

Multi-Label Generalized Zero Shot Learning for the Classification of Disease in Chest Radiographs

Nasir Hayat

Engineering Division

NYU Abu Dhabi, Abu Dhabi, UAE

NH2218@NYU.EDU

Hazem Lashen

Engineering Division

NYU Abu Dhabi, Abu Dhabi, UAE

HL3372@NYU.EDU

Farah E. Shamout

Engineering Division

NYU Abu Dhabi, Abu Dhabi, UAE

FS999@NYU.EDU

Abstract

Despite the success of deep neural networks in chest X-ray (CXR) diagnosis, supervised learning only allows the prediction of disease classes that were seen during training. At inference, these networks cannot predict an unseen disease class. Incorporating a new class requires the collection of labeled data, which is not a trivial task, especially for less frequently-occurring diseases. As a result, it becomes inconceivable to build a model that can diagnose all possible disease classes. Here, we propose a multi-label generalized zero shot learning (CXR-ML-GZSL) network that can simultaneously predict multiple seen and unseen diseases in CXR images. Given an input image, CXR-ML-GZSL learns a visual representation guided by the input’s corresponding semantics extracted from a rich medical text corpus. Towards this ambitious goal, we propose to map both visual and semantic modalities to a latent feature space using a novel learning objective. The objective ensures that (i) the most relevant labels for the query image are ranked higher than irrelevant labels, (ii) the network learns a visual representation that is aligned with its semantics in the latent feature space, and (iii) the mapped semantics preserve their original inter-class representation. The network is end-to-end trainable and requires no independent pre-training for the offline feature extractor. Experiments on the NIH Chest X-ray dataset show that our network outperforms two strong baselines in terms of recall, precision, f1 score, and area under the receiver operating characteristic curve. Our code is publicly available at: <https://github.com/nyuad-cai/CXR-ML-GZSL.git>

1. Introduction

Deep learning has enabled the development of computer aided diagnostic systems that can classify diseases in medical images with human level precision (Qin et al., 2018; Rajpurkar et al., 2018; Hwang et al., 2019). One limitation of these networks is that they require a large amount of annotated data for training, which is often a tedious and expensive process that requires expert knowledge. Along with this, it becomes even more challenging to collect a significant amount of data for naturally occurring rare diseases or novel outbreaks, such as COVID-19 (Rezaei and Shahidi, 2020). Therefore, it becomes impractical to annotate a

substantial amount of training data for all possible diseases in order to train a deep learning network. Due to the data limitations, at inference, a deep learning network is unable to classify disease classes that were unseen during training (Larochelle et al., 2008). However, radiologists can recognize previously unseen diseases by relying on the disease attributes learned through medical literature.

Zero shot learning (ZSL) has the potential to mimic the radiologist’s behavior by recognizing unseen diseases using other sources of information (Larochelle et al., 2008). It is arguably the most extreme case of learning with minimal supervision. In brief, given a query image, ZSL methods find the correspondence between the visual representation of the image and its semantic representation (Xian et al., 2017). ZSL has achieved impressive results on natural images (Rahman et al., 2020). Nevertheless, most of the proposed methods assign only a single label to the each query image, and in many implementations, the assigned label can only be drawn from the seen classes only (Larochelle et al., 2008; Long et al., 2017). In medical imaging tasks, such as the classification of disease in chest X-ray images, an image may contain more than one disease and the labels could be from either the seen or unseen classes (Wang et al., 2017). This limits the direct application of existing single-label ZSL methods for medical multi-label classification tasks. On the other hand, *Multi-label ZSL* (ML-ZSL) allows the assignment of multiple labels to each image. Under Multi-label *Generalized ZSL* (ML-GZSL), the goal is to assign multiple labels to a query image that could be in both seen and unseen classes (Scheirer et al., 2013). Existing ML-GZSL approaches developed for natural images perform the nearest neighbor search in the semantic space:

$$P_x^c = hf_x, Wi, \quad (1)$$

where h, i is the cosine similarity function, f_x are the visual features of a query image x , and W are the semantic embeddings for all possible classes. Equation 1 computes the similarity scores between the visual features and the semantic embeddings of each possible class. The similarity score indicates the relevance of the label to the image. More appropriately, these approaches perform the nearest neighbor search by mapping the aggregated visual features to the semantic space (Rahman et al., 2020; Huynh and Elhamifar, 2020). However, these methods extract a fixed visual representation of the image from a pre-trained visual encoder or a detection network. Additionally, projecting these extracted visual features to the semantic space shrinks the diversity of the visual information, which gives rise to inherent limitations such as the hubness problem (Dinu et al., 2015).

To overcome the aforementioned challenges, we propose the CXR-ML-GZSL network for the classification of disease in chest X-rays (commonly referred to as CXR). In summary, our network consists of a visual feature encoder and two projection modules for each modality, which are the visual features and semantic embeddings. These modules project the independent modalities to the common latent space where the correspondence among them can be established. We evaluate our proposed method on the NIH chest X-ray dataset (Wang et al., 2017) and the results demonstrate that CXR-ML-GZSL outperforms baseline models. In this work, we present the following three contributions:

- From a technical perspective, we design an end-to-end trainable network that jointly learns the visual representation and aligns it with the semantic representation. Our design does not require any offline training for the visual feature encoder.

