

## RESEARCH ARTICLE

# Deep learning models for COVID-19 chest x-ray classification: Preventing shortcut learning using feature disentanglement

Anusua Trivedi<sup>1,6</sup> , Caleb Robinson<sup>1</sup> , Marian Blazes<sup>3</sup>, Anthony Ortiz<sup>1</sup>, Jocelyn Desbiens<sup>2</sup>, Sunil Gupta<sup>2</sup>, Rahul Dodhia<sup>1</sup> , Pavan K. Bhatraju<sup>5,7</sup>, W. Conrad Liles<sup>5</sup>, Jayashree Kalpathy-Cramer<sup>4</sup>, Aaron Y. Lee<sup>3\*</sup> , Juan M. Lavista Ferres<sup>1\*</sup>

**1** AI for Good Research Lab, Microsoft, Redmond, WA, United States of America, **2** Intelligent Retinal Imaging Systems, Pensacola, FL, United States of America, **3** Department of Ophthalmology, University of Washington, Seattle, WA, United States of America, **4** Neuroscience, Massachusetts General Hospital, Boston, MA, United States of America, **5** Department of Medicine and Sepsis Center of Research Excellence, University of Washington (SCORE-UW), Seattle, WA, United States of America, **6** Flipkart US R&D, Seattle, WA, United States of America, **7** Division of Pulmonary, Critical Care and Sleep Medicine, University of Washington, Seattle, WA, United States of America

 These authors contributed equally to this work.

\* [leeay@uw.edu](mailto:leeay@uw.edu) (AYL); [jlavista@microsoft.com](mailto:jlavista@microsoft.com) (JMLF)



## OPEN ACCESS

**Citation:** Trivedi A, Robinson C, Blazes M, Ortiz A, Desbiens J, Gupta S, et al. (2022) Deep learning models for COVID-19 chest x-ray classification: Preventing shortcut learning using feature disentanglement. PLoS ONE 17(10): e0274098. <https://doi.org/10.1371/journal.pone.0274098>

**Editor:** Nguyen Quoc Khanh Le, Taipei Medical University, TAIWAN

**Received:** August 10, 2021

**Accepted:** August 22, 2022

**Published:** October 6, 2022

**Copyright:** © 2022 Trivedi et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** All research data files are available from the COVID-Net database (<https://github.com/lindawangg/COVID-Net>). All software and algorithms developed have been open sourced and are available at <https://github.com/microsoft/xray-feature-disentanglement>. The COVIDx dataset is publicly available. The CT images in the China Consortium CC-CCII dataset are publicly available. Use of the CC-CCII x-ray image data for validation required approval by the Sun Yat-sen Memorial University IRB.

## Abstract

In response to the COVID-19 global pandemic, recent research has proposed creating deep learning based models that use chest radiographs (CXRs) in a variety of clinical tasks to help manage the crisis. However, the size of existing datasets of CXRs from COVID-19+ patients are relatively small, and researchers often pool CXR data from multiple sources, for example, using different x-ray machines in various patient populations under different clinical scenarios. Deep learning models trained on such datasets have been shown to overfit to erroneous features instead of learning pulmonary characteristics in a phenomenon known as shortcut learning. We propose adding feature disentanglement to the training process. This technique forces the models to identify pulmonary features from the images and penalizes them for learning features that can discriminate between the original datasets that the images come from. We find that models trained in this way indeed have better generalization performance on unseen data; in the best case we found that it improved AUC by 0.13 on held out data. We further find that this outperforms masking out non-lung parts of the CXRs and performing histogram equalization, both of which are recently proposed methods for removing biases in CXR datasets.

## Introduction

The Coronavirus Disease (COVID)-19 pandemic has exposed many vulnerabilities in public health systems around the world. Artificial intelligence (AI) is poised to address some of these challenges with potential applications including early disease detection, clinical management tools, disease modeling, and vaccine research [1]. Screening patients for COVID-19 based on

**Funding:** This research has been funded by National Institutes of Health Grants NIH K23DK116967 (awarded to PKB), NIH/NEI K23EY029246 (awarded to AYL), Latham Vision Innovation Award (awarded to AYL), the C. Dan and Irene Hunter Endowed Professorship (awarded to AYL), and an unrestricted grant from Research to Prevent Blindness (awarded to AYL). The Microsoft Corporation (microsoft.com) provided support in the form of salaries to CR, AT, AO, RD, and JMLF. IRIS (<https://retinalscreenings.com/>) provided support in the form of salaries to JD and SG. The funders did not have any additional role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript. The specific roles of these authors are articulated in the 'author contributions' section.

