

CheXpert++: Approximating the CheXpert Labeler for Speed, Differentiability, and Probabilistic Output

Matthew B.A. McDermott

MMD@MIT.EDU
CSAIL, MIT

Tzu Ming Harry Hsu

STMHARRY@MIT.EDU
CSAIL, MIT

Wei-Hung Weng

CKBJIMMY@MIT.EDU
CSAIL, MIT

Marzyeh Ghassemi

MARZYEH@CS.TORONTO.EDU
University of Toronto

Peter Szolovits

PSZ@MIT.EDU
CSAIL, MIT

Abstract

It is often infeasible or impossible to obtain ground truth labels for medical data. To circumvent this, one may build rule-based or other expert-knowledge driven labelers to ingest data and yield silver labels absent any ground-truth training data. One popular such labeler is CheXpert (Irvin et al., 2019), a labeler that produces diagnostic labels for chest X-ray radiology reports. CheXpert is very useful, but is relatively computationally slow, especially when integrated with end-to-end neural pipelines, is non-differentiable so can't be used in any applications that require gradients to flow through the labeler, and does not yield probabilistic outputs, which limits our ability to improve the quality of the silver labeler through techniques such as active learning.

In this work, we solve all three of these problems with CheXpert++, a BERT-based, high-fidelity approximation to CheXpert. CheXpert++ achieves 99.81% parity with CheXpert, which means it can be reliably used as a drop-in replacement for CheXpert, all while being significantly faster, fully differentiable, and probabilistic in output. Error analysis of CheXpert++ also demonstrates that CheXpert++ has a tendency to actually correct errors in the CheXpert labels, with CheXpert++ labels being more often preferred by a clinician over CheXpert labels (when they disagree) on all but one disease task. To further demonstrate the utility of these advantages in this model, we conduct a proof-of-concept active learning study, demonstrating we can improve accuracy on an expert labeled random subset of report sentences by approximately 8% over raw, unaltered CheXpert by using one-iteration of active-learning inspired re-training. These findings suggest that simple techniques in co-learning and active learning can yield high-quality labelers under minimal, and controllable human labeling demands.

1. Introduction

In many non-medical settings, annotating even large datasets is relatively feasible via crowd-sourcing. In medicine, however, producing high-quality labels requires human expertise, and is thus much more expensive and impractical. To compensate for these difficulties, in several recent large radiology datasets (Johnson et al., 2019; Irvin et al., 2019; Bustos et al., 2019; Wang et al., 2017b), researchers have relied on various rule-based, expert-defined labeling systems, most recently the CheXpert Labeler (Irvin et al., 2019), which labels chest X-rays by running their corresponding free-text reports through a rule-based NLP model to predict either positive, negative, or uncertain mentions of 14 radiographic diagnostic categories. CheXpert has been used to define classification labels for two recent chest X-ray datasets (Johnson et al., 2019; Irvin et al., 2019), and has been used to provide a metric for clinical accuracy in radiology report generation (Liu et al., 2019), all in the absence of any ground-truth labels for these tasks.

CheXpert builds on top of NegBio (Peng et al., 2017) which provides regular expression (regex) infrastructure for uncertainty and negation detection. As a result of being a regex-based system, CheXpert is not differentiable, and its outputs are binary labels, not continuous probabilities. CheXpert can also only be run on the CPU, rendering it relatively slow compared to GPU-optimized systems like neural network models. CheXpert has been extremely valuable to the research community, but these three issues (lack of differentiability, probabilistic output, and runtime speed) induce significant hurdles for several exciting research directions.

For example, Liu et al. attempt to use CheXpert to enforce that automatically generated free-text radiology reports are clinically accurate by comparing their generated text to the ground truth in the CheXpert label space. However, as CheXpert is non-differentiable, they were forced to use a reinforcement learning policy gradient solution to optimize through the discrete labeler, a process which induced significant computational cost and, likely added additional instability to the results (Liu et al., 2019). In a different direction, several notable researchers have also questioned the validity of CheXpert’s labels (Oakden-Rayner, 2017), a finding we reproduce here. Were CheXpert to output probabilistic labels, we could attempt to improve these labels by employing active learning strategies leveraging the model’s output uncertainty as a signal of which inputs most required human relabeling, or we could attempt to compensate for these label errors by using a small set of human labels to re-calibrate the model¹ and then weighting our downstream use of these labels by their confidence scores. As it stands, we can do none of these things.

In this work, we solve these problem via CheXpert++, a Bidirectional Encoder Representations from Transformers (BERT; Devlin et al. (2018)) model trained to yield an extremely high-fidelity approximation to CheXpert. CheXpert++ is differentiable, yields probabilistic outputs, and is fully GPU-ready. CheXpert++ is able to match CheXpert labels 99.81% of the time on held-out radiology report sentences, so it can serve as a viable drop-in replacement for CheXpert, all while running approximately 1.8 times faster. Additionally, it is fully differentiable, and thus can be integrated natively into downstream neural pipelines. Finally, it yields probabilistic outputs, and therefore can be used with active learning systems. Unexpectedly, we also find that in the 0.07% of cases where the two systems disagree, CheXpert++ labels are actually more often preferred over CheXpert labels by a clinician, suggesting some inductive bias may be enabling CheXpert++ to correct errors in its own training labels.

