VisualCheXbert: Addressing the Discrepancy Between Radiology Report Labels and Image Labels

Saahil Jain*
saahil.jain@cs.stanford.edu
Stanford University
USA

Akshay Smit*
akshaysm@stanford.edu
Stanford University
USA

Steven QH Truong
VinBrain
Vietnam

Chanh DT Nguyen
Minh-Thanh Huynh
VinBrain
Vietnam

Mudit Jain
USA

Victoria A. Young
Stanford University
USA

Andrew Y. Ng
ang@cs.stanford.edu
Stanford University
USA

Matthew P. Lungren†
mlungren@stanford.edu
Stanford University
USA

Pranav Rajpurkar†
pranavsr@cs.stanford.edu
Stanford University
USA

Figure 1: The VisualCheXbert training procedure. VisualCheXbert uses a biomedically-pretrained BERT model to directly map from a radiology report to the labels obtained by a radiologist interpreting the associated X-ray image. The training procedure for VisualCheXbert is supervised by a computer vision model trained to detect medical conditions from chest X-ray images.

ABSTRACT

Automatic extraction of medical conditions from free-text radiology reports is critical for supervising computer vision models to

*Equal Contribution

†Equal Contribution

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ACM CHIL '21, April 8–10, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8359-2/21/04.

<https://doi.org/10.1145/3450439.3451862>

interpret medical images. In this work, we show that radiologists labeling reports significantly disagree with radiologists labeling corresponding chest X-ray images, which reduces the quality of report labels as proxies for image labels. We develop and evaluate methods to produce labels from radiology reports that have better agreement with radiologists labeling images. Our best performing method, called VisualCheXbert, uses a biomedically-pretrained BERT model to directly map from a radiology report to the image labels, with a supervisory signal determined by a computer vision model trained to detect medical conditions from chest X-ray images. We find that VisualCheXbert outperforms an approach using an existing radiology report labeler by an average F1 score of 0.14 (95% CI 0.12, 0.17). We also find that VisualCheXbert better agrees

with radiologists labeling chest X-ray images than do radiologists labeling the corresponding radiology reports by an average F1 score across several medical conditions of between 0.12 (95% CI 0.09, 0.15) and 0.21 (95% CI 0.18, 0.24).

CCS CONCEPTS

• **Computing methodologies** → **Natural language processing**; **Information extraction**.

KEYWORDS

natural language processing, BERT, medical report labeling, chest X-ray diagnosis

ACM Reference Format:

Saahil Jain, Akshay Smit, Steven QH Truong, Chanh DT Nguyen, Minh Thanh Huynh, Mudit Jain, Victoria A. Young, Andrew Y. Ng, Matthew P. Lungren, and Pranav Rajpurkar. 2021. VisualCheXbert: Addressing the Discrepancy Between Radiology Report Labels and Image Labels. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '21)*, April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3450439.3451862>

1 INTRODUCTION

Because manually annotating a large number of medical images is costly [1, 4, 6, 11, 17, 27, 31], an appealing solution is the use of automatic labelers to extract labels from medical text reports that accompany the images. On the task of chest X-ray interpretation, high-performing vision models have been successfully trained [22, 24–26, 30, 33] on large, publicly available chest X-ray datasets [13, 14, 23, 32] labeled by automated radiology report labelers [13, 16, 20, 28]. However, training these vision models on labels obtained from reports assumes that the report labels are good proxies for image labels. Prior work has found that report labels may not accurately reflect the visual content of medical images [18, 19, 29].

We investigate this assumption in the setting of automated chest X-ray labeling and develop methods to produce labels from radiology reports that better agree with radiologists labeling the corresponding X-ray images. Our primary contributions are:

- (1) We quantify the agreement between radiologists labeling reports and radiologists labeling images across several medical conditions. We find that there is significant disagreement between board-certified radiologists when labeling a chest X-ray image and when labeling the corresponding radiology report.
- (2) Upon board-certified radiologist review of examples of disagreements between radiologists labeling reports and radiologists labeling images, we find various reasons for disagreement related to (a) label hierarchy relationships, (b) access to clinical history, (c) the use of the *Impression* and *Findings* section of radiology reports, and (d) the inherent noise of the labeling task.
- (3) We find many significant relationships between presence of conditions labeled using reports and presence of conditions labeled using images. We report and clinically interpret various radiology report labels that increase (or decrease) the odds of particular conditions in an image with statistical significance.

- (4) We learn to map textual radiology reports directly to the X-ray image labels. Our best performing method, called *VisualCheXbert*, uses a biomedically-pretrained BERT model to directly map from a radiology report to the image labels. We find that *VisualCheXbert* better agrees with radiologists labeling chest X-ray images than do radiologists labeling the corresponding radiology reports by an average F1 score across several medical conditions of between 0.12 (95% CI 0.09, 0.15) and 0.21 (95% CI 0.18, 0.24). We also find that *VisualCheXbert* outperforms an approach using the CheXpert radiology report labeler [13] by an average F1 score of 0.14 (95% CI 0.12, 0.17).

We expect that our methods of addressing the discrepancy between medical report labels and image labels are broadly useful across the medical domain and may facilitate the development of improved medical imaging models.

2 DATA

We made use of two large publicly available datasets of chest X-rays: CheXpert [13] and MIMIC-CXR [14]. For both datasets, we use the *Impression* section of the radiology reports, which summarizes the key findings in the radiographic study. Each of the X-rays in these datasets was labeled for 14 commonly occurring medical conditions. CheXpert consists of 224,316 chest radiographs, with labels generated from the corresponding radiology report impression by the automatic, rules-based CheXpert labeler. Given a radiology report impression as input, the CheXpert labeler labels each medical condition (except “No Finding”) as “positive”, “negative”, “uncertain” or “blank”. A “blank” label is produced by the CheXpert labeler if the condition was not mentioned at all in the report impression. If the condition was mentioned but its presence was negated, a “negative” label is produced. If the condition was mentioned but its presence was uncertain, an “uncertain” label is produced. For “No Finding”, the CheXpert labeler only produces “positive” or “blank” labels. “No Finding” is only labeled as “positive” if no medical abnormality whatsoever was mentioned in the report impression. The MIMIC-CXR dataset consists of 377,110 chest X-rays and their corresponding radiology reports, and it has also been labeled by the CheXpert labeler.

The CheXpert dataset contains a separate set of 200 chest X-ray studies called the “CheXpert validation set” and another set of 500 chest X-ray studies called the “CheXpert test set”. The CheXpert validation set is labeled by the majority vote of 3 board-certified radiologists examining the X-ray images and labeling each of the 14 conditions as “positive” or “negative”, similar to the image ground truth on the CheXpert test set, which is described below. No radiologist report labels are obtained for the validation set.

The CheXpert test set, which was collected by Irvin et al. [13], is labeled by radiologists in two distinct ways:

Image ground truth. 5 board-certified radiologists looked at each X-ray image and labeled each of the 14 conditions as “positive” or “negative”. The final label is their majority vote. These radiologists only observed the X-ray images and did not have access to the radiology report or patients’ historical records at the time of image labeling.