**Competing interests:** The authors have read the journal's policy and have the following competing interests to declare: AYL reports support from the US Food and Drug Administration, grants from Santen, Regeneron, Carl Zeiss Meditec, and Novartis, and personal fees from Genentech, Roche, Verana Health, and Johnson and Johnson, outside of the submitted work. This article does not reflect the opinions of the Food and Drug Administration. CR, AT, AO, RD, and JMLF are paid employees of the Microsoft Corporation. JMLF also receives personal fees from Singularity University as a speaker. JD and SG are paid employees of IRIS. In addition, SG is a consultant/advisor for Alcon Laboratories, Allergan, Inc., Andrews Institute, GENENTECH, Novartis, Alcon Pharmaceuticals, Regeneron Pharmaceuticals, Inc., Roche Diagnostics, and Spark Therapeutics, Inc. as well as an equity owner in Intelligent Retinal Imaging Systems (IRIS), Retina Specialty Institute, and USRetina. This does not alter the authors' adherence to PLOS ONE policies on sharing data and materials. There are no patents, products in development or marketed products associated with this research to declare. MB, PKB, WCL, and JKC have no relevant financial or non-financial interests to disclose.

chest radiograph (CXR) imaging is one potential AI application, and deep learning models have been developed to distinguish COVID-19 pneumonia from normal findings and even rule out other types of pneumonia with high accuracy [2–4]. Due to the increasing practice of reporting results prior to peer review in the setting of a novel pandemic, many additional models have been reported as preprint papers [5–11].

The main challenge with training these deep learning models has been the shortage of COVID-19 CXR imaging data. Several public datasets are available, but most are small (100–200 patients). A recent review assessed many AI diagnostic and prognostic models for COVID-19 using the prediction model risk of bias assessment tool (PROBAST) and found that all but one of the 51 studies were at high risk for analysis bias, mainly overfitting due to small sample size [12]. Larger databases have typically pooled images from different patient populations at multiple research sites and hospitals, or even from different countries [7, 13, 14]. These unbalanced datasets often contain images with source-specific identifiers (antero-posterior versus posteroanterior positioning, imaging device type, image size, etc) that the models can misidentify as relevant features.

The most common approach for developing COVID-19 CXR diagnostic models is to use transfer and representational learning techniques to adapt classification models that were trained on labeled CXR datasets (to classify pre-COVID-19 pulmonary findings), or on unrelated image datasets, to the new task [2, 4, 15]. However, due to the problems mentioned above—small datasets or larger datasets with a latent imbalance—leads to learned models that suffer from a phenomenon summarized as “shortcut learning” [16]. Here, the models will identify dataset-specific features rather than pulmonary specific features in order to distinguish between classes, and the reported performance of these models may not be generalizable to other patient populations. Indeed, recent research proves exactly this [17] and further shows that some shortcuts can persist in external datasets. The problem now becomes, “how to train highly-parameterized deep learning models on publicly available data, without learning shortcuts that are present in such data?” A recent study proposed two preprocessing techniques to solve this problem: histogram equalization, to correct for issues such as contrast differences between images, and lung masking to remove potential source-specific information located outside of the lung area [18].

In this paper, we sought to evaluate the use of feature disentanglement [19], a multi-task training approach for deep neural networks, to prevent shortcut learning in the context of classification of COVID-19 CXR images. Models trained with this method are forced to learn features that can identify COVID-19+ CXRs and are penalized for learning features that can identify what sub-population a CXR is from. We further sought to compare feature disentanglement to the previously proposed histogram equalization and lung masking methods. Finally, we sought to test all three methods on the ultimate goal of improving model generalization performance on unseen data. We have released source code to reproduce the models and results found in this paper: <https://github.com/microsoft/xray-feature-disentanglement>.

## Methods

### Datasets

We used two CXR datasets: the open-source COVIDx dataset [20] and a dataset received from the China Consortium of Chest CT CC-CCII [21]. These are both image classification datasets that share the same three *class* labels for CXR images: “normal” (Normal)—collected from patients without pneumonia, “common pneumonia” (Pneumonia) CXRs—collected from patients with pneumonias other than from COVID-19, and “novel COVID-19 pneumonia” (COVID-19+). We treated the source dataset of each image as its *domain* label. Specifically,

**Table 1. Dataset overview.** Counts of disease label type per dataset. The COVIDx dataset is made up of 5 sub-datasets and the CC-CCII dataset is used as a held-out test set.