1. Recalibration, being only an alignment between 1D inputs and 1D outputs, would require far fewer labels than would producing a set of labels for a new CheXpert

Building on that finding and to demonstrate the utility of this approximator, we additionally build a proof-of-concept active learning system and show that after only one epoch of re-training under solicited clinician annotations, CheXpert++ is able to improve accuracy by 8% in matching true clinician-defined gold standard labels over CheXpert. Note that in a concordant work, released while this paper was under review, Smit et al. has also found that mixed training a BERT based model on CheXpert labels and human labels can offer improvements to overall labeling performance. Whereas our active learning study is a proof-of-concept, their work uses a larger set of gold-standard annotations and should be regarded as a more full-fledged model aimed at improving labeler performance alone (Smit et al., 2020).

In the rest of this work, we will first briefly outline the original CheXpert labeler and the construction, training, and performance of our approximation model CheXpert++, both as an approximation (e.g., in terms of fidelity to CheXpert) and in comparison to CheXpert via clinician evaluation. Next, we walk through in detail some projected use cases of this approximation model that would not be possible with the original CheXpert labeler, to better justify why we believe it is sensible to create such an approximation in the first place, focusing specifically on our proof-of-concept active learning system and the labeling improvements it imparts. Finally, we close with future work and concluding thoughts.

Generalizable Insights about Machine Learning in the Context of Healthcare

Technically, this paper is (intentionally) quite simple. It uses existing models and algorithms, applying them in relatively standard ways. Our goal is not to demonstrate a new method or procedure for collecting labels or for modeling clinical text. Instead, our goal is two-fold:

1. To provide as a resource a drop-in replacement for CheXpert that solves the 3 major hurdles identified above which inhibit novel research directions. Based on our own collective research history, we are confident this solution will prove useful to the community.
2. To challenge the assumption that moving beyond silver labels is generally prohibitive, even if it requires additional annotations.

This second point is our primary source of generalizable insight. Too often in machine learning for healthcare do we settle for sub-par labels before exploring procedures to intelligently and efficiently increase label quality, such as, for example, active learning. Evidence of this is the numerous works using the CheXpert labeler over the MIMIC-CXR dataset despite its known inadequacies.² In this work, with only roughly 3 hours of labeling time by a clinical coauthor and a single pass of fine-tuning, we were able to improve over the raw performance of CheXpert by approximately 8% (averaged across all 14 CheXpert tasks) on a disjoint (also clinician-labeled) test set. The ease with which we obtain this gain is suggestive that these problems are not as insurmountable as they appear, and that machine learning for health, as a field, would be well served investing more in co-learning, model-approximation, and active-learning techniques to efficiently and rapidly improve label quality.

2. To be clear, the authors of this work are not disparaging the CheXpert effort – CheXpert has been extremely valuable to the community in enabling new kinds of research and facilitating new data release. Any purported flaws are opportunities to further benefit the field, not negative reflections on the original effort or execution.

2. From CheXpert to CheXpert++

CheXpert Model Details CheXpert (Irvin et al., 2019) processes a span of text, and produces classification predictions across 14 different disease categories: *Support Devices*, *Airspace Opacity*, *Atelectasis*, *Cardiomegaly*, *Consolidation*, *Edema*, *Enlarged Cardiomediastinum*, *Fracture*, *Lung Lesion*, *No Finding*, *Pleural Effusion*, *Pleural Other*, *Pneumonia*, and *Pneumothorax*. For each disease category, CheXpert classifies the text as either containing *no mention*, an *uncertain mention* (indicating that the physician is uncertain), a *negative mention*, or a *positive mention* of the disease category.

CheXpert is an extension upon NegBio (Peng et al., 2017), a rule-based algorithm for detecting the 14 label categories used with the NIH Chest X-Ray 14 dataset (Wang et al., 2017b). Being largely rule-based, CheXpert is non-differentiable and yields only predictions, not probabilities. Additionally, in practice, we have found it relatively slow and only suitable for offline labeling, even when run with many parallel processes.

NegBio (in particular the labels assigned to the NIH Chest X-Ray 14 dataset) have been previously criticized for their inaccuracy (Oakden-Rayner, 2017), so a valid question is whether these weaknesses extend to CheXpert. Irvin et al. (2019) addressed these concerns by reporting comparisons to expert extracted labels for 1000 held-out radiology reports, finding strong performance for Mention F1 (0.948), and weaker but still strong performance for Negation F1 (0.899) and Uncertainty F1 (0.770). However, when tested on 687 reports from the (different) MIMIC-CXR dataset, Johnson et al. (2019) found significant reductions in performance across all three tasks (0.874, 0.565, 0.470, respectively).

