

MedShift: identifying shift data for medical dataset curation

Xiaoyuan Guo^a, Judy Wawira Gichoya^{b,c}, Hari Trivedi^{b,c}, Saptarshi Purkayastha^d, Imon Banerjee^{e,f,*}

^aDepartment of Computer Science, Emory University, GA, USA

^bSchool of Medicine, Emory University, GA, USA

^cDepartment of Radiology and Imaging Sciences, Emory University, GA, USA

^dIndiana University-Purdue University Indianapolis, School of Informatics and Computing, IN, USA

^eDepartment of Radiology, Mayo Clinic, Phoenix, AZ, USA

^fSchool of Computing and Augmented Intelligence, Arizona State University, AZ, USA

Abstract

Automated dataset curation in the medical domain has long been demanding as AI technologies are often hungry for annotated data. To curate a high-quality dataset, identifying data variance between the internal and external sources is a fundamental and crucial step as the data distributions from different sources can vary significantly and thus affect the performance of the AI models. However, methods to detect shift or variance in data have not been significantly researched. Challenges to this are the lack of effective approaches to learn dense representation of a dataset by capturing its semantics and difficulties of sharing private data across medical institutions. To overcome the problems, we propose a unified pipeline called *MedShift* to detect the top-level shift samples and thus facilitate the medical curation. Given an internal dataset A as the base source, we first train anomaly detectors for each class of dataset A to learn internal distributions in an unsupervised way. Second, without exchanging data across sources, we run the trained anomaly detectors on an external dataset B for each class. The data samples with high anomaly scores are identified as shift data. To quantify the *shiftiness* of the external dataset, we cluster B 's data into groups class-wise based on the obtained scores. We then train a multi-class classifier on A and measure the *shiftiness* with the classifier's performance variance on B by gradually dropping the group with the largest anomaly score for each class. Additionally, we adapt a dataset quality metric to help inspect the distribution differences for multiple medical sources. We verify the efficacy of *MedShift* with musculoskeletal radiographs (MURA) and chest X-rays datasets from more than one external source. Experiments show our proposed shift data detection pipeline can be beneficial for medical centers to curate high-quality datasets more efficiently. An interface introduction video to visualize our results is available at <https://youtu.be/V3BF0P1sxQE>.

Keywords: Dataset curation, Medical shift data, Anomaly detection, OOD detection, MURA X-ray, Chest X-ray

1. Introduction

*Corresponding author

Email addresses: xiaoyuan.guo@emory.edu (Xiaoyuan Guo), judywawira@emory.edu (Judy Wawira Gichoya), hari.trivedi@emory.edu (Hari Trivedi), saptpurk@iupui.edu (Saptarshi Purkayastha), banerjee.imon@mayo.edu (Imon Banerjee)

Supervised deep learning has been promising in solving various medical image-related tasks, and often requires well-annotated datasets for training, which highly drives the generation of medical datasets by research institutions and hospitals. Many of them have established or plan to establish research data curation services. When building large data collections for usage in training and

validation of machine learning, merely collecting a lot of data is not enough Van Ooijen (2019); Yamoah et al. (2019). It is essential that the quality of the data is sufficient for the intended application in order to obtain valid results Van Ooijen (2019). The medical datasets from different institutions can be heterogeneous and with distribution shifts. Models trained on an internal dataset A from a specific institute may show degraded performance on an external dataset B from other sources due to the possible noisy data, distribution shift and poor-quality data, which are called *shift data* in this paper. Dataset/Distribution shift is a common problem in predictive modelling and present in most practical applications, for reasons ranging from the bias introduced by experimental design the irreproducibility of the testing conditions at training time Quiñonero-Candela et al. (2009); Wang et al. (2021), of which imbalanced data, domain shift, source component shift, may be the most common forms Storkey (2009). The shift data introduces out-of-distribution (OOD) in the dataset, and should account for the performance dropping of well-trained models. Thus, identifying the shift data is crucial for cleaning the datasets and helpful in enhancing the model’s generalization with future training. Unfortunately, it still lacks an effective way to identify the difference for a bunch of datasets from the same medical domain. The main challenge lies in the inaccessibility to external medical datasets. Privacy concerns around sharing personally identifiable information are a major barrier to data sharing in medical research Schütte et al. (2021). To address these privacy concerns, there has been an impressive number of large-scale research collaborations to pool and curate de-identified medical data for open-source research purposes Clark et al. (2013). Nevertheless, most medical data is still isolated and locally stored in hospitals and laboratories due to the worries associated with sharing patient data Van Panhuis et al. (2014). Therefore, an efficient way of external dataset curation/cleaning without sharing data is desired.

To overcome the obstacle, we propose *MedShift*, a pipeline for identifying shift data, which takes advantage of the accessible models trained on the internal dataset to gain the in-distribution knowledge. As observed by Ref. Rabanser et al. (2018), domain-discriminating approaches tend to be helpful for characterizing shifts qualitatively and determining if they are harmful. Therefore,

we utilize unsupervised anomaly detectors to learn the “normality” of in-domain features. Suppose the internal dataset has multiple classes, the feature representation of each class is learnt by an OOD detector. Without sharing the internal dataset with others, the shift data is theoretically under-represented and should be detected by the accessible anomaly detectors as outliers from the external datasets. Since the supervised deep learning suffers from the performance dropping when facing the distribution/dataset shifting, especially when training data and test data are from two sources, two intuitions for example, the *shiftiness* of the identified data can be reflected via the performance variance of a well-trained model. Instead of checking the shift sample one by one, *MedShift* quantifies the *shiftiness* for each class in small groups. Based on the assigned anomaly scores, each class of the external datasets is clustered into multiple groups. Data samples with similar qualities will be grouped together. A multi-class classifier is then trained on the internal dataset and evaluated on the external datasets. Each group of each class in external datasets is gradually dropped in the decreasing order of anomaly scores. Meanwhile, the classification performance on the updated external data is recorded. The corresponding variation in performance, hence, reflects the significance of the distribution shift based on the fact that subtle changes in data distribution may affect the performance of well-trained classifiers. Additionally, we adopt a dataset quality metric (OTDD Alvarez Melis and Fusi (2020)) for helping facilitate the comparison of differences among a series of datasets coming from the same medical domain. We summarize our contributions as follows:

1. We propose an automatic pipeline of identifying shift data for medical data curation applications and evaluating the significance of shift data without sharing data between the internal and external organizations;
2. We employ two unsupervised anomaly detectors to learn the internal distribution and identify samples showing the significant *shiftiness* for external datasets, and compared their performance;
3. We quantify the effects of the shift data by training a multi-class classifier that learns internal domain knowledge and evaluating the classification performance for each sub-group of each class in external domains after dropping the shift data;

4. We adapt a data quality metric to quantify the dissimilarity between the internal and external datasets;
5. We experiment on two pairs of representative medical datasets and show effective qualitative and quantitative results, which prove the usefulness of the suggested pipeline for future medical dataset curation.

2. Methodology

In Section 2.1 and 2.2, we formulate the dataset shift identification problem and introduce the necessary notations. Then, we propose and illustrate the pipeline of shift identification in Section 2.3; we further dive deep in the *shiftiness* evaluation in Section 2.4. To complement, we introduce the details of our anomaly detection architecture used for *MedShift* pipeline in Section 2.5. Additionally we introduce the dataset quality measurement in Section 2.6.

2.1. Problem Statement

In view of the fact that the digital healthcare research is hugely limited by the data sharing and privacy issues because of the regulation imposed by Health Insurance Portability and Accountability Act (HIPPA), *MedShift* aims to overcome the barrier by exploiting the advantage of sharing data quality evaluation models across the organizations and inspects the *shiftiness* of external datasets based on the learnt internal domain.

2.2. Formulation and Notation

Given two datasets D_A and D_B of the same medical domain with the same classes (say c_1, c_2, \dots, c_n , n is the total number of classes) from two intuitions A and B (e.g., a chest X-ray dataset from Emory University D_A and a chest X-ray dataset from Stanford University D_B), let D_A be the internal dataset and D_B be the external dataset. Dataset distribution shift is termed the situation where $P_{D_A}(Y|X) = P_{D_B}(Y|X)$ but $P_{D_A}(X) \neq P_{D_B}(X)$, where Y and X represent the class labels and input data respectively.

Suppose we are given i.i.d. internal data $\{X_{c_i}^A\}_{i=1}^n$ with n classes, and input samples $\{x_j^{Ac_i}\}_{j=1}^{N_{c_i}^A} \subset X_{c_i}^A$ ($N_{c_i}^A$ is the sample number of dataset A 's class c_i) from the internal input distribution, and i.i.d. external data $\{X_{c_i}^B\}_{i=1}^n$ and input samples $\{x_j^{Bc_i}\}_{j=1}^{N_{c_i}^B} \subset X_{c_i}^B$ ($N_{c_i}^B$ is the sample number of dataset

B 's class c_i) from external distribution, the detection of class-wise distribution shift for dataset D_B based on D_A is to identify the anomalous samples $\bar{X}_{c_i}^B \subseteq X_{c_i}^B$. Take D_A class data as in-distribution (ID) data and train machine learning models (e.g. classification models), the models can learn the distribution of D_A 's classes and make predictions $P(y_{c_i}^A|x_{c_i}^A)$ for some targets $y_{c_i}^A$ given data samples $x_{c_i}^A$ for class c_i . Theoretically, given the target model trained on the ID data $X_{c_i}^A$, the predictions over set $X_{c_i}^B - \bar{X}_{c_i}^B$ should produce more relevant results than on the whole set $X_{c_i}^B$.

2.3. Shift Identification

In this section, we introduce the methodology for identification of image data distribution shift to discriminate the poor-quality, noisy and under-represented samples from the external data in an automatic way. The pipeline is built on top of the anomaly detection architecture to leverage the anomaly score as illustrated in the framework in Fig. 1, which involves two separate phases - internal training and test phase. An interesting challenge of shift identification is that the anomaly detectors should be able to identify unknown anomalous patterns of an external dataset without including any anomalous data samples in training since in the real situation, exchanging healthcare data among institutions and manually identifying noisy or anomalous data are not trivial tasks.

During the training phase, only internal data samples and the anomaly detection models (see introductions in Sec. 2.5) are involved. As shown in the left blue part of Fig. 1, a set of anomaly detectors \mathcal{F} s for each targeted categories of D_A are trained on the internal dataset in an unsupervised fashion, considering the unavailability of external data sources. Each class will then obtain a unique OOD detector \mathcal{F}_c . The anomaly detector learns to assign each data item with a specific anomaly score, a higher score means more possibility of being an anomalous data. Notably, the anomaly detectors are trained with accessible internal data, and then shared with the external validation sites.

In the test phase, no internal data will be shared but the trained anomaly detector model with shift identification capability will be exchanged. As represented with pink figures and dotted flows in Fig. 1, each trained anomaly detector is evaluated on each corresponding class of dataset D_B and assigns anomaly scores for the external dataset. To prepare for the *shiftiness* quantification in

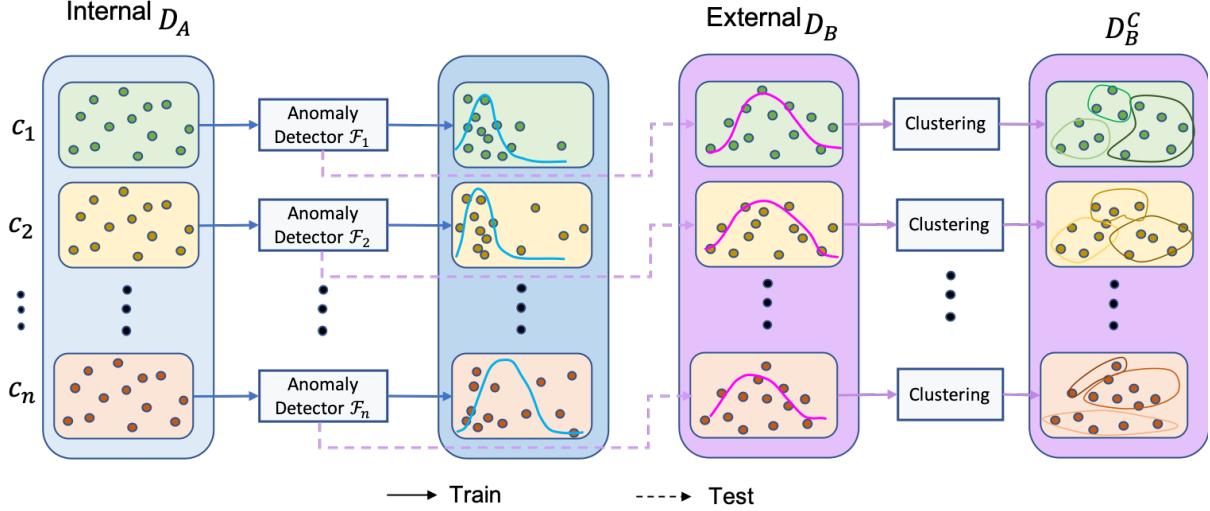


Figure 1: Shift data identification pipeline

Sec. 2.4, an unsupervised clustering algorithm is subsequently applied to each class and clusters the data items into k groups based on the learnt anomaly scores. For each class, the optimal number of cluster k is determined by the Elbow Method. As observed during our experiments, data collected from the same source usually presents similar distributions. Therefore, we keep k as the same across all the classes.

2.4. Shiftness Quantification

The above pipeline can be applied to detect the shift data and assign each data with an anomaly score for indicating its contribution to the dataset shift. Nonetheless, the *shiftness* of the identified data is not simple and straightforward to evaluate in relation with the targeted task. We suggest evaluating them in groups. As prepared in the first stage of the whole pipeline, the clustering has split each class of dataset D_B into multiple groups according to the anomaly scores. For simplicity, we assume that each class has k groups. To evaluate the significance of detected outliers, we train a multi-class classifier \mathcal{G} for D_A and test on D_B . As presented in Fig. 2, we gradually drop one group that has the largest anomaly scores among current groups for each class until only one group

remains. The corresponding class-wise classification performance is recorded. The performance variation thus is an indicator of the *shiftness* of the specific group.

2.5. Anomaly Detection

As claimed in Sec. 2.3, we propose to utilize anomaly detection models to not only identify distribution shifts in the external dataset but also automated cleaning of the external data without any data sharing. First, we briefly describe our anomaly detection model - *Cascade Variational autoencoder-based Anomaly Detector (CVAD)* Guo et al. (2021) used in *MedShift*, which was previously been tested on both generic and medical image datasets. As shown in Fig. 3, CVAD is a self-supervised variational autoencoder-based anomaly detection model which combines latent representation at multiple scales using the cascade architecture of variational autoencoders and thus, can reconstruct the in-distribution image I with high quality. Both the original image I and the reconstruction I' are then fed into a binary discriminator D to separate the synthetic data from the in-distribution ones. The anomaly score includes two parts: the reconstruction error S_{rec} in the first stage and the probability of being the anomaly class S_{dis} in the second stage. To adapt the application

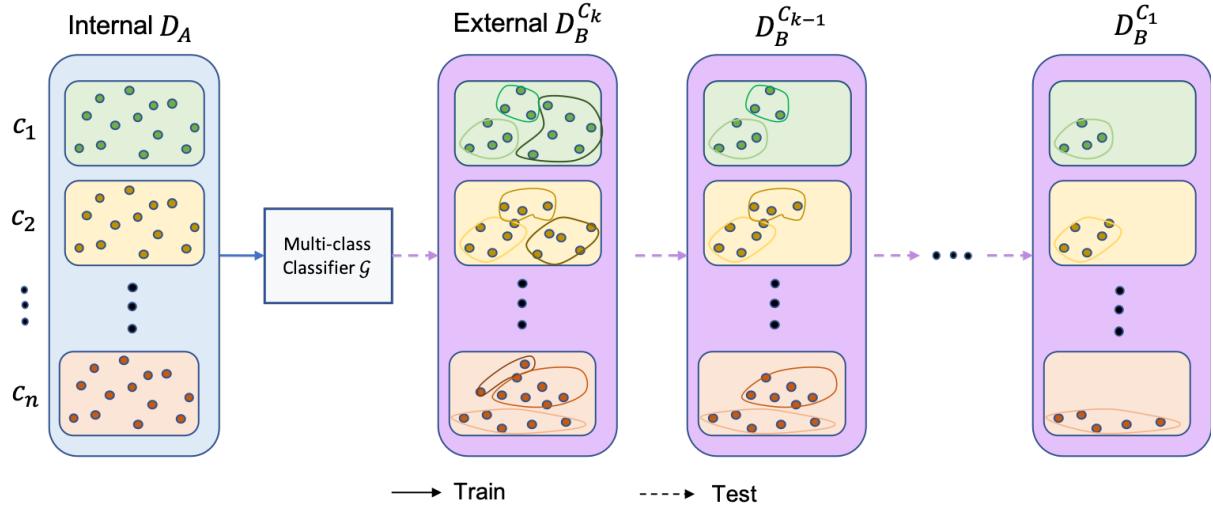
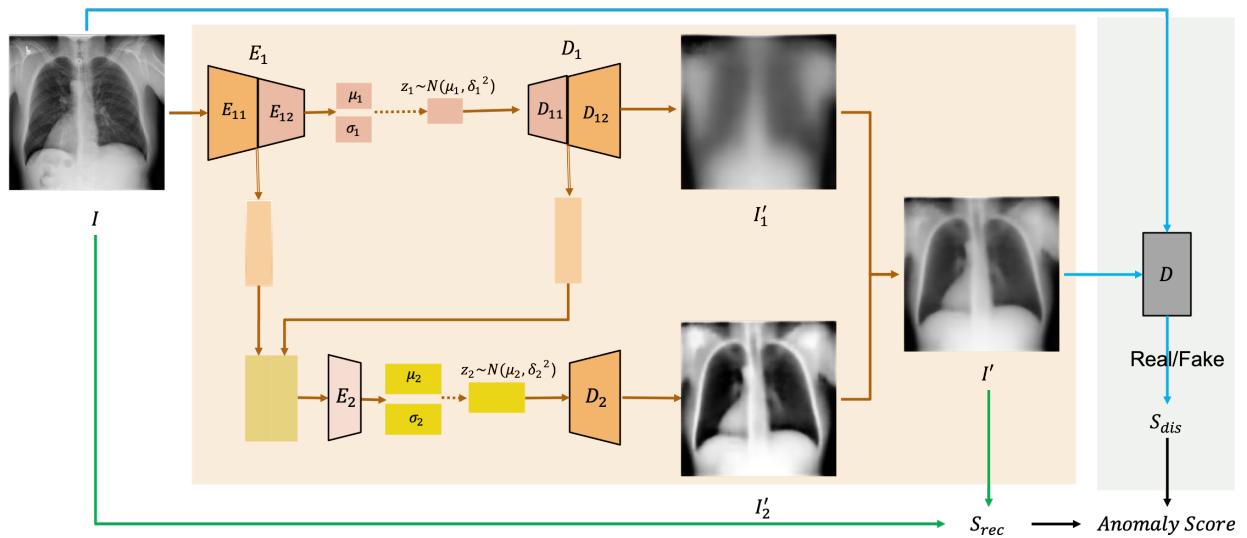


Figure 2: Shiftness quantification pipeline



of detecting abnormal data for multiple unknown external sources, we modified that anomaly score computation by simply adding the two parts together $S = S_{rec} + S_{dis}$. This gives us the advantage that when dealing with heavy noisy data, the reconstruction error will be the dominant indicator for *shiftiness*; when facing the hard distinguished cases the class probability plays the decision role.