	Normal	Pneumonia	COVID-19+
<b>COVIDx (totals)</b>	8,851	6,040	421
COVID-CHESTXRAY	–	28	289
Fig 1	–	–	35
ACTUALMED	–	–	51
SIRM	–	–	46
RSNA	8,851	6,012	–
<b>CC-CCII</b>	11,604	18,236	1,690

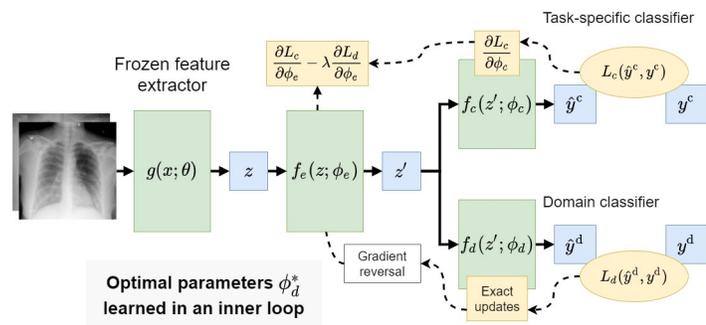
<https://doi.org/10.1371/journal.pone.0274098.t001>

there are five domain labels: one for each of the sub-datasets in the COVIDx dataset. See Table 1 for a breakdown of the number of class labels of each type per dataset. Below we describe each dataset in more detail.

The COVIDx dataset is a conglomeration of CXR samples from five other datasets: the COVID-CHESTXRAY dataset (also referred to as the COHEN dataset in related work) [22], the Fig 1 COVID-19 CXR dataset [23], the ACTUALMED COVID-19 CXR dataset [24], the Kaggle RSNA Pneumonia Detection Challenge dataset [25], and the SIRM samples from the Kaggle COVID-19 radiography database [26]. The samples are divided among three classes: “normal” control images, non COVID-19 pneumonia images, and COVID-19+ images. We used the scripts provided on the COVIDx GitHub repository to create the dataset. We then merge the training and test sets and filter out all but the first sample for each patient. The label counts in Table 1 for this dataset therefore equal the number of patients.

The CC-CCII dataset is a large CT dataset encompassing CT images from retrospective cohorts from the China Consortium of Chest CT Image Investigation (CFC-CCII) [21]. This datasets has samples from Sun Yat-sen Memorial Hospital and Third Affiliated Hospital of Sun Yat-sen University, The First Affiliated Hospital of Anhui Medical University, West China Hospital, Guangzhou Medical University First Affiliated Hospital, Nanjing Renmin Hospital, Yichang Central People’s Hospital, and Renmin Hospital of Wuhan University. The dataset we use in this paper consists of CXRs taken from a subset of the patients included in the CT dataset. These images have been labeled with the same classes as in COVIDx.

**Ethics statement.** The COVIDx dataset is publicly available. The CC-CCII xray images were collected by Dr. Jun Shen from Sun Yat-sen Memorial Hospital. All patient information



**Fig 1. Overview of the feature disentanglement modeling approach.** We propose to learn a model that simultaneously predicts the class label and domain label for a given CXR image. The parameters of the model are updated to extract representations that contain information about the class label but not about domain label.

<https://doi.org/10.1371/journal.pone.0274098.g001>

for the CC-CCII dataset are de-identified and anonymized before the authors received the data, and consent was obtained from all participants. The usage of the x-ray images in the CC-CCII dataset for our validation studies have been approved by the Sun Yat-sen Memorial University IRB. The CT images in the CC-CCII are publicly available. IRB approval was obtained before the start of the study.

### Model training

Formally, we trained a model,  $f(g(x; \theta); \phi)$ , where  $x$  is a CXR. This model is decomposed into a feature-extractor,  $g(x; \theta) = z$ , parameterized by  $\theta$ , and classifier,  $f(z; \phi)$ , parameterized by  $\phi$ . The feature-extractor is responsible for creating an embedding,  $z$ , for a given CXR, and the classifier is responsible for predicting the label from a given embedding. This representation is helpful in transfer learning settings where frozen feature-extractor models (i.e. models with fixed  $\theta$ ) that have been pre-trained on large CXR datasets can be used to reduce overfitting when relatively few labeled samples are available.