- Additionally, we propose a novel learning objective for the CXR-ML-GZSL to optimize the learned representations. The objective function simultaneously learns the predictive distribution, ensures that the visual representations are well centered around the class semantics in the latent embedding space, and enforces a constraint to preserve the original representation of class semantics in the latent space.
- From a medical imaging perspective, to the best of our knowledge, we are the first to propose a ML-GZSL framework for the classification of disease in chest X-rays.

We summarize related work in Section 2, the methodology in Section 3, the experiments in Section 4, the results in Section 5, the discussion in Section 6, and finally, the conclusion in Section 7.

Generalizable Insights about ML in the Context of Healthcare

Conventional multi-label chest X-ray classification models are constrained by the availability of data and its annotations. Our work overcomes the challenge of collecting large-scale annotated datasets by leveraging the use of rich medical literature, since it is the main knowledge source for all the discovered diseases by the medical community. This highlights the role of multi-modal learning in healthcare applications. Although we focus on chest X-rays, the network design can be potentially generalized to any medical imaging task, since the semantic encoder is task-agnostic. Improving the diagnosis of unseen diseases at inference has the potential to save patient lives.

2. Related work

2.1. Inductive & transductive ZSL

ZSL classifies previously unseen classes during training by transferring knowledge from seen classes. It relies on class semantics to bridge the gap between seen and unseen classes. The semantics are either obtained through manually annotated class attributes and textual descriptions embedded in a high-dimensional space (Annadani and Biswas, 2018; Zhang et al., 2016), or by extracting the label’s semantic vectors using Word2vec or Glove (Mikolov et al., 2013; Elhoseiny et al., 2017). Previously proposed ZSL training methods can be categorized into *inductive ZSL*, which strictly trains on the data related to seen classes only (Xian et al., 2017), and *transductive ZSL*, which assumes that unlabeled visual examples for unseen classes are also available during training (Ali Cheraghian and Petersson, 2019). However, transductive ZSL violates the assumption that unseen classes are not seen during training, as it is impractical to assume that visual data of unseen classes is available during training. Therefore, addressing the ZSL under the inductive setting would be more realistic and renders a practical solution for the classification of chest X-rays.

2.2. Multi-label generalized zero shot learning

Few of the existing work address multi-label classification with natural images. Lee et al. (2018) constructed structured knowledge graphs by exploiting the semantic relationships in WordNet (Miller, 1995) to assign labels for unseen classes. A recent work, deep0tag

(Rahman et al., 2020), extracts a bag of localized patches using a pre-trained detection network, aggregates these proposals, and then maps them in the semantic space to find the correspondence between seen and unseen classes. This approach however, requires a large-scale dataset with bounding box annotations to pre-train FasterRCNN (Ren et al., 2015; Hayat et al., 2021) for meaningful proposal generation. Huynh and Elhamifar (2020) proposed a shared multi-attention model to learn class agnostic multi-attention features by extracting a set of cropped regions’ features using a pre-trained CNN. These attention features are projected to the semantic space to find the relevance within labels. Although most of these proposed approaches perform reasonably well on natural images, they are either not directly applicable to chest X-ray classification or rely on offline visual feature extractors trained on ImageNet (Russakovsky et al., 2015). As shown in the work by Raghu et al. (2019), transferring knowledge from ImageNet provides little to no gain for radiograph diagnosis. As such, it is equally important to learn a meaningful visual representation for the visual encoder of CXR-ML-GZSL. Here, we propose to jointly learn the visual representation and align these representation with the semantics in an end-to-end network.

2.3. Deep learning for chest radiographs

Multi-label classification assumes that multiple labels could appear in a single image. This is quite common for the classification of chest X-rays, where each image could contain multiple diseases. Previously, Rajpurkar et al. (2017) trained a Densenet-121 (Huang et al., 2017) by formulating a standard multi-label classification problem. However, their approach does not capture meaningful correlations among the different classes. Guan et al. (2018) used an attention mechanism by cropping the regions of interest extracted from attention heat maps provided by a global image encoder, in order to train a local network. They eventually fuse the features after the last pooling layers from both the global and local branches to assign classification scores. Yao et al. (2018) used a DenseNet to encode a global representation of the X-rays and adopted a long-short term memory network to learn inter-dependencies among classes for better diagnosis (Hochreiter and Schmidhuber, 1997). Despite their promising performance, these methods rely on large amounts of labeled data. Additionally, at inference, they can only predict classes that were seen during training.

There is very limited work on using ZSL for chest radiography. Recently, Paul et al. (2021) proposed a GZSL framework using a two-branch autoencoder that combines external knowledge from three textual resources: chest X-ray reports, CT reports, and manually defined visual traits. Their design has a few limitations. First, the network predicts a single label, although the datasets have multiple labels per image. Second, it is transductive and hence assumes that the reports and unlabeled images for unseen cases are available during training, which contradicts the unseen assumption. Third, their semantic representation is learned from radiology and CT reports of already observed cases and is guided by manually annotated class attributes. Therefore, due to this data dependency on reports and manual annotations, its application is confined towards a closed set of seen and already observed unlabeled unseen classes. Unlike the work by Paul et al. (2021), our network can assign multiple labels to a query image and can generalize to an open set of unseen classes since it does not rely on existing reports and manual annotations.

