

CheXRelNet: An Anatomy-Aware Model for Tracking Longitudinal Relationships between Chest X-Rays

Gaurang Karwande ^{*1}, Amarachi B. Mbakwe ^{*1}, Joy T. Wu^{2,3}, Leo A. Celi^{4,5,6}, Mehdi Moradi⁷, and Ismini Lourentzou¹

¹ Virginia Tech {gaurangajitk,bmamarachi,ilourentzou}@vt.edu

² Stanford Medicine joytywu@stanford.edu

³ IBM Research

⁴ Institute for Medical Engineering and Science, Massachusetts Institute of Technology

⁵ Beth Israel Deaconess Medical Center

⁶ Harvard T.H. Chan School of Public Health lceli@mit.edu

⁷ University of British Columbia moradi@ece.ubc.ca

Abstract. Despite the progress in utilizing deep learning to automate chest radiograph interpretation and disease diagnosis tasks, change between sequential Chest X-rays (CXRs) has received limited attention. Monitoring the progression of pathologies that are visualized through chest imaging poses several challenges in anatomical motion estimation and image registration, *i.e.*, spatially aligning the two images and modeling temporal dynamics in change detection. In this work, we propose **CheXRelNet**, a neural model that can track longitudinal pathology change relations between two CXRs. **CheXRelNet** incorporates local and global visual features, utilizes inter-image and intra-image anatomical information, and learns dependencies between anatomical region attributes, to accurately predict disease change for a pair of CXRs. Experimental results on the Chest ImaGenome dataset show increased downstream performance compared to baselines. Code is available at <https://github.com/PLAN-Lab/CheXRelNet>.

Keywords: Graph Attention Networks · CXR Graph Representations · Chest X-Ray Comparison Relations · Longitudinal CXR Relationships

1 Introduction

Medical imaging research has experienced tremendous growth over the past years, spurred by continuous AI advancements [24,16,11], and in particular in the development of specialized digital devices and neural medical imaging architectures. Chest radiography is one of the most performed diagnostic examinations worldwide. The demand for chest radiography has increased the radiologists' workload. As manually interpreting Chest X-rays (CXRs) and radiology reports

^{*} Equal Contribution, with authors listed in alphabetical order.

can be time-consuming, these challenges contribute to the delays in detecting findings and providing exemplary patient clinical management plans. Though there has been substantial progress in radiology such as disease diagnostics [21,5,19,12], medical image segmentation [15,22,9,20], *etc.*, more complex reasoning tasks remain fairly unexplored. For example, despite significant progress in the application of machine learning in chest radiograph medical diagnosis, detecting longitudinal change between CXRs has attracted limited attention. Yet, understanding whether the patient’s condition has deteriorated or improved is crucial to guide the physician’s decision-making and determine the patient’s clinical management.

Automating this process is a challenging task. At times, differences between x-rays might go undetected, hindering early detection of disease progression that requires an immediate change in treatment plans. Previous work tackles change between longitudinal patient visits and evaluates the severity of diseases at each time point on a continuous scope on osteoarthritis in knee radiographs and retinopathy of prematurity in retinal photographs [14]. Other works target longitudinal disease tracking and outcome prediction severity for COVID-19 pulmonary diseases [13], by calculating a severity score for pulmonary x-rays via computing the Euclidean distance between each of the normal images and the image of interest. In addition, geometric correlation maps have been used to study the CXR longitudinal change detection problem [17], in which feature maps are extracted from CXR pairs and their matching scores are used to generate a geometric correlation map that can detect map-specific patterns showing lesion change. However, these works rely on global image information. To the best of our knowledge, no prior work considers capturing correlations among anatomical regions and findings when modeling change between CXRs. Yet, localizing pathologies to anatomy is critical for the radiologists’ reasoning and reporting process, where correlations between image findings and anatomical regions can help narrow down potential diagnoses.

The development of imaging models that track progress or retrogression between CXRs findings or diseases remains still an open issue. Therefore, in this work, we propose **CheXRelNet**, an anatomy-aware neural model that utilizes information from anatomical regions, learns their intra-image and inter-image dependencies with a graph attention network, and combines the localized region features with global image-level features to accurately capture anatomical location semantics for each finding when performing longitudinal relation comparison between CXR exams for a variety of anatomical findings.