CheXpert++ Model Details CheXpert++ is a BERT based model, initialized from the pre-trained clinical BERT model (Alsentzer et al., 2019), and followed by a multi-task prediction head. We trained CheXpert++ at the per patient and report level over the MIMIC-CXR dataset (Johnson et al., 2019), using an 80/10 train/test data split (602,855 train sentences and 75,748/29,166 test sentences/not-in-train test sentences). Note two nuances to these data; first, that while we do split the data by unique report, due to the frequency of the use of templates in this modality, some individual sentences (which is the level at which we model) are shared across all splits. To account for this, we analyzed our results both on the overall test set and the unseen-sentences level (in which we simply removed the sentences that were also seen in training), finding nearly identical performance in both. Here, to be conservative, we report performances on the unseen-sentences test set, which had minutely worse numbers (overall parity on the full test set was 99.93%, vs. 99.81% on the unseen-sentences set) and imposes a slightly stronger generalization requirement at the expense of imposing some domain shift between the train and test sentences. Secondly, note that the first model trained yielded such strong performance that no hyperparameter tuning was necessary, so while we originally had an additional 10% separated for validation, we did not use it at all in this work. The pre-trained BERT model was fine-tuned using 4 GPUs (GTX TITAN X) for 5 epochs using a batch size of 32 and an initial learning rate of 5×10^{-5} . Full training code and the pre-trained model are available.³

CheXpert++ Model Results

Overall Performance CheXpert++ approximates CheXpert with extremely high fidelity, matching CheXpert labels 99.81% of the time on a held out set. This number is an average accuracy (match %) over all 4 classes for all 14 labels. Each label individually also obtains over 99.4% match

3. <https://github.com/mmcdermott/chexpertplusplus>

Table 1: Failure-to-match rate (i.e., how often does CheXpert++ disagree with CheXpert) for each task, for both a majority class classifier and CheXpert++. Smaller is better.

Task	Majority Class	CheXpert++
No Finding	33.55	0.58
Enlarged Cardiomeastinum	11.70	0.20
Cardiomegaly	9.42	0.21
Lung Lesion	0.96	0.03
Airspace Opacity	8.11	0.19
Edema	6.13	0.13
Consolidation	7.12	0.19
Pneumonia	3.33	0.23
Atelectasis	5.63	0.26
Pneumothorax	13.43	0.14
Pleural Effusion	16.19	0.25
Pleural Other	0.44	0.03
Fracture	1.31	0.05
Support Devices	8.98	0.15

accuracy, and all tasks save “No Finding” obtain over 99.7% match accuracy. Especially when considering that CheXpert itself makes errors at rates far greater than this failure-to-match rate (as evaluated by [Johnson et al. \(2019\)](#); [Irvin et al. \(2019\)](#)), these performance numbers can justify CheXpert++ as a direct replacement for CheXpert. Some readers may question why we don’t just train CheXpert++ directly on the underlying data backing CheXpert. The answer is simple: *such data do not exist*. CheXpert is an expert-defined rule-based model, built and validated on separate, non-public datasets that we do not have access to. Using CheXpert as our training target is therefore very appropriate — by matching at this rate we can confidently use CheXpert++ as a drop in replacement, thereby solving our 3 main pain points, with no significant risk of disagreement all without requiring *any* human-generated labels of any kind. Later in this work, we will discuss a proof-of-concept active learning system where we use an additional, human annotated set we construct to improve CheXpert++ over CheXpert, but up to this point we rely on no human labels whatsoever.

Per-task Performance Breakdown For more detailed performance numbers, Table 1 lists the failure-to-match rate for all 14 labeling tasks, comparing majority class assignment performance (e.g., if one just predicted CheXpert’s most frequent class, how often would this class match the CheXpert labels) to CheXpert++. We also break down performance across all labels for all tasks separately with per-task confusion matrices in the Appendix, Figure 2. We see some patterns emerge among our (very small) number of mis-classifications — in particular, we very rarely mistake a “No Mention” label for a different label, whereas mis-classifications within the three other kinds of mentions are much more common.

Error Analysis In addition to the analyses above, we also had a clinician colleague examine a random subset of up to approximately 46 discrepant examples per task (some tasks had fewer than 46 discrepancies total across the approximately 30,000 held out sentences, as noted in the

Table 2: Fraction of labels preferred by a clinician coauthor when CheXpert and CheXpert++ disagree out of N randomly chosen discrepant sentences. Note that N is capped by the number of disagreements, which in several cases (denoted by a *) our annotator simply annotated all errors across the entire set. Larger means more preferred/better.