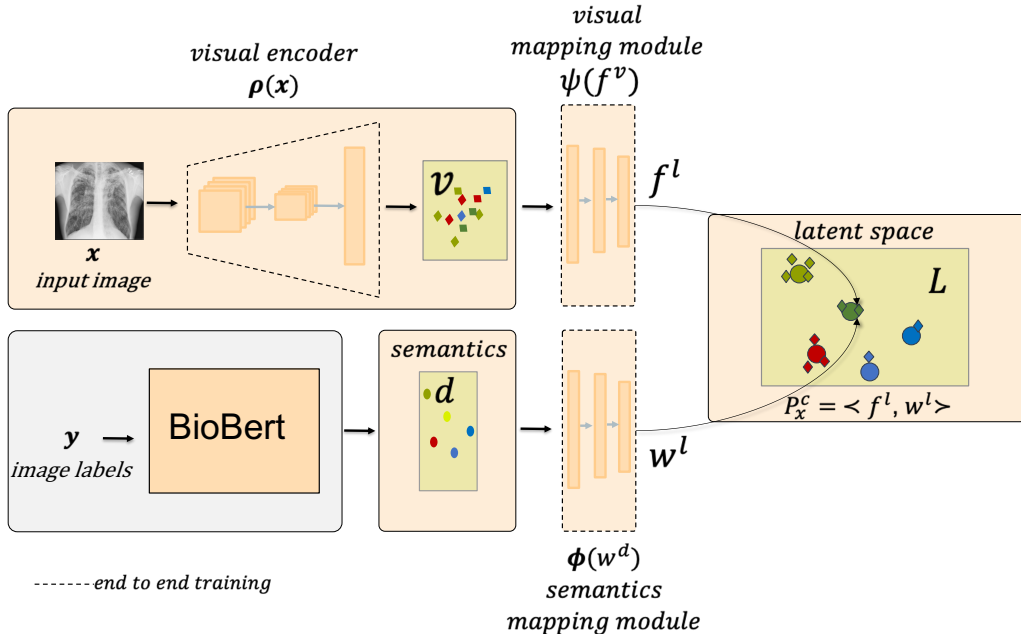


Figure 1: **Overview of our CXR-ML-GZSL model for learning visual representations of chest X-rays.** An overview of our network for chest X-ray images. It includes a trainable visual encoder and v - and d -dimensional visual and semantics spaces, respectively. For an input image x and its labels y , the network learns a visual representation guided by semantics extracted by BioBERT. We perform end-to-end training of the visual encoder, visual mapping module, and the semantics mapping module, as indicated by the black dashed line.

3. Methodology

3.1. Problem formulation

Consider a set X^S that contains the training images of seen classes only. Each $x \in X^S$ is associated with a set of labels \mathbf{y}_x , where $y_x^i \in \{0, 1\}_{i=1}^S$ and one indicates the presence of the i^{th} disease among the S seen classes during training. We denote the corresponding seen and unseen classes as $\mathcal{Y}^S = \{1, \dots, S\}$ and $\mathcal{Y}^U = \{S+1, \dots, C\}$, respectively, where C is the total number of seen and unseen classes. Note here that both of these sets are disjoint such that $\mathcal{Y}^S \cap \mathcal{Y}^U = \emptyset$. The training objective is to learn the visual representation of x guided by the semantics of its labels \mathbf{y}_x . At inference, given a test image x_{test} , the goal is to predict $\mathbf{y}_{x_{\text{test}}}$, where $y_{x_{\text{test}}}^i \in [0, 1]_{i=1}^C$. In the following sections, we describe the proposed network’s architecture and the learning objective.

3.2. Network architecture

An overview of the network architecture is visualized in Figure 1. It consists of a trainable visual encoder, fixed semantic encoder, and alignment modules. We provide more details about each component and the loss function in this section.

Visual encoder: To learn the visual representation, we first define a visual encoder $\rho(x) : \mathbb{R}^{w \times h \times c} \rightarrow \mathbb{R}^v$ which computes f^v , a v -dimensional visual representation of the input image

x . Subsequently, a sub-module, referred to as the *visual mapping module*, processes the visual representation $(f^V) : \mathbb{R}^V \rightarrow \mathbb{R}^l$ and maps it to a latent manifold l , which is jointly learned with the semantic information.

Semantics: Let the semantic embeddings for seen classes be $W^S = \{w_1, w_2, \dots, w_S\}$, where w_i is a d -dimensional semantic representation for a class label $i \in \mathcal{Y}^S$, and it is extracted from the penultimate layer of BioBert (Lee et al., 2019) for all the trainable labels. To the best of our knowledge, our work is the first to adopt BioBert for ZSL problems in healthcare. Our preference to use BioBert was because of the fact that it learns effective contextualized word embeddings specific to the biomedical corpora.