The contributions of this work are summarized as follows: 1) we introduce **CheXRelNet**, an anatomy-aware model for tracking longitudinal relations between CXRs. The proposed model utilizes both local and global anatomical information to output accurate localized comparisons between two sequential CXR examinations, 2) we propose a graph construction to capture correlations between anatomical regions from a pair of CXRs, and 3) we conduct experimental analysis to demonstrate that our proposed **CheXRelNet** model outperforms baselines.




Fig. 1: Graph Construction overview. Detected anatomical regions of interest (ROIs) are fed into a ResNet101 pretrained autoencoder to extract their corresponding visual features, formulating initial node representations $\mathcal{V}_i, \mathcal{V}'_i$ for CXR graphs $\mathcal{G}_i = (\mathcal{V}_i, \mathcal{E})$ and $\mathcal{G}'_i = (\mathcal{V}'_i, \mathcal{E})$. Here, $\mathcal{E}$ is constructed based on intra-image region-disease co-occurrence. Moreover, the nodes of the two graphs are connected via a set $\hat{\mathcal{E}}$ of directed edges indicating inter-image relations.

Finally, 4) we perform transfer learning experiments to test the generalization capabilities of our model across pathologies.

2 Methodology

Let $\mathcal{C} = \{(x_i, x'_i)\}_{i=1}^N$ be the set of CXR image pairs. Each image x_i has k anatomical regions. In addition, each image is associated with a set of labels $\mathcal{Y}_i = \{y_{i,m}\}_{m=1}^M$, $y_{i,m} \in \{0, 1\}$ indicating whether the label for pathology m appears in image x_i or not, and each pair (x_i, x'_i) is associated with a set of labels $\mathcal{Z} = \{z_{i,m}\}_{m=1}^M$, $z_{i,m} \in \{0, 1\}$ indicating whether the pathology m appearing in the image pair has improved or worsened. The goal is to design a model that compares the two images and predicts their labels as accurately as possible for an unseen image pair (x, x') and a wide range of pathologies. This is achieved by utilizing (i) the correlation among anatomical region features from the images x_i, x'_i , i.e., $R = f(x)$ and $R' = f(x')$, $R, R' \in \mathbb{R}^{k \times d}$, where k is the number of anatomical regions, each embedded into a row vector with dimensionality d (extracted by a pretrained feature extractor f) and (ii) the correlation among anatomical regions between the two images in the pair.

Given the initial training set of anatomical region representations $\{(R_i, R'_i)\}_{i=1}^N$, we define a normalized adjacency matrix $A \in \mathbb{R}^{2k \times 2k}$ that captures intra-image and inter-image region correlations. The intra-image correlations corresponding to the two $k \times k$ diagonal blocks of A are constructed based on the region-disease co-occurrence [1], i.e., the number of times two anatomical regions co-occur with the same disease or finding in the set of images $R_i, R'_i, i = 1, \dots, N$. Each

of these $k \times k$ co-occurrence blocks can be computed via the Jaccard similarity

$$J(r_s, r_t) = \frac{1}{M} \sum_{m=1}^M \frac{|\mathcal{Y}_{s,m} \cap \mathcal{Y}_{t,m}|}{|\mathcal{Y}_{s,m} \cup \mathcal{Y}_{t,m}|}. \quad (1)$$

Here, r_s represents an anatomical region, $\mathcal{Y}_s^m$ is the set of disease labels for region r_s and pathology m across all images and $\cap, \cup$ denote the intersection and union over multi-sets. To overcome the shortcomings of the label co-occurrence construction tendency to overfit the training data, a filtering threshold τ is adopted, *i.e.*,

$$A_{st} = \begin{cases} 1 & \text{if } J(R_s, R_t) \geq \tau \\ 0 & \text{if } J(R_s, R_t) < \tau \end{cases}. \quad (2)$$

Here, A_{st} corresponds to an element of one of the two diagonal blocks. We further note that the two diagonal blocks are identical.

