CheXternal: Generalization of Deep Learning Models for Chest X-ray Interpretation to Photos of Chest X-rays and External Clinical Settings

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

Anirudh Joshi*
anirudhjoshi@cs.stanford.edu
Stanford University
USA

Anuj Pareek*
anujpare@cs.stanford.edu
Stanford University
USA

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

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

Figure 1: We measured the diagnostic performance for 8 different chest X-ray models when applied to (1) smartphone photos of chest X-rays and (2) external datasets without any finetuning. All models were developed by different groups and submitted to the CheXpert challenge, and re-applied to test datasets without further tuning.

ABSTRACT

Recent advances in training deep learning models have demonstrated the potential to provide accurate chest X-ray interpretation and increase access to radiology expertise. However, poor generalization due to data distribution shifts in clinical settings is a key

barrier to implementation. In this study, we measured the diagnostic performance for 8 different chest X-ray models when applied to (1) smartphone photos of chest X-rays and (2) external datasets without any finetuning. All models were developed by different groups and submitted to the CheXpert challenge, and re-applied to test datasets without further tuning. We found that (1) on photos of chest X-rays, all 8 models experienced a statistically significant drop in task performance, but only 3 performed significantly worse than radiologists on average, and (2) on the external set, none of the models performed statistically significantly worse than radiologists, and five models performed statistically significantly better than radiologists. Our results demonstrate that some chest X-ray models, under clinically relevant distribution shifts, were comparable to radiologists while other models were not. Future work should investigate aspects of model training procedures and dataset collection

* Authors contributed equally to this research.

that influence generalization in the presence of data distribution shifts.

CCS CONCEPTS

• **Applied computing** → **Health informatics**; • **Computing methodologies** → **Image representations**.

KEYWORDS

Generalizability, Distribution Shifts, Chest X-ray Interpretation, Radiology, Clinical Deployment

ACM Reference Format:

Pranav Rajpurkar, Anirudh Joshi, Anuj Pareek, Andrew Y. Ng, and Matthew P. Lungren. 2021. CheXternal: Generalization of Deep Learning Models for Chest X-ray Interpretation to Photos of Chest X-rays and External Clinical Settings. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '21)*, April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3450439.3451876>

1 INTRODUCTION

Chest X-rays are the most common imaging examination in the world, critical for diagnosis and management of many diseases. With over 2 billion chest X-rays performed globally annually, many clinics in both developing and developed countries lack sufficient trained radiologists to perform timely X-ray interpretation. Automating cognitive tasks in medical imaging interpretation with deep learning models could improve access, efficiency, and augment existing workflows [18, 20, 22, 27]. However, poor generalization due to data distribution shifts in clinical settings is a key barrier to implementation.

First, a major obstacle to clinical adoption of such technologies is in model deployment, an effort often frustrated by vast heterogeneity of clinical workflows across the world [14]. Chest X-ray models are developed and validated using digital X-rays with many deployment solutions relying on heavily integrated yet often disparate infrastructures [1, 12, 13, 17, 21, 25, 26]. One appealing solution to scaled deployment across disparate clinical frameworks is to leverage the ubiquity of smartphones. Interpretation of medical imaging via cell phone photography is an existing “store-and-forward telemedicine” approach in which one or more photos of medical imaging are captured and sent as email attachments or instant messages by practitioners to obtain second opinions from specialists in routine clinical care [7, 31]. Smartphone photographs have been shown to be of sufficient diagnostic quality to allow for medical interpretation, thus leveraging deep learning models in automated interpretation of photos of medical imaging examinations may serve as an infrastructure agnostic approach to deployment, particularly in resource limited settings. However, significant technical barriers exist in automated interpretation of photos of chest X-rays. Photographs of X-rays introduce visual artifacts which are not commonly found in digital X-rays, such as altered viewing angles, variable lighting conditions, glare, moiré, rotations, translations, and blur [19]. These artifacts have been shown to reduce algorithm performance when input images are perceived through a camera [16]. The extent to which such artifacts reduces the performance of chest X-ray models has not been well investigated.

Figure 2: MCC differences of 8 chest X-ray models on different pathologies between photos of the X-rays and the original X-rays with 95% confidence intervals.