Table 1: Agreement between radiologists looking at reports and radiologists looking at the corresponding X-ray images. The high and low scores are obtained by mapping uncertain labels in the radiologist report labels to the image ground truth labels and the opposite of the image ground truth labels respectively.

Condition (n = # positive)	Low F1	High F1	Low Kappa	High Kappa
Atelectasis (n=153)	0.230	0.595	-0.014	0.457
Cardiomegaly (n=151)	0.422	0.463	0.290	0.344
Edema (n=78)	0.453	0.581	0.335	0.492
Pleural Effusion (n=104)	0.638	0.710	0.511	0.613
Enlarged Cardiom. (n=253)	0.089	0.208	-0.053	0.097
Lung Opacity (n=264)	0.683	0.686	0.401	0.405
Support Devices (n=261)	0.863	0.863	0.737	0.737
No Finding (n=62)	0.381	0.381	0.292	0.292
Average	0.470	0.561	0.312	0.430
Weighted Average	0.492	0.575	0.320	0.427

Radiologist report labels. A board-certified radiologist looked at each radiology report impression corresponding to the X-rays and labeled each of the 14 conditions as being “positive”, “negative”, “uncertain”, or “blank”. This radiologist did not observe any X-ray images. A condition was labeled as “blank” if it was not at all mentioned in the report impression. If the condition was mentioned but its presence in the chest X-ray was negated, then the condition was labeled as “negative”. If the condition was mentioned but its presence was uncertain, it was labeled as “uncertain”.

3 EVALUATION

We only evaluate our models on medical conditions for which at least 50 out of the 500 chest X-ray studies in the CheXpert test set were marked positive by the radiologists labeling the X-ray images (image ground truth). These conditions, which we refer to as the *evaluation conditions*, are: Atelectasis, Cardiomegaly, Edema, Pleural Effusion, Enlarged Cardiomediastinum, Lung Opacity, Support Devices, and No Finding. We evaluate models using the average and weighted average of the F1 score across conditions on the CheXpert test set with the image ground truth. To compute the weighted average, each condition is weighted by the portion of positive labels for that condition in the CheXpert test set.

4 EXPERIMENTS

4.1 Do radiologists labeling reports agree with radiologists labeling X-ray images?

We first investigate the extent of the disagreement between board-certified radiologists when labeling a chest X-ray image and when labeling the corresponding radiology report.

4.1.1 Method. We compute the level of agreement between radiologists labeling X-ray images and radiologists labeling the corresponding radiology reports on the CheXpert test set. The CheXpert test set contains a set of labels from radiologists labeling X-ray images as well as another set of labels from radiologists labeling the corresponding radiology reports. Using the labels from X-ray images as the ground truth, we compute Cohen’s Kappa [5] as well as the F1 score to measure the agreement between these two sets of

labels. To compare the radiologist report labels to the image ground truth labels, we convert the radiologist report labels to binary labels as follows. We map the blank labels produced for the radiology report to negative labels. We map uncertain labels to either the image ground truth label or the opposite of the image ground truth label, and we record the results for both these strategies to obtain “Low F1”, “High F1”, “Low Kappa”, and “High Kappa” scores. The low and high scores represent the most pessimistic and optimistic mapping of the uncertainty labels.

4.1.2 Results. We find that there is significant disagreement, which is indicated by low Kappa and F1 scores for almost all conditions evaluated. For example, Enlarged Cardiomediastinum and No Finding have a relatively small “High Kappa” score of 0.097 and 0.292 and a “High F1” score of 0.208 and 0.381, indicating high levels of disagreement even when assuming the most optimistic mapping of the uncertainty labels. Atelectasis, Cardiomegaly, Edema, Pleural Effusion, and Lung Opacity also have a low “High Kappa” score of 0.457, 0.344, 0.492, 0.613, and 0.405 respectively and a “High F1” score of 0.595, 0.463, 0.581, 0.710, and 0.686 respectively. Support Devices has the highest Kappa score, with a “High Kappa” of 0.737, and the highest F1 score, with a “High F1” of 0.863. The average Kappa score is between 0.312 and 0.430, and the average F1 score is between 0.470 and 0.561. The low and high F1 / Kappa scores for the evaluation conditions are shown in Table 1.

4.2 Why do radiologists labeling reports disagree with radiologists labeling X-ray images?

We investigate why there is disagreement between board-certified radiologists when labeling a chest X-ray image and when labeling the corresponding radiology report.

4.2.1 Method. A board-certified radiologist was given access to the chest X-ray image, the full radiology report, the radiology report impression section, the image ground truth across all conditions, and the radiologist report labels across all conditions for each of the 500 examples in the CheXpert test set. The radiologist then explained examples where radiologists labeling reports disagree with radiologists labeling X-ray images. We also calculated the counts of disagreements between radiologists labeling reports and radiologists labeling X-ray images for each condition on the CheXpert test set. A board-certified radiologist explained why there were large numbers of disagreements on certain conditions.

4.2.2 Results. We find various reasons why radiologists labeling reports might disagree with radiologists labeling images. First, there is a difference between the setup of the report labeling and image labeling tasks related to the label hierarchy. On the report labeling task on the CheXpert test set, radiologists were instructed to label only the most specific condition as positive and leave parent conditions blank. For example, although Lung Opacity is a parent condition of Edema, a radiologist marking a report as positive for Edema would leave Lung Opacity blank. Blank report labels are typically mapped to negative image labels. However, radiologists labeling images label each condition as positive or negative independent of the presence of other conditions. Second, radiologists labeling reports have access to clinical report history, which biases

Table 2: Clinical explanations of disagreements between radiologists looking at reports and radiologists looking at images on the CheXpert test set. Given access to the X-ray image, the full radiology report, the radiology report impression, the radiology report labels, and the image ground truth, a board-certified radiologist explained disagreements between radiologist report labels and the image ground truth. We show select examples with explanations in this table.