The inter-image correlations correspond to the two off-diagonal $k \times k$ blocks of the adjacency matrix A , and they are chosen to indicate a relationship between the same anatomical regions of every pair of images. More precisely, we set $A_{st} = \mathbb{1}\{t = s + k\}$ for $s = 1, \dots, k$. The rationale of this adjacency matrix definition is that A will be associated with every pair (x_i, x'_i) and will capture useful inter-image correlations and local intra-image region-level correlations. More precisely, the upper $k \times k$ diagonal block is associated with image x_i , forming a graph $G_i = (\mathcal{V}_i, \mathcal{E})$ with nodes being the vector representations of the k anatomical regions of image x_i . Similarly, the lower $k \times k$ diagonal block is associated with image x'_i , forming a graph $G'_i = (\mathcal{V}'_i, \mathcal{E})$ as before. Finally, the upper $k \times k$ off-diagonal block indicates a set of edges $\tilde{\mathcal{E}}$ between the same regions of images x_i, x'_i . This graph construction is also depicted in Figure 1.

To capture global and local dependencies between anatomical regions, we utilize a graph attention network (GAT) [25] $Z_i = g(R_i, A) \in \mathbb{R}^{k \times d}$ to update R_i as follows:

$$R_i^{(t+1)} = \alpha_{i,i}^{(t)} W_1^{(t)} R_i^{(t)} + \sum_{j \in \mathcal{N}(i)} \alpha_{i,j}^{(t)} W_1^{(t)} R_j^{(t)}, \quad (3)$$

where $W_1 \in \mathbb{R}^{d \times d}$ is a learned weight matrix, $\mathcal{N}(i)$ denotes the neighborhood of x_i , t is the number of stacked GAT layers, and $\alpha_{i,j}$ are the attention coefficients computed as

$$\alpha_{i,j}^{(t)} = \frac{\exp \left(\text{LeakyReLU} \left(\mathbf{a}^\top \left[W_1^{(t)} R_i^{(t)}; W_1^{(t)} R_j^{(t)} \right] \right) \right)}{\sum_{k \in \mathcal{N}(i) \cup \{i\}} \exp \left(\text{LeakyReLU} \left(\mathbf{a}^\top \left[W_1^{(t)} R_i^{(t)}; W_1^{(t)} R_k^{(t)} \right] \right) \right)} \quad (4)$$

Here, $\mathbf{a}$ is a learned weight vector, and $;$ denotes concatenation. The final region representations are computed by a weighted combination of the neighbor




Fig. 2: Classification module. The constructed graph, *i.e.*, the anatomical regions of interest (ROIs) vector representations and the corresponding adjacency matrix, is passed through a graph attention network that learns ROI inter-dependencies, essentially capturing local ROI information. Global image-level representations extracted from a pretrained ResNet101 autoencoder model are concatenated with the ROI learned representations and are passed through a final dense classification layer. The model is trained end-to-end with a cross-entropy classification loss.

vector representations, scaled by their attention scores

$$R_i^{(t+1)} = \phi \left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(t)} R_j^{(t)} \right), \quad (5)$$

where $\phi(\cdot)$ is a non-linear transformation. Given the past history of the patient, a medical expert has enough information to direct the majority of their focus on a particular region within a CXR. In Figure 1, the node highlighted in red corresponds to the physician-designated focus region $k^* \in [1, k]$ for the particular CXR examination. We extract the node embedding corresponding to the focus region of x'_i for each CXR image pair and forward this embedding to the final dense classification layer. Specifically, for a focus-region $k^* \in [1, k]$, the extracted node embedding $R'_i \in \mathbb{R}^d$ is given by,

$$R'_i = R_i^{(t+1)} \mathbb{1}\{k = k^*\}. \quad (6)$$

To capture global image-level information, each image in pair (x_i, x'_i) is encoded into two d -dimensional vectors by utilizing the pretrained feature extractor f , *i.e.*, $Q_i = f(x_i)$ and $Q'_i = f(x'_i)$, $Q_i, Q'_i \in \mathbb{R}^d$. The final prediction is computed via

$$\hat{y} = [R'_i; Q_i; Q'_i] W_2^T, \quad (7)$$

where $;$ denotes the concatenation of the local region-level and global image-level features, $W_2 \in \mathbb{R}^{3d \times M}$ is a fully connected layer that obtains the label predictions.

Table 1: Dataset Characteristics. # Image Pairs (number of comparison CXR pairs) and # Bboxes (number of bounding boxes) and # Training Pairs (number of training comparison CXR pairs) per pathology label. Each pathology is indexed with a pathology ID (first column).