We define, $(w^d) : \mathbb{R}^d \rightarrow \mathbb{R}^l$, referred to as the *semantics mapping module*, which learns to map a d -dimensional semantic embedding to the joint latent manifold l .

With respect to the proposed architecture, we redefine Equation 1 as:

$$P_x^S = h((x)), (W^S)^T i, \quad (2)$$

where P_x^S represents the relevance scores of the training labels, i.e. $P_x^S = \{p_1, p_2, \dots, p_S\}$. Those relevance scores represent the similarity between the visual input and each possible label.

Training objective: We formulate a training objective to jointly optimize the parameters of the network as follows:

$$\min_{\phi, \rho, \psi} L = L_{rank} + \gamma_1 L_{align} + \gamma_2 L_{con}, \quad (3)$$

where γ_1 and γ_2 are the regularization parameters for the loss components L_{align} and L_{con} , respectively. In the next sections, we formally define each of the loss components.

3.3. L_{rank} Ranking loss of relevance scores

During training, the network assigns the relevance scores $P_x^S = \{p_1, p_2, \dots, p_S\}$ for each of the S seen classes for an input image x . Given the ground-truth labels \mathbf{y}_x , where $y_x^i \in \{0, 1\}$, we denote Y_p as the set of positive ground-truth labels (i.e., diseases that are present in the image $y_x^i = 1$) and Y_n as the set of negative ground-truth labels (i.e., diseases that are not present in the image $y_x^i = 0$). Let us assume that p_{y_p} and p_{y_n} are the computed scores for a positive label and a negative label, respectively. Since this is a multi-label classification task, ideally, we wish to meet two conditions. First, p_{y_p} should be higher than p_{y_n} , and second, the difference between p_{y_p} and p_{y_n} should be at least a margin value of δ . Therefore, we formulate an image-level penalty score using the margin-based ranking loss as follows:

$$L(P_x^S, \mathbf{y}_x) = \frac{1}{S} \sum_{y_p \in Y_p} \sum_{y_n \in Y_n} \max(\delta + (p_{y_n} - p_{y_p}), 0). \quad (4)$$

In the above formulation, the loss is 0 if $p_{y_p} > p_{y_n}$ by a margin of at least δ . It would inflict a penalty of at least δ when $p_{y_p} \leq p_{y_n}$. For all training samples, we average the image-level penalty scores over all images:

$$L_{rank} = \frac{1}{N} \sum_{x \in \mathcal{X}^s} L(P_x^S, \mathbf{y}_x), \quad (5)$$

where N is the total number of images.

3.4. L_{align} Alignment loss of visual-semantic representations

To align the visual representations with their relevant semantic representations during training, we formulate an alignment loss for the two modalities as follows:

$$L_{align} = \frac{1}{N} \left(1 - \sum_{x \in \mathcal{X}^s} h \left(\langle f(x), w_x \rangle \right) \right), \quad (6)$$

where w_x is the corresponding semantic embeddings for the input image x and h, i defines the cosine similarity function. In case multiple labels are present in the image, then their semantic embeddings are averaged in w_x . The averaging allows the alignment of the visual representation with its semantics in the multi-label setting.

3.5. L_{con} Semantics inter-class consistency regularizer

The semantic and visual representations stem from two separate modalities. The semantic representations are learned through medical textual literature, while the visual representations are learned from chest radiograph. To bridge the gap between both modalities, we learn two mapping functions that project the representations to a common domain. While the visual representations are fine-tuned during training, the semantic representations are fixed after being extracted from the natural language encoder. Mapping the semantic representations to the latent space may lead to a loss in inter-class relevance. Hence, we aim to ensure that the semantic manifold remains consistent throughout the projection process by leveraging the semantic inter-class relationship. To achieve this, we design an L_1 regularization term within the learning objective which considers the similarity between the classes between the original space and the projected space, such as:

$$L_{con} = \sum_{w_i \in \mathcal{W}} \sum_{\substack{w_j \in \mathcal{W} \\ j \neq i}} k |h(w_i) - h(w_j)| \quad (7)$$

where w_i and w_j are the original semantic representations for two classes and $h(\cdot)$ computes their projected representations. Ideally, the cosine similarity between the two classes in the original space should be the same in the projected space, and hence the loss would be 0.

3.6. Inference

After training the network, we obtain the fine-tuned image encoder $f(x)$ and mapping modules f^v and w^d , which project both the encoded visual features and the semantic embeddings to the joint latent space, respectively. At inference, for a test image $x \in \mathcal{X}^C$, we update \mathcal{W} in Equation 2 to include both seen and unseen semantic embeddings. Contrary to conventional multi-label classification, after a simple modification of \mathcal{W} , we can get the prediction scores $P_x^C = \{p_1, p_2, \dots, p_C\}$ for a set of C classes comprising both seen and unseen, such that:

$$P_x^C = h \left(\langle f(x), w^C \rangle \right). \quad (8)$$

In Equation 8, \mathcal{W}^C represents the updated \mathcal{W} and includes the semantic embeddings of both seen and unseen classes.