As this method poses no assumption on the input data and the applied situations, we utilize this anomaly detection architecture in our pipeline called *MedShift_w_CVAD* across all the experiments. Apart from our anomaly detection model *CVAD*, we also test other anomaly detector f-AnoGAN Schlegl et al. (2019) in *MedShift* for comparison (*MedShift_w_f-AnoGAN* in short).

2.6. Dataset Quality Measurement

To further quantify the efficacy of identifying the shift data among external datasets, we measure the quality of external datasets compared to the internal dataset and observe the difference after removing the shift data from the external sources in an iterative fashion. We apply the Optimal Transport Dataset Distance Alvarez Melis and Fusi (2020) (OTDD) measure to calculating similarities, or distances, between classification datasets. It relies on optimal transport Villani (2009), which is a flexible geometric method for comparing probability distributions, and can be used to compare any two datasets, regardless of whether their label sets are directly comparable. Formally, the optimal transport dataset distance is defined as:

$$OTDD(\mathbf{D}_A, \mathbf{D}_B) = \min_{\pi \in \Pi(P_A, P_B)} \int_{\mathbf{Z} \times \mathbf{Z}} d(z, z') d\pi(z, z') \quad (1)$$

, of which

$$d(z, z') = (d(x, x')^2 + W_2(P_y, P_{y'})^2)^{\frac{1}{2}} \quad (2)$$

, where D_A , D_B are the two datasets, x, x' and y, y' are their samples and labels respectively, W_p denotes the p-Wasserstein distance. Please refer Ref. Alvarez Melis and Fusi (2020) for more details.

3. Experiments

3.1. Datasets

There are two categories of medical datasets used in this paper: (1) *Musculoskeletal radiographs* - Emory MURA dataset (internal) and Stanford MURA dataset Rajpurkar et al. (2017) (external); (2) *Chest radiographs* - Emory Chest X-rays (internal, Emory_CXR in short), CheXpert dataset Irvin et al. (2019) (external_1) and MIMIC dataset Johnson et al. (2019) (external_2).

MURA (musculoskeletal radiographs) is a large dataset of bone X-rays. Each MURA dataset has seven classes, *XR_HAND*, *XR_FORARM*, *XR_FIGER*, *XR_SHOULDER*, *XR_ELBOU*, *XR_WRIST*, *XR_HUMERUS*. Image examples are illustrated in Fig. 4a for each class. To demonstrate the effectiveness of detecting shift data, we have Emory MURA and Stanford MURA datasets as a pair and treat Emory MURA as the internal dataset with Stanford MURA as the external one. More class-wise details of the datasets are presented in the upper of Table. 1.

For chest X-ray, we used three datasets - Emory_CXR (199,029 training and 12,873 test images retrieved from Emory Healthcare system), CheXpert and MIMIC datasets. The bottom part of Table. 1 shows the details of the three datasets. The chest X-ray datasets have 14 classes (or diagnosis) in total. The classes are *No Finding*, *Enlarged Cardiomediastinum*, *Cardiomegaly*, *Lung Lesion*, *Lung Opacity*, *Edema*, *Consolidation*, *Pneumonia*, *Atelectasis*, *Pneumothorax*, *Pleural Effusion*, *Pleural Other*, *Fracture*, *Support Devices*. Image examples are displayed in Fig. 4b. Different from the MURA dataset where class labels are mutually exclusive, each chest X-ray data may have multiple common diagnoses.

3.2. Implementation Details

We implement the pipeline using Pytorch 1.5.0, Python 3.7.3 and Cuda compilation tools V10.0.130 on a machine with 4 NVIDIA RTX A6000 GPUs with 48 GB memory. More details about the training of anomaly detectors and classifiers are introduced below.

3.2.1. Anomaly Detectors

We resize all the medical images to 256×256×channel for simplicity considering the irregular image sizes. To train CVAD, we use the Adam optimizer with a batch size

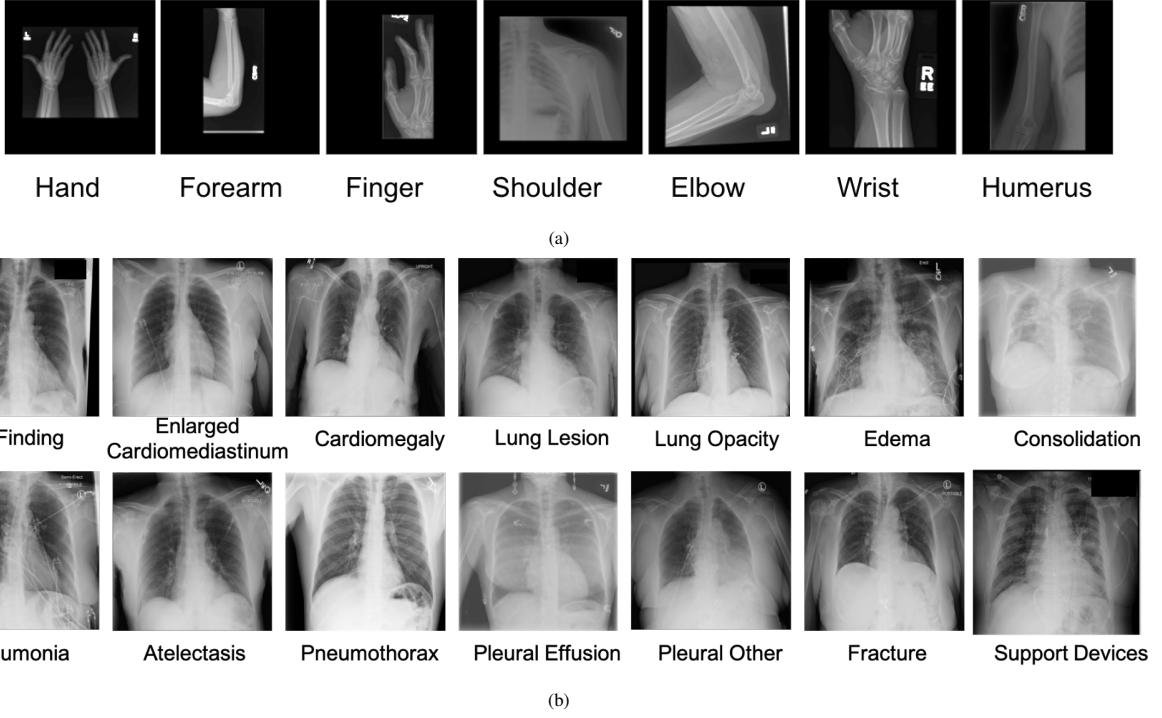


Figure 4: Sample images from the datasets: (a) MURA examples for each anatomical joint class. (Intensity contrasts are changed for better visualization); (b) chest X-ray examples for each class. (Image are resized for better visualization);

	HAND		FOREARM			FINGER		SHOULDER			ELBOW			WRIST		HUMERUS	
Emory_MURA	2,473	(21.33%)		368	(3.17%)		368	(3.17%)	3,451	(29.77%)	1,521	(13.12%)	2,858	(24.65%)	553	(4.77%)	
Stanford_MURA	3,851	(17.94%)		1,097	(5.11%)		3,660	(17.05%)	5,621	(26.18%)	2,397	(11.16%)	3,993	(18.60%)	852	(3.97%)	
	No Finding	Enlarged Cardiome diastinum	Cardiomegaly	Lung Lesion	Edema	Consolidation	Pneu- monia	Atele- ctasis	Pneu- thorax	Pleural Effusion	Pleural Other	Fracture	Support Devices				
Emory_CXR (train)	57,973 (11.35%)	7,825 (1.53%)	27,019 (5.29%)	6,157 (1.21%)	64,439 (12.62%)	22,540 (4.41%)	6,906 (1.35%)	9,188 (1.80%)	66,150 (12.95%)	11,550 (2.26%)	51,828 (10.15%)	2,325 (0.46%)	2,114 (0.41%)	174,768 (34.22%)			
Emory_CXR (test)	7,962 (30.44%)	523 (2.00%)	1,256 (4.80%)	397 (1.52%)	2,141 (8.18%)	475 (1.82%)	151 (0.58%)	439 (1.68%)	1,684 (6.44%)	150 (0.57%)	711 (2.72%)	98 (0.37%)	177 (0.68%)	9,995 (38.21%)			
CheXpert	22,381 (4.34%)	10,798 (2.09%)	27,000 (5.24%)	9,186 (1.78%)	105,581 (20.48%)	52,246 (10.13%)	14,783 (2.87%)	6,039 (1.17%)	33,376 (6.47%)	19,448 (3.77%)	86,187 (16.72%)	3,523 (0.68%)	9,040 (1.75%)	116,001 (22.50%)			
MIMIC	143,352 (22.62%)	84,073 (13.26%)	76,957 (12.14%)	76,423 (12.06%)	65,047 (10.26%)	64,346 (10.15%)	36,564 (5.77%)	26,222 (4.14%)	14,675 (2.32%)	14,257 (2.25%)	10,801 (1.70%)	10,042 (1.58%)	7,605 (1.20%)	3,460 (0.55%)			

Table 1: Dataset details, with total image number and the percentage (in brackets) of each class presented. Upper part of the table present the MURA datasets and the lower is for Chest X-ray datasets.

of 256 and 2,048 for MURA and chest X-ray dataset, respectively; we set the learning rate of 1×10^{-5} and 1×10^{-3} for the generator and the discriminator of CVAD, respectively; we train the generator with 250-500 epochs and the discriminator with 10-20 epochs.

To train f-AnoGAN, we use the default Adam optimizer with a learning rate of 2×10^{-4} and the same batch sizes as CVAD for the corresponding datasets; we run the generative adversarial training for 1000-1500 epochs and the encoder training for 300-500 epochs.

3.2.2. Multi-class Classifiers

To quantify the *shiftiness* of each clustered group for each class of external dataset D_B , we first train a multi-class classifier \mathcal{G} for the internal dataset D_A . The classifier learns the class latent features of the internal domain and is able to predict class labels for test data. For MURA data, we train ResNet152 He et al. (2016) on the Emory MURA dataset with the publicly available pre-trained weights as initialization. We optimize the classifier using the Adam optimizer with a batch size of 512, a learning rate of 1×10^{-3} for 50 epochs. For chest X-ray data, we utilize the model proposed by Ref. Yuan et al. (2021), which originally aims for multi-label classification of the CheXpert dataset, and modifies it for the Emory_CXR 14-class classification task. Following the same implementations in Ref. Yuan et al. (2021), we use DenseNet121 Huang et al. (2017) as the feature extraction backbone and initialize it with the public pretrained model weights. We train the classifier with a batch size of 256 for 20 epochs. The corresponding classification performances, including the *Precision*, *Recall*, *F1-score* and *AUC* score are reported in Sec. 3.3.3.

3.3. Results

In this section, we evaluate the performance of our pipeline on three objectives - (i) shift data identification, (ii) shift data partition and (iii) shift data significance evaluation. To increase readability, one representative class is selected for explanation and more results about other classes are supplemented in Appendix A and Appendix B.

3.3.1. Shift Identification with Anomaly Detection

In the process of identifying the shift data from the external source, each class of the internal dataset will obtain

its own anomaly detector. Figure 5 presents the anomaly score distributions of the representative class from both MURA and Chest X-ray obtained by *MedShift_w_CVAD* architecture. The X-axis represents the anomaly score and Y-axis stands for the number of images that have anomaly scores in the corresponding range. In both cases, Emory data is considered as internal data.

For **MURA** dataset, the anomaly score distribution for ***XR_HAND*** is shown in the left of Fig. 5, with the blue curve for Emory ***XR_HAND*** and the orange distribution curve for Stanford ***XR_HAND*** data. As can be observed, the peaks of the two distributions are clearly separated, the Stanford data generally gets higher OOD scores than the internal Emory data. The difference between the internal and external anomaly score distributions can be easily observed. The closer and more similar the two distributions are, the less shift the external dataset has.

The similar phenomenon can also be seen in chest X-ray data when being tested on two external datasets. For **chest X-ray** dataset, the OOD detection for *Consolidation* is shown in the right of Fig. 5, with the blue histogram and curve for internal Emory_CXR dataset, the orange for CheXpert dataset and the green for MIMIC dataset. The differences in the distributions reflect how different the external chest X-ray data is from the internal domain. MIMIC *Consolidation* has less overlapping with the internal Emory_CXR compared with the CheXpert distribution, which indicates that MIMIC contains more shift *Consolidation* data than the CheXpert dataset.

3.3.2. Shift Data Clustering Results

In this section, we showcase the clustering results based on anomaly scores for both MURA and chest X-ray datasets. Specifically, Stanford MURA dataset, CheXpert and MIMIC data are clustered into different groups according to their anomaly scores obtained in the previous step. The selection of group numbers is decided by the Elbow distortion curves, please refer to Appendix A.3 and Appendix A.4 for the curve details.

An ***XR_HAND*** example of **MURA** dataset is shown in the left of Fig. 6. There are 5 cluster groups in total, with each row representing one cluster. The groups are sorted in ascending order, namely, the top row is with the lowest anomaly scores and the bottom has the largest anomaly scores. For better understanding, their corresponding scores are labelled on top of each example

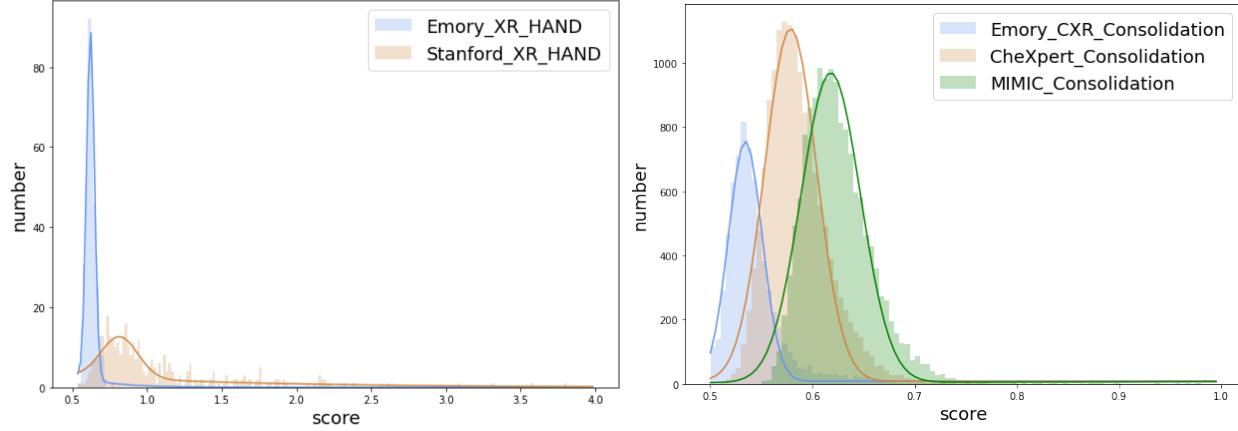


Figure 5: *MedShift-w_CVAD* example results of shift identification with anomaly detection - (left) anomaly score distributions on MURA *HAND*; (right) anomaly score distributions for chest X-ray *Consolidation*. Distributions are truncated on samples with large anomaly scores for better visualization.

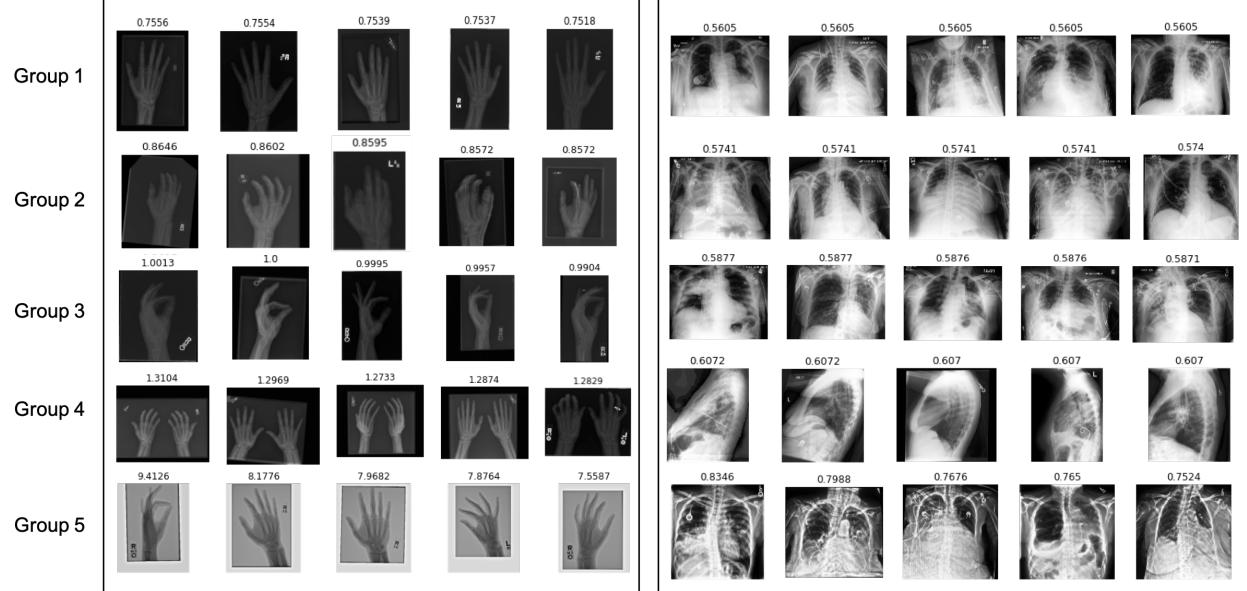


Figure 6: *MedShift-w_CVAD* examples of clustering results - (left) clustering results on Stanford_MURA *HAND* data; and (right) clustering results on CheXpert *Consolidation* data. Each row represents one group with five example images. The groups are illustrated in ascending order based on the anomaly scores from top to bottom. The corresponding anomaly score is on top of each image.

item. As can be observed, the hand data gradually shows more and more variations in terms of image quality, positioning, and noise, as the anomaly score becomes large, especially when comparing the group 1 (first row with lowest anomaly score) to group 5 (last row with highest anomaly score). The variance exhibiting in the abnormal data indicates the existence of distribution shift in the external dataset. Nonetheless, the significance of the detected under-represented/shift data samples in affecting deep learning models' prediction/classification remains to be explored. Similarly, an example of **chest X-ray Consolidation** is presented in the right of Fig. 6. Following the same arrange order, the difference for each group can be clearly captured.

3.3.3. Classification Results for Shiftness Evaluation

As introduced in Sec. 3.2.2, a multi-class classifier has to be trained on the internal dataset to quantify the effect of removing the *shiftness* of external datasets for the two targeted classification tasks. In this section, we report the classification training and testing performance on the internal dataset, and the performance on the external datasets after dropping the highest anomaly score group gradually. The external group-wise *shiftness* is thus revealed by the performance variation. An evident decrease suggests a significant distribution shift in the dropped group. For comparison, we report the classification outcomes on external dataset based on the clustering results obtained with both anomaly scores computed with CVAD Guo et al. (2021) and f-AnoGAN Schlegl et al. (2019) architectures.