Task	N	CheXpert	CheXpert++	Both Wrong
No Finding	44	25%	39%	36%
Enlarged Cardiomeastinum	43	33%	56%	12%
Cardiomegaly	44	23%	61%	14%
Lung Lesion	9*	44%	33%	22%
Airspace Opacity	46	26%	61%	13%
Edema	38*	29%	66%	5%
Consolidation	46	30%	65%	4%
Pneumonia	45	22%	62%	16%
Atelectasis	46	24%	70%	7%
Pneumothorax	41	22%	76%	2%
Pleural Effusion	45	31%	58%	9%
Pleural Other	8*	38%	13%	50%
Fracture	14*	21%	57%	21%
Support Devices	45	29%	62%	7%
Micro Average		27%	60%	12%

table)⁴. Perhaps surprisingly, we found that in a majority of cases (59%), said clinician preferred CheXpert++ labels over those of CheXpert, despite the fact that CheXpert++ was trained to match CheXpert exactly. In 13% of cases, both labels were deemed incorrect, and in only 28% of cases was CheXpert preferred over CheXpert++. A per-task breakdown of these results is in Table 2. We find that CheXpert++ is preferred in all tasks save “Lung Lesion” (where CheXpert labels are preferred) and “Pleural Other,” where most often both are incorrect.

Runtime CheXpert++ is also faster than CheXpert. To label all MIMIC-CXR (Johnson et al., 2019) sentences, CheXpert (using 32 processes) takes approximately 2.75 hours, whereas CheXpert++ (using 1 GPU) takes only approximately 1.53 hours. We believe CheXpert++ could also be optimized further, via neural network acceleration (e.g., Lebedev and Lempitsky (2018); Cheng et al. (2018)).

3. Projected Use Cases of CheXpert++

In addition to the raw speed improvement, which may justify its use directly, we envision several specific use-cases where having a differentiable, probabilistic model in place of CheXpert could be highly beneficial. We outline these at a high-level here, then explore a proof-of-concept system performing active learning using CheXpert++ in Section 4.

Active Learning/Weak Supervision The automatic labeling tool used for Chest X-ray 14 (Wang et al., 2017b) has been criticized for poor accuracy historically (Oakden-Rayner, 2017). The CheXpert system, which builds on this system, may also suffer from similar issues. Given the cost of acquiring

4. Annotator was presented with a sentence, both labels (blinded), and indicated whether they preferred one label to the other, found both labels wrong, or were unsure (used 3 times of 528 total sentences annotated).

additional labels in this domain, techniques like active learning or weak supervision are extremely appealing. Prior works have found both sets of techniques helpful in processing hard-to-label data (Wang et al., 2017a; Wang and Poon, 2018; Halpern et al., 2016; Ratner et al., 2017). Specifically within radiology report classification, Nguyen and Patrick (2014) study found dramatic annotation savings through active learning approaches. However, many active learning approaches require a notion of *confidence* from the underlying classifier, which CheXpert alone cannot provide. As CheXpert++ is fully probabilistic, it can be used natively with such active learning systems. We demonstrate an active learning proof-of-concept system with CheXpert++ in Section 4.

Downstream Project Use Beyond its immediate utility in labeling, other works could also consider using the clinical labels provided by CheXpert or CheXpert++ to enhance their radiograph processing systems in other ways. In our own prior work, for example, used CheXpert to assess the clinical accuracy of machine generated free-text reports for a radiograph. In order to get around the non-differentiable CheXpert labeler, we used a reinforcement learning (RL) policy gradient framework. As CheXpert++ is fully differentiable, works like this could instead adopt CheXpert++ in an end-to-end setting without any RL component, offering both a more direct objective function and a technically simpler solution (Liu et al., 2019). Wang et al. (2018) also included a text-based disease classification branch in their joint embedding learning network. With a pre-trained CheXpert++ in place for such models, we save efforts in learning the text branch. As CheXpert++ is also much faster and can be run on a GPU, it also would yield a much more computationally performant integration than CheXpert does. This ability may help replace some of the more contrived methods of ensuring “clinically accurate” report generation that have been used, such as generating reports based on a training dataset composed exclusively of pre-filled templates (Gale et al., 2018).

4. Active Learning Proof-of-Concept System

In this section, we describe a proof-of-concept active learning system using CheXpert++, which ultimately improved ground-truth labeling performance by approximately 8% (on average across all 14 tasks) as compared to CheXpert. While this system is not a full-fledged replacement labeler, it demonstrates that the active learning strategies enabled by CheXpert++ can facilitate efficient, significant improvement in label quality. In this section, we will first detail how we structured our active learning experiments, then we will detail the results and implications of our findings.

4.1. Active Learning Experiment Setup

There are several key components to our active learning experiments. First, our new source of ground truth for evaluating the system against CheXpert; second, how we selected which examples to relabel for fine-tuning the CheXpert++ system; and, finally, third, how we actually did the re-labeling and re-training.

Held-out Set Construction As stated previously, by default we have no gold-standard labels on these sentences; just the CheXpert defined silver labels. Thus, to enable evaluating our active learning system, we curated a set of 540 sentences for relabeling by 2 clinician colleagues in the following way.