Pathology ID	Description	# Image Pairs	# Bboxes	# Training Pairs
D1	Lung Opacity	32,524	455,336	22,620
D2	Pleural Effusion	13,122	183,708	9,192
D3	Atelectasis	9,660	135,240	6,922
D4	Enlarged Cardiac Silhouette	1,958	3,916	1,384
D5	Pulmonary Edema/Hazy Opacity	12,090	169,260	8,424
D6	Pneumothorax	2,728	38,192	1,930
D7	Consolidation	3,332	46,648	2,310
D8	Fluid Overload/Heart Failure	674	9,436	132
D9	Pneumonia	3,814	53,396	2,590
All 9 Pathologies	Total	79,902	1,095,132	55,504

The network is trained with a multi-label cross-entropy classification loss

$$L = \frac{1}{N} \sum_{i=1}^N \sum_{m=1}^M z_{i,m} \log(\sigma(\hat{z}_{i,m})) + (1 - z_{i,m}) \log(1 - \sigma(\hat{z}_{i,m})), \quad (8)$$

where σ is the sigmoid function and $\{\hat{y}_i^m, y_i^m\} \in \mathbb{R}^M$ are the model prediction and the ground truth for example x_i , respectively. Figure 2 presents an overview of the model architecture.

3 Experiments

3.1 Dataset

The proposed **CheXRelNet** model is trained and evaluated on the CHEST IMAGENOME dataset [29]. This dataset was generated by locally labeling 242,072 frontal MIMIC-CXR [7] (AP or PA view) automatically through a combination of rule-based text analysis and atlas-based bounding box extraction techniques [28,30]. CHEST IMAGENOME represents the connections of each CXR annotation as an anatomy-centered scene graph, following a radiologist-constructed CXR ontology. The dataset contains 1,256 combinations of relation annotations between 29 CXR anatomical locations and their attributes structured as one scene-graph per image, and about 670,000 localized comparison relations between the anatomical locations across sequential exams. In this work, we utilize the localized comparison relations data that involves cross-image relations for the 9 pathologies. Each comparison relation in the CHEST IMAGENOME dataset consists of the DICOM identifiers of the two CXRs being compared, the particular pathological finding observed in those two CXRs, the anatomical region of interest on which the radiologist’s comparison is focused, and the corresponding comparison label. In addition to comparison relations, the CHEST IMAGENOME dataset also

Table 2: Comparison against baselines (accuracy).

Method	D1	D2	D3	D4	D5	D6	D7	D8	D9	All
Local	0.59	0.53	0.60	0.47	0.56	0.46	0.61	0.47	0.63	0.60
Global	0.67	0.69	0.64	0.74	0.71	0.50	0.65	0.69	0.67	0.67
CheXRelNet	0.67	0.68	0.66	0.75	0.71	0.52	0.67	0.73	0.67	0.68

provides bounding box information for extracting individual anatomical regions from the CXRs, viz. ‘Left Lung’, ‘Cardiac Silhouette’, *etc.* For each of the 242,072 frontal MIMIC-CXRs a list of anatomical regions (bboxes) is provided, as well as the corresponding euclidean coordinates for each bounding box. We utilize these coordinates to crop different anatomical regions within a CXR. There are a total of 122,444 unique comparisons in the dataset, of which 79,902 have at least one of the nine selected pathology labels, in addition to regions detected by the object detection pipeline and the overall comparison relation. For each image, except for those with the pathology label as ‘Enlarged Cardiac Silhouette’, 7 of the most frequently occurring anatomical regions were extracted. For the pathology label ‘Enlarged Cardiac Silhouette’, the dataset provides only one corresponding bounding box. Table 1 shows high-level data statistics.

3.2 Baselines

We compare the **CheXRelNet** model against the following baselines: 1) **Local** model: we utilize a previously proposed siamese network trained on cropped ROIs, encoded with a pre-trained ResNet101 autoencoder and passed through a dense layer and a final classification layer [29]. This model essentially only looks at the corresponding anatomical regions and considers neither global information nor intra-region dependencies. 2) **Global** model: we also design a siamese architecture that encodes the entire CXR as opposed to only the cropped ROIs in the Local model. Apart from the input being a full image rather than an ROI, the model architecture is the same as the Local model. Hence, this baseline incorporates the global information but does not take into consideration the anatomical region of interest nor explicitly models inter-region dependencies. These two siamese models serve as baseline methods to contrast the effectiveness of **CheXRelNet**, which not only is location-aware but can also explicitly model both inter-region and intra-image CXR dependencies.