Table. 2 shows the classification results for the **MURA** data, including the test results of Emory MURA and evaluation on Stanford MURA groups. Both the class-wise and average performances are reported, including *Precision*, *Recall*, *F1-score* and *AUC* scores. As the classification is evaluated in the order of *TOP_k*, *TOP_k-2*, ..., *TOP_1* order, which is *TOP_5*, *TOP_4*, *TOP_3*, *TOP_2*, *TOP_1* for our experiments, meaning that we gradually drop the group that with the highest anomaly scores and evaluate the classification performance on the remaining data. There are five groups being clustered for each class. Therefore, the *TOP 5* clusters constitute the whole external dataset and the corresponding classification results for CVAD version and f-AnoGAN version are the same. For simplicity, only one version is present (see Ta-

ble. 2 *Row Stanford_MURA_TOP 5*). The total number of images being evaluated on is listed in the row *#images* for each class. The amount of data samples in the dropped group is the number difference between the adjacent groups. Take *XR_HAND* for example, group 5 of *MedShift_w_CVAD* has 753 samples by calculating the difference of total image number of *TOP 5* clusters (3851) and *TOP 4* clusters (3098), (i.e., $753 = 3851 - 3098$) and group 5 of *MedShift_w_f-AnoGAN* has 13 samples ($13 = 3851 - 3838$). As can be observed in the table, the classifier's predictions become more and more accurate as the groups are discarded gradually based on their anomaly score order. Look into the *AUC* scores of *XR_HAND* from *TOP 5* to *TOP 1*, the values of both CVAD and f-AnoGAN are growing, which means the removed group contains data with certain *shiftness* and will affect the in-domain model's ability. The extent of *shiftness* can be inferred via the change of classification measurements for a notable improvement indicates a severe shifting exists in the dropped group. Although the same trend is noted for both CVAD and f-AnoGAN versions in general, the CVAD version can get more increase after expelling the most anomalous group than the f-AnoGAN version, which demonstrates the effectiveness of our *MedShift* framework in determining shift data among external datasets.

Aside from the quantification results above, we show the confusion matrices in Fig. 7 to report more details about the classification performance. Each confusion matrix represents a particular situation of using different cluster data during evaluation. To be clear, we use the clustering results based on the pipeline with our CVAD architecture. From left to right, top to bottom of Fig. 7, the confusion matrices are for *TOP 5*, *TOP 4*, *TOP 3*, *TOP 2*, *TOP 1* clusters, respectively. Since the groups that are dropped contain shift data, the classification accuracy is gradually rising after removing the shift data items, which can be observed from the confusion matrix differences.

Furthermore, to demonstrate our pipeline's capability of separating shift data that are with high anomaly scores, we random sample the same number of images as the *MedShift_w_CVAD*'s *TOP 1* cluster of Stanford MURA dataset for each class and run the classifier to evaluate the classification accuracy. The random sampling confusion matrix is displayed in Fig 7f. Compared to the confusion matrix of *TOP 1* in Fig. 7e, there are more misclassified

Table 2: MURA classification class-wise results with CVAD (left) and f-AnoGAN (right). The best classification values are in bold for each method.

Dataset	Metric	HAND	FOREARM	SHOULDER	FINGER	ELBOW	WRIST	HUMERUS	Average	
									Macro	Weighted
Emory_test	#images	495	74	691	74	305	572	111	2,322	
	Precision	0.842	0.704	0.979	0.312	0.929	0.957	0.875	0.800	0.903
	Recall	0.970	0.770	0.999	0.068	0.862	0.942	0.820	0.776	0.915
	F1-score	0.901	0.735	0.989	0.111	0.895	0.950	0.847	0.775	0.905
Stanford_TOP 5	#images	3,851	1,097	5,621	3,660	2,397	3,993	852	21,471	
	Precision	0.921	0.758	0.977	0.765	0.695	0.380	0.395	0.699	0.754
	Recall	0.450	0.160	0.746	0.188	0.701	0.983	0.664	0.556	0.604
	F1-score	0.605	0.264	0.846	0.301	0.698	0.548	0.496	0.537	0.594
Stanford_TOP 4	AUC	0.721	0.578	0.870	0.588	0.831	0.808	0.811	0.902	0.915
	#images	3,098 / 3,838	880 / 1,091	4,499 / 5,584	2,904 / 3,658	1,923 / 2,387	3,182 / 3,933	686 / 848	17,172 / 21,339	
	Precision	0.921 / 0.921	0.758 / 0.758	0.986 / 0.978	0.768 / 0.765	0.695 / 0.695	0.426 / 0.379	0.545 / 0.404	0.728 / 0.700	0.772 / 0.755
	Recall	0.558 / 0.452	0.195 / 0.160	0.827 / 0.750	0.233 / 0.188	0.777 / 0.704	0.990 / 0.987	0.691 / 0.665	0.610 / 0.558	0.665 / 0.605
Stanford_TOP 3	F1-score	0.695 / 0.606	0.311 / 0.265	0.899 / 0.849	0.358 / 0.302	0.734 / 0.700	0.596 / 0.548	0.609 / 0.503	0.600 / 0.539	0.654 / 0.596
	AUC	0.774 / 0.722	0.596 / 0.579	0.911 / 0.872	0.609 / 0.588	0.867 / 0.832	0.843 / 0.811	0.833 / 0.812	0.938 / 0.903	0.949 / 0.916
	#images	2,331 / 3,814	661 / 1,079	3,368 / 5,419	2,159 / 3,648	1,443 / 2,367	2,380 / 3,066	517 / 808	12,859 / 20,201	
	Precision	0.913 / 0.920	0.759 / 0.758	0.986 / 0.981	0.789 / 0.765	0.690 / 0.698	0.471 / 0.329	0.589 / 0.414	0.743 / 0.695	0.784 / 0.764
Stanford_TOP 2	Recall	0.661 / 0.455	0.253 / 0.162	0.821 / 0.769	0.279 / 0.188	0.831 / 0.710	0.988 / 0.991	0.747 / 0.666	0.654 / 0.563	0.701 / 0.595
	F1-score	0.767 / 0.609	0.379 / 0.267	0.896 / 0.862	0.413 / 0.302	0.754 / 0.704	0.638 / 0.494	0.659 / 0.511	0.644 / 0.535	0.692 / 0.593
	AUC	0.823 / 0.723	0.624 / 0.580	0.908 / 0.882	0.632 / 0.588	0.891 / 0.834	0.867 / 0.814	0.862 / 0.813	0.953 / 0.905	0.961 / 0.916
	#images	1,553 / 3,761	440 / 1,068	2,234 / 3,839	1,429 / 3,483	959 / 2,335	1,587 / 2,048	345 / 717	8,547 / 17,251	
Stanford_TOP 1	Precision	0.894 / 0.921	0.763 / 0.771	0.984 / 0.979	0.801 / 0.765	0.666 / 0.724	0.520 / 0.262	0.592 / 0.483	0.746 / 0.701	0.788 / 0.770
	Recall	0.748 / 0.461	0.359 / 0.164	0.795 / 0.818	0.324 / 0.195	0.842 / 0.719	0.986 / 0.991	0.754 / 0.658	0.687 / 0.572	0.724 / 0.575
	F1-score	0.815 / 0.614	0.488 / 0.270	0.879 / 0.891	0.461 / 0.311	0.744 / 0.722	0.681 / 0.414	0.663 / 0.557	0.676 / 0.540	0.717 / 0.582
	AUC	0.864 / 0.725	0.677 / 0.580	0.895 / 0.907	0.654 / 0.590	0.894 / 0.838	0.889 / 0.808	0.866 / 0.814	0.963 / 0.904	0.968 / 0.908
	#images	773 / 3,697	219 / 1,042	1,110 / 1,921	711 / 2,417	477 / 2,236	795 / 1,023	172 / 463	4,257 / 12,799	
	Precision	0.855 / 0.925	0.779 / 0.799	0.989 / 0.966	0.816 / 0.751	0.612 / 0.767	0.575 / 0.184	0.559 / 0.467	0.741 / 0.694	0.788 / 0.784
	Recall	0.814 / 0.468	0.434 / 0.168	0.730 / 0.791	0.368 / 0.257	0.881 / 0.743	0.974 / 0.997	0.738 / 0.590	0.705 / 0.573	0.732 / 0.547
	F1-score	0.834 / 0.622	0.557 / 0.278	0.840 / 0.869	0.508 / 0.382	0.722 / 0.755	0.723 / 0.310	0.637 / 0.521	0.689 / 0.534	0.726 / 0.580
	AUC	0.892 / 0.726	0.714 / 0.582	0.863 / 0.893	0.676 / 0.618	0.905 / 0.848	0.904 / 0.806	0.857 / 0.782	0.969 / 0.909	0.971 / 0.907

cases, especially for *XR_HAND* class, where there are 629 samples correctly predicted and 144 wrong predictions after applying our pipeline whereas 354 correct predictions and 419 misclassified cases for the random sampling situation. The improvement of classification accuracy manifests that our *MedShift* can identify the shift data that will degrade the performance of an in-domain model. The same observation holds for the **chest X-ray** experiments, please refer to Appendix B.3 for the classification details.

3.3.4. Dataset Quality Measurement Results

We report the Stanford **MURA** dataset quality calculated via the OTDD metric proposed above in the left of Fig. 8. We respectively evaluate the quality for TOP_5, TOP_4, TOP_3, TOP_2, TOP_1 cases as indicated by the X-axis values of the plots. To compare, we test our pipeline with both CVAD and f-AnoGAN anomaly detection architectures. As can be seen, the distance between Stanford MURA and Emory MURA datasets is decreasing when the anomalous groups with shift data are removed gradually. Nevertheless, our CVAD version (in

blue) shortens the distance more and faster than the f-AnoGAN (in orange) version. And the general external dataset quality achieves the best when it is composed by the group with the lowest anomaly scores, which follows the same conclusion as the average classification performance in Tab. 2.

For the reason that the OTDD method computes the distance values with label-data pairs, it was not designed for multi-label datasets. To adapt for the **chest X-ray** scenario, we report the class quality instead of the whole dataset. In the right of Fig. 8, we show CheXpert and MIMIC Consolidation class quality obtained by both the CVAD and f-AnoGAN versions. Generally, the distances between the internal and external are shortened in a limited way with *MedShift_w-CVAD* model, but the distance values are enlarged by the f-AnoGAN version. Since the distance represents the dissimilarity between the evaluated dataset pair, an increase of distance indicates a failure of identifying shift data in the external domain. Here, the CVAD version shows better performance than the *MedShift_w-f-AnoGAN* model.

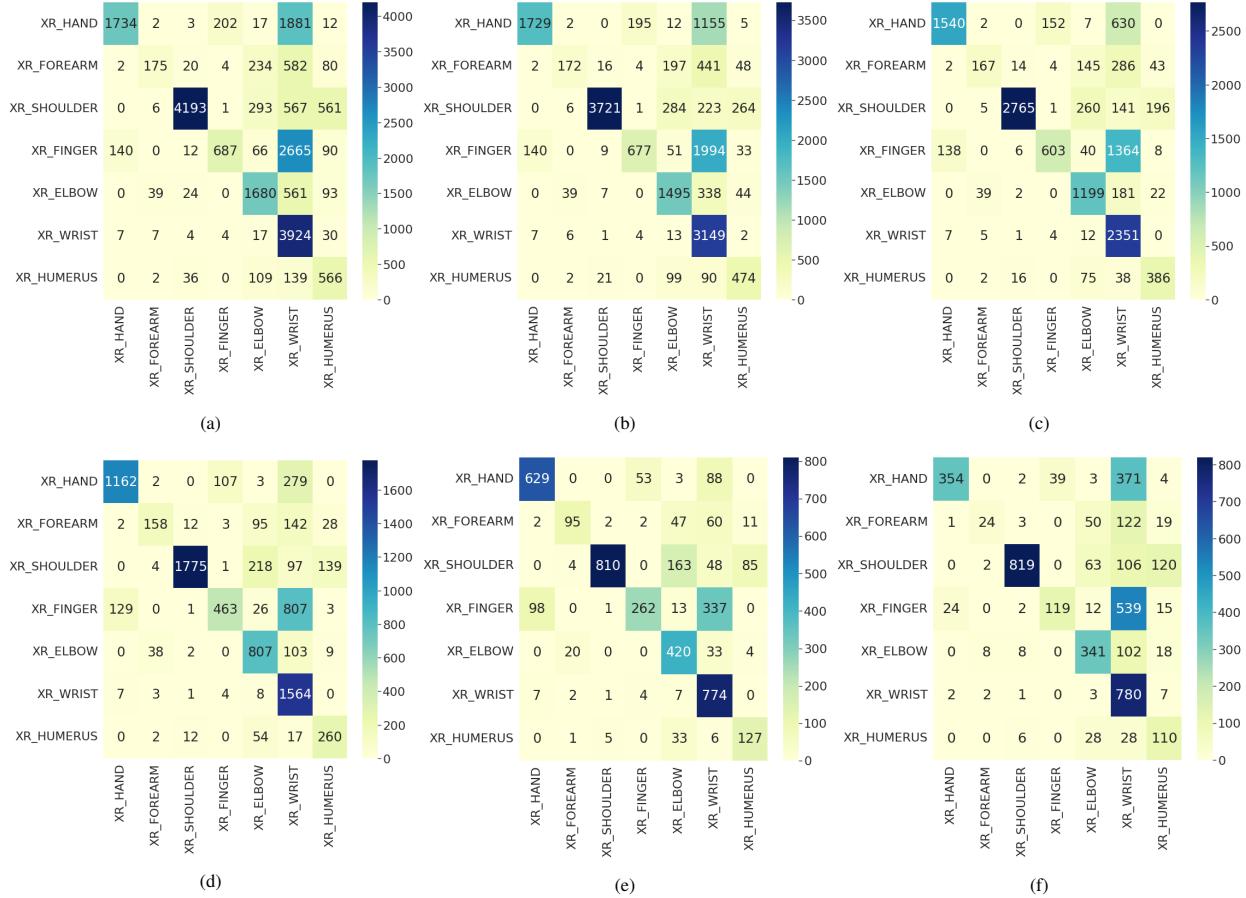


Figure 7: *MedShift-w-CVAD* confusion matrices for Stanford MURA with different numbers of clusters - (a) confusion matrix of TOP 5 clusters; (b) confusion matrix of TOP 4 clusters; (c) confusion matrix of TOP 3 clusters; (d) confusion matrix of TOP 2 clusters; (e) confusion matrix of TOP 1 cluster; (f) confusion sampling with the same number of images as the TOP 1 cluster.

Moreover, an increase of distance is also an indicator of stop sign for detecting shift data of a well-performed shift identification model. From the anomaly score distribution plots of Fig. 8, it is clear that external MURA *HAND* has more variance than the external chest X-ray *Consolidation* data. Thus, shift data identification is relatively difficult for the chest X-ray dataset, and the quality improvement is limited when little *shiftness* exists in the external dataset. Depending on the quality expectations, users can decide to remain the original *Consolidation* class or remove one or two top groups from *Consolidation*. Due to the space limitation, only one chest X-ray class case is

illustrated, please refer to Appendix B.4 for more class quality results.

4. Discussion

In this paper, we have designed an automated pipeline - *MedShift*, for medical dataset curation based on anomaly score. Under-the-hood, *MedShift* identifies image data distribution shift based on anomaly detection and unsupervised clustering to discriminate the poor-quality, noisy and under-represented samples from the external data. The anomaly detection architecture involves two separate

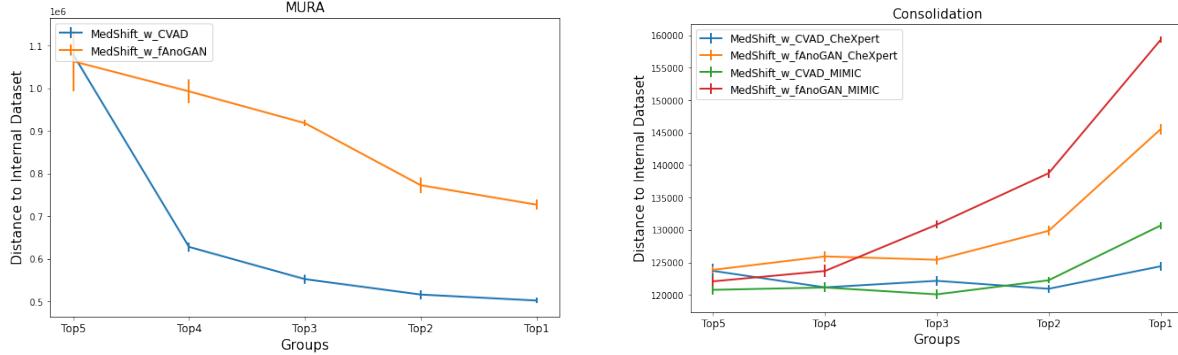


Figure 8: Dataset quality measurement results - (left) Stanford MURA whole dataset’s quality; (right) CheXpert and MIMIC Consolidation class quality. X-axis values represent situations of the groups in use, and Y-axis values indicate the distance between the internal and external datasets (the lower the better). Distance mean and stdev values of ten rounds of evaluations are present in the plots.

implementation phases - (1) internal training - time consuming and needs to be trained for each targeted class label, and (2) test phase - quick, only forward pass which needs minimal data pre-processing and cleaning from the external sites. Once trained, the anomaly detectors should be able to identify unknown anomalous patterns from an external dataset without ever seeing such anomalous data examples in training. This quality makes the proposed pipeline particularly suitable for medical image dataset curation since exchanging healthcare data among institutions and manually identifying noisy or anomalous data are both extremely challenging in the current healthcare situation.

Our pipeline is flexible towards the particular anomaly detector architectures. We evaluated two use-cases - diagnosis from chest X-ray and classifying anatomical joints from MURA and applied two different anomaly detectors CVAD and f-AnoGAN. Even though our CVAD version efficiently shortens the data quality matrix (OTDD) faster than f-AnoGAN and reaches convergence for the shift data removal by dropping lower number of cases from external data, the targeted final classification performance stays similar for both architectures.

Our experiments showed that being trained only on internal Emory datasets, deep learning models classification accuracy is gradually rising on the external dataset after removing the shift data items via *MedShift* and ultimately achieved performance close to the internal data. The improvement of classification accuracy represents the

fact that the *MedShift* can identify relevant shift data that will degrade the performance of an in-domain model and be able to reproduce the internal performance on unseen external data. Moreover, the brief cluster exploration on the external dataset showed that higher anomaly cluster groups contain more variations in terms of image quality, positioning, noise, and the pipeline correctly identified the shift data. As an immediate future study, we plan to conduct a reader study with expert radiologists to interactively evaluate the proposed platform and quantify the performance based on user-feedback matrices.

In its current state, the proposed pipeline *MedShift* can serve domain-specific quality checks and derive powerful and actionable insights. The suggested workflow will be beneficial in future non-shareable healthcare collaboration where the *MedShift* pipeline will be set up as a browser-based service within the local firewall for automated dataset curation with multi-class labels.