Report Impression and Labels	Clinical Explanation
<p>1. single ap upright view of the chest showing a mildly increased opacity at the left lung base that could represent atelectasis versus consolidation.</p> <p style="text-align: right;"><i>Cardiomegaly</i></p> <p style="text-align: right;">Radiologist Report Label: Negative Image Ground Truth: Positive</p>	<p>The radiologist looking at the report marks Cardiomegaly as negative as it is not mentioned in the report. Since the image is a Intensive Care Unit (ICU) film and cardiomegaly is not a clinically relevant condition for the population selected for in ICU films, the presence of cardiomegaly was never mentioned in the report, resulting in the discrepancy between radiologists looking at the report and radiologists looking at the image.</p>
<p>1. pulmonary vascular congestion. left lower lobe opacity compatible with atelectasis and/or consolidation.</p> <p style="text-align: right;"><i>Cardiomegaly</i></p> <p style="text-align: right;">Radiologist Report Label: Negative Image Ground Truth: Positive</p>	<p>Although cardiomegaly was mentioned in the radiology report "Findings" section, cardiomegaly was not mentioned in the report "Impression". Since the radiologist looking at the report only had access to the "Impression" section, they labeled Cardiomegaly as negative when it was actually present in the image.</p>
<p>1. decreased pulmonary edema. stable bilateral pleural effusions and bibasilar atelectasis.</p> <p style="text-align: right;"><i>Edema</i></p> <p style="text-align: right;">Radiologist Report Label: Positive Image Ground Truth: Negative</p>	<p>The phrase "decreased pulmonary edema" shows that the radiologist writing the report had relevant clinical context, as the edema has "decreased" compared to a previous report or image. However, the radiologist looking at the image does not have this clinical context, resulting in a discrepancy.</p>
<p>1. single frontal radiograph of the chest is limited secondary to poor inspiration and rotation.</p> <p>2. cardiac silhouette is partially obscured secondary to rotation. lungs demonstrate bibasilar opacities, likely reflecting atelectasis. possible small right pleural effusion. no pneumothorax.</p> <p>3. visualized osseous structures and soft tissues unremarkable.</p> <p style="text-align: right;"><i>Pleural Effusion</i></p> <p style="text-align: right;">Radiologist Report Label: Positive Image Ground Truth: Negative</p>	<p>The phrase "possible small right pleural effusion" indicates the uncertainty regarding the presence of pleural effusion. This natural uncertainty may explain the disagreement between radiologists looking at the image and radiologists looking at the report. On review, it was noted that pleural effusion was borderline in this example.</p>
<p>1. crowding of the pulmonary vasculature. cannot exclude mild interstitial pulmonary edema.</p> <p>2. no focal air space consolidation. the cardiomedastinal silhouette appears grossly within normal limits.</p> <p style="text-align: right;"><i>Pleural Effusion</i></p> <p style="text-align: right;">Radiologist Report Label: Negative Image Ground Truth: Positive</p>	<p>Upon review by a board-certified radiologist, there was an error in the radiology report, which did not mention the presence of pleural effusion. The error in the report itself may explain the disagreement between the image and report labels.</p>

radiologists towards reporting certain conditions in reports while a radiologist labeling the image may not observe the condition on the image. Busby et al. [3] explain biases from clinical history in terms of framing bias, where the presentation of the clinical history can lead to different diagnostic conclusions, and attribution bias, where information in the clinical history can lead to different diagnostic conclusions. Third, radiologists labeling reports were only given access to the report impression section when labeling the CheXpert test set. Sometimes, conditions are mentioned in the *Findings* section of the report but not mentioned in the *Impression* section. This results in more negative labels when radiologists looked at

reports. For chest CT scan reports, Gershanik et al. [10] also find that a condition mentioned in the *Findings* section is not always mentioned in the *Impression* section of the report. Fourth, labeling images and reports is inherently noisy to a certain extent, resulting in disagreement. Drivers of noise include mistakes on the part of radiologists labeling reports and radiologists labeling images, uncertainty regarding the presence of a condition based on an image or report, and different thresholds for diagnosing conditions as positive among radiologists. Brady et al. [2] describe additional factors that contribute to discrepancies in radiologist interpretations,

including radiologist specific causes of error like under reading as well as system issues like excess workload.

Next, we explain the counts of the largest disagreements between radiologists labeling reports and radiologists labeling images. Out of the 500 examples on the CheXpert test set, there were 223 examples where the image was labeled positive while the report was labeled negative for Enlarged Cardiomeastinum. We hypothesize that this results from the difference in the task setup related to the label hierarchy. Since Enlarged Cardiomeastinum is a parent condition of Cardiomegaly, radiologists labeling reports were instructed to leave Enlarged Cardiomeastinum blank if they labeled Cardiomegaly positive. There were 101 examples where the image was labeled positive while the report was labeled negative for Cardiomegaly. Diagnosis of cardiomegaly on chest radiographs can depend on patient positioning and clinical history. Further, particularly in the ICU setting in which multiple consecutive radiographs are taken, cardiomegaly is not consistently described in the report even when present unless a clinically significant change is observed (i.e. pericardial effusion). There were 100 examples where the image was labeled positive while the report was labeled negative for Lung Opacity. We hypothesize that this results from the difference in task setup related to label hierarchy, as Lung Opacity is a parent condition. Further, particularly in the setting of atelectasis, lung opacity may not have risen to clinical relevance for the reporting radiologist despite being seen on the isolated imaging task. There were 65 examples where the image was labeled negative while the report was labeled positive for Pleural Effusion. We hypothesize that this partially results from both the variant thresholds for diagnosis of pleural effusion among radiologists and the clinical setting in which the reporting radiologist has access to prior films. It was common to see the report state "decreased" or "trace residual" effusion due to the context of prior imaging on that patient. However, in the isolated image labeling task, the perceived likelihood of the condition fell below the threshold of a board-certified radiologist. There were 49 examples where the image was labeled negative, while the report was labeled positive for Edema. Similar to the effusion example, clinical context and prior imaging played a role in these discrepancies as, again, diagnoses were carried forward from prior studies and language such as "some residual" or "nearly resolved" in the report were used to indicate the presence of edema based on the clinical context. However, when labeling the corresponding image in isolation, the presence of edema fell below the threshold of a board-certified radiologist. Table 2 contains specific examples of these disagreements with clinical explanations. Table 3 shows the counts of disagreements between radiologists labeling reports and radiologists labeling images by condition.

4.3 Are there significant relationships between conditions labeled from reports and conditions labeled from images?

To determine whether there are significant relationships between conditions labeled from reports and conditions labeled from images, we learn a mapping from the output of radiologists labeling reports to the output of radiologists labeling images. We then analyze the significant relationships implied by this mapping from a clinical perspective.

Table 3: Counts of disagreements by condition between radiologists labeling reports and radiologists labeling the corresponding X-ray images on the CheXpert test set. The first column reports the number of times the image ground truth was positive, while the radiologist report label was negative. The second column reports the number of times the image ground truth was negative, while the radiologist report label was positive.

Condition	Positive on Image Negative on Report	Negative on Image Positive on Report
No Finding	38	40
Enlarged Cardiom.	223	5
Cardiomegaly	101	15
Lung Opacity	100	50
Lung Lesion	2	12
Edema	26	49
Consolidation	16	17
Pneumonia	6	5
Atelectasis	75	31
Pneumothorax	1	13
Pleural Effusion	11	65
Pleural Other	3	15
Fracture	3	21
Support Devices	53	13

4.3.1 Method. We train logistic regression models to map the radiologist report labels for all conditions to the image ground truth for each of the evaluation conditions. We quantitatively measure the relationship between the radiologist report labels and the image ground truth by obtaining odds ratios from the coefficients of these logistic regression models. We review the odds ratios from these models with a board-certified radiologist to understand how particular radiologist report labels might clinically change the odds of image labels.

4.3.2 Training details. We one-hot encode the radiologist report labels and provide these binary variables as inputs to a logistic regression model. For example, the "Atelectasis Positive" variable is 1 if the radiologist labels Atelectasis as positive on the report and 0 otherwise. Similarly, the "Atelectasis Negative" variable is 1 if the radiologist labels Atelectasis as negative on the report and 0 otherwise. The same logic applies to the "Atelectasis Uncertain" variable as well as the other variables for each condition. We then train the logistic regression model with L1 regularization ($\alpha = 0.5$) on the CheXpert test set using the one-hot encoded radiologist report labels (for all conditions) as input and the image ground truth for a condition as output. In total, we train different logistic regression models to map the radiologist report labels to binary image labels for each of the 8 evaluation conditions. We compute odds ratios by exponentiating the coefficients of the logistic regression models.