We fit  $f(z; \phi)$  over a dataset,  $\mathcal{D} = (x_i, y_i^{\text{class}}, y_i^{\text{domain}})_{i=1}^N$ , where each image contains a *class* label and a *domain* label. The *class* labels are clinically relevant; for example, they describe whether a CXR is from a normal (Normal) patient, a patient with common pneumonia (Pneumonia), or from a patient that may exhibit novel COVID-19 pneumonia (COVID-19+). The *domain* labels are *not* clinically relevant. They encode information about how the sample was collected, such as the clinical site and patient cohorts. In our experiments the *domain* labels encode which dataset a sample originally came from. An ideal classifier would not be sensitive to how a sample was collected but would rely on pulmonary features present in the different types of *class* labels. However, for the reasons we outline in the Introduction, classifiers can overfit to spurious signals in CXRs leading to poor generalization when applied to new imagery.

We fit  $f(z; \phi)$  in two optimization settings: a straightforward baseline setting and a feature disentanglement setting in which we force the classifier to learn data representations that are not useful in predicting the *domain* labels, while remaining useful for predicting the *class* labels:

**Baseline.** In this setting, we learned  $\phi^*$  by minimizing a negative log-likelihood loss,  $L(\hat{y}^{\text{class}}, y^{\text{class}})$ , between the class predictions,  $\hat{y}^{\text{class}}$ , and the *class* labels,  $y^{\text{class}}$ , over  $\mathcal{D}$  in a standard setup:

$$\phi^* = \arg \min_{\phi} \sum_{x, y^{\text{class}} \in \mathcal{D}} L(f(g(x); \phi), y^{\text{class}}) \tag{1}$$

Our hypothesis was that, in this setting, even simple classifiers will overfit to spurious features and exhibit poor generalization performance.

**Feature disentanglement.** In this setting we assumed that we can further decompose  $f$  into a feature extractor,  $f_e(z; \phi_e) = z'$ , and two classification heads,  $f_c(z'; \phi_c) = \hat{y}^{\text{class}}$  and  $f_d(z'; \phi_d) = \hat{y}^{\text{domain}}$ —see Fig 1 for an overview of this setup. We then follow the methodology proposed in [19] for training  $f$  in a multi-task setting to transform the initial embedding  $z$  into a compressed form  $z'$  that is useful for predicting the class label and not useful for predicting the domain label for a given CXR. Formally, we defined the empirical error with a *class* loss as above, and an additional *domain* loss:

$$E(\phi_e, \phi_c, \phi_d) = \sum_{x, y^{\text{class}} \in \mathcal{D}} L_c(f_c(f_e(z; \phi_e); \phi_c), y^{\text{class}}) - \lambda \sum_{x, y^{\text{domain}} \in \mathcal{D}} L_d(f_d(f_e(z; \phi_e); \phi_d), y^{\text{domain}}) \tag{2}$$

Note that the error is defined as the difference between the summed class loss and summed domain loss. Larger loss values in the second term will contribute to the overall minimization of the total error which influence the model to learn a compressed representation that is *not* predictive of the domain labels.  $\lambda$  is a hyperparameter that controls the relative influence of the domain loss on the error. We optimized for parameters,  $\phi$ , according to:

$$(\phi_e^*, \phi_c^*) = \arg \min_{\phi_e, \phi_c} E(\phi_e, \phi_c, \phi_d) \quad (3)$$

$$\phi_d^* = \arg \max_{\phi_d} E(\phi_e, \phi_c, \phi_d) \quad (4)$$

As described in [19], this can be solved by iterating using gradient-based optimization with an additional “gradient reversal layer” inserted between  $f_e$  and  $f_d$  in the structure of the model (see Fig 1). We considered a modification wherein after each training epoch we optimize for  $\phi_d^*$ , over the whole training set, for the current, fixed values of  $\phi_e$ . Put simply, we trained  $f_d(z; \phi_d)$  to convergence at the end of each epoch. As a consequence, the updates to  $\phi_e$  in the next epoch of training are done with respect to the domain loss computed at a local minimum for the domain classifier. We found empirically that this helped reduce variance in results between training runs. This modification is cheap when the parameters of the original feature extractor remain frozen and the intermediate representations,  $z$ , are cached—which was always the case in our experiments.