Appendix A. MURA Results

Appendix A.1. Anomaly Score Distribution Results with CVAD

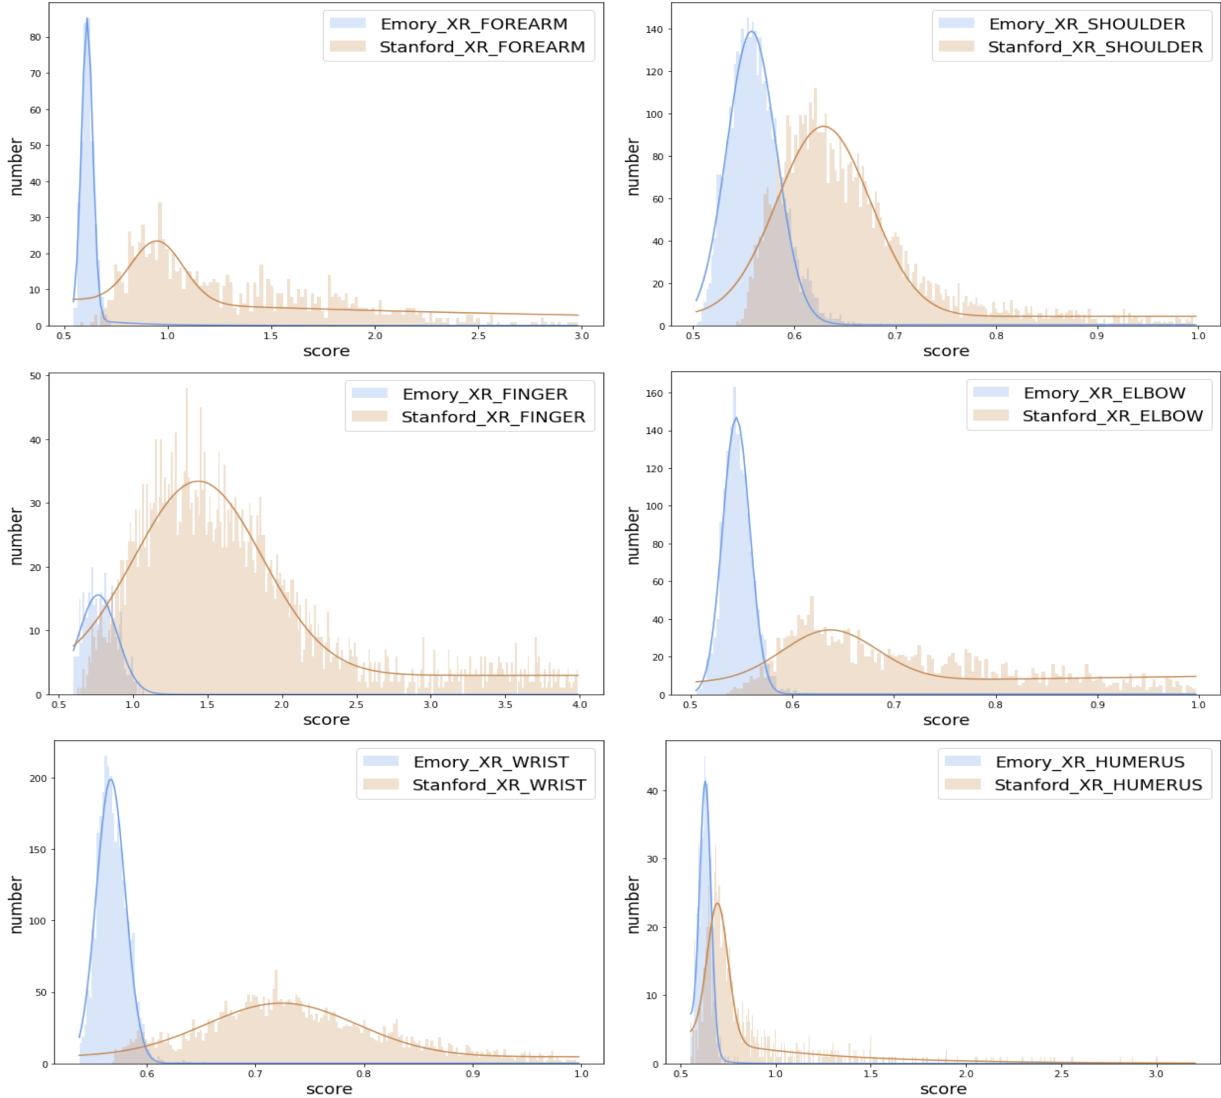


Figure A.9: *MedShift_w_CVAD*'s MURA anomaly score distribution results for classes *FOREARM*, *SHOULDER*, *FINGER*, *ELBOW*, *WRIST* and *HUMERUS* from left to right, top to bottom.

Appendix A.2. Anomaly Score Distribution Results with f -AnoGAN

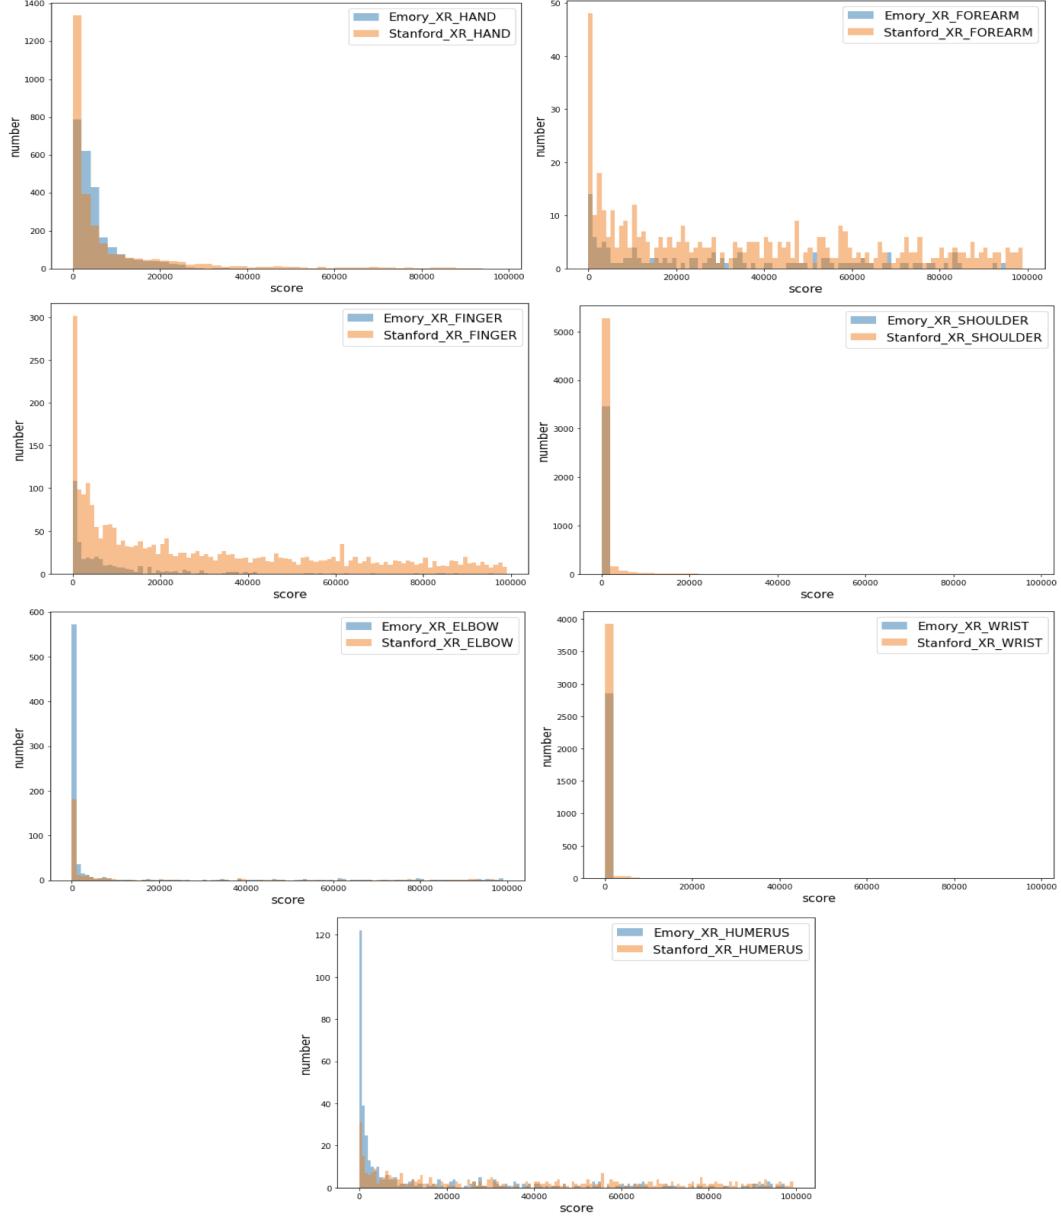


Figure A.10: *MedShift-w-f*-AnoGAN's MURA anomaly score distribution results for classes *HAND*, *FOREARM*, *SHOULDER*, *FINGER*, *ELBOW*, *WRIST* and *HUMERUS* from left to right, top to bottom.

Appendix A.3. MedShift_w_CVAD Elbow Distortion Curve Results

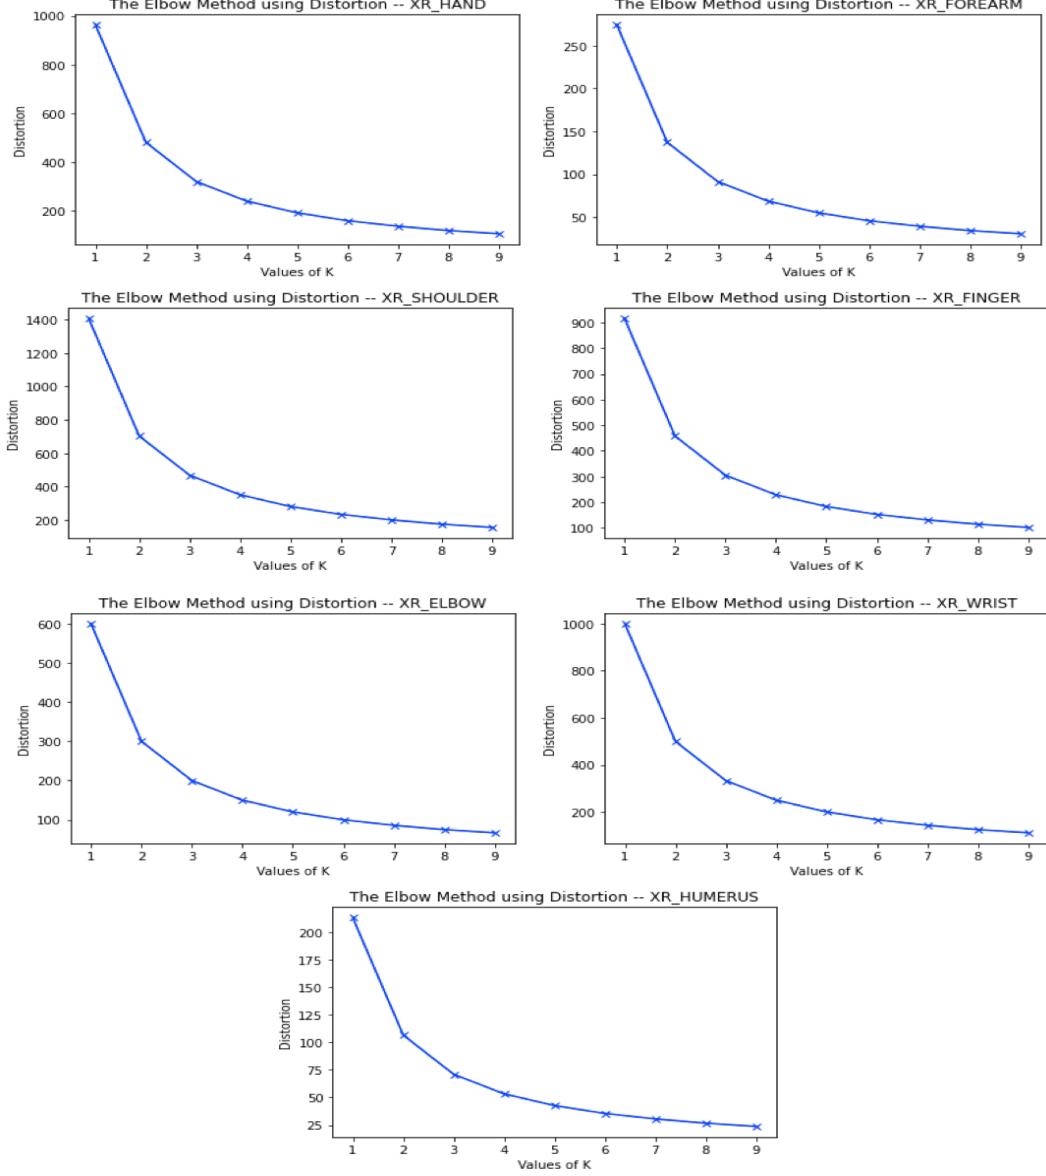


Figure A.11: *MedShift_w_CVAD*'s elbow distortion curves on Stanford MURA dataset. From left to right, top to bottom, there are plots for *XR_HAND*, *XR_FOREARM*, *XR_SHOULDER*, *XR_FINGER*, *XR_ELBO*, *XR_WRIST*, *XR_HUMERUS*, respectively. X-axis values represent the selection of K , the number of groups to be clustered into, and Y-axis values indicate the distortion.

Appendix A.4. *MedShift_w_fAnoGAN* Elbow Distortion Curve Results

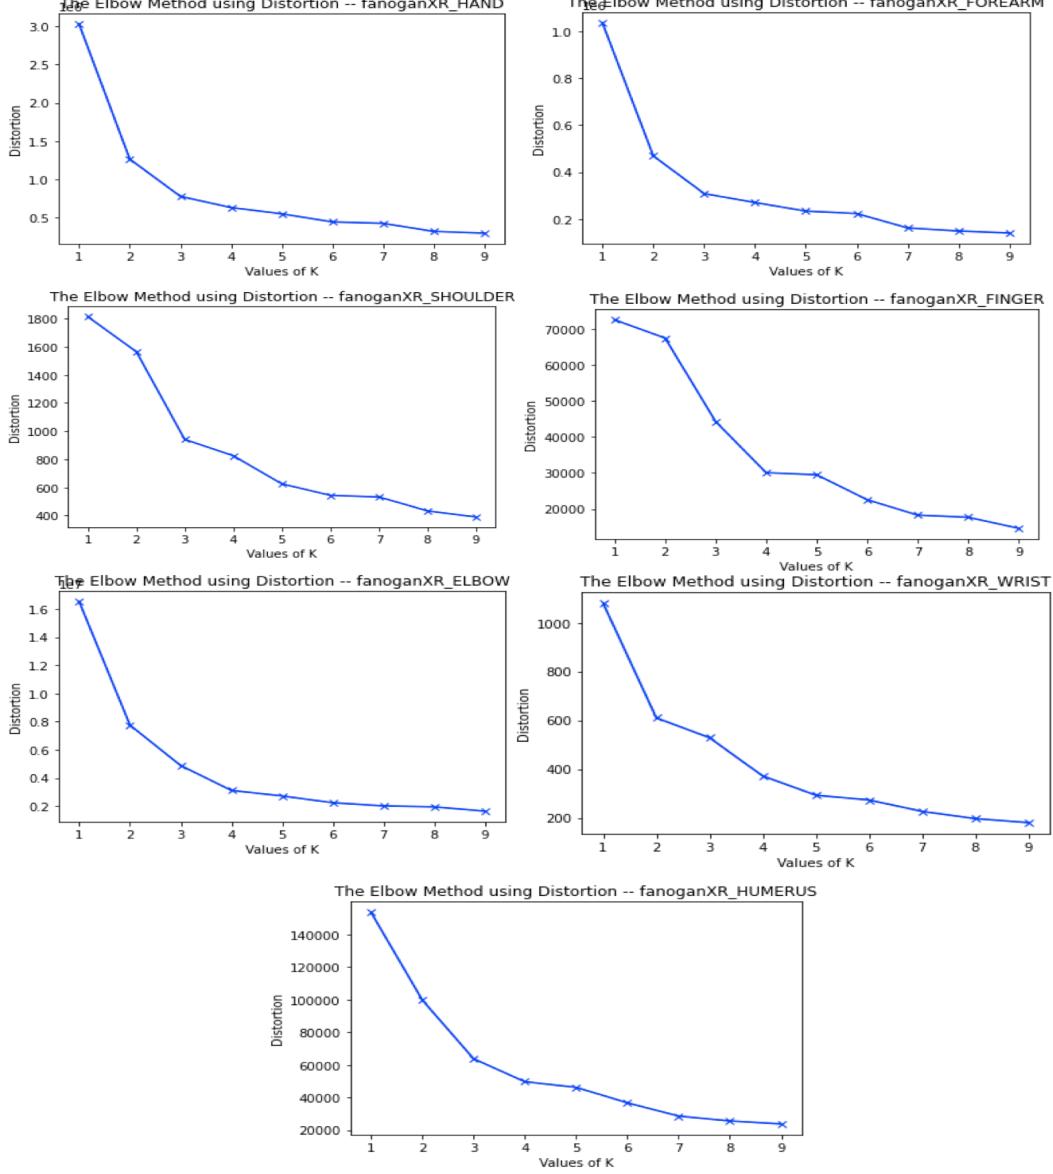


Figure A.12: *MedShift_w_fAnoGAN*'s elbow distortion curves on Stanford MURA dataset. From left to right, top to bottom, there are plots for *XR_HAND*, *XR_FOREARM*, *XR_SHOULDER*, *XR_FINGER*, *XR_ELBOW*, *XR_WRIST*, *XR_HUMERUS*, respectively. X-axis values represent the selection of K , the number of groups to be clustered into, and Y-axis values indicate the distortion.