4.3.3 Results. After training the logistic regression models, we find that particular radiology report labels increased (or decreased) the odds of particular conditions in an image with statistical significance ($P < 0.05$). As expected, we find that radiology report labels associated with a condition increase the odds of that same

condition in the image; for example, a Cardiomegaly positive report label increases the odds of Cardiomegaly in the image. We also find that the regression model corrects for label hierarchy. A Cardiomegaly positive report label increases the odds of Enlarged Cardiomegistinum (the parent of Cardiomegaly) on the image by 9.6 times. We similarly observe the model correcting for the label hierarchy of Lung Opacity. Radiology report labels of Edema positive, Consolidation positive, and Atelectasis positive, which all correspond to child conditions of Lung Opacity, increase the odds of Lung Opacity. We also find that the model maps particular uncertainties in report labels to the presence of a condition in the image. For example, Atelectasis uncertain report labels and Edema uncertain report labels increase the odds of Lung Opacity by 2.9 and 7.9 times respectively.

Next, we find that the model maps positive report labels to the presence of other conditions in the image. A Pleural Effusion positive report label increases the odds of Lung Opacity by 4.4 times. We hypothesize that this results from co-occurrence between Pleural Effusion and child conditions of Lung Opacity such as Atelectasis and Edema. Pleural effusion physiologically often leads to adjacent lung collapse, atelectasis, and is often seen in physiologic fluid overload conditions, edema. We find that an Atelectasis positive report label decreases the odds of Support Devices in the image by 0.28 times. On the patient population who have support devices, many of whom are in an Intensive Care Unit (ICU) setting, it is not clinically useful for radiologists to comment on the presence of atelectasis on reports, as they would rather focus on more clinically relevant changes. This may explain the mechanism by which the presence of atelectasis in a report signals that there are no support devices in the image. We find that a Fracture positive report label decreases the odds of Support Devices by 0.17 times. We hypothesize that this results from a negative co-occurrence between Fractures and Support Devices, as the two observations select for different patient populations: X-rays for fractures are often done in the Emergency Department (ED) or other outpatient settings rather than the ICU setting. We find that an Edema positive report label increases the odds of Enlarged Cardiomegistinum on the image by 2.1 times. This may be explained by the fact that Edema and Enlarged Cardiomegistinum often co-occur in a clinical setting, as they can both be caused by congestive heart failure. Lastly, we find that a Support Devices positive report label decreases the odds of No Finding in the image by 0.03 times. This may be explained by the fact that patients with support devices are usually in the ICU setting and sick with other pathologies. We visualize these statistically significant odds ratios for each type of radiologist report label (such as "Atelectasis Negative") as a factor for the presence of an evaluation condition in the X-ray image in Figure 2.

4.4 Can we naively map labels obtained from reports to X-ray image labels?

We map the output of an automated radiology report labeler to X-ray image labels using simple uncertainty handling strategies.

4.4.1 Method. For a baseline approach, we naively map labels obtained from running the CheXpert labeler on the radiology report impressions to X-ray image labels. The CheXpert labeler is an automatic, rules-based radiology report labeler [13]. The labels

Table 4: F1 scores obtained by the Zero-One and LogReg baselines, evaluated on the CheXpert test set. The weighted average is weighted by prevalence (n = # positive).

Condition (n = # positive)	Zero-One Baseline	LogReg Baseline
Atelectasis (n=153)	0.52	0.63
Cardiomegaly (n=151)	0.46	0.56
Edema (n=78)	0.53	0.47
Pleural Effusion (n=104)	0.65	0.65
Enlarged Cardiom. (n=253)	0.20	0.67
Lung Opacity (n=264)	0.69	0.81
Support Devices (n=261)	0.85	0.84
No Finding (n=62)	0.39	0.55
Average	0.54	0.65
Weighted Average	0.56	0.70

produced by the CheXpert labeler include 4 classes per medical condition (positive, negative, uncertain, and blank). Since the image ground truth only has positive or negative labels per condition, we must map the labels produced by the CheXpert labeler to binary labels. We map the blank labels produced by the CheXpert labeler to negative labels. We do not change the positive and negative labels produced by the CheXpert labeler. To handle the uncertain labels, we use the two common uncertainty handling strategies in Irvin et al. [13]: we map the uncertain labels to either all negative labels (zeros-uncertainty handling strategy) or all positive labels (ones-uncertainty handling strategy). We record the F1 score from the better performing strategy on the CheXpert test set, using as ground truth the labels provided by radiologists labeling X-ray images (image ground truth). We refer to this method as the *Zero-One Baseline*. Since we only report the maximum of the zeros-uncertainty handling strategy and the ones-uncertainty handling strategy, the F1 scores for the Zero-One Baseline represent the most optimistic global mapping of the uncertainty labels for this method.

4.4.2 Results. We find that the average and weighted average F1 scores across the evaluation conditions for the Zero-One Baseline are 0.54 and 0.56 respectively, which are in between the average / weighted average "Low F1" and "High F1" scores for radiologists labeling reports (see Table 1). This indicates that the Zero-One Baseline is not strictly better or worse than radiologists labeling reports, who we previously show to have poor agreement with radiologists labeling images. The Zero-One Baseline F1 scores for Atelectasis, Cardiomegaly, Edema, Pleural Effusion, and Enlarged Cardiomegistinum are 0.52, 0.46, 0.53, 0.65, and 0.20 respectively, which are all between the respective "Low F1" and "High F1" scores for radiologists labeling reports. The Zero-One Baseline F1 scores for Lung Opacity and No Finding are 0.69 and 0.39 respectively, which are slightly higher (~ 0.01 difference) than the respective "High F1" scores for radiologists labeling reports. Similarly, the Zero-One Baseline F1 score for Support Devices is 0.39, which is slightly lower (~ 0.01 difference) than the "Low F1" Support Devices score for radiologists labeling reports. The F1 scores for the Zero-One Baseline across the evaluation conditions are shown in Table 4.

Figure 2: Odds ratios for radiologist report labels as factors for the presence of a condition in the X-ray image. We map the radiologist report labels across all conditions to the image ground truth using a logistic regression model. We obtain odds ratios for the input variables, which are the one-hot encoded radiologist report labels, and only display odds ratios for which the corresponding P value (two-sided t test) is less than 0.05.

Table 5: F1 scores for BERT+Thresholding and BERT+LogReg trained on the MIMIC-CXR and CheXpert datasets. We refer to the BERT+Thresholding method on the MIMIC-CXR dataset as VisualCheXbert. The models here are evaluated on the CheXpert test set.