A second major obstacle to clinical adoption of chest X-ray models is that clinical deployment requires models trained on data from one institution to generalize to data from another institution [2, 14]. Early work has shown that chest X-ray models may not generalize well when externally validated on data from a different institution and are possibly vulnerable to distribution shift stemming from change in patient population or rely on non-medically relevant cues between institutions [33]. However, the difference in diagnostic performance of more recent chest X-ray models to external datasets has not been investigated.

We measured the diagnostic performance for 8 different chest X-ray models when applied to (1) photos of chest X-rays, and (2) chest X-rays obtained at a different institution. Specifically, we applied these models to a dataset of smartphone photos of 668 X-rays from 500 patients, and a set of 420 frontal chest X-rays from the ChestXray-14 dataset collected at the National Institutes of Health Clinical Center [32]. All models were developed by different groups and submitted to the CheXpert challenge, a large public competition

Figure 3: MCC differences of the same models on photos of chest X-rays compared to radiologist performance with 95% confidence intervals.

for digital chest X-ray analysis [10]. Models were evaluated on their diagnostic performance in binary classification, as measured by Matthew’s Correlation Coefficient (MCC) [3], on the following pathologies selected in Irvin et al. [10]: atelectasis, cardiomegaly, consolidation, edema, and pleural effusion [10].

We found that:

- (1) In comparison of model performance on digital chest X-rays to photos, all 8 models experienced a statistically significant drop in task performance on photos with an average drop of 0.036 MCC. In comparison of performance of models on photos compared to radiologist performance, three out of eight models performed significantly worse than radiologists on average, and the other five had no significant difference.
- (2) On the external set (NIH), none of the models performed statistically significantly worse than radiologists. On average over the pathologies, five models performed significantly better than radiologists. On specific pathologies (consolidation, cardiomegaly, edema, and atelectasis), there were some

Figure 4: Comparison of the average AUC of 8 individual models on photos of chest X-rays compared to on standard images

models that achieved significantly better performance than radiologists.

Our systematic examination of the generalization capabilities of existing models can be extended to other tasks in medical AI, and provide a framework for tracking technical readiness towards clinical translation.

2 METHODS

2.1 Photos of Chest X-rays

We collected a test set of photos of chest x-rays, described in Phillips et al. [19]. In this set, chest X-rays from each CheXpert test study were displayed on a non-diagnostic computer monitor. Chest X-rays were displayed in full screen on a computer monitor with 1920×1080 screen resolution and a black background. A physician was instructed to capture the photos, keeping the mobile camera stable and center the lung fields in the camera view. A time-restriction of 5 seconds per image was imposed to simulate a busy healthcare environment. Subsequent inspection of photos showed that they were taken with slightly varying angles; some photos included artefacts such as Moiré patterns and subtle screen glare. Photos were labeled using the ground truth for the corresponding digital X-ray image. The reference standard on this set was determined using a majority vote of 5 board-certified radiologists. Three separate board-certified radiologists were used for the comparison against the models and all radiologists used the original chest X-ray images for making their diagnoses, rather than the photos.

Metric	Comparison	Average	Pleural Effusion	Edema	Atelectasis	Consolidation	Cardiomegaly
AUC	Photos	0.856 (0.840,0.869)	0.950 (0.932,0.968)	0.917 (0.884,0.943)	0.882 (0.856,0.912)	0.914 (0.865,0.946)	0.921 (0.900,0.940)
	Standard	0.871 (0.855,0.883)	0.960 (0.944,0.975)	0.926 (0.892,0.950)	0.885 (0.858,0.910)	0.918 (0.879,0.948)	0.934 (0.914,0.951)
	Standard - Photos	0.016 (0.012,0.019)	0.011 (0.004,0.019)	0.009 (0.001,0.018)	0.003 (-0.006,0.013)	0.005 (-0.009,0.016)	0.013 (0.006,0.023)
MCC	Photos	0.534 (0.507,0.559)	0.571 (0.526,0.631)	0.556 (0.481,0.639)	0.574 (0.505,0.634)	0.316 (0.246,0.386)	0.580 (0.522,0.630)
	Standard	0.570 (0.543,0.599)	0.621 (0.575,0.670)	0.550 (0.474,0.637)	0.587 (0.529,0.640)	0.336 (0.264,0.418)	0.643 (0.584,0.695)
	Standard - Photos	0.036 (0.024,0.048)	0.049 (0.020,0.070)	-0.006 (-0.039,0.033)	0.012 (-0.016,0.041)	0.020 (-0.011,0.047)	0.063 (0.036,0.084)

Table 1: AUC and MCC performance of models and radiologists on the standard X-rays and the photos of chest X-rays, with 95% confidence intervals.