Appendix A.5. Clustering Results with CVAD

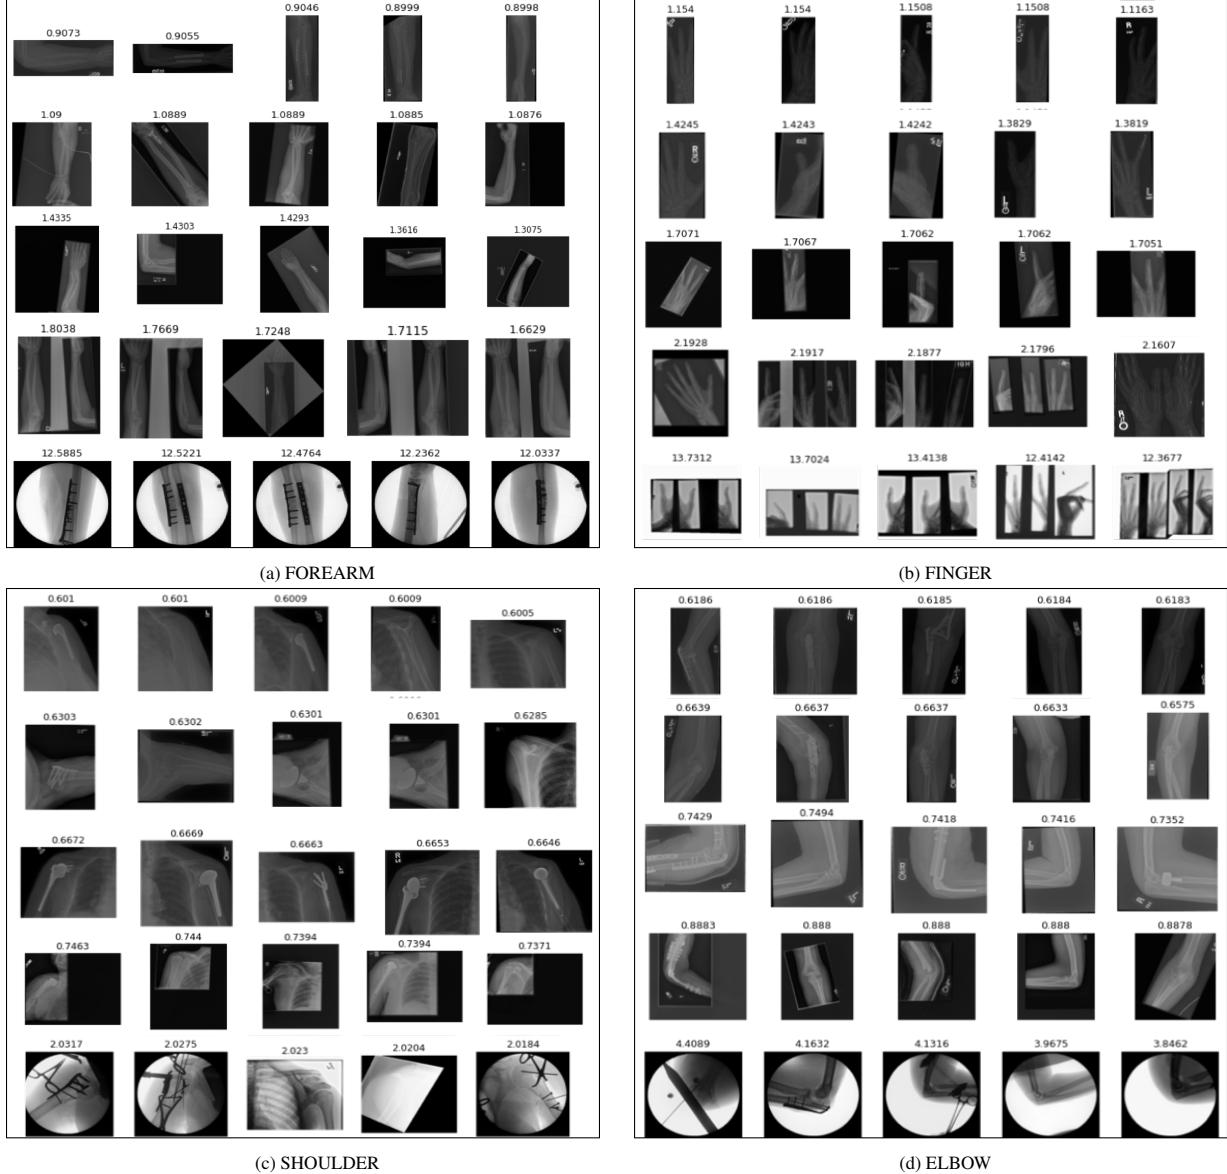
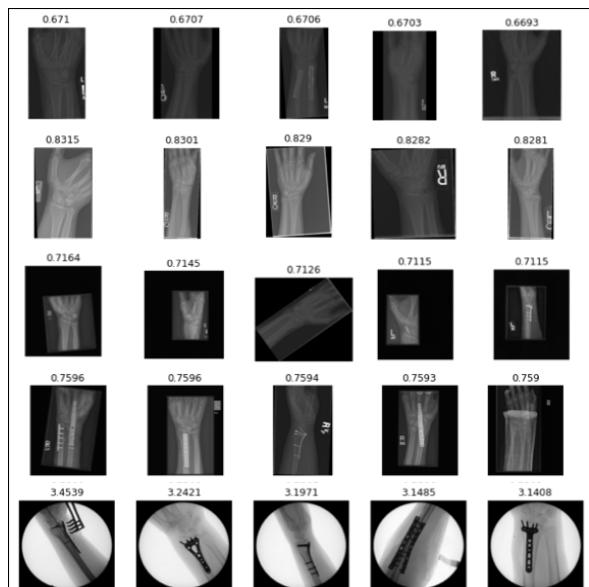
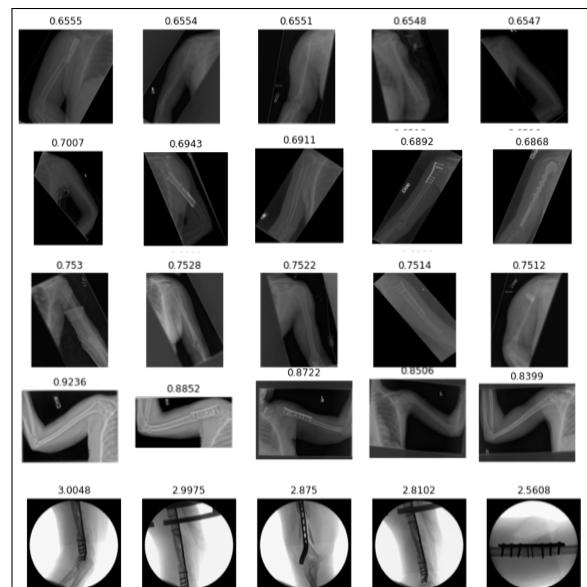


Figure A.13: *MedShift_w_CVAD*'s clustering results of Stanford MURA classes (a) *FOREARM*; (b) *FINGER*; (c) *SHOULDER*; and (d) *ELBOW*. Each row represents one group with five example images. The groups are illustrated in ascending order based on the anomaly scores from top to bottom. The corresponding anomaly score is on top of each image.



(a) WRIST



(b) HUMERUS

Figure A.14: More *MedShift-w-CVAD*'s clustering results of Stanford MURA classes (a) *WRIST*; (b) *HUMERUS* following the same arrangement style of Fig. A.13.

Appendix A.6. Clustering Results with *f*-*AnoGAN*

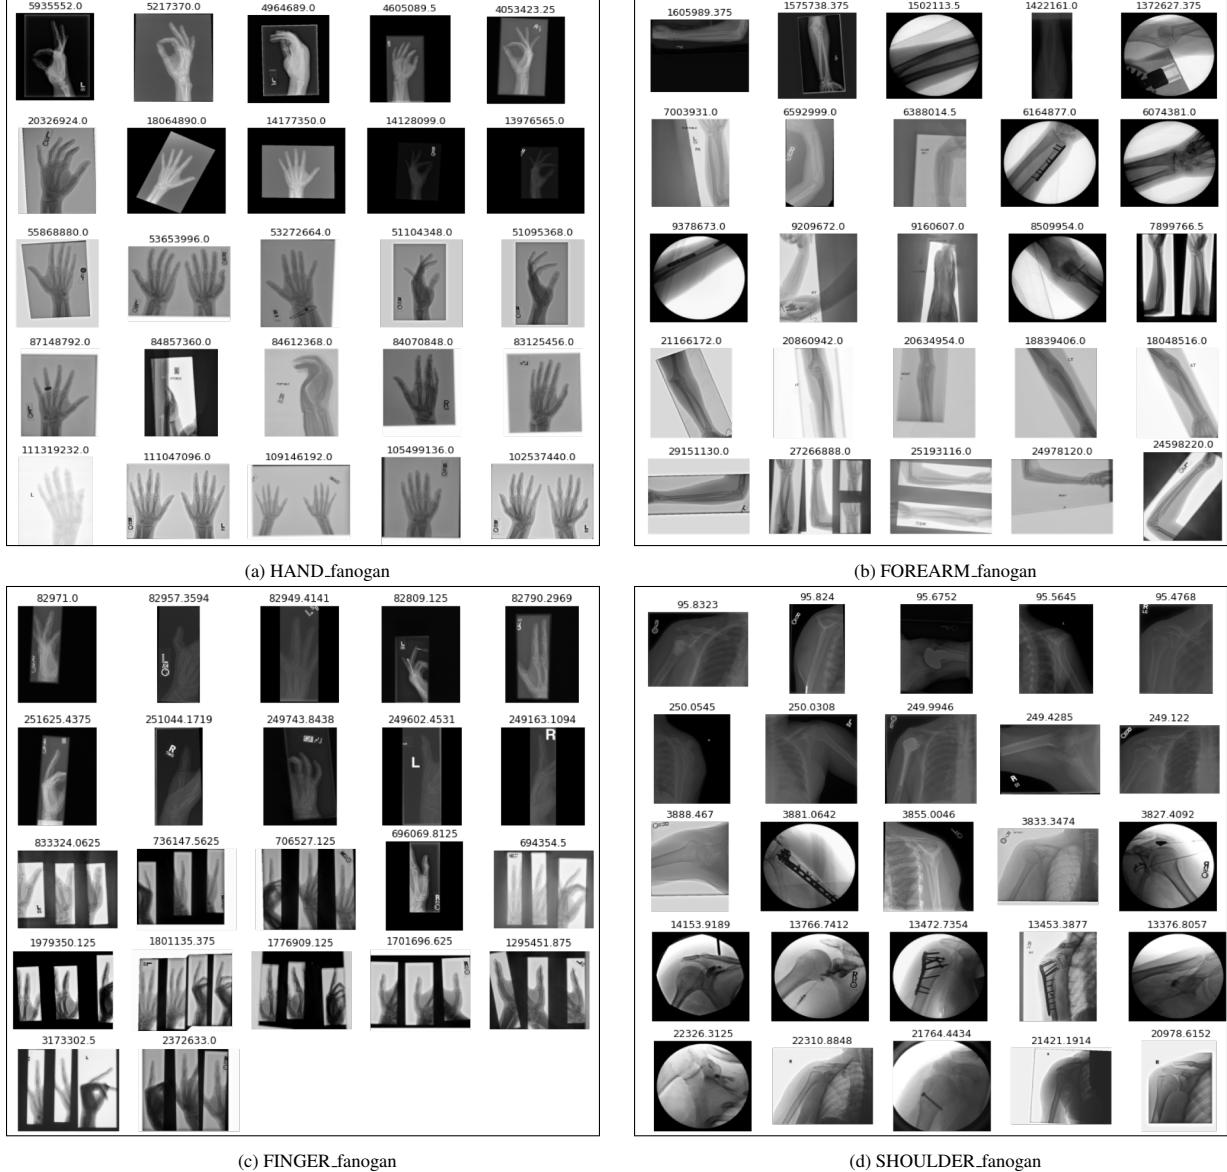
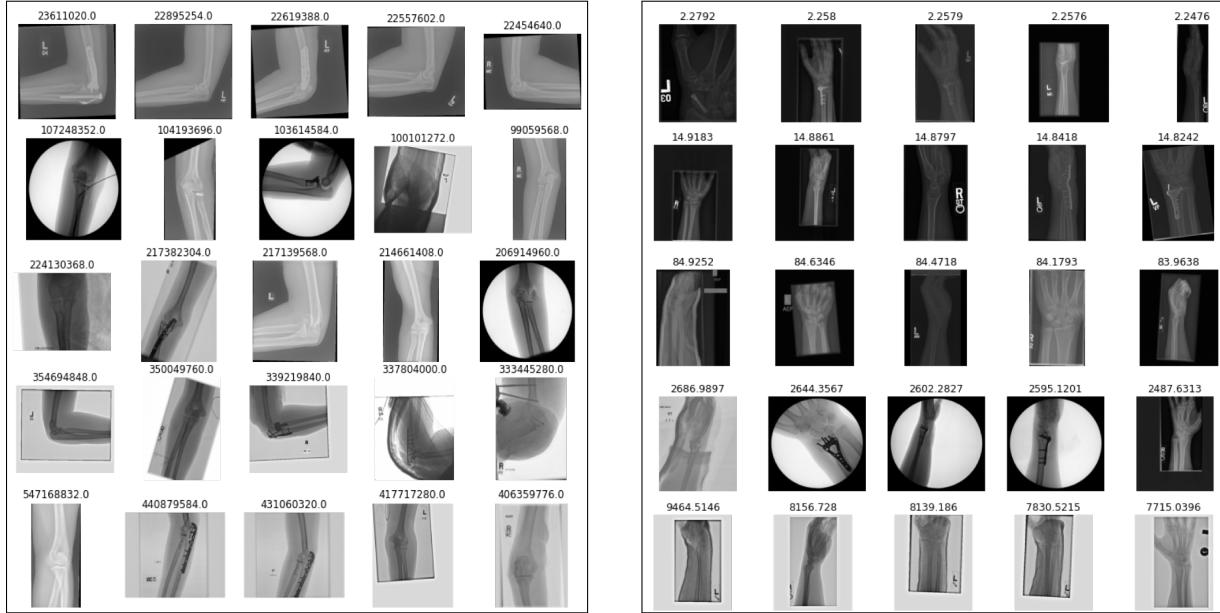


Figure A.15: *MedShift-w-f*-*AnoGAN*'s clustering results of Stanford MURA classes (a) *HAND*; (b) *FOREARM*; (c) *FINGER* and (d) *SHOULDER*. Each row represents one group with five example images. The groups are illustrated in ascending order based on the anomaly scores from top to bottom. The corresponding anomaly score is on top of each image.



(a) ELBOW_fanogan

(b) WRIST_fanogan

(c) HUMERUS_fanogan

Figure A.16: More *MedShift_w-f-AnoGAN*'s clustering results of Stanford MURA classes (a) *ELBOW*; (b) *WRIST* and (c) *HUMERUS* following the same arrangement style of Fig. A.13.

Appendix B. Chest X-ray Results

Appendix B.1. Anomaly Score Distribution Results

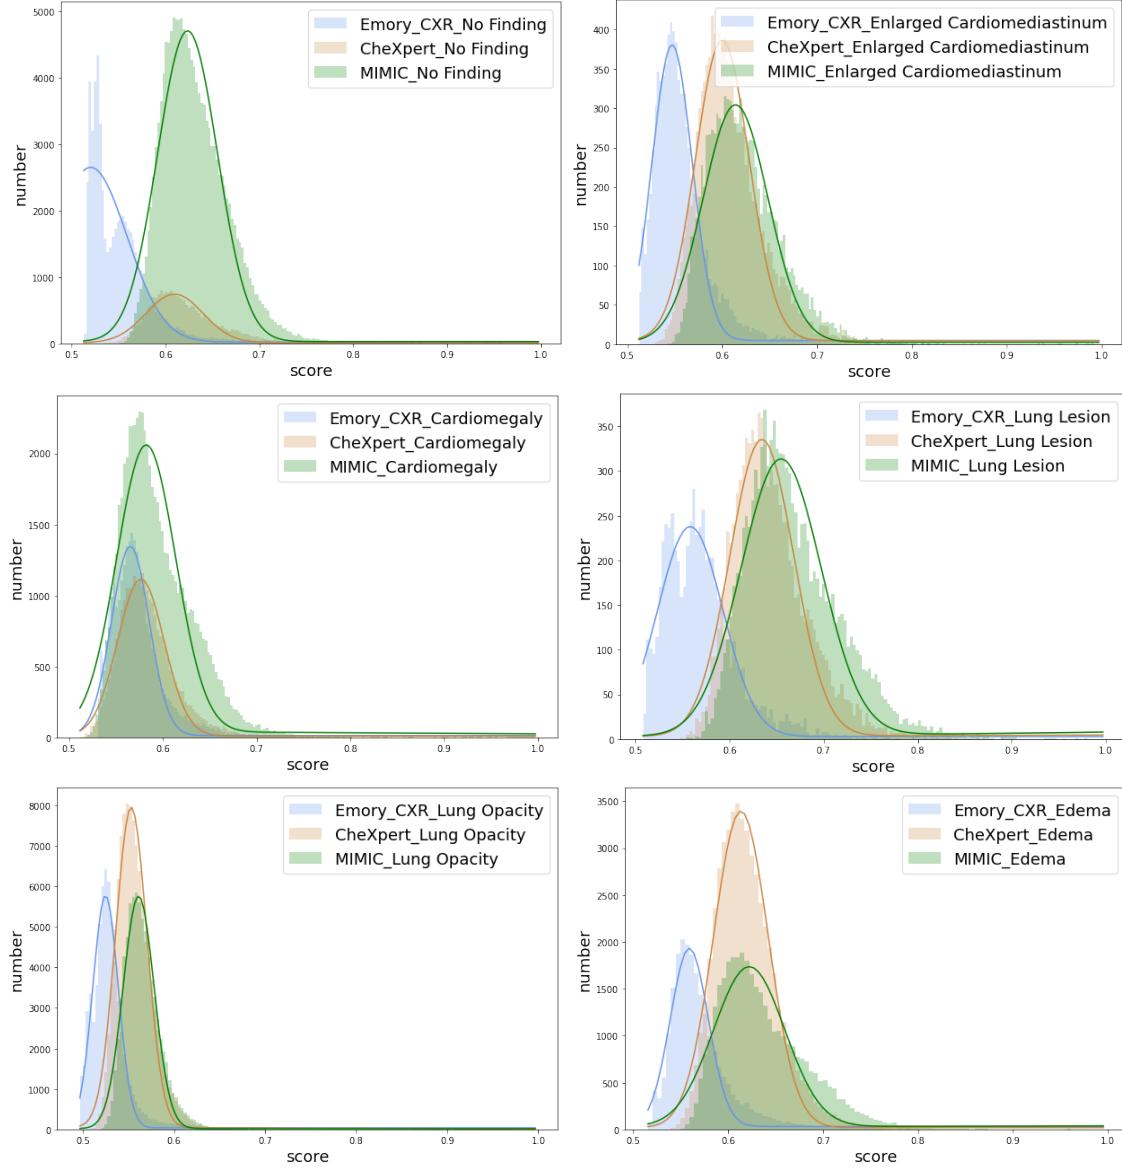


Figure B.17: *MedShift-w-CVAD* class-wise chest X-ray anomaly score distributions for Emory_CXR (blue), CheXpert (orange) and MIMIC (green) datasets. Distributions may be truncated for better visualization.