Condition (n = # positive)	BERT+Thresholding on MIMIC-CXR, DenseNet Labels	BERT+LogReg on MIMIC-CXR, DenseNet Labels (VisualCheXbert)	BERT+Thresholding on CheXpert, DenseNet Labels	BERT+LogReg on CheXpert, DenseNet Labels
Atelectasis (n=153)	0.65	0.64	0.67	0.66
Cardiomegaly (n=151)	0.53	0.62	0.61	0.61
Edema (n=78)	0.55	0.54	0.49	0.53
Pleural Effusion (n=104)	0.64	0.65	0.57	0.67
Enlarged Cardiom. (n=253)	0.44	0.73	0.60	0.70
Lung Opacity (n=264)	0.81	0.83	0.70	0.83
Support Devices (n=261)	0.85	0.87	0.80	0.84
No Finding (n=62)	0.44	0.54	0.46	0.52
Average	0.61	0.68	0.61	0.67
Weighted Average	0.65	0.73	0.65	0.72

4.5 Can we learn to map labels obtained from reports to X-ray image labels?

We map the output of an automated radiology report labeler to X-ray image labels, similarly to how we previously map the output of radiologists labeling reports to the output of radiologists labeling images. Previous work by Dunnmon et al. [8] showed that labels obtained from noisy labeling functions on radiology reports can be

mapped to labels that are of similar quality to image labels produced by radiologists for the simpler task of classifying X-rays as normal or abnormal.

4.5.1 Method. This approach, motivated by a prior experiment in which we map radiologist report labels to image labels, improves upon the naive uncertainty mapping strategy used in the Zero-One Baseline. As before, we obtain report labels by running the

CheXpert labeler on radiology report impressions. For each of the evaluation conditions, we train a logistic regression model that maps the CheXpert labeler's output on a radiology report impression to a positive or negative label for the target condition. This approach makes use of the automated report labels for all 14 conditions to predict the label for each target condition. We refer to this approach as the *LogReg Baseline*.

4.5.2 Training details. We one-hot encode the report labels outputted by the CheXpert labeler and provide these binary variables as inputs to a logistic regression model. We train a logistic regression model with $L2$ regularization ($C = 1.0$) and a max iteration of 500 using the one-hot encoded report labels (for all conditions) as input and the image ground truth for a condition as output. The class weights are the inverse prevalence of the respective class in the training set. We use a leave-one-out cross-validation strategy to train and validate the logistic regression model on the CheXpert test dataset. For each of the 8 evaluation conditions, we train different logistic regression models to map the labels produced by the CheXpert labeler to binary image labels.

4.5.3 Results. We find that the LogReg Baseline approach improves upon the Zero-One Baseline for most conditions. Compared to the Zero-One Baseline, the LogReg Baseline increases the average F1 score from 0.54 to 0.65 and the weighted average F1 score from 0.56 to 0.70. The LogReg Baseline increases the F1 score compared to the Zero-One Baseline from 0.52 to 0.63 for Atelectasis, 0.46 to 0.56 for Cardiomegaly, 0.20 to 0.67 for Enlarged Cardiomediastinum, 0.69 to 0.81 for Lung Opacity, and 0.39 to 0.55 for No Finding. However, the LogReg Baseline decreases the F1 scores compared to the Zero-One Baseline from 0.53 to 0.47 for Edema and 0.85 to 0.84 for Support Devices. For Pleural Effusion, both the LogReg Baseline and the Zero-One Baseline have an F1 score of 0.65. Although the LogReg Baseline is not better than the Zero-One Baseline for all conditions, these results suggest that a learned mapping from radiologist report labels to X-ray image labels can outperform naively mapping all uncertain labels to positive or negative for most conditions. The F1 scores obtained by the LogReg Baseline, along with head-to-head comparisons to the Zero-One Baseline, are shown in Table 4.

4.6 Can we learn to map the text reports directly to the X-ray image labels?

Previously, we mapped the output of an existing automated report labeler, which takes text reports as input, to X-ray image labels. We now map the textual radiology report directly to the X-ray image labels.

4.6.1 Method. We develop a deep learning model that maps a radiology report directly to the corresponding X-ray image labels.

Since it is too expensive to obtain labels from radiologists for hundreds of thousands of X-ray images to supervise our model, we instead train a single DenseNet model [12] to detect medical conditions from chest X-ray images, as is described by Irvin et al. [13], and we use this computer vision model as a proxy for a radiologist labeling an X-ray image. We use the DenseNet model to output probabilities for each of the 14 conditions for all X-rays in the MIMIC-CXR dataset and the CheXpert training dataset. To obtain the output of the vision model on the MIMIC-CXR dataset,

we train the DenseNet on the CheXpert training dataset. Similarly, to obtain the output of the vision model on the CheXpert training dataset, we train the DenseNet on the MIMIC-CXR dataset. We find that the DenseNet trained on the CheXpert training set has an AUROC of 0.875 on the CheXpert test set across all conditions, and the DenseNet trained on the MIMIC-CXR dataset has an AUROC of 0.883 on the CheXpert test set across all conditions.

We then use the probabilities outputted from these computer vision models as ground truth to fine-tune a BERT-base model. We train one BERT model using the MIMIC-CXR dataset and one using the CheXpert training dataset. The BERT model takes a tokenized radiology report impression from the MIMIC-CXR or CheXpert dataset as input and is trained to output the labels produced by the DenseNet model. We feed the BERT model's output corresponding to the *[CLS]* token into linear heads (one head for each medical condition) to produce scores for each medical condition. We use the cross-entropy loss to fine-tune BERT. The BERT model is initialized with biomedically pretrained weights produced by Peng et al. [21]. This model training process is shown in Figure 1.

After training the BERT model, we map the outputs of BERT, which are probabilities, to positive or negative labels for each condition. To do so, we try two different methods. Our first method uses optimal probability thresholds to convert the BERT outputs to binary labels. We calculate optimal thresholds by finding the threshold for each condition that maximizes Youden's index [34] (the sum of sensitivity and specificity minus one) on the CheXpert validation dataset. We refer to this approach as *BERT+Thresholding*. Our second method trains a logistic regression model to map the output of BERT across all 14 conditions to a positive or negative label for the target condition. We refer to this approach as *BERT+LogReg*. Ultimately, we develop four different models by using both methods on outputs from a BERT model trained on the MIMIC-CXR dataset and a BERT model trained on the CheXpert training dataset. The four resulting models are called BERT+Thresholding on MIMIC-CXR, BERT+LogReg on MIMIC-CXR, BERT+Thresholding on CheXpert, and BERT+LogReg on CheXpert. We refer to the BERT+LogReg model trained on the MIMIC-CXR dataset with labels provided by the DenseNet model, which is our best performing approach, as *VisualCheXbert*.

4.6.2 Training details. We train the BERT model on 3 TITAN-XP GPUs using the Adam optimizer [15] with a learning rate of 2×10^{-5} , following Devlin et al. [7] for fine-tuning tasks. We use a random 85%-15% training-validation split, as in Smit et al. [28]. The BERT model is trained until convergence. We use a batch size of 18 radiology report impressions. For the BERT+LogReg approach, the logistic regression model uses $L2$ regularization ($C = 1.0$) and a max iteration of 500. Similar to the LogReg Baseline, the class weights are the inverse prevalence of the respective class in the training set, and we use a leave-one-out cross-validation strategy to train and test the logistic regression model on the CheXpert test dataset. We train different logistic regression models to map the probabilities outputted by the BERT model to the binary image labels for each of the 8 evaluation conditions.