First, for each CheXpert (task, label) pair (e.g., “Edema, no mention”, or “Fracture, positive mention”), we randomly selected ten sentences from the dataset which were assigned that label

on that task by CheXpert. This amounted to 40 sentences for all tasks save “No Finding” (which only permits “Positive Mention” and “Negative Mention”) which received 20 sentences. These 540 sentences served as our held out set, and each clinician annotated all of them, though each sentence was only annotated on its specific task (e.g., “Edema”), to minimize annotation effort. The reason to do this tiered selection is to ensure that our held-out, golden-labeled subset spanned the breadth of conditions present in the reports as much as possible. Given performance numbers of CheXpert on this held-out set, however, we could use the overall propensity of these labels for their specific tasks to calculate the overall expected match rate, so we do not lose any information by doing this tiered selection process, and we gain significant diversity across tasks and expected labels.

When we compute CheXpert’s raw match-rate against these labels, the results are not compelling; CheXpert is accurate only 69.8%/71.6% of the time on our held-out set according to each of our annotators individually. This echoes prior concerns about the validity of these automatic labelers (Oakden-Rayner, 2017). Note that the two annotators also disagreed non-trivially often (their agreement rate was 78.2%), but they agree significantly higher with each other than either alone agrees with CheXpert.

Sample Selection In the interest of providing a simple proof-of-concept for the utility of CheXpert++’s probabilistic outputs towards active-learning, we use a very straightforward selection process here. Namely, we selected, per task, the top 100 most uncertain samples under CheXpert++’s probabilistic output within all reports in our validation set,⁵ and asked our more senior annotator⁶ to (again in a per-task fashion) label all of these sentences, a process that took just over 3 hours. This resulted in 1086 new labeled sentences to use in training, each with only a subset of measured labels.⁷ As we only used one annotator to do this re-training, we also evaluate with that annotator’s held-out set to minimize sources of variance, though performance improvements (albeit much reduced) were observed on average in the other annotator as well despite the previously mentioned inter-annotator variance.

Retraining Procedure We fine-tuned CheXpert++ for one additional epoch only over these 1086 sentences using 3 GPUs (note that using fewer GPUs would also work – the system does not require significant GPU memory), only applying losses for labels that were actually set by our clinician annotator. This process took less than 30 seconds given how small the dataset was. Traditionally, in a full active learning pipeline, we would then solicit another round of “most-uncertain-samples” and repeat this procedure until performance reached sufficient level. Here, for our proof-of-concept, one epoch was enough to establish a significant performance gain.

4.2. Active Learning Results

After this single epoch, performance improved on our new, gold-standard held-out set by approximately 8% (averaged across all tasks) relative to both CheXpert++ with no active-learning and CheXpert. Figure 1 shows the difference in accuracy of both CheXpert++ raw and CheXpert++ post active learning vs. the raw CheXpert. We can see that even this single epoch of active-learning

5. Note that, as sentences are duplicated within reports, this includes a subset of sentences that were in our original training set, but was fully deduplicated from the sentences in our annotated held out set

6. We used just one annotator for the training set given the significantly increased burden

7. Note that this number is lower than the expected 1400 as many sentences are repeated across many radiology reports, and we eliminated these from our setup.

fine-tuning significantly improves performance on the majority of tasks, some quite significant margins.