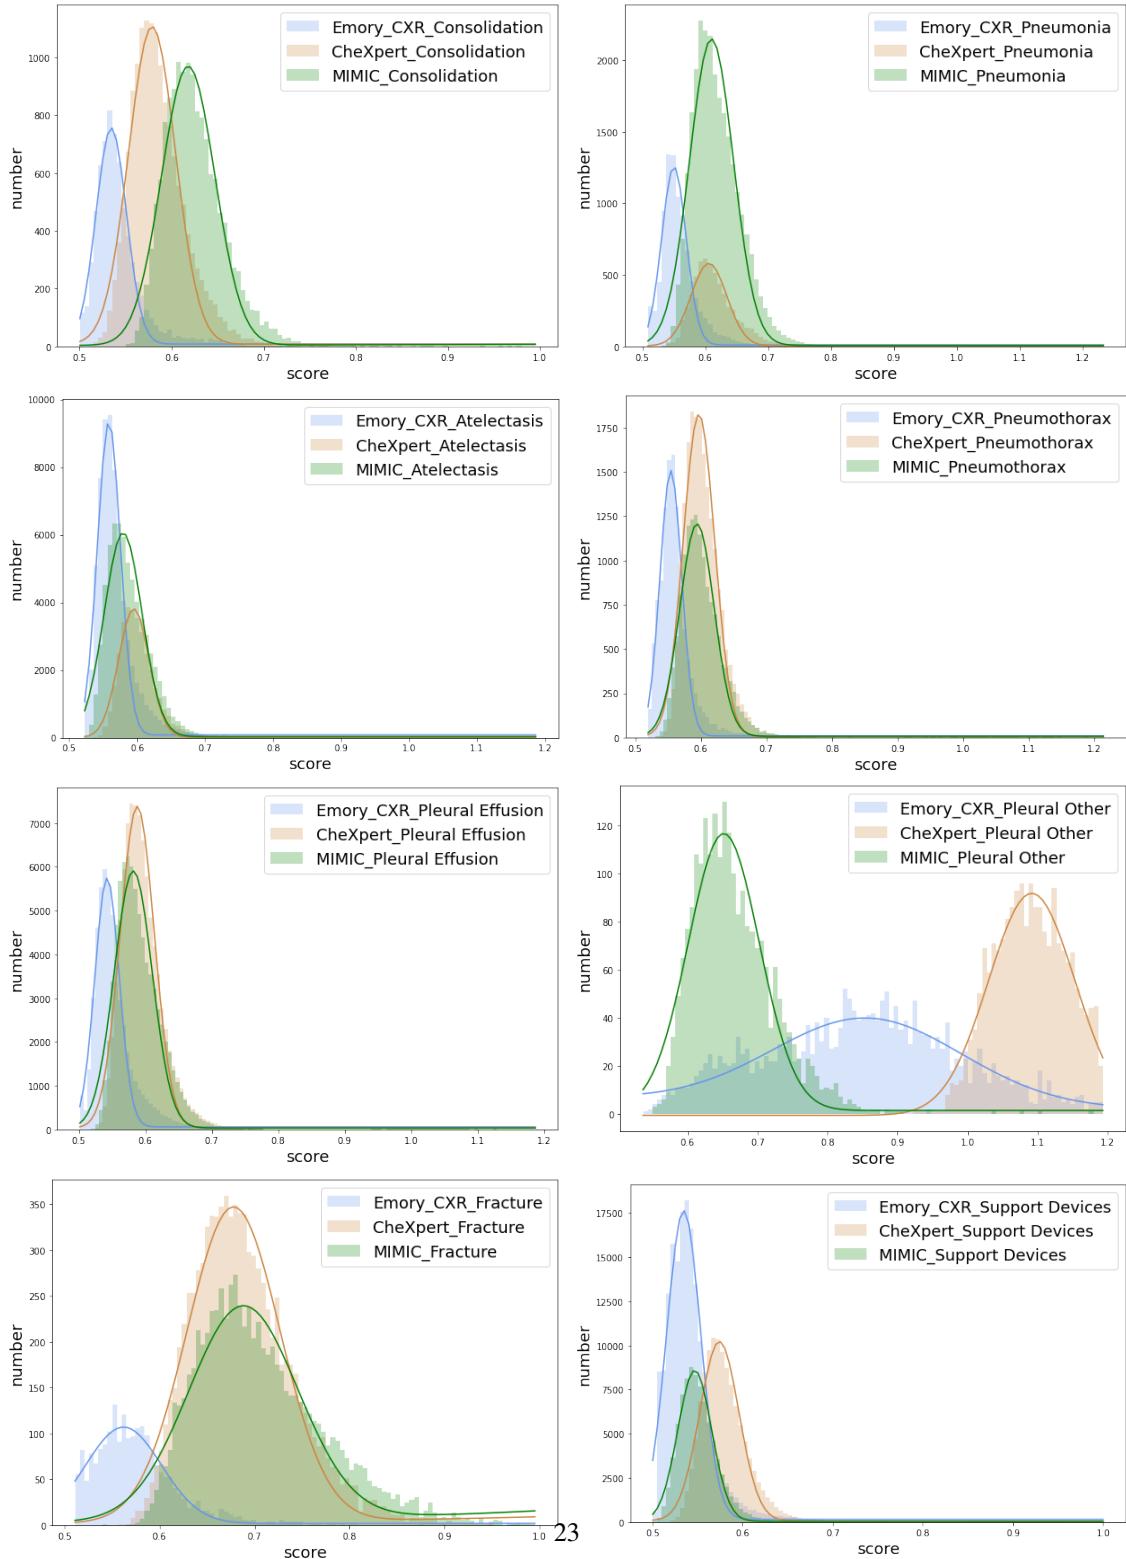


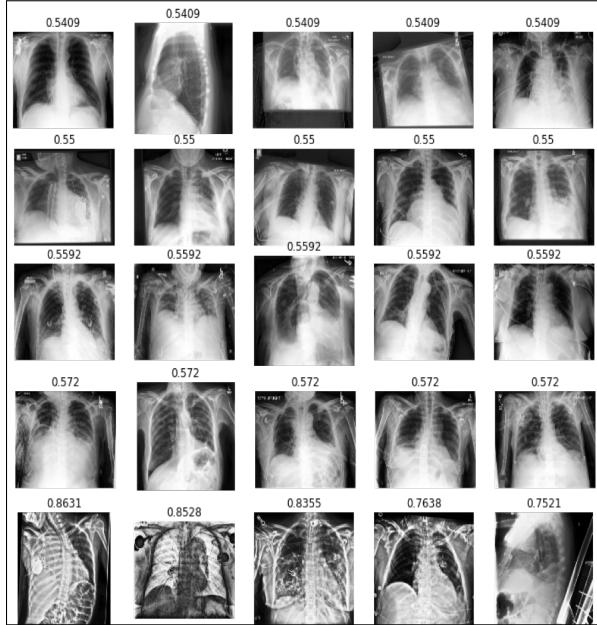
Figure B.18: More *MedShift.w_CVAD* class-wise chest X-ray anomaly score distributions for Emory_CXR (blue), CheXpert (orange) and MIMIC (green) datasets. Distributions may be truncated for better visualization.

Appendix B.2. Clustering Results

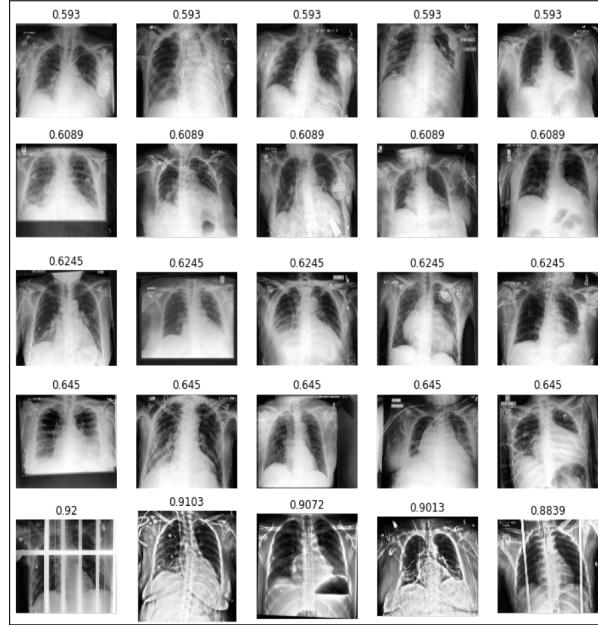
Appendix B.2.1. Clustering Results for CheXpert



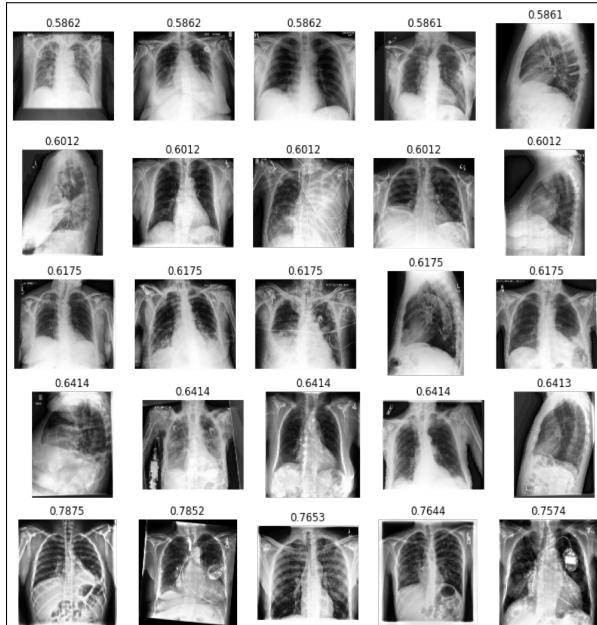
Figure B.19: *MedShift_w_CVAD* CheXpert clustering results for classes (a) *No Finding*; (b) *Enlarge Cardiomediastinum*; (c) *Cardiomegaly* and (d) *Lung Lesion* following the same arrangement style of Fig. A.13.



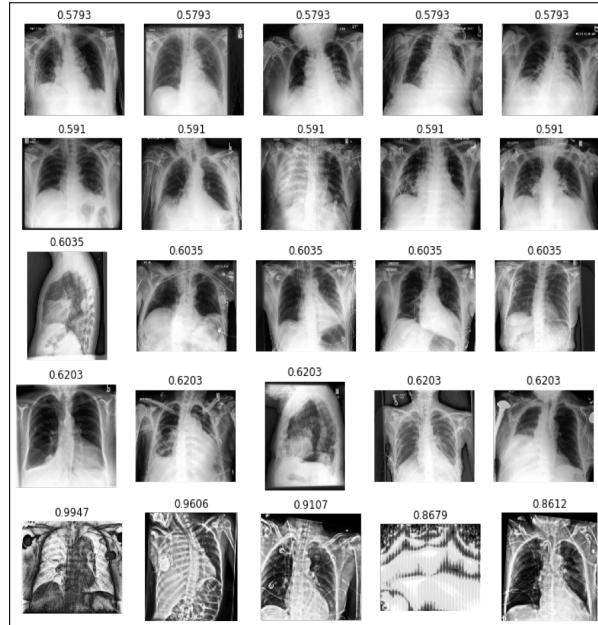
(a) Lung Opacity



(b) Edema

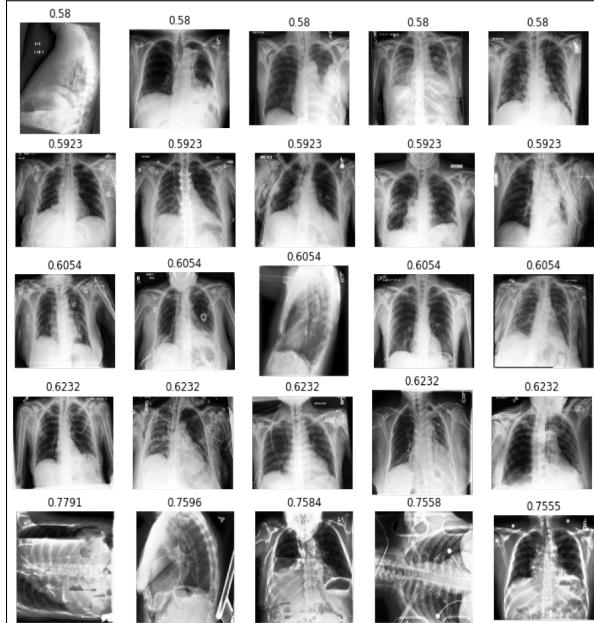


(c) Pneumonia

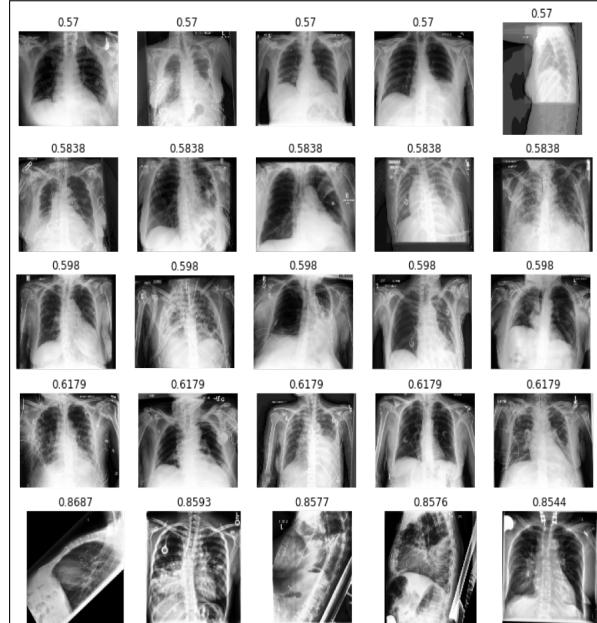


(d) Atelectasis

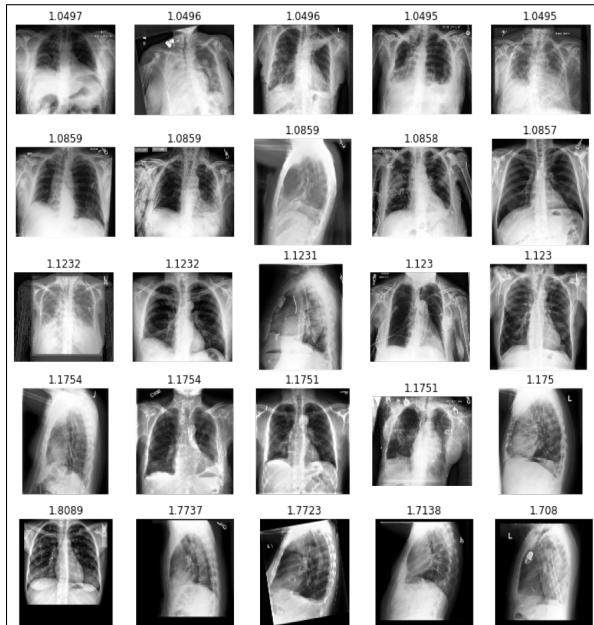
Figure B.20: More *MedShift_w-CVAD* CheXpert clustering results for classes (a) *Lung Opacity*; (b) *Edema*; (c) *Pneumonia* and (d) *Atelectasis* following the same arrangement style of Fig. A.13.



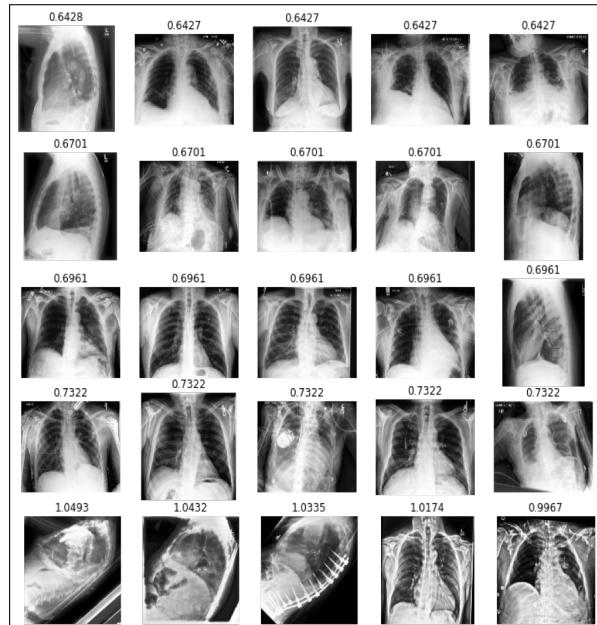
(a) Pneumothorax



(b) Pleural Effusion

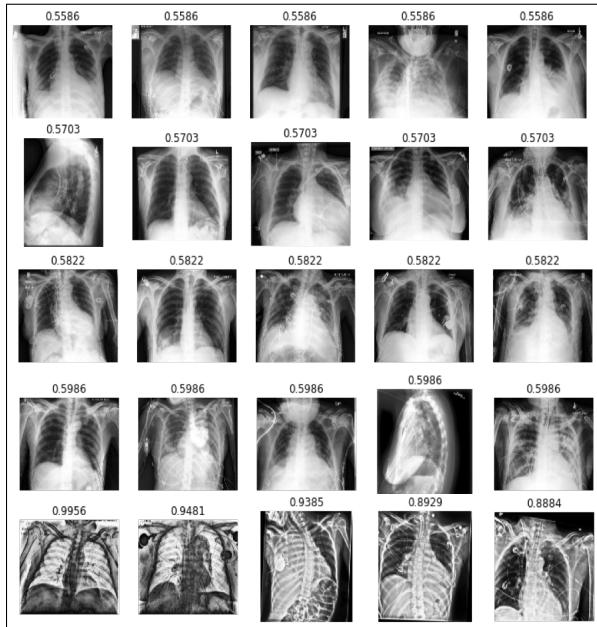


(c) Pleural Other



(d) Fracture

Figure B.21: More *MedShift_w_CVAD* CheXpert clustering results for classes (a) *Pneumothorax*; (b) *Pleural Effusion*; (c) *Pleural Other* and (d) *Fracture* following the same arrangement style of Fig. A.13.



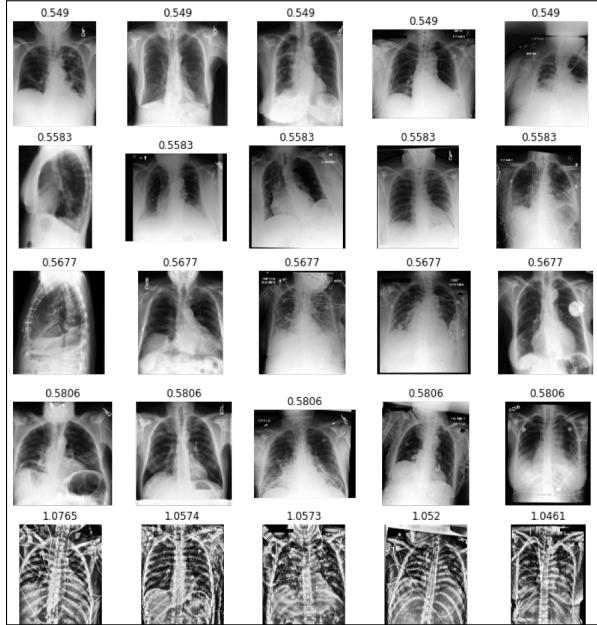
(a) Support Devices

Figure B.22: More *MedShift_w_CVAD* CheXpert clustering results for classes (a) *Support Devices* following the same arrangement style of Fig. A.13.