Comparison	Average	Pleural Effusion	Edema	Atelectasis	Consolidation	Cardiomegaly
Photos	0.534 (0.507,0.559)	0.571 (0.526,0.631)	0.556 (0.481,0.639)	0.574 (0.505,0.634)	0.316 (0.246,0.386)	0.580 (0.522,0.630)
Radiologists	0.568 (0.542,0.597)	0.671 (0.618,0.727)	0.507 (0.431,0.570)	0.548 (0.496,0.606)	0.359 (0.262,0.444)	0.566 (0.511,0.620)
Radiologists - Photos	0.035 (0.009,0.065)	0.099 (0.056,0.145)	-0.049 (-0.136,0.029)	-0.027 (-0.086,0.050)	0.042 (-0.056,0.124)	-0.014 (-0.069,0.029)

Table 2: MCC performance of models on the photos of chest X-rays, radiologist performance, and their difference, with 95% confidence intervals.

2.2 Running Models on New Test Sets

CheXpert used a hidden test set for official evaluation of models. Teams submitted their executable code, which was then run on a test set that was not publicly readable to preserve the integrity of the test results. We made use of the CodaLab platform to re-run these chest X-ray models by substituting the hidden CheXpert test set with the datasets used in this study.

2.3 Evaluation Metrics

Our primary evaluation metric was Matthew’s Correlation Coefficient (MCC), a statistical rate which produces a high score only if the prediction obtained good results in all of the four confusion matrix categories (true positives, false negatives, true negatives, and false positives); MCC is proportionally both to the size of positive elements and the size of negative elements in the dataset [3].

We reported the average MCC of 8 models for five pathologies, namely atelectasis, cardiomegaly, consolidation, edema, and pleural effusion. Additionally, in experiments comparing the models on standard chest X-rays to photos of chest X-rays, we reported the AUC and MCC of the models. In experiments comparing models to board-certified radiologists, we reported the difference in MCC for each of the five pathologies.

3 RESULTS

3.1 Model Performance on Photos of Chest X-rays vs Original X-rays.

3.1.1 Performance Drop In Application To Photos. In comparison of model performance on digital chest X-rays to photos, all eight models experienced a statistically significant drop in task performance on photos with an average drop of 0.036 MCC (95% CI 0.024, 0.048) (See Figure 2, Table 1). All models had a statistically significant drop on at least one of the pathologies between native digital image to photos. One model had a statistically significant drop in

performance on three pathologies: pleural effusion, edema, and consolidation. Two models had a significant drop on two pathologies: one on pleural effusion and edema, and the other on pleural effusion and cardiomegaly. The cardiomegaly and pleural effusion tasks led to decreased performance in five and four models respectively.

3.1.2 Performance on Photos In Comparison to Radiologist Performance on Standard Images. In comparison of performance of models on photos compared to radiologist performance, three out of eight models performed significantly worse than radiologists on average, and the other five had no significant difference (see Figure 3). On specific pathologies, there were some models that had a significantly higher performance than radiologists: two models on cardiomegaly, and one model on edema. Conversely, there were some models that had a significantly lower performance than radiologists: two models on cardiomegaly, and one model on consolidation. The pathology with the greatest number of models that had a significantly lower performance than radiologists was pleural effusion (seven models).

3.1.3 Performance drop in context of radiologist performance. Our results demonstrated that while most models experienced a significant drop in performance when applied to photos of chest X-rays compared to the native digital image, their performance was nonetheless largely equivalent to radiologist performance. We found that although there were thirteen times that models had a statistically significant drop in performance on photos on the different pathologies, the models had significantly lower performance than radiologists only 6 of those 13 times. Comparison to radiologist performance provides context in regard to clinical applicability: several models remained comparable to radiologist performance standard despite decreased performance on photos. Further investigation could be directed towards understanding how different model training procedures may affect model generalization to photos of chest X-rays, and understanding etiologies behind trends for changes in performance for specific pathologies or specific artifacts.

Figure 5: MCC differences in performance of models on the CheXpert test set, with 95% confidence intervals (higher than 0 is in favor of the models being better).