## Lung masking and histogram equalization preprocessing

Motivated by the same points that we are [18], proposed two preprocessing methods for eliminating signal in CXR imagery that a classifier might overfit to: lung masking and histogram equalization. Lung masking uses a model to remove the non-lung parts of CXR images with the assumption that features from the imagery surrounding the lungs should not be used in predictive models of lung diseases. Histogram equalization is a common technique in computer vision used to improve the contrast of grayscale images. If histogram equalization is applied to each image in a dataset of CXR image, then a classification model will not be able to exploit relative differences in the contrast of images in its decisions.

These methods can be used in combination with our proposed method for feature disentanglement. In our experiments we used the recent SOTA lung VAE model [27] to create lung masks and implementation of histogram equalization from OpenCV [28].

## Experimental setup

Our method depends on a feature extractor model,  $g(x; \theta)$ , to create an initial set of embeddings. We considered three different pre-trained CNNs for this role:

**Torchxrayvision pretrained DenseNet121** We used the DenseNet121 model [29] weights released in the torchxrayvision package [30, 31]. This model has been trained on CXR imagery, collected before the COVID-19 pandemic, with multi-task pulmonary disease labels from different datasets. We extracted a 1,024 dimensional feature vector for a given CXR by applying global average pooling after the last convolutional layer in the DenseNet.

**ImageNet pretrained DenseNet121** We also used the DenseNet121 model [29] weights from the PyTorch torchvision library [32]. This model has been trained on ImageNet, and the *only* difference in implementation from the torchxrayvision version is that the first convolutional layer operates over 3 channel images instead of 1 channel images. We duplicated

input values across the 3 input channels in order to run this model on CXR imagery. We extract features from this DenseNet in the same way as above.

**COVID-Net** We used the COVID-Net model and the “COVIDNet-CXR Large” pretrained weights from the same repository that proposed the COVIDx dataset [20]. We extract a 2,048 dimensional feature vector for an input CXR by applying global average pooling after the last convolutional layer in the COVID-Net architecture. This is in contrast with the methodology from [20], who flatten the representation after the last convolutional layer to get a 460,800 dimensional representation.

In all experiments we reported average and standard deviation results from 5×10-fold cross validation. In the experiments testing whether we can identify domain labels from pre-trained model representations, we stratified by the domain label in order to ensure that there are samples from each dataset in training and testing splits. In the experiments where we measure domain generalization, we report results on the out-of-sample CC-CCII data, we used the out-of-sample data as an additional test set in the cross-validation folds used to test in-sample performance. Specifically, we split the training set into 10-folds, trained a model on data from 9 folds, then tested on the held-out fold, as well as the entire out-of-sample dataset.

Our classification problems are all multi-class; we reported an unweighted average of the area under the ROC curve (AUC) calculated individually for each class in a one-versus-rest manner and average per class accuracy (ACC). Both methods are not sensitive to the class imbalance which we observe in both our *class* and *domain* tasks.

In the experiments where we measure domain generalization of the **Baseline** models, i.e. the models trained without feature disentanglement, we fit an additional logistic regression model on the domain labels, from the representations generated by  $g(x; \theta)$ . This additional step happens after fully training the  $f(g(x); \phi)$  and will allow us to measure how much information about the domain labels the encoder model uses.

When using feature disentanglement, we set  $f_e(z)$  as a fully connected model with the following structure: Dense(256) → BatchNorm → Dense(64) → BatchNorm, where both dense layers are followed by a ReLU nonlinearity. We set both  $f_c(z')$  and  $f_d(z')$  as logistic regression layers. We trained with the AdamW optimizer using AMSGrad and an initial learning rate of 0.001. We divided the learning rate by a factor of 10 if the validation *class* loss has stagnated for over 10 epochs and stopped training either the third time this happens, or after 200 total epochs.

We also decayed  $\lambda$  throughout training; we set  $\lambda_0 = 10$  and use the following update rule evaluated each epoch (where  $t$  is the epoch) to calculate  $\lambda$ :

$$\lambda \leftarrow \frac{\lambda_0}{1 + (\lambda_0 - 1) \left(\frac{t}{200}\right)^3}$$

This moves  $\lambda$  from  $\lambda_0$  to 1 over the course of the maximum 200 epochs of training. We did not thoroughly test this schedule against other choices. We simply aimed for the domain loss to guide training in the initial epochs of training, and for the class loss to converge to a local minimum in the later epochs of training.