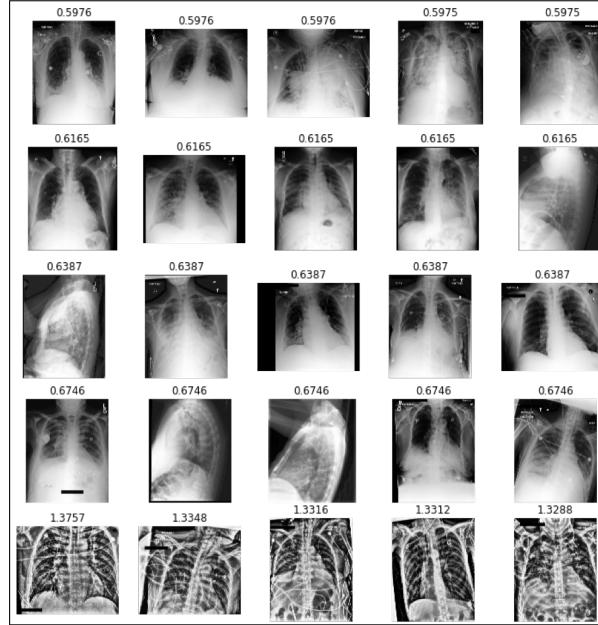
Appendix B.2.2. Clustering Results for MIMIC



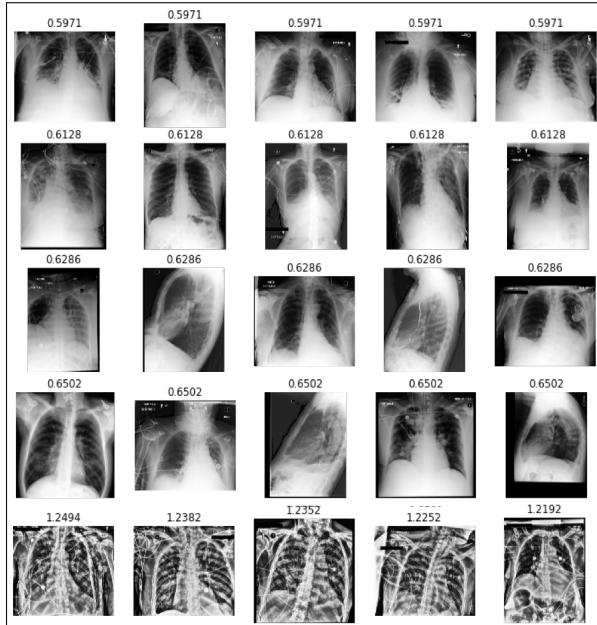
Figure B.23: `MedShift_w_CVAD` MIMIC clustering results for classes (a) *No Finding*; (b) *Enlarge Cardiomedastinum*; (c) *Cardiomegaly* and (d) *Lung Lesion* following the same arrangement style of Fig. A.13.



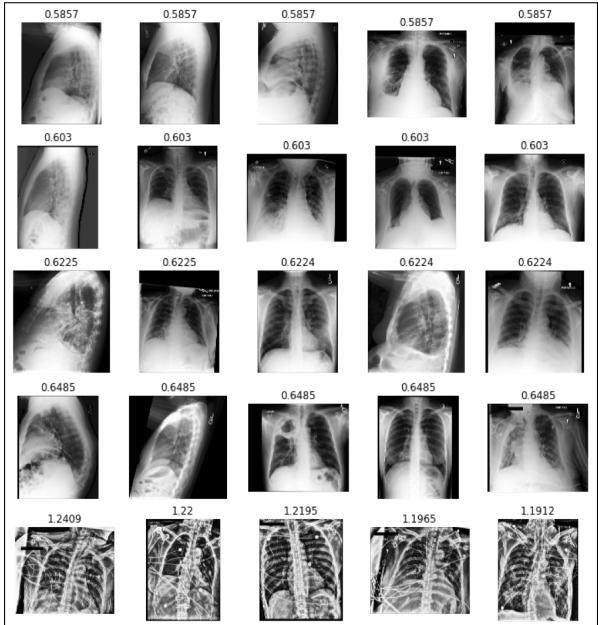
(a) Lung Opacity



(b) Edema

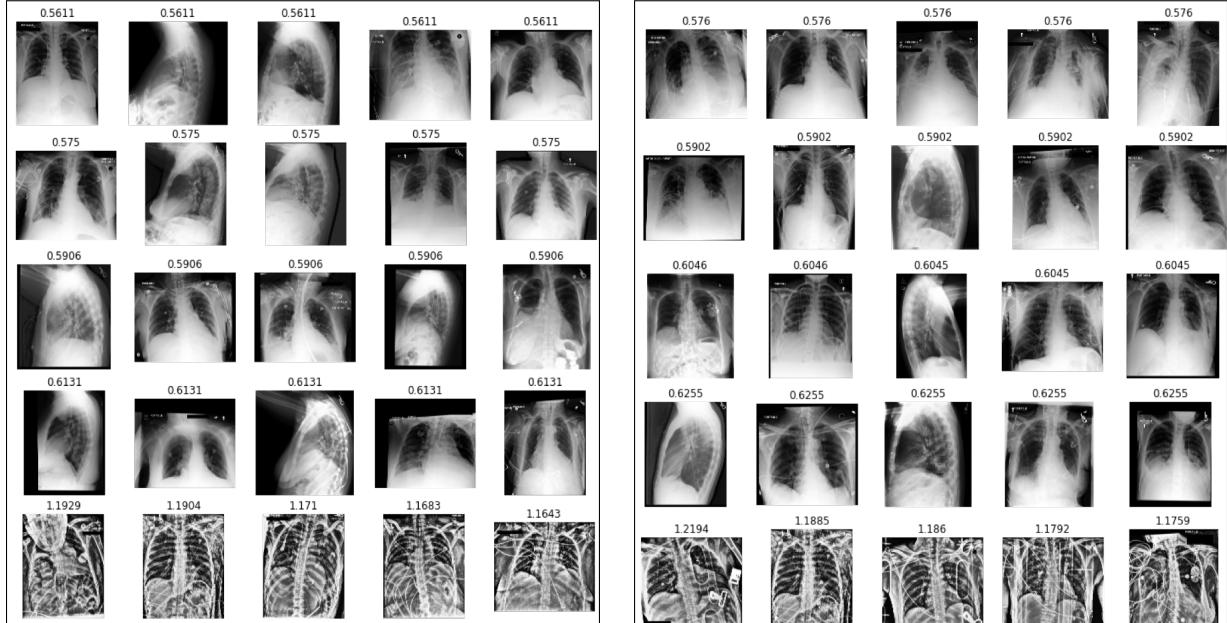


(c) Consolidation



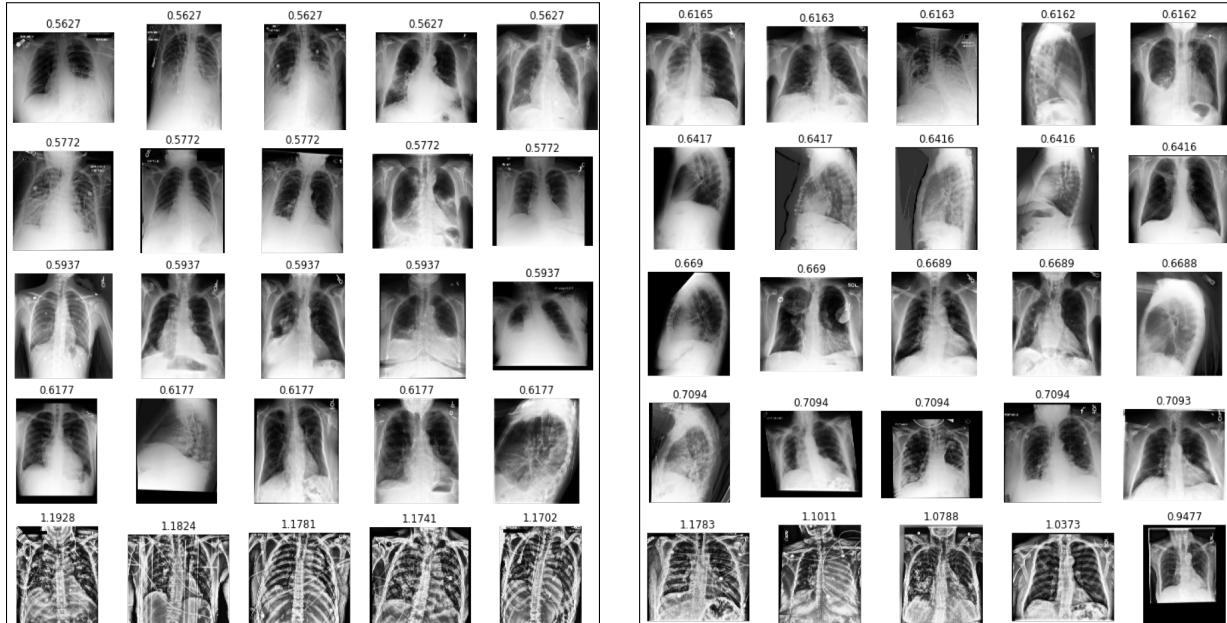
(d) Pneumonia

Figure B.24: *MedShift_w_CVAD* MIMIC clustering results for classes (a) *Lung Opacity*; (b) *Edema*; (c) *Consolidation* and (d) *Pneumonia* following the same arrangement style of Fig. A.13.



(a) Atelectasis

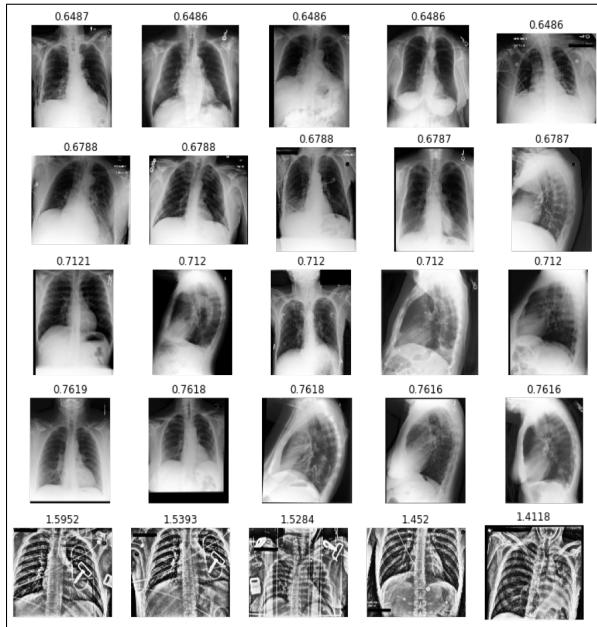
(b) Pneumothorax



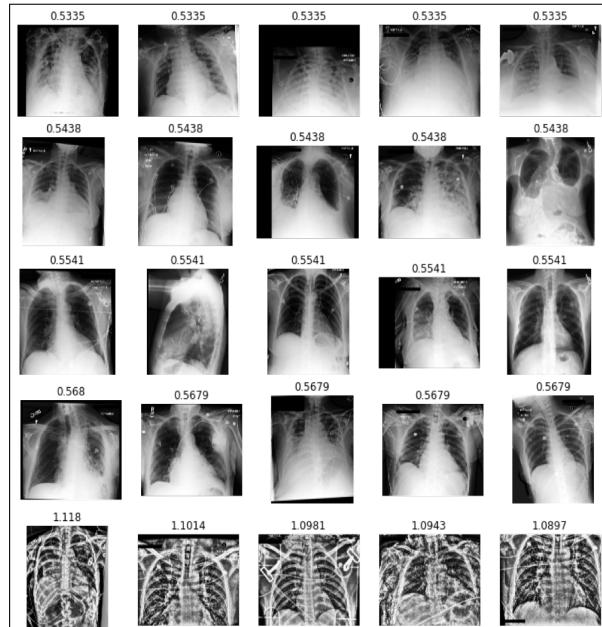
(c) Pleural Effusion

(d) Pleural Other

Figure B.25: *MedShift_w_CVAD* MIMIC clustering results for classes (a) Atelectasis; (b) Pneumothorax; (c) Pleural Effusion and (d) Pleural Other following the same arrangement style of Fig. A.13.



(a) Fracture



(b) Support Devices

Figure B.26: *MedShift_w_CVAD* MIMIC clustering results for classes (a) *Fracture*; (b) *Support Devices* following the same arrangement style of Fig. A.13.

Appendix B.3. Classification Results

Table B.3: Chest X-ray data classification class-wise results with CVAD (left) and f-AnoGAN (right) architectures

Dataset	Metrics	Performance Metrics										AVERAGE								
		No Finding	Enlarged Cardiomediastinum	Cardiomegaly	Lung Lesion	Lung Opacity	Edema	Consolidation	Pneumonia	Atelectasis	Pneumothorax	Pleural Effusion	Pleural Other	Fracture	Support Devices	Micro	Macro	Weighted	Samples	
Emory_CXR	#images	7,936	532	1,256	397	2,141	830	151	439	2,215	150	711	98	177	9,994	0.338	0.705	0.633	661	
	Precision	0.753	0.082	0.385	0.090	0.531	0.416	0.107	0.153	0.577	0.066	0.516	0.066	0.058	0.93	0.489	0.338	0.705	0.633	
	Recall	0.960	0.232	0.697	0.446	0.588	0.627	0.437	0.490	0.580	0.433	0.686	0.520	0.277	0.109	0.514	0.506	0.514	0.537	
	F1-score	0.844	0.121	0.496	0.150	0.558	0.500	0.171	0.233	0.579	0.115	0.589	0.118	0.096	0.195	0.502	0.340	0.477	0.54	
	AUC	0.724	0.561	0.788	0.652	0.742	0.783	0.697	0.697	0.743	0.681	0.824	0.732	0.607	0.540	0.710	0.698	0.719	0.661	
TOP 5	#images	22,381	10,798	27,000	9,186	10,5581	52,246	14,783	6,039	33,376	19,448	86,187	3,523	9,040	116,001	212,273				
	ChexNet	0.356	0.098	0.319	0.081	0.581	0.462	0.122	0.055	0.196	0.118	0.623	0.033	0.091	0.670	0.343	0.272	0.479	0.351	
	Recall	0.684	0.073	0.610	0.610	0.832	0.625	0.556	0.495	0.818	0.884	0.835	0.461	0.173	0.845	0.752	0.607	0.752	0.722	
	F1-score	0.468	0.084	0.419	0.143	0.684	0.531	0.200	0.098	0.316	0.208	0.714	0.120	0.074	0.471	0.342	0.566	0.446	0.446	
	AUC	0.754	0.513	0.635	0.513	0.665	0.625	0.063	0.063	0.693	0.067	0.533	0.623	0.533	0.664	0.724	0.709	0.724	0.709	
TOP 4	#images	14,520	10,042	64,346	10,801	76,423	26,564	14,675	26,232	65,047	14,257	70,557	3,469	7,605	84,073	173,664 / 88,757				
	ChexNet	0.674	0.040	0.412	0.060	0.318	0.358	0.096	0.133	0.315	0.084	0.564	0.028	0.036	0.502	0.315	0.259	0.433	0.4	
	Recall	0.754	0.078	0.440	0.539	0.740	0.619	0.543	0.428	0.754	0.719	0.700	0.379	0.158	0.739	0.661	0.542	0.661	0.661	
	F1-score	0.712	0.053	0.425	0.108	0.445	0.454	0.164	0.202	0.445	0.151	0.624	0.052	0.058	0.598	0.427	0.321	0.504	0.464	
	AUC	0.757	0.515	0.656	0.639	0.764	0.661	0.602	0.693	0.786	0.631	0.759	0.731	0.667	0.727	0.712	0.661	0.727	0.712	
TOP 3	#images	17,834 / 4,845	8,899 / 4,282	22,846 / 11,259	7,518	21,609	88,278 / 39,924	45,239 / 21,838	4,934 / 21,074	28,414 / 13,719	16,509 / 6,620	73,595 / 35,301	3,141 / 1,079	7,464	3,325	97,280 / 45,373	173,664 / 88,757			
	ChexNet	0.382 / 0.327	0.099 / 0.112	0.315 / 0.369	0.085 / 0.078	0.463 / 0.523	0.119 / 0.138	0.054 / 0.059	0.199 / 0.234	0.115 / 0.120	0.626 / 0.690	0.037 / 0.031	0.098 / 0.120	0.670 / 0.750	0.351 / 0.395	0.275 / 0.300	0.483 / 0.545	0.358 / 0.393	0.358 / 0.393	
	Recall	0.640 / 0.638	0.071 / 0.066	0.624 / 0.631	0.060 / 0.054	0.848 / 0.853	0.647 / 0.660	0.553 / 0.562	0.505 / 0.519	0.833 / 0.848	0.884 / 0.902	0.855 / 0.888	0.469 / 0.450	0.157 / 0.146	0.864 / 0.863	0.765 / 0.772	0.611 / 0.614	0.765 / 0.772	0.737 / 0.737	
	F1-score	0.478 / 0.433	0.083 / 0.088	0.419 / 0.466	0.150 / 0.138	0.691 / 0.734	0.196 / 0.222	0.098 / 0.109	0.321 / 0.366	0.204 / 0.212	0.723 / 0.767	0.069 / 0.079	0.121 / 0.132	0.755 / 0.808	0.481 / 0.522	0.346 / 0.363	0.573 / 0.616	0.457 / 0.499	0.457 / 0.499	
	AUC	0.760 / 0.769	0.518 / 0.515	0.770 / 0.710	0.655 / 0.659	0.691 / 0.693	0.622 / 0.618	0.389 / 0.573	0.667 / 0.605	0.740 / 0.770	0.623 / 0.654	0.532 / 0.535	0.665 / 0.664	0.535 / 0.535	0.729 / 0.730	0.640 / 0.633	0.715 / 0.722	0.661 / 0.661	0.720 / 0.720	
TOP 2	#images	115,452	10,000	58,303	5,638	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	10,000	
	ChexNet	0.455 / 0.544	0.042 / 0.047	0.415 / 0.403	0.062 / 0.062	0.323 / 0.330	0.359 / 0.422	0.077 / 0.135	0.133 / 0.163	0.323 / 0.407	0.044 / 0.047	0.505 / 0.653	0.039 / 0.036	0.039 / 0.038	0.506 / 0.643	0.319 / 0.370	0.269 / 0.304	0.437 / 0.476	0.399 / 0.424	
	Recall	0.749 / 0.726	0.082 / 0.070	0.451 / 0.472	0.540 / 0.510	0.739 / 0.786	0.645 / 0.675	0.542 / 0.556	0.431 / 0.471	0.756 / 0.809	0.704 / 0.715	0.716 / 0.750	0.381 / 0.400	0.158 / 0.140	0.749 / 0.744	0.666 / 0.681	0.546 / 0.560	0.666 / 0.681	0.663 / 0.671	
	F1-score	0.710 / 0.677	0.056 / 0.063	0.431 / 0.487	0.111 / 0.111	0.449 / 0.518	0.461 / 0.519	0.165 / 0.217	0.203 / 0.242	0.453 / 0.541	0.150 / 0.186	0.632 / 0.668	0.055 / 0.072	0.058 / 0.058	0.604 / 0.698	0.431 / 0.483	0.324 / 0.364	0.507 / 0.543	0.464 / 0.484	
	AUC	0.755 / 0.782	0.513 / 0.505	0.650 / 0.630	0.638 / 0.631	0.651 / 0.653	0.745 / 0.744	0.661 / 0.646	0.694 / 0.675	0.683 / 0.693	0.729 / 0.737	0.603 / 0.654	0.653 / 0.654	0.727 / 0.726	0.665 / 0.661	0.724 / 0.720	0.710 / 0.700	0.727 / 0.720		
TOP 1	#images	13,296 / 5,030	6,986 / 4,554	18,241 / 11,723	5,773 / 2,766	68,914 / 41,990	38,897 / 22,654	9,479 / 6,706	3,781 / 2,455	22,396 / 13,716	11,697 / 7,688	58,410 / 36,948	2,498 / 1,249	5,796 / 3,788	58,410 / 45,873	133,375 / 77,585	173,664 / 88,757			
	ChexNet	0.394 / 0.302	0.104 / 0.113	0.315 / 0.367	0.089 / 0.083	0.586 / 0.638	0.462 / 0.524	0.118 / 0.133	0.054 / 0.066	0.202 / 0.229	0.114 / 0.132	0.630 / 0.680	0.041 / 0.033	0.109 / 0.115	0.667 / 0.762	0.357 / 0.392	0.277 / 0.309	0.485 / 0.539	0.364 / 0.393	
	Recall	0.607 / 0.623	0.075 / 0.076	0.638 / 0.628	0.582 / 0.597	0.856 / 0.853	0.669 / 0.664	0.548 / 0.552	0.515 / 0.514	0.842 / 0.844	0.876 / 0.899	0.865 / 0.861	0.441 / 0.441	0.130 / 0.132	0.875 / 0.861	0.772 / 0.766	0.610 / 0.609	0.772 / 0.766	0.745 / 0.744	
	F1-score	0.478 / 0.407	0.087 / 0.084	0.421 / 0.463	0.154 / 0.146	0.694 / 0.730	0.547 / 0.584	0.194 / 0.214	0.097 / 0.114	0.326 / 0.360	0.202 / 0.231	0.729 / 0.756	0.075 / 0.061	0.113 / 0.123	0.707 / 0.709	0.488 / 0.519	0.348 / 0.363	0.575 / 0.612	0.464 / 0.494	
	AUC	0.752 / 0.759	0.520 / 0.516	0.709 / 0.711	0.656 / 0.660	0.691 / 0.697	0.617 / 0.616	0.625 / 0.628	0.585 / 0.577	0.610 / 0.607	0.735 / 0.725	0.629 / 0.609	0.538 / 0.538	0.653 / 0.666	0.732 / 0.732	0.637 / 0.637	0.617 / 0.617	0.709 / 0.720		
TOP 2	#images	86,000 / 25,331	6,641 / 4,228	41,690 / 24,530	6,906 / 4,500	48,974 / 27,226	24,134 / 17,784	6,268 / 6,248	16,540 / 9,331	43,116 / 26,506	18,590 / 5,729	50,435 / 30,506	2,282 / 1,999	4,785 / 2,888	50,435 / 45,373	126,263 / 65,373	173,664 / 88,757			
	MIMIC	0.681 / 0.680	0.070 / 0.070	0.645 / 0.646	0.474 / 0.475	0.600 / 0.600	0.677 / 0.681	0.541 / 0.544	0.432 / 0.461	0.766 / 0.813	0.638 / 0.679	0.677 / 0.709	0.404 / 0.404	0.149 / 0.148	0.626 / 0.641	0.326 / 0.371	0.263 / 0.320	0.435 / 0.464	0.404 / 0.441	
	Recall	0.741 / 0.712	0.090 / 0.070	0.472 / 0.466	0.528 / 0.523	0.747 / 0.758	0.647 / 0.681	0.541 / 0.544	0.432 / 0.461	0.766 / 0.813	0.638 / 0.679	0.749 / 0.770	0.379 / 0.379	0.148 / 0.148	0.626 / 0.641	0.326 / 0.371	0.263 / 0.320	0.435 / 0.464	0.404 / 0.441	
	F1-score	0.710 / 0.662	0.060 / 0.064	0.437 / 0.479	0.114 / 0.123	0.452 / 0.516	0.469 / 0.515	0.166 / 0.212	0.204 / 0.245	0.461 / 0.537	0.153 / 0.206	0.642 / 0.691	0.060 / 0.070	0.059 / 0.076	0.611 / 0.698	0.430 / 0.489	0.324 / 0.364	0.511 / 0.535	0.469 / 0.484	
	AUC	0.757 / 0.773	0.515 / 0.502	0.656 / 0.647	0.649 / 0.642	0.650 / 0.632	0.736 / 0.747	0.664 / 0.664	0.603 / 0.659	0.691 / 0.670	0.633 / 0.630	0.532 / 0.535	0.648 / 0.666	0.602 / 0.602	0.786 / 0.786	0.324 / 0.324	0.263 / 0.303	0.453 / 0.453	0.404 / 0.441	
TOP 1	#images	8,808 / 5,029	4,911 / 4,668	13,019 / 11,951	3,927 / 4,193	47,978 / 43,264	25,529 / 22,804	3,804 / 3,673	9,390 / 9,632	3,138 / 3,254	2,215 / 2,210	20,003 / 37,642	889 / 1,777	2,049 / 2,049	3,268 / 3,268	26,981 / 50,665	47,417 / 28,008			
	ChexNet	0.406 / 0.270	0.109 / 0.133	0.313 / 0.360	0.092 / 0.111	0.590 / 0.632	0.459 / 0.517	0.115 / 0.130	0.053 / 0.067	0.204 / 0.227	0.115 / 0.142	0.634 / 0.676	0.044 / 0.044	0.106 / 0.104	0.665 / 0.755	0.364 / 0.389	0.279 / 0.295	0.487 / 0.529	0.370 / 0.391	
	Recall	0.559 / 0.603	0.079 / 0.057	0.465 / 0.511	0.554 / 0.603	0.865 / 0.848	0.697 / 0.659	0.544 / 0.548	0.524 / 0.519	0.845 / 0.839	0.963 / 0.892	0.876 / 0.843	0.484 / 0.452	0.159 / 0.133	0.573 / 0.573	0.779 / 0.766	0.609 / 0.609	0.776 / 0.772	0.754 / 0.754	
	F1-score	0.477 / 0.373	0.092 / 0.072	0.425 / 0.453	0.158 / 0.177	0.701 / 0.725	0.554 / 0.579	0.190 / 0.210	0.096 / 0.118	0.329 / 0.357	0.205 / 0.245	0.735 / 0.750	0.080 / 0.080	0.116 / 0.116	0.496 / 0.509	0.496 / 0.515	0.375 / 0.375	0.478 / 0.483	0.478 / 0.483	
	AUC	0.711 / 0.641	0.063 / 0.053	0.436 / 0.447	0.123 / 0.138	0.454 / 0.516	0.479 / 0.512	0.170 / 0.210	0.025 / 0.025	0.470 / 0.530	0.613 / 0.628	0.653 / 0.680	0.688 / 0.688	0.081 / 0.071	0.619 / 0.694	0.487 / 0.487	0.334 / 0.363	0.564 / 0.526	0.480 / 0.487	
TOP 1	#images	4,390 / 4,940	2,587 / 4,050	7,357 / 11,902	2,000 / 5,579	25,328 / 44,160	19,477 / 22,913	3,241 / 3,241	6,694 / 6,694	5,655 / 11,006	16,002 / 20,590	2,901 / 6,644	7,148 / 9,632	1,797 / 1,797	4,744 / 28,008	47,434 / 97,347	173,664 / 88,757			
	ChexNet	0.422 / 0.249	0.116 / 0.095	0.319 / 0.358	0															

