i01.3301590
- Johnson AEW, Pollard TJ, Berkowitz SJ, Greenbaum NR, Lungren MP, Deng CY, et al. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Sci Data*. (2019) 6:317. doi: 10.1038/s41597-019-0322-0
- Karagyris A, Kashyap S, Lourentzou I, Wu JT, Sharma A, Tong M, et al. Creation and validation of a chest x-ray dataset with eye-tracking and report dictation for AI development. *Sci Data*. (2021) 8:92. doi: 10.1038/s41597-021-00863-5
- Le Meur O, Baccino T. Methods for comparing scanpaths, saliency maps: strengths, weaknesses. *Behav Res Methods*. (2013) 45(1):251–66. doi: 10.3758/s13428-012-0226-9
- Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In: *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22–29, 2017*. IEEE Computer Society (2017). p. 618–626. doi: 10.1109/ICCV.2017.74
- He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. In: *2016 IEEE Conference on Computer Vision, Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27–30, 2016*. IEEE Computer Society (2016). p. 770–778. doi: 10.1109/CVPR.2016.90
- Ioffe S, Szegedy C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: Bach FR, Blei DM, editors. *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6–11 July 2015, JMLR Workshop, Conference Proceedings*. (JMLR.org) (2015). vol. 37, p. 448–456
- Bigolin Lanfredi R, Zhang M, Auffermann W, Chan J, Duong P, Srikumar V, et al. REFLACX: Reports and eye-tracking data for localization of abnormalities in chest x-rays (2021). doi:10.13026/E0DJ-8498
- Bigolin Lanfredi R, Zhang M, Auffermann WF, Chan J, Duong PT, Srikumar V, et al. Reflax, a dataset of reports and eye-tracking data for localization of abnormalities in chest x-rays. *Sci Data*. (2022) 9:350. doi: 10.1038/s41597-022-01441-z
- Goldberger AL, Amaral LAN, Glass L, Hausdorff JM, Ivanov PC, Mark RG, et al. PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation*. (2000) 101:e215–20. doi: 10.1161/01.CIR.101.23.e215
- Johnson A, Lungren M, Peng Y, Lu Z, Mark R, Berkowitz S, et al. MIMIC-CXR-JPG- chest radiographs with structured labels (version 2.0.0) (2019). doi:10.13026/8360-t248
- Johnson AEW, Pollard TJ, Berkowitz SJ, Greenbaum NR, Lungren MP, Deng C, et al. MIMIC-CXR-JPG: A large publicly available database of labeled chest radiographs. CoRR[Preprint] abs/1901.07042 (2019).
- Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. PyTorch: An imperative style, high-performance deep learning library. In: Wallach HM, Larochelle H, Beygelzimer A, d'Alché-Buc F, Fox EB, Garnett R, editors. *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada* (2019). p. 8024–8035.
- Reddi SJ, Kale S, Kumar S. On the convergence of adam and beyond. In: *International Conference on Learning Representations* (2018).
- Stember JN, Celik H, Gutman D, Swinburne N, Young R, Eskreis-Winkler S, et al. Integrating eye tracking and speech recognition accurately annotates MR brain images for deep learning: proof of principle. *Radiol Artificial Intell*. (2021) 3:e200047. doi: 10.1148/ryai.2020200047
- Saab K, Hooper SM, Sohoni NS, Parmar J, Pogatchnik B, Wu S, et al. Observational supervision for medical image classification using gaze data. In: *Medical Image Computing and Computer Assisted Intervention - MICCAI 2021*. Springer International Publishing (2021). p. 603–614. doi: 10.1007/978-3-030-87196-3_56
- Li L, Xu M, Wang X, Jiang L, Liu H. Attention based glaucoma detection: a large-scale database and CNN model. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019*. Computer Vision Foundation/IEEE (2019). p. 10571–10580. doi: 10.1109/CVPR.2019.01082
- Wang S, Ouyang X, Liu T, Wang Q, Shen D. Follow my eye: using gaze to supervise computer-aided diagnosis. *IEEE Trans Med Imaging*. (2022) 41:1688–98. doi: 10.1109/TMI.2022.3146973
- Agrawal M, Hegselmann S, Lang H, Kim Y, Sontag DA. Large language models are few-shot clinical information extractors. In: Goldberg Y, Kozareva Z, Zhang Y, editors. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7–11, 2022*. Association for Computational Linguistics (2022). p. 1998–2022.
- Bigolin Lanfredi R, Schroeder JD, Tasdizen T. Localization supervision of chest x-ray classifiers using label-specific eye-tracking annotation. CoRR [Preprint] abs/2207.09771 (2022). doi: 10.48550/arXiv.2207.09771