4.6.3 Results. We compare the performance of the different BERT approaches on the CheXpert test set. First, we find that on most conditions, BERT+LogReg outperforms BERT+Thresholding. This

Table 6: Improvement in F1 score obtained by VisualCheXbert, evaluated on the CheXpert test set and reported with 95% confidence intervals. The left-most column shows the improvement over the Zero-One Baseline. The middle column shows the improvement over the radiologist report labels with uncertain mapped to the image ground truth label. The right-most column shows the improvement over the radiologist report labels with uncertain mapped to the opposite of image ground truth label.

Condition (n = # positive)	Improvement over Zero-One Baseline	Improvement over Higher Radiologist Score	Improvement over Lower Radiologist Score
Atelectasis (n=153)	0.12 (0.04, 0.20)	0.04 (-0.04, 0.12)	0.41 (0.32, 0.49)
Cardiomegaly (n=151)	0.16 (0.07, 0.25)	0.15 (0.07, 0.25)	0.20 (0.11, 0.28)
Edema (n=78)	0.01 (-0.05, 0.07)	-0.04 (-0.09, 0.02)	0.09 (0.02, 0.17)
Pleural Effusion (n=104)	-0.01 (-0.04, 0.03)	-0.06 (-0.10, -0.02)	0.01 (-0.03, 0.05)
Enlarged Cardiom. (n=253)	0.53 (0.46, 0.60)	0.52 (0.44, 0.60)	0.64 (0.57, 0.71)
Lung Opacity (n=264)	0.14 (0.09, 0.20)	0.15 (0.09, 0.20)	0.15 (0.10, 0.20)
Support Devices (n=261)	0.02 (-0.01, 0.06)	0.01 (-0.02, 0.04)	0.01 (-0.02, 0.04)
No Finding (n=62)	0.15 (0.05, 0.26)	0.16 (0.05, 0.28)	0.16 (0.05, 0.28)
Average	0.14 (0.12, 0.17)	0.12 (0.09, 0.15)	0.21 (0.18, 0.24)
Weighted Average	0.17 (0.15, 0.20)	0.15 (0.13, 0.18)	0.24 (0.21, 0.26)

finding holds true on both the CheXpert and MIMIC-CXR datasets. Second, we find that despite being trained on datasets from different institutions, the models trained on MIMIC-CXR and CheXpert datasets perform similarly. This indicates that the BERT model trained on radiology report impressions from the MIMIC-CXR distribution (Beth Israel Deaconess Medical Center Emergency Department between 2011–2016) [14] can perform as well as a model trained on radiology report impressions from the CheXpert distribution (Stanford Hospital between 2002–2017) [13], even when both models are evaluated on a test set from the CheXpert distribution. Since we obtain a slightly higher average and weighted average F1 using the MIMIC-CXR dataset, we use BERT trained on MIMIC-CXR in our final approach called VisualCheXbert. The performance of the BERT approaches is shown in Table 5.

Next, we compare VisualCheXbert to the Zero-One Baseline. When comparing VisualCheXbert to the Zero-One Baseline as well as the higher and lower scores of radiologists labeling reports described below, we report the improvements by computing the paired differences in F1 scores on 1000 bootstrap replicates and providing the mean difference along with a 95% two-sided confidence interval [9]. Overall, VisualCheXbert improves the average F1 and weighted average F1 over the Zero-One Baseline with statistical significance, increasing the average F1 score by 0.14 (95% CI 0.12, 0.17) and the weighted average F1 score by 0.17 (95% CI 0.15, 0.20). We find that VisualCheXbert obtains a statistically significant improvement over the Zero-One Baseline on most conditions. VisualCheXbert increases the F1 score on Enlarged Cardiomediastinum, Cardiomegaly, No Finding, Lung Opacity, and Atelectasis compared to the Zero-One Baseline by 0.53 (95% CI 0.46, 0.60), 0.16 (95% CI 0.07, 0.25), 0.15 (95% CI 0.05, 0.26), 0.14 (95% CI 0.09, 0.20), and 0.12 (95% CI 0.04, 0.20), respectively. VisualCheXbert obtains similar performance (no statistically significant difference) to the Zero-One Baseline on the rest of the conditions, which are Edema, Pleural Effusion, and Support Devices, with improvements of 0.01 (95% CI -0.05, 0.07), -0.01 (95% CI -0.04, 0.03), and 0.02 (95% CI -0.01, 0.06), respectively.

Lastly, we compare the F1 scores for VisualCheXbert to the higher and lower scores of radiologists labeling reports. The higher scores for radiologists labeling reports are obtained by mapping the uncertain radiologist report labels to the image ground truth label, while the lower scores for radiologists labeling reports are obtained by mapping the uncertain radiologist report labels to the opposite of the ground truth. Overall, VisualCheXbert obtains a statistically significant improvement over both the higher and lower radiologist scores, increasing the average F1 score by 0.12 (95% CI 0.09, 0.15) over the higher radiologist score and 0.21 (95% CI 0.18, 0.24) over the lower radiologist score and increasing the weighted average F1 score by 0.15 (95% CI 0.13, 0.18) over the higher radiologist score and 0.24 (95% CI 0.21, 0.26) over the lower radiologist score. Statistically significant improvements over the higher radiologist score are observed for Cardiomegaly (0.15 [95% CI 0.07, 0.25]), Enlarged Cardiomediastinum (0.52 [95% CI 0.44, 0.60]), Lung Opacity (0.15 [95% CI 0.09, 0.20]), and No Finding (0.16 [95% CI 0.05, 0.28]). VisualCheXbert performs similarly (no statistically significant difference) to the higher radiologist score on Atelectasis (0.04 [95% CI -0.04, 0.12]), Edema (-0.04 [95% CI -0.09, 0.02]), and Support Devices (0.01 [95% CI -0.02, 0.04]). VisualCheXbert performs slightly worse than the higher radiologist score on one condition, which is Pleural Effusion (-0.06 [95% CI -0.10, -0.02]). VisualCheXbert observes considerable, statistically significant improvements compared to the lower radiologist score on all but two conditions. There is no statistically significant difference between VisualCheXbert and the lower radiologist score on these two conditions, which are Pleural Effusion (0.01 [95% CI -0.03, 0.05]) and Support Devices (0.01 [95% CI -0.02, 0.04]). We show the improvements obtained by VisualCheXbert over the Zero-One Baseline and the improvements over radiologists labeling reports in Table 6.

5 LIMITATIONS

Our work has the following limitations. First, our study only made use of the *Impression* section of the radiology reports, which is a summary of the radiology report. Prior work regarding automated

chest X-ray labeling has also extensively used the impression section in radiology reports [13, 14, 32]. However, conditions are sometimes mentioned in the *Findings* section of the report but not in the *Impression* section. As a result, negative and blank labels are more frequent when using the *Impression* section, and this could increase the disparity between labels extracted from the impression and the corresponding chest X-ray image labels. Second, the VisualCheXbert model has a maximum input size of 512 tokens. In practice, only 3 of the report impressions in the entire CheXpert dataset were longer than this limit. Third, the CheXpert test set, on which we evaluated our models, consists of 500 radiology studies and is therefore limited in size. As a result, some of the medical conditions contained very few positive examples; we only evaluated our models on conditions for which at least 10% of the examples in the CheXpert test set were positive. Using a larger test set would allow evaluation on rarer conditions. Fourth, our models are evaluated on chest X-rays from a single institution. Further evaluation on data from other institutions could be used to evaluate the generalizability of our models.