4. Experimental settings

4.1. Dataset

To evaluate the proposed network CXR-ML-GZSL, we ran the experiments on the publicly available NIH Chest X-ray dataset (Wang et al., 2017). The dataset contains 112,120 frontal X-ray images from across 30,805 unique patients. We split the dataset randomly to obtain a training set (70%), validation set (10%), and test set (20%). Each image is associated with 14 possible classes. We include all 14 classes and perform a random split of seen (10 in total) and unseen (4 in total) classes. The seen classes were Atelectasis, Effusion, Infiltration, Mass, Nodule, Pneumothorax, Consolidation, Cardiomegaly, Pleural Thickening, and Hernia, and the unseen classes were Edema, Pneumonia, Emphysema, and Fibrosis. We strictly exclude all images from the training set that are associated with any unseen class label, in accordance with the inductive setting. The final training set contained 30,758 images, the validation set contained 4,474 images, and the test set contained 10,510 images.

4.2. Model training and selection

In this section, we describe the experimental setting for CXR-ML-GZSL. To encode the visual information, our method is designed to work with any state-of-the-art image convolutional neural network. We ran all of our experiments with Densenet-121, since it has shown the best results for chest X-ray classification (Rajpurkar et al., 2017). To do so, we removed the final classification layer and used the resulting network as the visual encoder (x) that yields the visual representation $f^v \in \mathbb{R}^{1024}$.

The visual mapping module was parameterized as a three layer feed-forward neural network formally defined as $f^v \rightarrow \text{fc1} \rightarrow \text{Relu} \rightarrow \text{fc2} \rightarrow \text{Relu} \rightarrow \text{fc3} \rightarrow f^l$, where fc1 is a fully connected layer which learns the weight matrix $\mathbf{W}_{\text{fc1}} \in \mathbb{R}^{1024 \times 512}$ and a bias vector $\mathbf{b}_{\text{fc1}} \in \mathbb{R}^{512}$. The subsequent layer learns the weight matrix $\mathbf{W}_{\text{fc2}} \in \mathbb{R}^{512 \times 256}$ and a bias vector $\mathbf{b}_{\text{fc2}} \in \mathbb{R}^{256}$. The last layer learns the weight matrix $\mathbf{W}_{\text{fc3}} \in \mathbb{R}^{256 \times 128}$ and a bias vector $\mathbf{b}_{\text{fc3}} \in \mathbb{R}^{128}$, which is eventually projected into the joint latent space with the semantic embeddings. The semantic mapping module follows a symmetric architecture to visual mapping module. Our design of the mapping modules enables the network to work with embeddings extracted from different architectures, which are eventually mapped to a common space irrespective of the original dimensions.

We trained the network using the Adam optimizer (Kingma and Ba, 2017) for a total of 100 epochs while reducing the initial learning rate (lr) by a factor of 0.01 once the validation loss stagnates for 10 epochs. It took about 8 hours to train a single model on an NVIDIA Quadro RTX 6000 GPU. Through out our experiments, we set $\delta = 0.5$ in Equation 4. To optimize the other hyperparameters (γ_1, γ_2 in Equation 3, and the learning rate), we ran multiple experiments by randomly selecting $\gamma \in \{0.1, 0.01, 0.05\}$ and $lr \in \{0.0001, 0.00005, 0.00001\}$ and then chose the best performing model on the validation set in terms of the mean harmonic AUROC. The grid search ranges were selected based on the results of preliminary experiments. Our codebase was developed using the Pytorch deep learning library (Paszke et al., 2019).

4.3. Performance metrics

We used evaluation metrics that are commonly reported in ML-GZSL methods (Lee et al., 2018; Huynh and Elhamifar, 2020). We computed overall precision, recall and f1 scores for the top k predictions where $k \in \{2, 3\}$ for GZSL. We selected a smaller value of k since our dataset has fewer number of classes compared to natural imaging datasets (Chua et al., July 8-10, 2009; Kuznetsova et al., 2020). Additionally, we report the average area under the receiving operating characteristic curve (AUROC) for seen and unseen classes and its harmonic mean, since computing recall for top k predictions may not be a sufficient indicator of class-wise performance. It is important to note that the harmonic mean measures the inherent bias of GZSL methods towards the seen classes.

Table 1: **Performance evaluation of the network compared to state-of-the-art baselines on the test set.** Multi-label GZSL performance on the NIH chest X-ray dataset in terms of recall (r), precision (p) and $f1$ scores at the top k predictions. We report the **AUROC** for seen (**S**) classes, unseen (**U**) classes, and their harmonic mean (**H**). Our proposed end-to-end trained network CXR-ML-ZSL is denoted as " OUR_{e2e} ". The numbers in **bold** are the best values in the respective column.

Method	k=2			k=3			AUROC		
	r@k	p@k	f1@k	r@k	p@k	f1@k	S	U	H
<i>LESA</i> (2020)	0.14	0.10	0.03	0.21	0.11	0.05	0.51	0.50	0.50
<i>MLZSL</i> (2018)	0.20	0.19	0.16	0.30	0.17	0.20	0.72	0.54	0.62
<i>OUR_{e2e}</i>	0.36	0.33	0.32	0.47	0.28	0.34	0.79	0.66	0.72

Table 2: **Class-wise AUROC comparison across all disease classes in the test set.** In this table, we show the performance AUROC across all seen and unseen classes. Our proposed method (OUR_{e2e}) shows significant gains for a number of individual classes. Compared with the second best method (MLZSL), our method shows an AUROC gain of 22.2% for the mean performance for unseen classes. The *italicized* classes are unseen and the **bold** numbers represent the best AUROC in the respective column.