3.3 Implementation Details

To train each model, we use the train/validation/testing splits and detected ROIs provided by CHEST IMAGENOME. For each image within the comparison pair, we crop the image ROIs and resize them to 224×224 . Each cropped ROI is then embedded into a vector with 2048 dimensions, by utilizing a pre-trained ResNet101 autoencoder [3] that is trained on several medical imaging datasets,

e.g., NIH, CheXpert, and MIMIC datasets, *etc.* [6,7,26]. The same autoencoder is utilized for the baseline models. The co-occurrence matrix threshold is set to 0.5. Our model is a 2-layer graph neural network with 2048 and 1024 neurons per layer, in the first and second layers respectively. There are 5 and 3 multi-head-attentions in each respective layer. The output from the graph attention network is concatenated with the global information and then passed through two dense layers of sizes 768 and 128 respectively. We train the network using Adam [10] optimizer for 200 epochs, with a $0.8e^{-3}$ initial learning rate set [27] and a batch size of 32. To avoid overfitting, we utilize early stopping with 11 patience and gradient clipping that is set to 0.1. In addition, we use 0.5 Dropout [23] and a learning rate decay factor of 0.3 with the patience threshold set to 4. The model is implemented by utilizing the PyTorch [18] and pytorch-geometric [4] deep learning frameworks. The evaluation metric is accuracy and results are reported over six experimental trials.

3.4 Experimental Results

Results are summarized in Table 2. **CheXRelNet** achieves a mean accuracy of 0.683 (SD=0.0024), while the Local model has 0.602 mean accuracy (SD=0.0059) and the Global model has 0.672 mean accuracy (SD=0.0046) over six trials. We observe that the Local model is generally underperforming, and it is most likely limited because it focuses on a specific anatomical region and completely neglects the global information. In contrast, radiologists often take into consideration more than one anatomical region when drawing inferences from CXRs. The Global model is a lot more effective than the Local one, and incorporating global information boosts the prediction accuracy. Yet, the Global model is also limited as it focuses on the entire image but fails to consider the relationships among anatomical regions. We additionally perform statistical significance tests, *i.e.*, an unpaired t-test ($p = 0.049$) and a one-tailed t-test ($p = 0.018$) comparing **CheXRelNet** and the Global baseline. These t-test results verify that the **CheXRelNet** and Global baseline predictions follow distinct distributions and that the improvement in accuracy is significant at $p < 0.05$. Additional experiments w.r.t. model architectures and model capacity (number of parameters) are reported in the supplementary. Overall, **CheXRelNet** improves upon the Global model’s prediction accuracy by modeling the inter-image and intra-image region correlations and attending to the anatomical regions of interest.

We also perform a transfer learning experiment wherein we train **CheXRelNet** on a set of diseases and test performance on a different set of diseases [2,31]. Specifically, we train **CheXRelNet** on a subset of the data with ‘Pneumothorax’, ‘Consolidation’, ‘Fluid Overload/Heart Failure’, ‘Pneumonia’ (D6-D9) pathologies, and test on the following pathology labels that are unseen during training: ‘Lung Opacity’, ‘Pleural Effusion’, ‘Atelectasis’, ‘Enlarged Cardiac Silhouette’, and ‘Pulmonary Edema/Hazy Opacity’ (D1-D5). Results are reported in Table 3. We perform this experiment on individual unseen pathology labels as well as on sets of multiple unseen pathology labels. We observe that our model can generalize well to unseen pathology labels. We can attribute this to the incorporation

Table 3: Transfer learning evaluation against baselines (accuracy). Models are trained on D6-D9 and tested on unseen pathologies (D1-D5). SetA consists of unseen pathologies {D1, D2}. SetB consists of unseen pathology labels, {D3, D4}. Set C consists of all unseen pathology labels {D1,D2,D3,D4,D5}.

Method	D1	D2	D3	D4	D5	SetA	SetB	SetC
Local	0.56	0.49	0.54	0.49	0.55	0.54	0.55	0.54
Global	0.61	0.63	0.60	0.65	0.63	0.61	0.63	0.62
CheXRelNet (ours)	0.64	0.60	0.61	0.68	0.67	0.63	0.64	0.64

of both local and global information during training. The model is learning associations between different anatomical regions and therefore can identify complex bio-markers associated with the progression of pathologies.