Finally, all of our experiments are fully reproducible by setting the seed of the random number generators in all component software packages. We varied this seed over the 5x restarts in our experiments.

**Table 2. Results showing how well logistic regression classifiers can identify which sub-dataset a CXR is from within the COVIDx dataset, and how well classifiers can identify which dataset a “COVID-19+” CXR is from across both the COVIDx and CC-CCII datasets.** We report AUC values as averages of the one-vs-all binary AUCs between all classes, and accuracy (ACC) as the average accuracy over all classes. We observe that the representations generated by the classifiers, even from masked/equalized inputs, contain enough information to accurately identify the sources of the imagery in both cases.

Unmasked Images	Embedding size	COVIDx datasets		All COVID-19+ samples	
		AUC	ACC	AUC	ACC
Pixel intensity histogram	256	0.75 ± 0.03	0.36 ± 0.04	0.83 ± 0.03	0.47 ± 0.07
Torchxrayvision embedding	1024	0.93 ± 0.03	0.54 ± 0.09	0.93 ± 0.03	0.59 ± 0.07
ImageNet embedding	1024	0.95 ± 0.02	0.53 ± 0.07	0.97 ± 0.01	0.66 ± 0.09
COVID-Net embedding	2048	0.89 ± 0.05	0.39 ± 0.05	0.91 ± 0.03	0.45 ± 0.04
Masked/equalized Images	Embedding size	AUC	ACC	AUC	ACC
Pixel intensity histogram	256	0.57 ± 0.06	0.27 ± 0.06	0.75 ± 0.02	0.25 ± 0.04
Torchxrayvision embedding	1024	0.76 ± 0.05	0.35 ± 0.07	0.88 ± 0.02	0.42 ± 0.08
ImageNet embedding	1024	0.84 ± 0.02	0.34 ± 0.04	0.91 ± 0.02	0.50 ± 0.10
COVID-Net embedding	2048	0.71 ± 0.04	0.26 ± 0.04	0.85 ± 0.04	0.35 ± 0.05

<https://doi.org/10.1371/journal.pone.0274098.t002>

## Experiments and results

### Identifying domain labels from pre-trained model representations

In Table 2 we show that all of our existing models extract representations that contain enough information, even when run on masked/equalized imagery, for discriminating between domain labels throughout the COVIDx dataset (i.e. determining which sub-dataset from COVIDx a sample belongs to), and domain labels throughout *only* COVID-19+ samples (i.e. determining where a COVID-19+ sample comes from out of the sub-datasets in COVIDx or the CC-CCII dataset). For example, we find that a DenseNet121 pre-trained on ImageNet extracts features from unmasked CXR images that are sufficient to train a logistic regression model that can identify the source dataset among samples from the COVIDx dataset with a held-out average AUC of  $0.95 \pm 0.02$ . The same embeddings can be used to train a logistic regression model that can identify the source dataset given COVID-19+ sample, among the constituent datasets of COVIDx and CC-CCII with a held out AUC of  $0.97 \pm 0.01$ . While the preprocessing steps help to reduce this performance, the linear models still perform much better than random guessing, suggesting that overfitting to domain signals in the embedded representations is trivial.

### Class and domain performance

In Table 3 we show the within-dataset and generalization performance of models,  $f(z; \phi)$ , trained on top of the different feature extractors. Similar to Table 2, the models that are trained in a standard way have high performance on the domain task, and correspondingly high performance on the actual task labels. At the same time, the poor performance of these models on the out-of-sample dataset, CC-CCII, shows that they are overfitting to the training dataset. For example, a model trained on top of torchxrayvision unmasked image embeddings performs similar to random guessing—random guessing performance is 0.33 average accuracy, while this model gets 0.34 average accuracy. Lung masking and histogram equalization improves generalization performance in all cases, however we hypothesize that this may throw out signals that are relevant to the task.