Method	seen mean	unseen mean	atelectasis	cardiomegaly	effusion	infiltration	mass	nodule	pneumothorax	consolidation	pleural thickening	hernia	<i>pneumonia</i>	<i>edema</i>	<i>emphysema</i>	<i>fibrosis</i>
<i>LESA</i> (2020)	0.51	0.52	0.54	0.51	0.48	0.50	0.52	0.50	0.51	0.54	0.52	0.47	0.50	0.54	0.50	0.53
<i>MLZSL</i> (2018)	0.72	0.54	0.70	0.76	0.76	0.64	0.67	0.69	0.76	0.68	0.66	0.90	0.52	0.51	0.58	0.54
<i>OUR_{e2e}</i>	0.79	0.66	0.76	0.90	0.83	0.70	0.80	0.75	0.83	0.69	0.72	0.90	0.62	0.67	0.74	0.60

5. Results

5.1. Comparison to baseline models

We compared the performance of proposed approach (OUR_{e2e}), with two state-of-the-art ML-GZSL methods: *LESA* (Huynh and Elhamifar, 2020) and *MLZSL* (Lee et al., 2018). Table 1 summarizes the performance results on the test set. The results show that our method performs best across all metrics, reaching an AUROC of 0.66 on the unseen classes and a harmonic AUROC of 0.72 across both seen and unseen classes. *LESA* achieved the




Figure 2: **Qualitative results for the top three predictions per image within the test set.** The labels in green, orange and red represent true positives, false negatives, and false positives respectively. The green and orange labels combined represent the ground-truth labels.

worst performance, while MLZSL achieved the second best performance. In particular, our method outperforms MLZSL by a significant margin, with a percentage change of 73.68% in terms of precision@2, for example.

Table 2 compares the AUROC values for each class against the state-of-the-art methods. Our method achieves the best performance across all seen classes compared to the baseline methods, except for Hernia, where it achieves a comparable performance to MLZSL (0.90 AUROC). Across the unseen classes, it achieves the best AUROC performance compared to both baseline methods.

Figure 2 shows example predictions for 9 test chest X-ray images. We select the top three predictions on each test image. We observe that our method is still able to predict the unseen classes even when the number of ground-truth labels is high, such as greater than 3. This demonstrates the effectiveness of the proposed method under challenging settings of simultaneously predicting multiple seen and unseen classes.

5.2. Ablation studies

We performed two ablation studies using the validation set. In all ablation studies, we set the initial learning rate to 0.0001, $\gamma_1 = 0.01$, and $\gamma_2 = 0.01$. Table 3 shows the AUROC

Table 3: **Ablation study for the loss function using the validation set.** The individual contribution of each loss component towards the performance in terms of AUROC of multi-label generalized ZSL. We progressively integrate each component, improving the performance across unseen classes, while the performance of seen classes remains persistent. The **bold** numbers represent the best AUROC in the respective column.

L_{rank}	L_{align}	L_{con}	Seen	Unseen	Harmonic Mean
✓			0.790	0.574	0.665
✓	✓		0.791	0.601	0.683
✓	✓	✓	0.783	0.663	0.718

Table 4: **Comparison of training strategies on the validation set.** We compare different training strategies and report the AUROC for our proposed method compared to offline visual encoders trained on ImageNet and the NIH chest X-ray datasets. End-to-end training achieves the best performance across the seen and unseen classes. The **bold** numbers represent the best AUROC in the respective column.

Training method	Seen	Unseen	Harmonic Mean
Pre-training with Imagenet	0.721	0.608	0.660
Pre-training with NIH	0.771	0.609	0.681
End-to-end training	0.783	0.663	0.718

for seen and unseen classes and their harmonic mean with different formulations of the objective function. In particular, we ran experiments to evaluate the individual contribution of incorporating the alignment loss and the semantic consistency regularizer. While the performance of seen classes remains consistent through out the experiments (0.783-0.791 AUROC), an improvement in the unseen AUROC can be seen when incorporating the alignment loss L_{align} , which ensures that the learned visual representations are centered around the class semantics. Additionally, preserving a consistent semantic manifold using the L_{con} regularizer contributed significantly to improving the performance among the unseen classes. Therefore, we used the loss component that incorporated all of the ranking loss, alignment loss, and the semantics regularizer to train the final model.

We also investigated the impact of using an offline visual encoder compared to our proposed end-to-end training approach. We ran experiments by freezing the visual encoder pre-trained on ImageNet for the purpose of extracting the visual features. Then, we separately trained the visual encoder on the NIH Chest X-ray dataset with seen classes only and eventually used it to extract fixed visual features. Table 4 shows the performance of these approaches compared to our end-to-end training approach. It is interesting to note that the end-to-end approach performs reasonably well compared to the other approaches. This emphasizes the point that learning a well-aligned visual representation guided by the semantic embeddings helps in improving the performance, especially for unseen classes where the AUROC is boosted significantly.