6 CONCLUSION

We investigate the discrepancy between labels extracted from radiology reports and the X-ray image ground truth labels. We then develop and evaluate methods to address this discrepancy. In our work, we aim to answer the following questions.

Do radiologists labeling reports agree with radiologists labeling X-ray images? We find that there is significant disagreement between radiologists labeling reports and radiologists labeling images. On the CheXpert test set, we observe low Kappa scores for almost all conditions evaluated. The average Kappa across the evaluation conditions is between 0.312 and 0.430. These bounds are based on the most pessimistic mapping and most optimistic mapping of uncertain radiology report labels.

Why do radiologists labeling reports disagree with radiologists labeling X-ray images? Upon a board-certified radiologist review of examples of disagreements between radiologists labeling reports and radiologists labeling images, we find four main reasons for disagreement. First, on the CheXpert test set, radiologists labeling reports typically do not mark a parent condition as positive if a child condition is positive. An example of a parent and child condition would be Lung Opacity and Edema, respectively. Second, radiologists labeling reports have access to clinical report history, which biases their diagnoses compared to radiologists labeling images who do not have access to this information. Third, conditions are sometimes reported in the *Findings* section of radiology reports but not the *Impression* section of radiology reports. However, the *Impression* section of radiology reports is commonly used to label reports. This discrepancy can cause radiologists labeling reports to miss pathologies present on the X-ray image. Fourth, labeling images and reports is noisy to a certain extent due to factors such as human mistakes, uncertainty in both reports and images, and different thresholds for diagnosing conditions as positive among radiologists.

Are there significant relationships between conditions labeled from reports and conditions labeled from images? We find many significant relationships between conditions labeled from reports and

conditions labeled from images. We report and clinically interpret various radiology report labels that increase (or decrease) the odds of particular conditions in an image with statistical significance ($P < 0.05$). As expected, we find that positive report labels for a condition increase the odds of that condition in an image. We find that positive report labels for children of a condition increase the odds of the parent condition in an image, a phenomenon that is correcting for the label hierarchy. We find that particular uncertain report labels for a condition increase the odds of the condition (and/or its parent condition). We also find that positive report labels for certain conditions increase (or decrease) the odds of other conditions in the image. One example is that a positive Atelectasis report label decreases the odds of Support Devices in the X-ray image by 0.28 times. We explain potential mechanisms by which the presence of a condition in a report signals the presence (or absence) of another condition in the image.

Can we learn to map radiology reports directly to the X-ray image labels? We learn to map a textual radiology report directly to the X-ray image labels. We use a computer vision model trained to detect diseases from chest X-rays as a proxy for a radiologist labeling an X-ray image. Our final model, VisualCheXbert, uses a biomedically-pretrained BERT model that is supervised by the computer vision model. When evaluated on radiologist image labels on the CheXpert test set, VisualCheXbert increases the average F1 score across the evaluation conditions by between 0.12 (95% CI 0.09, 0.15) and 0.21 (95% CI 0.18, 0.24) compared to radiologists labeling reports. VisualCheXbert also increases the average F1 score by 0.14 (95% CI 0.12, 0.17) compared to a common approach that uses a previous rules-based radiology report labeler.

Given the considerable, statistically significant improvement obtained by VisualCheXbert over the approach using an existing radiology report labeler [13] when evaluated on the image ground truth, we hypothesize that VisualCheXbert's labels could be used to train better computer vision models for automated chest X-ray diagnosis.

ACKNOWLEDGMENTS

We would like to acknowledge the Stanford Machine Learning Group (stanfordmlgroup.github.io) and the Stanford Center for Artificial Intelligence in Medicine and Imaging ([AIMI.stanford.edu](https://aiml.stanford.edu)) for infrastructure support.

REFERENCES

- [1] Michael David Abramoff, Yiyue Lou, Ali Erginay, Warren Clarida, Ryan Amelon, James C Folk, and Meindert Niemeijer. 2016. Improved Automated Detection of Diabetic Retinopathy on a Publicly Available Dataset Through Integration of Deep Learning. *Invest Ophthalmol Vis Sci* 57, 13 (Oct 2016), 5200–5206. <https://doi.org/10.1167/iovs.16-19964>
- [2] Adrian Brady, Risteárd Ó Laoide, Peter McCarthy, and Ronan McDermott. 2012. Discrepancy and error in radiology: concepts, causes and consequences. *The Ulster medical journal* 81, 1 (2012), 3.
- [3] Lindsay P Busby, Jesse L Courtier, and Christine M Glastonbury. 2018. Bias in radiology: the how and why of misses and misinterpretations. *Radiographics* 38, 1 (2018), 236–247.
- [4] Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria de la Iglesia-Vaya. 2020. PadChest: A large chest x-ray image dataset with multi-label annotated reports. *Medical Image Analysis* 66 (Dec 2020), 101797. <https://doi.org/10.1016/j.media.2020.101797>
- [5] Jacob Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement* 20, 1 (2021/01/13 1960), 37–46. <https://doi.org/10.1177/001316446002000104>