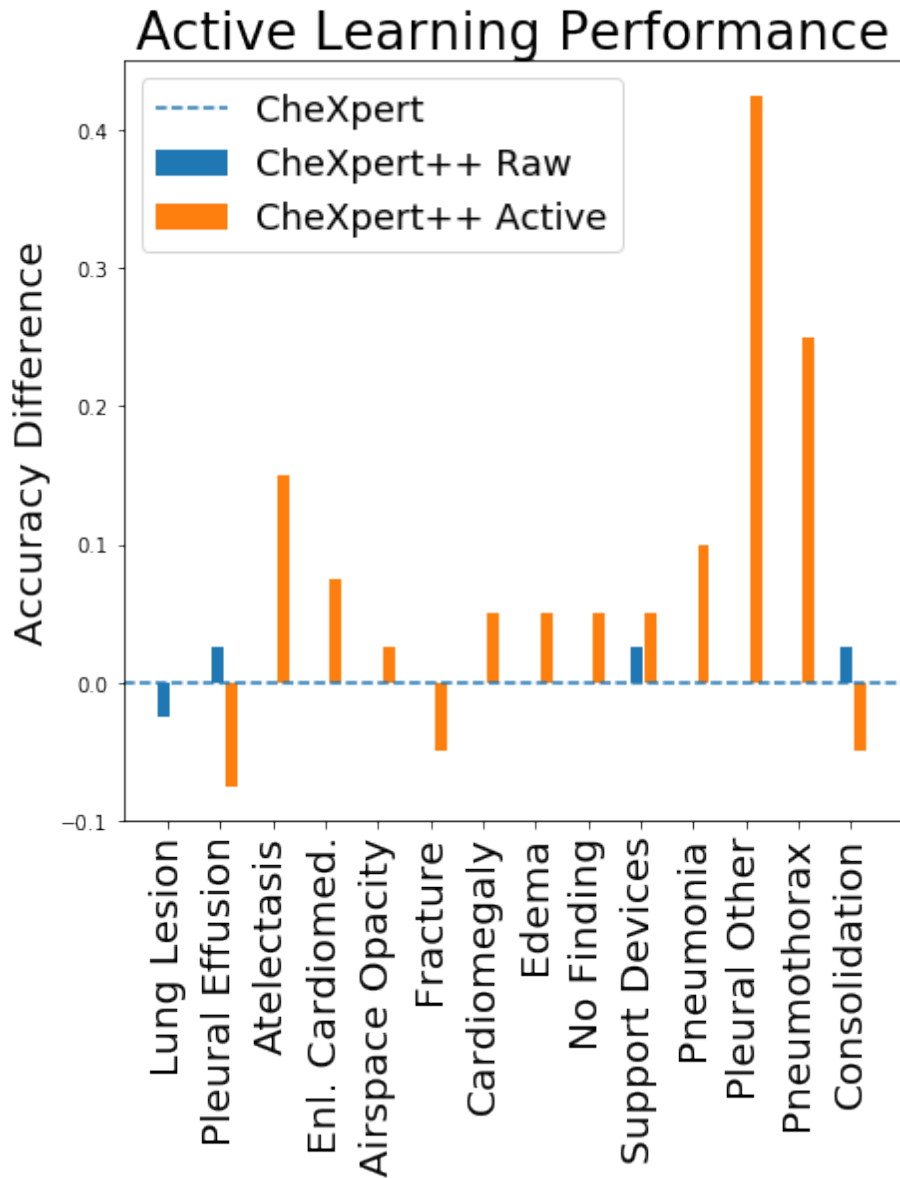


Figure 1: Accuracy on our clinician-labeled held out set of CheXpert++ both with and without active learning minus accuracy of CheXpert on the same held-out set. Average accuracy of CheXpert is 71.6%, average accuracy of CheXpert++ raw is 72.0%, average accuracy of CheXpert++ post active learning is 79.1%. Higher is better.

These results show that with only 3 hours of annotation effort and 1 epoch of fine-tuning, we are able to reap significant label quality improvements via this simple approach.

5. Discussion

In this section, we first discuss several key takeaways from the work, then offer commentary on the next steps for this project. We then progress to concluding thoughts in the next section.

5.1. Key Takeaways

This work demonstrates both the utility of CheXpert++ as a direct replacement for CheXpert in contexts where speed, differentiability, or probabilistic output matter, but also suggest that perhaps we need to place greater emphasis on strategies to intelligently and efficiently improve the quality of silver-labels in medicine. Many works have already used the CheXpert labels across different datasets, including (Wang et al., 2018; Liu et al., 2019; Seyyed-Kalantari et al., 2020), among others, despite known issues relating to label quality. Here, we demonstrate first through the power of the apparent inductive biases of CheXpert++, as it improves on CheXpert’s labels simply by training to match them, and second through our proof-of-concept active learning study, where we improve performance notably with minimal effort, that simple efforts could’ve appeased some of these label quality issues. This suggests we should place greater emphasis on similar strategies in general within machine learning for healthcare.

5.2. Future Work

CheXpert++ has many advantages over CheXpert, but also has several notable opportunities for future improvements. Firstly, it has only been tested on one dataset, MIMIC-CXR. Testing on other datasets would be important for users to have full confidence in using CheXpert++ on novel sources of radiology reports. Secondly, our model only operates at a per-sentence level—we have not investigated any methods for combining our per-sentence outputs into a global per-report set of labels, as CheXpert does. There are a number of strategies that could be explored in future work. Finally, our partial manual annotation studies suggests that CheXpert++’s “errors” are often corrections to CheXpert’s labels, and our active-learning proof-of-concept strongly suggests that a full active-learning effort could likely yield higher quality labels with minimal additional annotator effort. This warrants significant additional analysis; using a full active learning analysis, with more annotators and a more in-depth error resolution procedure to raise the annotator agreement rate would be significant improvements and help transition our proof-of-concept into a full fledged system. Note that Smit et al. (2020) perform a more involved annotation study in this vein in their (co-discovered) BERT based CheXpert improvement system, using 2 board certified radiologists and error resolution policies.

6. Conclusion

In this work, we presented CheXpert++, an extremely high fidelity neural approximation to the expert-labeling system CheXpert. We believe this work demonstrates that CheXpert++ can serve as a faster, fully differentiable drop-in replacement for CheXpert, thereby enabling native integrations with downstream neural pipelines such as that of Liu et al. (2019). Further, as it provides probabilistic output, it can be used with active learning algorithms to refine original label quality for classification purposes. These properties will render CheXpert++ useful in a variety of settings as the community continues to make increasing use of the chest X-ray modality. We demonstrate this likely utility both by outlining several strategies of projects that could benefit from CheXpert++ and via an

active-learning study which improved expert-defined gold standard performance by an average of 8% as compared to CheXpert across all 14 labeling tasks.