3.1.4 Implication. While using photos of chest X-rays to input into chest X-ray algorithms could enable any physician with a smartphone to get instant AI algorithm assistance, the performance of chest X-ray algorithms on photos of chest X-rays has not been thoroughly investigated. Several studies have highlighted the importance of generalizability of computer vision models with noise in [8]. Dodge and Karam [4] demonstrated that deep neural networks perform poorly compared to humans on image classification on distorted images. Geirhos et al. [6], Schmidt et al. [24] have found that convolutional neural networks trained on specific image corruptions did not generalize, and the error patterns of network and human predictions were not similar on noisy and elastically deformed images.

Figure 6: MCC differences in performance of the same models compared to another set of radiologists across the same pathologies on an external institution’s (NIH) data.

3.2 Comparison of Models and Radiologists on External Institution

We measured the change in diagnostic performance of the same eight chest X-ray models on chest X-rays obtained at a different institution. We applied these models, trained on the CheXpert dataset from the Stanford Hospital, to a set of 420 frontal chest X-rays labeled as part of Rajpurkar et al. [22]. These X-rays are sourced from the ChestXray-14 dataset collected at the National Institutes of Health Clinical Center [32], and sampled to contain at least 50 cases of each pathology according to the original labels provided in the dataset. The reference standard on this set (NIH) was determined using a majority vote of three cardiothoracic subspecialty radiologists; six board-certified radiologists were used for comparison against the models.

3.2.1 Performance on external institution in comparison to radiologists. On the external set (NIH), none of the models performed statistically significantly worse than radiologists (see Figure 6). On

Institution	Comparison	Average	Pleural Effusion	Edema	Atelectasis	Consolidation	Cardiomegaly
CheXpert	Radiologists	0.568 (0.542,0.597)	0.671 (0.618,0.727)	0.507 (0.431,0.570)	0.548 (0.496,0.606)	0.359 (0.262,0.444)	0.566 (0.511,0.620)
	Models	0.570 (0.543,0.599)	0.621 (0.575,0.670)	0.550 (0.474,0.637)	0.587 (0.529,0.640)	0.336 (0.264,0.418)	0.643 (0.584,0.695)
	Models - Radiologists	0.002 (-0.028,0.030)	-0.05 (-0.092,-0.007)	0.043 (-0.033,0.114)	0.039 (-0.029,0.106)	-0.022 (-0.104,0.076)	0.077 (0.040,0.135)
NIH	Radiologists	0.537 (0.515,0.555)	0.642 (0.590,0.690)	0.618 (0.549,0.669)	0.469 (0.423,0.515)	0.455 (0.385,0.509)	0.492 (0.443,0.530)
	Models	0.578 (0.551,0.601)	0.673 (0.605,0.734)	0.662 (0.582,0.742)	0.529 (0.454,0.595)	0.551 (0.499,0.623)	0.517 (0.466,0.567)
	Models - Radiologists	0.041 (0.010,0.072)	0.032 (-0.019,0.078)	0.044 (-0.028,0.124)	0.060 (-0.003,0.126)	0.096 (0.027,0.155)	0.025 (-0.028,0.078)

Table 3: MCC performance of models and radiologists on the CheXpert and NIH sets of chest X-rays, and their difference, with 95% confidence intervals.

Figure 7: Overall change in performance of models (blue) and radiologists (orange) across CheXpert and the external institution dataset (NIH).

average over the pathologies, five models performed significantly better than radiologists. On specific pathologies, there were some models that achieved significantly better performance than radiologists: six models on consolidation, three models on cardiomegaly, four on edema, and two on atelectasis, one on pleural effusion.

3.2.2 Implication. Our finding that these models perform comparably to or at a level exceeding radiologists differs from a previous study which reported that a chest X-ray model failed to generalize to new populations or institutions separate from the training data, relying on institution specific and/or confounding cues to infer the label of interest [33]. Our findings may be attributed to

the improvement in the generalizability of chest X-ray models owing to larger and higher-quality datasets that have been publicly released [10, 11]. Future work should investigate specific aspects of model training and dataset quality and size that lend to these differences, and whether self-supervised training procedures [28] increase generalizability across institutions.