4 Conclusion

CXRs are commonly repeatedly requested in the clinical workflow to assess for a myriad of attributes. Diagnosis and monitoring are typically performed through comparisons of sequential CXR images, both in in-patient and outpatient settings. Given a patient with two sequential CXR exams, the goal of this work is to automatically evaluate disease change. To this end, we describe a methodology for localized relation comparisons between CXR images. The proposed **CheXRelNet** fuses global image-level information, local intra-image region-level correlations, and inter-image correlations. Experimental results show that **CheXRelNet** outperforms baselines in both traditional and transfer learning settings. As a result, our method provides necessary components for monitoring the progression of pathologies that are visualized through chest imaging. In the future, we hope to expand our work to model disease progression among several sequential CXRs, incorporate additional temporal context information and physiological data [29,8] and account for the time interval variability found in longitudinal imaging records. Finally, future work can adapt the proposed methodology to other medical imaging tasks, and include interpretability mechanisms.

References

1. Agu, N.N., Wu, J.T., Chao, H., Lourentzou, I., Sharma, A., Moradi, M., Yan, P., Hendler, J.: Anaxnet: Anatomy aware multi-label finding classification in chest x-ray. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 804–813. Springer (2021)
2. Chao, W.L., Changpinyo, S., Gong, B., Sha, F.: An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 52–68. Springer (2016)

3. Cohen, J.P., Viviano, J.D., Bertin, P., Morrison, P., Torabian, P., Guarrera, M., Lungren, M.P., Chaudhari, A., Brooks, R., Hashir, M., Bertrand, H.: TorchXRyVision: A library of chest X-ray datasets and models. In: Medical Imaging with Deep Learning (2022), <https://github.com/mlmed/torchxrayvision>
4. Fey, M., Lenssen, J.E.: Fast graph representation learning with PyTorch Geometric. In: Proceedings of the ICLR Workshop on Representation Learning on Graphs and Manifolds (2019)
5. Guo, R., Passi, K., Jain, C.K.: Tuberculosis diagnostics and localization in chest x-rays via deep learning models. *Frontiers in Artificial Intelligence* p. 74 (2020)
6. Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghighi, B., Ball, R., Shpanskaya, K., et al.: Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). vol. 33, pp. 590–597 (2019)
7. Johnson, A.E., Pollard, T.J., Berkowitz, S.J., et al.: Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data* pp. 1–8 (2019)
8. Karargyris, A., Kashyap, S., Lourentzou, I., Wu, J.T., Sharma, A., Tong, M., Abedin, S., Beymer, D., Mukherjee, V., Krupinski, E.A., et al.: Creation and validation of a chest x-ray dataset with eye-tracking and report dictation for ai development. *Scientific data* **8**(1), 1–18 (2021)
9. Kim, M., Lee, B.D.: Automatic lung segmentation on chest x-rays using self-attention deep neural network. *Sensors* **21**(2), 369 (2021)
10. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proceedings of the International Conference on Learning Representations (ICLR) (2015)
11. Kiradloo, G.: Role of application of artificial intelligence and its importance in the healthcare industry. *International Journal of Advanced Research in Engineering and Technology* **9**(2) (2018)
12. Kobayashi, M., Ishioka, J., Matsuoka, Y., Fukuda, Y., Kohno, Y., Kawano, K., Morimoto, S., Muta, R., Fujiwara, M., Kawamura, N., et al.: Computer-aided diagnosis with a convolutional neural network algorithm for automated detection of urinary tract stones on plain x-ray. *BMC urology* **21**(1), 1–10 (2021)
13. Li, M.D., Arun, N.T., Gidwani, M., Chang, K., Deng, F., Little, B.P., Mendoza, D.P., Lang, M., Lee, S.I., O’Shea, A., et al.: Automated assessment and tracking of covid-19 pulmonary disease severity on chest radiographs using convolutional siamese neural networks. *Radiology: Artificial Intelligence* **2**(4) (2020)
14. Li, M.D., Chang, K., Bearce, B., Chang, C.Y., Huang, A.J., Campbell, J.P., Brown, J.M., Singh, P., Hoebel, K.V., Erdoğan, D., et al.: Siamese neural networks for continuous disease severity evaluation and change detection in medical imaging. *NPJ digital medicine* **3**(1), 1–9 (2020)
15. Maity, A., Nair, T.R., Mehta, S., Prakasam, P.: Automatic lung parenchyma segmentation using a deep convolutional neural network from chest x-rays. *Biomedical Signal Processing and Control* **73**, 103398 (2022)
16. Majkowska, A., Mittal, S., Steiner, D.F., Reicher, J.J., McKinney, S.M., Duggan, G.E., Eswaran, K., Cameron Chen, P.H., Liu, Y., Kalidindi, S.R., et al.: Chest radiograph interpretation with deep learning models: assessment with radiologist-adjudicated reference standards and population-adjusted evaluation. *Radiology* **294**(2), 421–431 (2020)
17. Oh, D.Y., Kim, J., Lee, K.J.: Longitudinal change detection on chest x-rays using geometric correlation maps. In: Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 748–756. Springer (2019)

18. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. In: *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*. pp. 8026–8037 (2019)
19. Rajpurkar, P., Irvin, J., Ball, R.L., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C.P., et al.: Deep learning for chest radiograph diagnosis: A retrospective comparison of the chexnext algorithm to practicing radiologists. *PLoS medicine* **15**(11), e1002686 (2018)
20. Reamaroon, N., Sjoding, M.W., Derksen, H., Sabeti, E., Gryak, J., Barbaro, R.P., Athey, B.D., Najarian, K.: Robust segmentation of lung in chest x-ray: applications in analysis of acute respiratory distress syndrome. *BMC Medical Imaging* **20**(1), 1–13 (2020)
21. Shelke, A., Inamdar, M., Shah, V., Tiwari, A., Hussain, A., Chafekar, T., Mehendale, N.: Chest x-ray classification using deep learning for automated covid-19 screening. *SN computer science* **2**(4), 1–9 (2021)
22. Souza, J.C., Diniz, J.O.B., Ferreira, J.L., da Silva, G.L.F., Silva, A.C., de Paiva, A.C.: An automatic method for lung segmentation and reconstruction in chest x-ray using deep neural networks. *Computer methods and programs in biomedicine* **177**, 285–296 (2019)
23. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* **15**(56), 1929–1958 (2014)
24. Tang, X.: The role of artificial intelligence in medical imaging research. *BJR Open* **2**(1) (2019)
25. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2018)
26. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: 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 (CVPR)*. pp. 2097–2106 (2017)
27. Wilson, D.R., Martinez, T.R.: The need for small learning rates on large problems. In: *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*. vol. 1, pp. 115–119. IEEE (2001)
28. Wu, J., Gur, Y., Karargyris, A., Syed, A.B., Boyko, O., Moradi, M., Syeda-Mahmood, T.: Automatic bounding box annotation of chest x-ray data for localization of abnormalities. In: *Proceedings of the 17th International Symposium on Biomedical Imaging (ISBI)*. pp. 799–803. IEEE (2020)
29. Wu, J.T., Agu, N.N., Lourentzou, I., Sharma, A., Paguio, J.A., Yao, J.S., Dee, E.C., Mitchell, W.G., Kashyap, S., Giovannini, A., et al.: Chest imagenome dataset for clinical reasoning. In: *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track* (2021)
30. Wu, J.T., Syed, A., Ahmad, H., et al.: Ai accelerated human-in-the-loop structuring of radiology reports. In: *Proceedings of the American Medical Informatics Association (AMIA) Annual Symposium* (2020)
31. Xian, Y., Schiele, B., Akata, Z.: Zero-shot learning-the good, the bad and the ugly. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4582–4591 (2017)