7. Acknowledgments

This work would not be possible without the aid of Dr. Catherine McDermott who, in addition to the critical support of coauthor Dr. Wei-Hung Weng in this regard, helped us annotate various examples between CheXpert++ and CheXpert. Additionally, this work is funded in part by National Institutes of Health: National Institutes of Mental Health grant P50-MH106933 as well as a Mitacs Globalink Research Award.

References

- Emily Alsentzer, John R. Murphy, Willie Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew B. A. McDermott. Publicly Available Clinical BERT Embeddings. *arXiv:1904.03323 [cs]*, April 2019. URL <http://arxiv.org/abs/1904.03323>. arXiv: 1904.03323.
- Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria de la Iglesia-Vayá. PadChest: A large chest x-ray image dataset with multi-label annotated reports. *arXiv:1901.07441 [cs, eess]*, January 2019. URL <http://arxiv.org/abs/1901.07441>. arXiv: 1901.07441.
- Y. Cheng, D. Wang, P. Zhou, and T. Zhang. Model Compression and Acceleration for Deep Neural Networks: The Principles, Progress, and Challenges. *IEEE Signal Processing Magazine*, 35(1): 126–136, January 2018. ISSN 1053-5888. doi: 10.1109/MSP.2017.2765695.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv:1810.04805 [cs]*, October 2018. URL <http://arxiv.org/abs/1810.04805>. arXiv: 1810.04805.
- William Gale, Luke Oakden-Rayner, Gustavo Carneiro, Andrew P. Bradley, and Lyle J. Palmer. Producing radiologist-quality reports for interpretable artificial intelligence. *arXiv:1806.00340 [cs]*, June 2018. URL <http://arxiv.org/abs/1806.00340>. arXiv: 1806.00340.
- Yoni Halpern, Steven Horng, Youngduck Choi, and David Sontag. Electronic medical record phenotyping using the anchor and learn framework. *Journal of the American Medical Informatics Association: JAMIA*, 23(4):731–740, 2016. ISSN 1527-974X. doi: 10.1093/jamia/ocw011.
- Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, 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. CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. *arXiv:1901.07031 [cs, eess]*, January 2019. URL <http://arxiv.org/abs/1901.07031>. arXiv: 1901.07031.
- Alistair E. W. Johnson, Tom J. Pollard, Seth J. Berkowitz, Nathaniel R. Greenbaum, Matthew P. Lungren, Chih-ying Deng, Roger G. Mark, and Steven Horng. MIMIC-CXR: A large publicly available database of labeled chest radiographs. *arXiv:1901.07042 [cs, eess]*, January 2019. URL <http://arxiv.org/abs/1901.07042>. arXiv: 1901.07042.

- V. Lebedev and V. Lempitsky. Speeding-up convolutional neural networks: A survey. *Bulletin of the Polish Academy of Sciences. Technical Sciences*, Vol. 66(nr 6), 2018. ISSN 0239-7528. doi: 10.24425/bpas.2018.125927. URL <http://yadda.icm.edu.pl/yadda/element/bwmeta1.element.baztech-3221666b-2abf-4665-a3a8-65b2cc86a375>.
- Guanxiong Liu, Tzu-Ming Harry Hsu, Matthew McDermott, Willie Boag, Wei-Hung Weng, Peter Szolovits, and Marzyeh Ghassemi. Clinically Accurate Chest X-Ray Report Generation. *arXiv:1904.02633 [cs]*, April 2019. URL <http://arxiv.org/abs/1904.02633>. arXiv: 1904.02633.
- Dung H M Nguyen and Jon D Patrick. Supervised machine learning and active learning in classification of radiology reports. *Journal of the American Medical Informatics Association : JAMIA*, 21(5):893–901, September 2014. ISSN 1067-5027. doi: 10.1136/amiajnl-2013-002516. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4147614/>.
- Luke Oakden-Rayner. Exploring the ChestXray14 dataset: problems, December 2017. URL <https://lukeoakdenrayner.wordpress.com/2017/12/18/the-chestxray14-dataset-problems/>.
- Yifan Peng, Xiaosong Wang, Le Lu, Mohammadhadi Bagheri, Ronald Summers, and Zhiyong Lu. NegBio: a high-performance tool for negation and uncertainty detection in radiology reports. *arXiv:1712.05898 [cs]*, December 2017. URL <http://arxiv.org/abs/1712.05898>. arXiv: 1712.05898.
- Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. Snorkel: Rapid training data creation with weak supervision. In *Proceedings of the VLDB Endowment. International Conference on Very Large Data Bases*, volume 11, page 269. NIH Public Access, 2017.
- Laleh Seyyed-Kalantari, Guanxiong Liu, Matthew McDermott, and Marzyeh Ghassemi. Chexclusion: Fairness gaps in deep chest x-ray classifiers. *arXiv preprint arXiv:2003.00827*, 2020.
- Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pareek, Andrew Y Ng, and Matthew P Lungren. Chexbert: Combining automatic labelers and expert annotations for accurate radiology report labeling using bert. *arXiv preprint arXiv:2004.09167*, 2020.
- Hai Wang and Hoifung Poon. Deep Probabilistic Logic: A Unifying Framework for Indirect Supervision. *arXiv:1808.08485 [cs]*, August 2018. URL <http://arxiv.org/abs/1808.08485>. arXiv: 1808.08485.
- Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. Cost-Effective Active Learning for Deep Image Classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(12):2591–2600, December 2017a. ISSN 1051-8215, 1558-2205. doi: 10.1109/TCSVT.2016.2589879. URL <http://arxiv.org/abs/1701.03551>. arXiv: 1701.03551.
- Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M. Summers. ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. *2017 IEEE Conference on Computer*