3.2.3 Performance change in context of radiologist performance. Comparing performances on the CheXpert and NIH test sets, we found that on the NIH data set, in 16 instances models had a significantly better performance than radiologists; on the internal CheXpert test set, we observed that in 6 instances, models had a significantly higher performance than radiologists (see Figure 5). This difference may be attributed to a variety of factors including the difference in prevalence of pathologies or the difficulty in identifying them in the external test set compared to the internal set. We are able to contextualize the generalization ability of models to external institutions by comparing their differences to a radiologist performance benchmark, rather than provide a comparison of their absolute performances, which would not control for these possible differences. For instance, when considering cardiomegaly (see Figure 7), we observe a drop in model performance, which in isolation would indicate poor generalizability. However, in light of a similar drop in radiologist performance, we may be able to attribute the difference to differences in difficulties between the two datasets.

4 DISCUSSION

The purpose of this work was to systematically address the key translation challenges for chest X-ray models in clinical application to common real-world scenarios. We found that several chest X-ray models had a drop in performance when applied to smartphone photos of chest X-rays, but even with this drop, some models still performed comparably to radiologists. We also found that when models were tested on an external institution’s data, they performed comparably to radiologists. In both forms of clinical distribution shifts we found that high-performance chest X-ray interpretation models trained on CheXpert produced clinically useful diagnostic performance.

Our work makes significant contributions over another investigation of chest X-ray models [23]. While their study considered the differences in AUC of models when applied to photos of X-rays, they did not (1) compare the resulting performances against radiologists, (2) investigate the drop in performances on specific tasks, or (3) analyze drops in performances of individual models across tasks. Finally, while they compared the performance of models to radiologists on an external dataset, they did not investigate the change in performance of models between the internal dataset and the external dataset.

Strengths of our study include our systematic investigation of generalization performance of several chest X-ray models developed by different teams. Limitations of our work include that our study is still retrospective in nature, and prospective studies would further advance understanding of generalization under distribution shifts. Our systematic examination of the generalization capabilities of existing models can be extended to other tasks in medical AI [5, 9, 15, 29, 30], and provide a framework for tracking technical readiness towards clinical translation.