6. Discussion

Recent advancements in the field of deep learning for medical imaging are highly dependent upon the availability of large-scale datasets. In this study, we present a promising approach to develop a multi-label diagnostic network that can classify unseen classes using generalized zero shot learning. Our CXR-ML-GZSL network leverages the contextualized semantics learned through rich medical literature and learns visual representations guided by these semantics, with a unique learning objective. We train our model for a set of seen classes and then evaluate it on a set of held-out seen and unseen classes using the NIH chest X-ray dataset. The assumption is that the unseen classes are never exposed to the model during training, in order to mimic a real-life scenario of classifying unseen “rare” diseases. The results show that the proposed network generalizes well to both seen and unseen classes and achieves impressive performance gains over previous state-of-the-art methods. Note that for clinical deployment, we suggest that the clinical practitioner is shown the list of diseases ranked from most-to-least likely based on the prediction scores. Clinicians can also define a threshold to convert the prediction scores (P_X^c) to binary using sensitivity and specificity analysis.

A recent work by Paul et al. (2021) proposed and evaluated a GZSL method for chest X-ray imaging. They include a subset of 9 classes out of 14 based on the availability of CT reports. Out of these, they randomly select 6 classes as seen and the remaining 3 are used as unseen classes. However, their assumption of the availability of unseen data during training (including CT reports, radiology reports & chest radiographs) violates the unseen characteristic for ZSL and may not give meaningful insight on the performance for unseen classes. Therefore, to fully assess the robustness of GZSL towards unseen classes, we strictly ensure that no auxiliary data related to unseen classes is used during the training. Due to evident differences in assumptions on the training settings, we are unable to directly compare the results between our work and the work of Paul et al. (2021).

Limitations

To evaluate the proposed network, we are restricted to using an existing dataset. This has several limitations. First, the dataset has a fewer number of classes compared to natural imaging benchmarks. To evaluate the robustness of our proposed method and future works, we emphasize on the need for creating a more challenging benchmark dataset with a sufficiently large number of classes. In addition, we randomly selected seen and unseen classes to evaluate the method for chest X-ray classification. Although this evaluation is based on the assumption that the four classes were unseen, the method needs to be further evaluated with other seen and unseen splits, potentially on other medical tasks using different types of imaging, using data from other institutions to assess for generalizability, and on standard computer vision benchmarks. Additionally, we perform hyperparameter tuning for the learning rate and γ parameters only. We expect the results to further improve if we tune other hyperparameters, like batch size and dimensions of mapping modules.

7. Conclusion

In this work, we propose a multi-label generalized zero shot learning (CXR-ML-GZSL) network for chest X-ray classification. Through our experimentation, we show that at inference, our network can simultaneously assign multiple labels both from seen and unseen classes. Since our network is strictly restricted to be trained on seen classes and no auxiliary information in the form of either unlabeled chest X-rays or clinical reports related to unseen classes are provided during training, we believe that CXR-ML-GZSL has substantial potential to diagnose unseen classes at inference especially for rare or emerging diseases that are challenged by the lack of annotated data.

References

- Dylan Campbell Ali Cheraghian, Shafin Rahman and Lars Petersson. Transductive zero-shot learning for 3d point cloud classification. *arXiv preprint arXiv:1912.07161*, 2019.
- Yashas Annadani and Soma Biswas. Preserving semantic relations for zero-shot learning. *CoRR*, abs/1803.03049, 2018. URL <http://arxiv.org/abs/1803.03049>.
- Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yan-Tao Zheng. Nus-wide: A real-world web image database from national university of singapore. In *Proc. of ACM Conf. on Image and Video Retrieval (CIVR’09)*, Santorini, Greece., July 8-10, 2009.
- Georgiana Dinu, Angeliki Lazaridou, and Marco Baroni. Improving zero-shot learning by mitigating the hubness problem. *arXiv preprint arXiv:1412.6568*, 2015.
- Mohamed Elhoseiny, Yizhe Zhu, Han Zhang, and Ahmed M. Elgammal. Link the head to the “beak”: Zero shot learning from noisy text description at part precision. *CoRR*, abs/1709.01148, 2017. URL <http://arxiv.org/abs/1709.01148>.
- Qingji Guan, Yaping Huang, Z. Zhong, Zhedong Zheng, L. Zheng, and Y. Yang. Diagnose like a radiologist: Attention guided convolutional neural network for thorax disease classification. *arXiv preprint arXiv:1801.09927*, abs/1801.09927, 2018.
- Nasir Hayat, Munawar Hayat, Shafin Rahman, Salman Khan, Syed Waqas Zamir, and Fahad Shahbaz Khan. Synthesizing the unseen for zero-shot object detection. In Hiroshi Ishikawa, Cheng-Lin Liu, Tomas Pajdla, and Jianbo Shi, editors, *Computer Vision – ACCV 2020*, pages 155–170, Cham, 2021. Springer International Publishing.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
- G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269, 2017. doi: 10.1109/CVPR.2017.243.