- [6] Dina Demner-Fushman, Marc D Kohli, Marc B Rosenman, Sonya E Shooshan, Laritza Rodriguez, Sameer Antani, George R Thoma, and Clement J McDonald. 2016. Preparing a collection of radiology examinations for distribution and retrieval. *J Am Med Inform Assoc* 23, 2 (Mar 2016), 304–310. <https://doi.org/10.1093/jamia/ocv080>
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805 [cs.CL]
- [8] Jared Dunnmon, Alexander Ratner, Nishith Khandwala, Khaled Saab, Matthew Markert, Hersh Sagreiya, Roger Goldman, Christopher Lee-Messer, Matthew Lungren, Daniel Rubin, and Christopher Ré. 2019. Cross-Modal Data Programming Enables Rapid Medical Machine Learning. arXiv:1903.11101 [cs.LG]
- [9] Bradley Efron and Robert Tibshirani. 1986. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical science* (1986), 54–75.
- [10] Esteban F Gershanik, Ronilda Lacson, and Ramin Khorasani. 2011. Critical finding capture in the impression section of radiology reports. In *AMIA Annual Symposium Proceedings*, Vol. 2011. American Medical Informatics Association, 465.
- [11] Varun Gulshan, Lily Peng, Marc Coram, Martin C. Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, Ramasamy Kim, Rajiv Raman, Philip C. Nelson, Jessica L. Mega, and Dale R. Webster. 2016. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. *JAMA* 316, 22 (1/13/2021 2016), 2402–2410. <https://doi.org/10.1001/jama.2016.17216>
- [12] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. 2018. Densely Connected Convolutional Networks. arXiv:1608.06993 [cs.CV]
- [13] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn Ball, Katie Shpanskaya, Jayne Seekins, David A. Mong, Safwan S. Halabi, Jesse K. Sandberg, Ricky Jones, David B. Larson, Curtis P. Langlotz, Bhavik N. Patel, Matthew P. Lungren, and Andrew Y. Ng. 2019. CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. arXiv:1901.07031 [cs.CV]
- [14] Alistair E. W. Johnson, Tom J. Pollard, Nathaniel R. Greenbaum, Matthew P. Lungren, Chih ying Deng, Yifan Peng, Zhiyong Lu, Roger G. Mark, Seth J. Berkowitz, and Steven Horng. 2019. MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs. arXiv:1901.07042 [cs.CV]
- [15] Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs.LG]
- [16] Matthew B. A. McDermott, Tzu Ming Harry Hsu, Wei-Hung Weng, Marzyeh Ghassemi, and Peter Szolovits. 2020. CheXpert++: Approximating the CheXpert labeler for Speed, Differentiability, and Probabilistic Output. arXiv:2006.15229 [cs.LG]
- [17] Ha Q. Nguyen, Khanh Lam, Linh T. Le, Hieu H. Pham, Dat Q. Tran, Dung B. Nguyen, Dung D. Le, Chi M. Pham, Hang T. T. Tong, Diep H. Dinh, Cuong D. Do, Luu T. Doan, Cuong N. Nguyen, Binh T. Nguyen, Que V. Nguyen, Au D. Hoang, Hien N. Phan, Anh T. Nguyen, Phuong H. Ho, Dat T. Ngo, Nghia T. Nguyen, Nhan T. Nguyen, Minh Dao, and Van Vu. 2021. VinDr-CXR: An open dataset of chest X-rays with radiologist's annotations. arXiv:2012.15029 [eess.IV]
- [18] Luke Oakden-Rayner. 2019. Exploring large scale public medical image datasets. arXiv:1907.12720 [eess.IV]
- [19] Tobi Olatunji, Li Yao, Ben Covington, Alexander Rhodes, and Anthony Upton. 2019. Caveats in Generating Medical Imaging Labels from Radiology Reports. arXiv:1905.02283 [cs.CL]
- [20] Yifan Peng, Xiaosong Wang, Le Lu, Mohammadhadi Bagheri, Ronald Summers, and Zhiyong Lu. 2017. NegBio: a high-performance tool for negation and uncertainty detection in radiology reports. arXiv:1712.05898 [cs.CL]
- [21] Yifan Peng, Shankai Yan, and Zhiyong Lu. 2019. Transfer Learning in Biomedical Natural Language Processing: An Evaluation of BERT and ELMo on Ten Benchmarking Datasets. arXiv:1906.05474 [cs.CL]
- [22] Hieu H. Pham, Tung T. Le, Dat Q. Tran, Dat T. Ngo, and Ha Q. Nguyen. 2020. Interpreting chest X-rays via CNNs that exploit hierarchical disease dependencies and uncertainty labels. arXiv:1911.06475 [eess.IV]
- [23] Nick A. Phillips, Pranav Rajpurkar, Mark Sabini, Rayan Krishnan, Sharon Zhou, Anuj Pareek, Nguyet Minh Phu, Chris Wang, Mudit Jain, Nguyen Duong Du, Steven QH Truong, Andrew Y. Ng, and Matthew P. Lungren. 2020. CheXphoto: 10,000+ Photos and Transformations of Chest X-rays for Benchmarking Deep Learning Robustness. arXiv:2007.06199 [eess.IV]
- [24] Pranav Rajpurkar, Anirudh Joshi, Anuj Pareek, Phil Chen, Amirhossein Kiani, Jeremy Irvin, Andrew Y. Ng, and Matthew P. Lungren. 2020. CheXpediton: Investigating Generalization Challenges for Translation of Chest X-Ray Algorithms to the Clinical Setting. arXiv:2002.11379 [eess.IV]
- [25] Pranav Rajpurkar, Anirudh Joshi, Anuj Pareek, Jeremy Irvin, Andrew Y. Ng, and Matthew Lungren. 2020. CheXphotogenic: Generalization of Deep Learning Models for Chest X-ray Interpretation to Photos of Chest X-rays. arXiv:2011.06129 [eess.IV]
- [26] Pranav Rajpurkar, Chloe O'Connell, Amit Schechter, Nishit Asnani, Jason Li, Amirhossein Kiani, Robyn L Ball, Marc Mendelson, Gary Maartens, Daniël J van Hoving, et al. 2020. CheXaid: deep learning assistance for physician diagnosis of tuberculosis using chest x-rays in patients with HIV. *NPJ digital medicine* 3, 1 (2020), 1–8.
- [27] George Shih, Carol C. Wu, Safwan S. Halabi, Marc D. Kohli, Luciano M. Prevedello, Tessa S. Cook, Arjun Sharma, Judith K. Amorosa, Veronica Arteaga, Maya Galperin-Aizenberg, Ritu R. Gill, Myrna C.B. Godoy, Stephen Hobbs, Jean Jeudy, Archana Laroia, Palmi N. Shah, Dharshan Vummidi, Kavitha Yaddanapudi, and Anouk Stein. 2019. Augmenting the National Institutes of Health Chest Radiograph Dataset with Expert Annotations of Possible Pneumonia. *Radiology: Artificial Intelligence* 1, 1 (2019), e180041. <https://doi.org/10.1148/ryai.2019180041> arXiv:https://doi.org/10.1148/ryai.2019180041
- [28] Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pareek, Andrew Y. Ng, and Matthew P. Lungren. 2020. CheXbert: Combining Automatic Labelers and Expert Annotations for Accurate Radiology Report Labeling Using BERT. arXiv:2004.09167 [cs.CL]
- [29] Siyi Tang, Amirata Ghorbani, Rikiya Yamashita, Sameer Rehman, Jared A. Dunnmon, James Zou, and Daniel L. Rubin. 2020. Data Valuation for Medical Imaging Using Shapley Value: Application on A Large-scale Chest X-ray Dataset. arXiv:2010.08006 [cs.LG]
- [30] Yu-Xing Tang, You-Bao Tang, Yifan Peng, Ke Yan, Mohammadhadi Bagheri, Bernadette A. Redd, Catherine J. Brandon, Zhiyong Lu, Mei Han, Jing Xiao, and Ronald M. Summers. 2020. Automated abnormality classification of chest radiographs using deep convolutional neural networks. *npj Digital Medicine* 3, 1 (2020), 70. <https://doi.org/10.1038/s41746-020-0273-z>
- [31] Linda Wang, Zhong Qiu Lin, and Alexander Wong. 2020. COVID-Net: a tailored deep convolutional neural network design for detection of COVID-19 cases from chest X-ray images. *Scientific Reports* 10, 1 (2020), 19549.
- [32] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers. 2017. 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*. 2097–2106.
- [33] Wenwu Ye, Jin Yao, Hui Xue, and Yi Li. 2020. Weakly Supervised Lesion Localization With Probabilistic-CAM Pooling. arXiv:2005.14480 [cs.CV]
- [34] William J Youden. 1950. Index for rating diagnostic tests. *Cancer* 3, 1 (1950), 32–35.