REFERENCES

- [1] Savvas Andronikou, Kieran McHugh, Nuraan Abdurahman, Bryan Khoury, Victor Mngomezulu, William E Brant, Ian Cowan, Mignon McCulloch, and Nathan Ford. 2011. Paediatric radiology seen from Africa. Part I: providing diagnostic imaging to a young population. *Pediatric radiology* 41, 7 (2011), 811–825.
- [2] David Chen, Sijia Liu, Paul Kingsbury, Sunghwan Sohn, Curtis B. Storie, Elizabeth B. Habermann, James M. Naessens, David W. Larson, and Hongfang Liu. 2019. Deep learning and alternative learning strategies for retrospective real-world clinical data. *npj Digital Medicine* 2, 1 (Dec. 2019), 43. <https://doi.org/10.1038/s41746-019-0122-0>
- [3] Davide Chicco and Giuseppe Jurman. 2020. The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC genomics* 21, 1 (2020), 6.
- [4] Samuel Dodge and Lina Karam. 2017. A Study and Comparison of Human and Deep Learning Recognition Performance under Visual Distortions. In *2017 26th International Conference on Computer Communication and Networks (ICCCN)*. 1–7. <https://doi.org/10.1109/ICCCN.2017.8038465>
- [5] Tony Duan, Pranav Rajpurkar, Dillon Laird, Andrew Y. Ng, and Sanjay Basu. 2019. Clinical Value of Predicting Individual Treatment Effects for Intensive Blood Pressure Therapy: A Machine Learning Experiment to Estimate Treatment Effects from Randomized Trial Data. *Circulation: Cardiovascular Quality and Outcomes* 12, 3 (March 2019). <https://doi.org/10.1161/CIRCOUTCOMES.118.005010>
- [6] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. 2019. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv:1811.12231 [cs, q-bio, stat]* (Jan. 2019).
- [7] Hans Goost, Johannes Witten, Andreas Heck, Dariusch R Hadizadeh, Oliver Weber, Ingo Gräff, Christof Burger, Mareen Montag, Felix Koerfer, and Korosh Kabir. 2012. Image and diagnosis quality of X-ray image transmission via cell phone camera: a project study evaluating quality and reliability. *PLoS One* 7, 10 (2012), e43402.
- [8] Dan Hendrycks and Thomas Dietterich. 2019. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. *arXiv:1903.12261 [cs, stat]* (March 2019).
- [9] Shih-Cheng Huang, Tanay Kothari, Imon Banerjee, Chris Chute, Robyn L Ball, Norah Borus, Andrew Huang, Bhavik N Patel, Pranav Rajpurkar, Jeremy Irvin, et al. 2020. PENet—A scalable deep-learning model for automated diagnosis of pulmonary embolism using volumetric CT imaging. *NPJ digital medicine* 3, 1 (2020), 1–9.
- [10] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana 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. *Proceedings of the AAAI Conference on Artificial Intelligence* 33 (July 2019), 590–597. <https://doi.org/10.1609/aaai.v33i01.3301590>
- [11] 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. 2019. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific Data* 6, 1 (Dec. 2019), 317. <https://doi.org/10.1038/s41597-019-0322-0>
- [12] K. Kallianos, J. Mongan, S. Antani, T. Henry, A. Taylor, J. Abuya, and M. Kohli. 2019. How far have we come? Artificial intelligence for chest radiograph interpretation. *Clinical Radiology* 74, 5 (May 2019), 338–345. <https://doi.org/10.1016/j.crad.2018.12.015>
- [13] Satyananda Kashyap, Mehdi Moradi, Alexandros Karargyris, Joy T. Wu, Michael Morris, Babak Saboury, Eliot Siegel, and Tanveer Syeda-Mahmood. 2019. Artificial intelligence for point of care radiograph quality assessment. In *Medical Imaging 2019: Computer-Aided Diagnosis*, Vol. 10950. International Society for Optics and Photonics, 109503K. <https://doi.org/10.1117/12.2513092>
- [14] Christopher J. Kelly, Alan Karthikesalingam, Mustafa Suleyman, Greg Corrado, and Dominic King. 2019. Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine* 17, 1 (Dec. 2019), 195. <https://doi.org/10.1186/s12916-019-1426-2>
- [15] Amirhossein Kiani, Bora Uyumazturk, Pranav Rajpurkar, Alex Wang, Rebecca Gao, Erik Jones, Yifan Yu, Curtis P. Langlotz, Robyn L. Ball, Thomas J. Montine, Brock A. Martin, Gerald J. Berry, Michael G. Ozawa, Florette K. Hazard, Ryanne A. Brown, Simon B. Chen, Mona Wood, Libby S. Allard, Lourdes Ylagan, Andrew Y. Ng, and Jeanne Shen. 2020. Impact of a deep learning assistant on the histopathologic classification of liver cancer. *npj Digital Medicine* 3, 1 (Dec. 2020), 23. <https://doi.org/10.1038/s41746-020-0232-8>
- [16] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. 2016. Adversarial examples in the physical world. *CoRR abs/1607.02533* (2016). eprint: 1607.02533.
- [17] Paras Lakhani and Baskaran Sundaram. 2017. Deep Learning at Chest Radiography: Automated Classification of Pulmonary Tuberculosis by Using Convolutional Neural Networks. *Radiology* 284, 2 (April 2017), 574–582. <https://doi.org/10.1148/radiol.2017162326>