- Dat Huynh and Ehsan Elhamifar. A shared multi-attention framework for multi-label zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Eui Jin Hwang, Ju Gang Nam, Woo Hyeon Lim, Sae Jin Park, Yun Soo Jeong, Ji Hee Kang, Eun Kyoung Hong, Taek Min Kim, Jin Mo Goo, Sunggyun Park, Ki Hwan Kim, and Chang Min Park. Deep learning for chest radiograph diagnosis in the emergency department. *Radiology*, 293(3):573–580, 2019. doi: 10.1148/radiol.2019191225. URL <https://doi.org/10.1148/radiol.2019191225>. PMID: 31638490.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
- Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, and et al. The open images dataset v4. *International Journal of Computer Vision*, 128(7):1956–1981, Mar 2020. ISSN 1573-1405. doi: 10.1007/s11263-020-01316-z. URL <http://dx.doi.org/10.1007/s11263-020-01316-z>.
- Hugo Larochelle, Dumitru Erhan, and Yoshua Bengio. Zero-data learning of new tasks. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 2, AAAI’08*, page 646–651. AAAI Press, 2008. ISBN 9781577353683.
- Chung-Wei Lee, Wei Fang, Chih-Kuan Yeh, and Yu-Chiang Frank Wang. Multi-label zero-shot learning with structured knowledge graphs. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1576–1585, June 2018. doi: 10.1109/CVPR.2018.00170.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, Sep 2019. ISSN 1460-2059. doi: 10.1093/bioinformatics/btz682. URL <http://dx.doi.org/10.1093/bioinformatics/btz682>.
- Yang Long, Li Liu, Ling Shao, Fumin Shen, Guiguang Ding, and Jungong Han. From zero-shot learning to conventional supervised classification: Unseen visual data synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL <https://proceedings.neurips.cc/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf>.
- George A. Miller. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41, Nov 1995. ISSN 0001-0782. doi: 10.1145/219717.219748. URL <https://doi.org/10.1145/219717.219748>.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Angshuman Paul, Thomas C Shen, Sungwon Lee, Niranjana Balachandar, Yifan Peng, Zhiyong Lu, and Ronald M Summers. Generalized zero-shot chest x-ray diagnosis through trait-guided multi-view semantic embedding with self-training. *IEEE transactions on medical imaging*, PP, February 2021. ISSN 0278-0062. doi: 10.1109/tmi.2021.3054817. URL <https://doi.org/10.1109/TMI.2021.3054817>.
- Chunli Qin, D. Yao, Yonghong Shi, and Z. Song. Computer-aided detection in chest radiography based on artificial intelligence: a survey. *BioMedical Engineering OnLine*, 17, 2018.
- Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging, 2019.
- S. Rahman, S. Khan, and N. Barnes. Deep0tag: Deep multiple instance learning for zero-shot image tagging. *IEEE Transactions on Multimedia*, 22(1):242–255, 2020. doi: 10.1109/TMM.2019.2924511.
- Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Yi Ding, Aarti Bagul, Curtis Langlotz, Katie S. Shpanskaya, Matthew P. Lungren, and Andrew Y. Ng. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *CoRR*, abs/1711.05225, 2017. URL <http://arxiv.org/abs/1711.05225>.
- Pranav Rajpurkar, Jeremy Irvin, Robyn Ball, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Bhavik Patel, Kristen Yeom, Katie Shpanskaya, Francis Blankenberg, Jayne Seekins, Timothy Amrhein, David Mong, Safwan Halabi, Evan Zucker, and Matthew Lungren. Deep learning for chest radiograph diagnosis: A retrospective comparison of the chexnext algorithm to practicing radiologists. *PLOS Medicine*, 15:e1002686, 11 2018. doi: 10.1371/journal.pmed.1002686.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, page 91–99, Cambridge, MA, USA, 2015. MIT Press.
- Mahdi Rezaei and Mahsa Shahidi. Zero-shot learning and its applications from autonomous vehicles to covid-19 diagnosis: A review. *Intelligence-Based Medicine*, 3-4:100005, 2020. ISSN 2666-5212. doi: <https://doi.org/10.1016/j.ibmed.2020.100005>. URL <https://www.sciencedirect.com/science/article/pii/S2666521220300053>.

- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- W. J. Scheirer, A. de Rezende Rocha, A. Sapkota, and T. E. Boult. Toward open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(7):1757–1772, 2013. doi: 10.1109/TPAMI.2012.256.
- Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald Summers. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3462–3471, 2017.
- Yongqin Xian, Christoph H. Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning - A comprehensive evaluation of the good, the bad and the ugly. *CoRR*, abs/1707.00600, 2017. URL <http://arxiv.org/abs/1707.00600>.
- Li Yao, Eric Poblenz, Dmitry Dagunts, Ben Covington, Devon Bernard, and Kevin Lyman. Learning to diagnose from scratch by exploiting dependencies among labels, 2018.
- Li Zhang, Tao Xiang, and Shaogang Gong. Learning a deep embedding model for zero-shot learning. *CoRR*, abs/1611.05088, 2016. URL <http://arxiv.org/abs/1611.05088>.