Vision and Pattern Recognition (CVPR), pages 3462–3471, July 2017b. doi: 10.1109/CVPR.2017.369. URL <http://arxiv.org/abs/1705.02315>. arXiv: 1705.02315.

Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, and Ronald M Summers. Tienet: Text-image embedding network for common thorax disease classification and reporting in chest x-rays. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9049–9058, 2018.

Appendix A. Further Experiments in CheXpert CheXpert++ Parity

Various Transformations of the Confusion Matrices for all Tasks

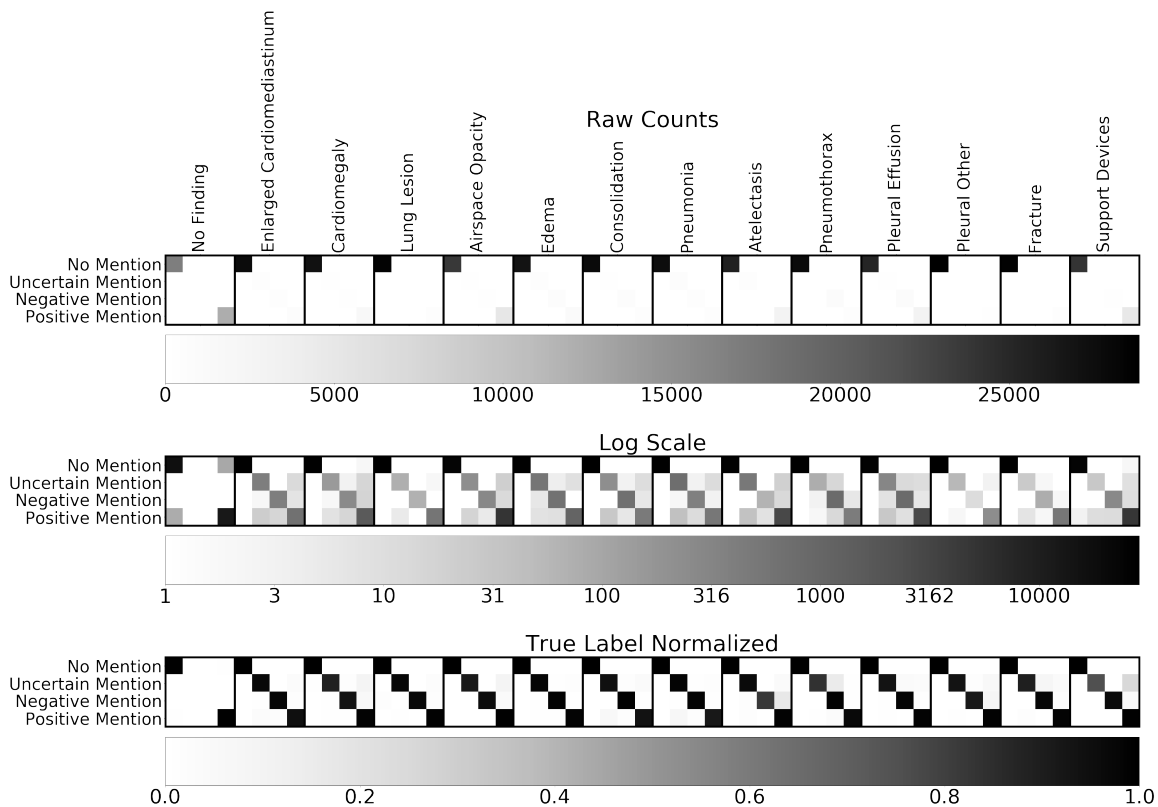


Figure 2: Confusion matrices of CheXpert++ against CheXpert in raw counts, log-scale counts, and reference-label (i.e., row) normalized ratio for all tasks. Note that each of the labeling tasks and the associated confusion matrices are independent. The vertical axis corresponds to the reference (CheXpert) label while the horizontal axis in each matrix corresponds to predictions from CheXpert++. Log scale is added to better illustrate the breakdown of the very small number of mis-classifications observed.