- [18] Ju Gang Nam, Sunggyun Park, Eui Jin Hwang, Jong Hyuk Lee, Kwang-Nam Jin, Kun Young Lim, Thienkai Huy Vu, Jae Ho Sohn, Sangheum Hwang, Jin Mo Goo, and others. 2018. Development and validation of deep learning-based automatic detection algorithm for malignant pulmonary nodules on chest radiographs. *Radiology* 290, 1 (2018), 218–228. Publisher: Radiological Society of North America.
- [19] Nick A. Phillips, Pranav Rajpurkar, Mark Sabini, Rayan Krishnan, Sharon Zhou, Anuj Pareek, Nguyet Minh Phu, Chris Wang, Andrew Y. Ng, and Matthew P. Lungren. 2020. CheXphoto: 10,000+ Smartphone Photos and Synthetic Photographic Transformations of Chest X-rays for Benchmarking Deep Learning Robustness. arXiv:2007.06199 [eess.IV]
- [20] Chunli Qin, Demin Yao, Yonghong Shi, and Zhijian Song. 2018. Computer-aided detection in chest radiography based on artificial intelligence: a survey. *BioMedical Engineering OnLine* 17, 1 (Aug. 2018), 113. <https://doi.org/10.1186/s12938-018-0544-y>
- [21] Zhi Zhen Qin, Melissa S. Sander, Bishwa Rai, Collins N. Titahong, Santat Sudrun-grot, Sylvain N. Laah, Lal Mani Adhikari, E. Jane Carter, Lekha Puri, Andrew J. Codlin, and Jacob Creswell. 2019. Using artificial intelligence to read chest radiographs for tuberculosis detection: A multi-site evaluation of the diagnostic accuracy of three deep learning systems. *Scientific Reports* 9, 1 (Oct. 2019), 1–10. <https://doi.org/10.1038/s41598-019-51503-3>
- [22] Pranav Rajpurkar, Jeremy Irvin, Robyn L. Ball, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis P. Langlotz, Bhavik N. Patel, Kristen W. Yeom, Katie Shpanskaya, Francis G. Blankenberg, Jayne Seekins, Timothy J. Amrhein, David A. Mong, Safwan S. Halabi, Evan J. Zucker, Andrew Y. Ng, and Matthew P. Lungren. 2018. Deep learning for chest radiograph diagnosis: A retrospective comparison of the CheXNeXt algorithm to practicing radiologists. *PLOS Medicine* 15, 11 (Nov. 2018), e1002686. <https://doi.org/10.1371/journal.pmed.1002686>
- [23] Pranav Rajpurkar, Anirudh Joshi, Anuj Pareek, Phil Chen, Amirhossein Kiani, Jeremy Irvin, Andrew Y. Ng, and Matthew P. Lungren. 2020. CheXpedition: Investigating Generalization Challenges for Translation of Chest X-Ray Algorithms to the Clinical Setting. arXiv:2002.11379 [eess.IV]
- [24] Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. 2018. Adversarially Robust Generalization Requires More Data. In *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.). Curran Associates, Inc., 5014–5026.
- [25] Adam B Schwartz, Gina S Siddiqui, John L Barbieri, Amana F Akhtar, Woojin K Kim, Ryan A Littman-Quinn, Emily S Conant, Narainder D Gupta, Bryan A Pukenas, Parvati H Ramchandani, and et al. 2014. The accuracy of mobile teleradiology in the evaluation of chest X-rays. *Journal of Telemedicine and Telecare* (Oct. 2014). Publisher: Journal of Telemedicine and Telecare.
- [26] 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 (Jan. 2019), e180041. <https://doi.org/10.1148/ryai.2019180041>
- [27] Ramandeep Singh, Mannudeep K. Kalra, Chayanin Nitiwarangkul, John A. Patti, Fatemeh Homayounieh, Atul Padole, Pooja Rao, Preetham Putha, Victorine V. Muse, Amita Sharma, and Subba R. Digumarthy. 2018. Deep learning in chest radiography: Detection of findings and presence of change. *PLoS ONE* 13, 10 (Oct. 2018). <https://doi.org/10.1371/journal.pone.0204155>
- [28] Hari Sowrirajan, Jingbo Yang, Andrew Y. Ng, and Pranav Rajpurkar. 2020. MoCo Pretraining Improves Representation and Transferability of Chest X-ray Models. arXiv:2010.05352 [cs.CV]
- [29] Eric J. Topol. 2019. High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine* 25, 1 (Jan. 2019), 44–56. <https://doi.org/10.1038/s41591-018-0300-7>
- [30] Maya Varma, Mandy Lu, Rachel Gardner, Jared Dunnmon, Nishith Khandwala, Pranav Rajpurkar, Jin Long, Christopher Beaulieu, Katie Shpanskaya, Li Fei-Fei, Matthew P. Lungren, and Bhavik N. Patel. 2019. Automated abnormality detection in lower extremity radiographs using deep learning. *Nature Machine Intelligence* 1, 12 (Dec. 2019), 578–583. <https://doi.org/10.1038/s42256-019-0126-0>
- [31] DJ Vassallo, PJ Buxton, JH Kilbey, and M Trasler. 1998. The first telemedicine link for the British Forces. *Journal of the Royal Army Medical Corps* 144, 3 (1998), 125–130.
- [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] John R. Zech, Marcus A. Badgeley, Manway Liu, Anthony B. Costa, Joseph J. Titano, and Eric Karl Oermann. 2018. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLOS Medicine* 15, 11 (Nov. 2018), e1002683. <https://doi.org/10.1371/journal.pmed.1002683>