Appendix B.4. Chest X-ray Quality Results

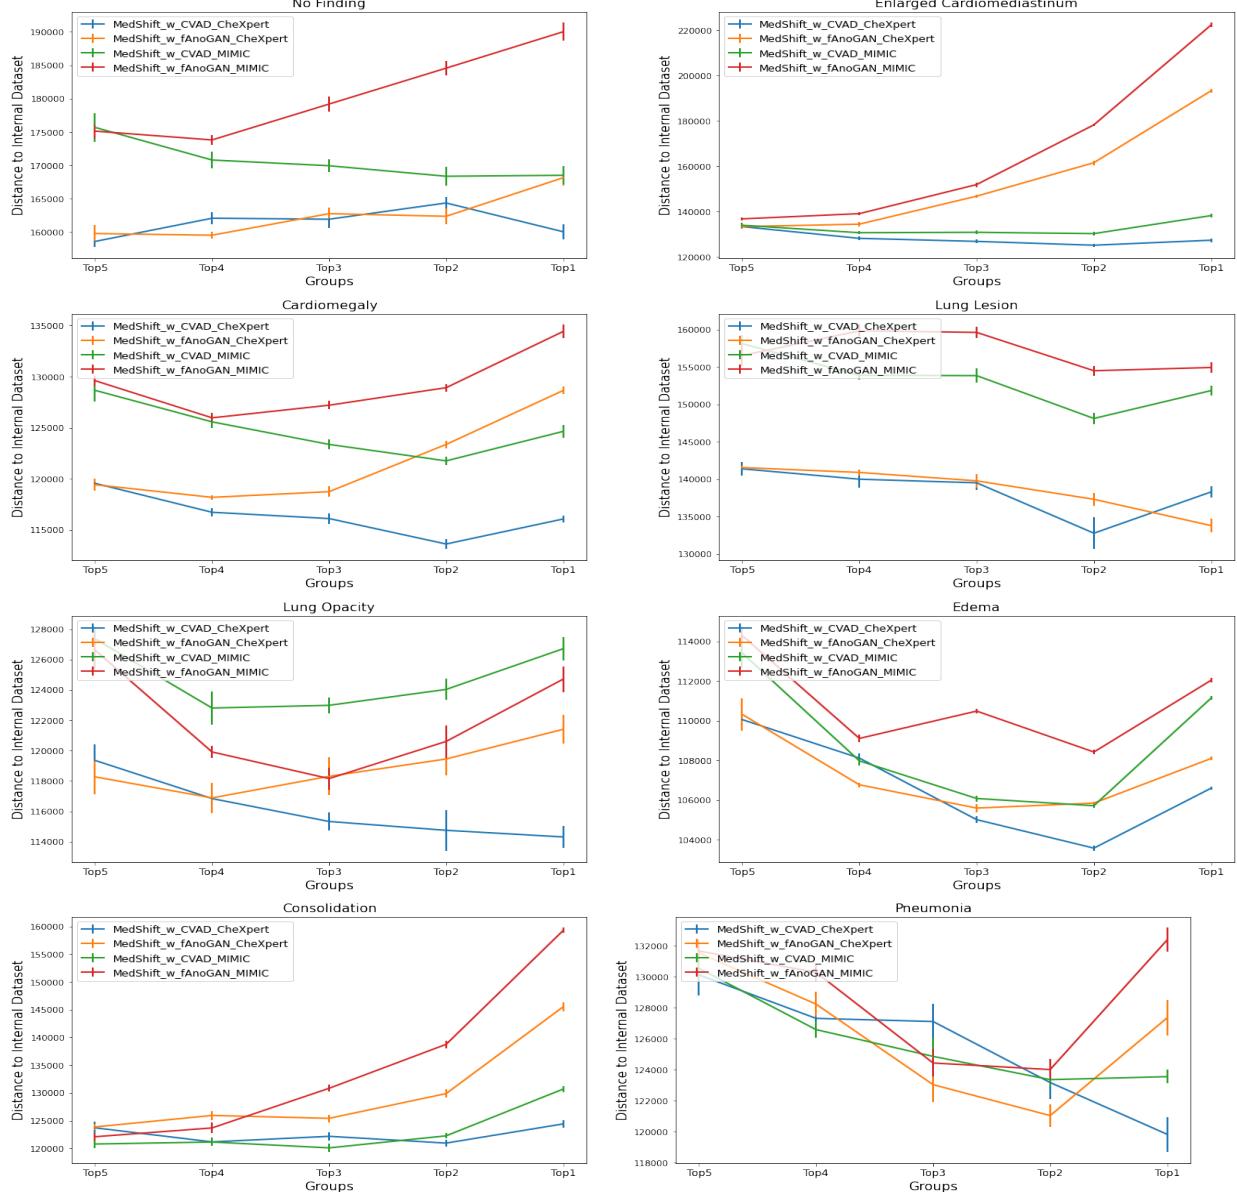


Figure B.27: Chest X-ray dataset quality measurement results. From left to right, top to bottom, there are plots for classes *No Finding*, *Enlarged Cardiomedastinum*, *Cardiomegaly*, *Lung Lesion*, *Lung Opacity*, *Edema*, and *Consolidation* respectively.

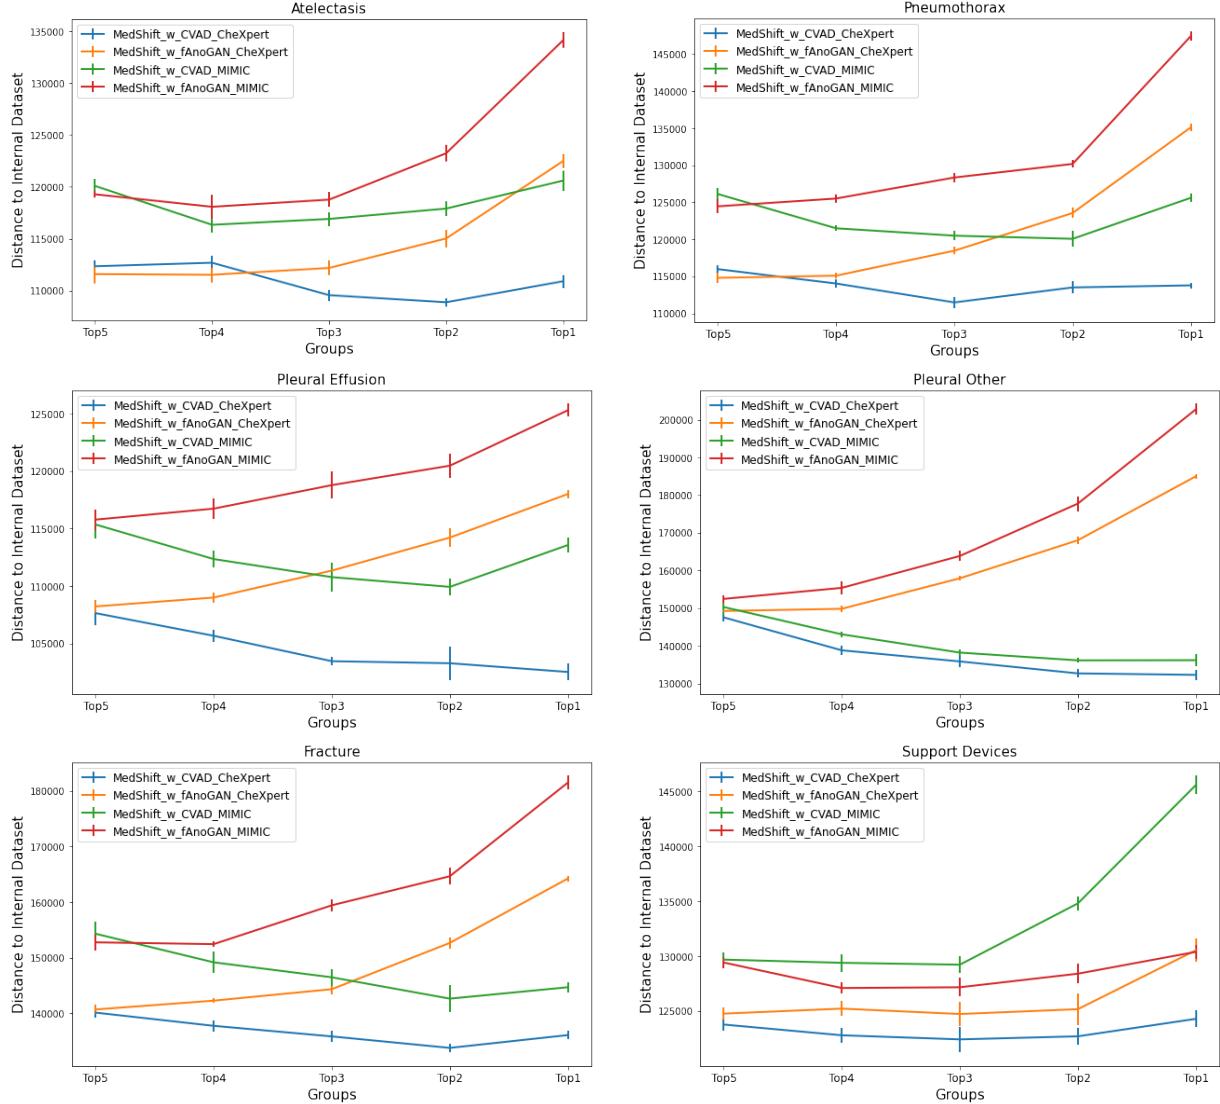


Figure B.28: More Chest X-ray dataset quality measurement results. From left to right, top to bottom, there are plots for classes *Atelectasis*, *Pneumothorax*, *Pleural Effusion*, *Pleural Other*, *Fracture* and *Support Devices* respectively.

References

Alvarez Melis, D., Fusi, N., 2020. Geometric dataset distances via optimal transport. *Advances in Neural Information Processing Systems* 33.

Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., et al., 2013. The cancer imaging archive (tcia): maintaining and operating a public information repository. *Journal of digital imaging* 26, 1045–1057.

Guo, X., Gichoya, J.W., Purkayastha, S., Banerjee, I., 2021. Cvad: A generic medical anomaly detector based on cascade vae. *arXiv preprint arXiv:2110.15811* .

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.

Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708.

Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., et al., 2019. CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison, in: *Proceedings of the AAAI conference on artificial intelligence*, pp. 590–597.

Johnson, A.E., Pollard, T.J., Berkowitz, S.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Mark, R.G., Horng, S., 2019. Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data* 6, 1–8.

Quiñonero-Candela, J., Sugiyama, M., Lawrence, N.D., Schwaighofer, A., 2009. *Dataset shift in machine learning*. Mit Press.

Rabanser, S., Günnemann, S., Lipton, Z.C., 2018. Failing loudly: An empirical study of methods for detecting dataset shift. *arXiv preprint arXiv:1810.11953* .

Rajpurkar, P., Irvin, J., Bagul, A., Ding, D., Duan, T., Mehta, H., Yang, B., Zhu, K., Laird, D., Ball, R.L., et al., 2017. Mura: Large dataset for abnormality detection in musculoskeletal radiographs. *arXiv preprint arXiv:1712.06957* .

Schlegl, T., Seeböck, P., Waldstein, S.M., Langs, G., Schmidt-Erfurth, U., 2019. f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis* 54, 30–44.

Schütte, A.D., Hetzel, J., Gatidis, S., Hepp, T., Dietz, B., Bauer, S., Schwab, P., 2021. Overcoming barriers to data sharing with medical image generation: a comprehensive evaluation. *NPJ digital medicine* 4.

Storkey, A., 2009. When training and test sets are different: characterizing learning transfer. *Dataset shift in machine learning* 30, 3–28.

Van Ooijen, P.M., 2019. Quality and curation of medical images and data, in: *Artificial intelligence in medical imaging*. Springer, pp. 247–255.

Van Panhuis, W.G., Paul, P., Emerson, C., Grefenstette, J., Wilder, R., Herbst, A.J., Heymann, D., Burke, D.S., 2014. A systematic review of barriers to data sharing in public health. *BMC public health* 14, 1–9.

Villani, C., 2009. *Optimal transport: old and new*. volume 338. Springer.

Wang, R., Chaudhari, P., Davatzikos, C., 2021. Embracing the disharmony in medical imaging: A simple and effective framework for domain adaptation. *arXiv preprint arXiv:2103.12857* .

Yamoah, G.G., Cao, L., Wu, C.W., Beekman, F.J., Vandeghinste, B., Mannheim, J.G., Rosenhain, S., Leonardic, K., Kiessling, F., Gremse, F., 2019. Data curation for preclinical and clinical multimodal imaging studies. *Molecular imaging and biology* 21, 1034–1043.

Yuan, Z., Yan, Y., Sonka, M., Yang, T., 2021. Large-scale robust deep auc maximization: A new surrogate loss and empirical studies on medical image classification, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3040–3049.