5 Supplementary Material

5.1 Qualitative Results

We visualize the model predictions for different pathologies. Figure 3 showcases an input image pair for the pathology label ‘Fluid Overload/ Heart Failure’ where there has been a worsening in the patient’s condition. For this particular pair, the anatomical region of interest (ROI) is ‘Cardiac Silhouette’, which is depicted with a red bounding box. Other anatomical regions that our model takes into consideration when making predictions are shown in green bounding boxes. The previous and current CXRs are named ‘CXr 1’ and ‘CXr 2’, respectively. Upon close inspection, we can see there are subtle changes within the ROI as well as in other parts of the CXR. There is increased haziness in the Left and Right Lungs, and minute changes in the Cardiac Silhouette. Similarly, Figure 4, depicts the input image pair for the pathology label ‘Pneumonia’, and the case where there has been an improvement in the patient’s condition. For this particular pair, the anatomical ROI is ‘Right Lower Lung Zone’. In addition to changes within the ROI, there are significant improvements in the regions ‘Left Mid Lung Zone’ and ‘Left Lower Lung Zone’. The Local model focuses only on the ROI, whereas the Global model focuses only on the entire image. Hence, both of these fail in making correct predictions over these images. Our **CheXRe1Net** model builds associations between various regions and hence is able to factor in the minute changes across the entire anatomy while making predictions.




Fig. 3: Qualitative results for pathology D8, class: Worsened

5.2 Pathology-specific Models

In this experiment, we train models that are specific to a given pathology label. Hence, unlike Table 2, where the model is trained jointly for all nine pathology labels and later tested over each individual pathology, here we train and test




Fig. 4: Qualitative results for pathology D9, class: Improved

Table 4: Pathology-specific comparison of CheXRelNet against baselines.

Method	D1	D2	D3	D4	D5	D6	D7	D8	D9	AVG
Local	0.63	0.55	0.59	0.62	0.68	0.53	0.60	0.45	0.63	0.59
Global	0.68	0.64	0.61	0.69	0.70	0.49	0.59	0.69	0.58	0.63
CheXRelNet (ours)	0.67	0.69	0.61	0.71	0.70	0.49	0.67	0.65	0.65	0.65

Table 5: Comparison against baselines with ‘no change’ samples.

Method	D1	D2	D3	D4	D5	D6	D7	D8	D9	All
Local	0.41	0.37	0.41	0.29	0.37	0.37	0.49	0.29	0.42	0.43
Global	0.45	0.47	0.44	0.48	0.48	0.36	0.47	0.50	0.43	0.45
CheXRelNet (ours)	0.49	0.47	0.44	0.49	0.49	0.36	0.47	0.44	0.47	0.47

our model on the same pathology label. Results are shown in Table 4. From these results, we can infer that CheXRelNet is comparable to or outperforms the Local and Global baselines for five out of nine pathologies. For the pathologies ‘Pleural Effusion’ (D2), ‘Consolidation’ (D7), and ‘Pneumonia’ (D9), the difference in accuracy is greater than or equal to 5%. Compared to the results in Table 2, another interesting observation is that when our model is trained for all pathologies, the classification accuracy is higher. We attribute this observation to the presence of a consistent pattern of disease progression across all pathologies. Hence, the model trained on all pathologies is able to generalize better and is more consistent than the one trained on a single pathology.

5.3 Multi-class Classification with ‘No Change’

Table 5 shows the results when we repeat the original experiment with samples wherein there is no change in the disease progression. Albeit the accuracy drops considerably for all models, CheXRelNet outperforms baselines. Further investigations are required to make the model robust to ‘no change’ samples.

Table 6: Ablation study on model structure and capacity.

Model	Local			Global			CheXRelNet			
	Type A	Type B	Type C	Type A	Type B	Type C	Type A	Type B	Type C	Type D
#Parameters (M)	25.6	34.7	41.4	25.6	34.7	41.4	38.6	27.8	54.3	28.9
Accuracy	0.60	0.64	0.63	0.67	0.67	0.67	0.68	0.68	0.68	0.67

5.4 Ablation Study: Model Architectures and Capacity

We perform an ablation study to investigate if there exists a correlation between the performance and the model capacity (number of trainable parameters). Results are shown in Table 6. For all three models, the architectures named *Type A* are the ones used throughout the study. The architectures named *Type B* and *Type C*, for the Local and Global baselines, have more dense layers and neurons. As for the graph models, *Type B*, *Type C* are the shallower and deeper versions of **CheXRelNet** having 1 and 3 Graph Attention (GAT) layers, respectively, whereas *Type A* is the actual **CheXRelNet** that has 2 GAT layers. *Type D* is the version of our **CheXRelNet** model without attention and in that we replace the GAT layers with simpler Graph Convolution layers. The number of trainable parameters and corresponding accuracy is reported in Table 6. We can infer that the model performance is less influenced by the model capacity and that the graph neural network is the prominent differentiating factor.