On the other hand, we observe that training  $f(z; \phi)$  with feature disentanglement results in better generalization performance with both unmasked or masked/equalized images in all cases. ImageNet embeddings of masked/equalized images and feature disentanglement training results in the best generalization performance on CC-CCII. Interestingly, training with

**Table 3. Results showing within-dataset class performance, within-dataset domain performance, and out-of-sample class performance from training models with the COVIDx dataset.** The task performance (task AUC and task accuracy) shows how well classifiers are able to distinguish between “Normal”, “Pneumonia”, and “COVID-19+” disease labels, while the domain performance (domain AUC) shows how well classifiers are able to distinguish which sub-dataset an image belongs to. We report AUC values as averages of the one-vs-all binary AUCs between all classes, and accuracy (ACC) as the average accuracy over all classes. In all cases class performance (both within-dataset and out-of-sample) is reported from the classifier trained on samples within-dataset, while domain performance is reported from an additional classifier trained to predict domain labels on top of the learned representations,  $z'$ , as a measure of how much domain information the representation contains. We observe that using feature disentanglement decreases within-dataset domain performance as expected, and increases out-of-sample class performance—i.e. improves generalization performance.

	COVIDx		CC-CCII	
	Task AUC	Domain AUC	Task AUC	Task ACC
<b>Torchxrayvision Embeddings</b>				
Unmasked	0.97 ± 0.01	0.94 ± 0.02	0.55 ± 0.03	0.34 ± 0.04
Masked/equalized	0.92 ± 0.01	0.85 ± 0.03	0.65 ± 0.02	0.42 ± 0.02
Unmasked + Disentanglement	0.90 ± 0.03	0.56 ± 0.07	0.68 ± 0.04	0.49 ± 0.02
Masked/equalized + Disentanglement	0.87 ± 0.02	0.53 ± 0.06	0.71 ± 0.03	0.47 ± 0.02
<b>ImageNet Embeddings</b>				
Unmasked	0.96 ± 0.01	0.97 ± 0.02	0.64 ± 0.02	0.37 ± 0.01
Masked/equalized	0.94 ± 0.01	0.88 ± 0.03	0.67 ± 0.03	0.41 ± 0.03
Unmasked + Disentanglement	0.85 ± 0.03	0.57 ± 0.07	0.73 ± 0.03	0.43 ± 0.02
Masked/equalized + Disentanglement	0.85 ± 0.03	0.57 ± 0.04	0.73 ± 0.03	0.46 ± 0.02
<b>COVID-Net Embeddings</b>				
Unmasked	0.96 ± 0.01	0.93 ± 0.02	0.59 ± 0.02	0.38 ± 0.01
Masked/equalized	0.89 ± 0.02	0.80 ± 0.04	0.63 ± 0.02	0.44 ± 0.02
Unmasked + Disentanglement	0.92 ± 0.02	0.53 ± 0.10	0.67 ± 0.03	0.41 ± 0.02
Masked/equalized + Disentanglement	0.83 ± 0.02	0.52 ± 0.07	0.62 ± 0.01	0.45 ± 0.00

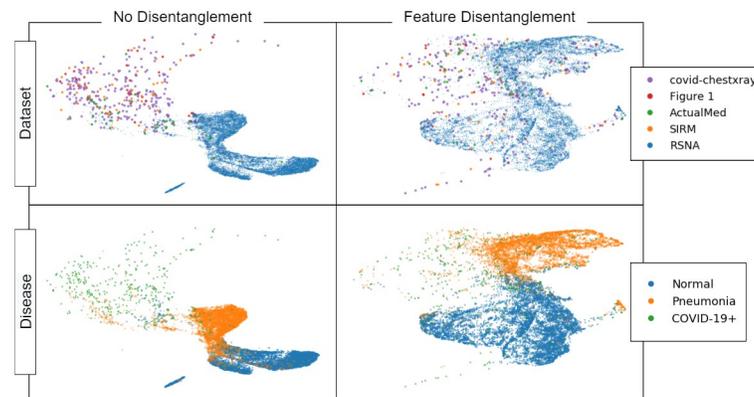
<https://doi.org/10.1371/journal.pone.0274098.t003>

feature disentanglement on masked/equalized images is not always better than training with feature disentanglement on unmasked images. This supports the idea that lung masking and histogram equalization might throw out relevant task signals. In all cases, training with feature disentanglement dramatically reduces the domain signal in learned representations. For example, the domain performance of the model tuned with torchxrayvision embeddings of unmasked images is 0.94 AUC, while using feature disentanglement reduces this to 0.56 AUC. The increase in performance from training with feature disentanglement is variable across the pre-trained model used to generate the initial representation. We observe the smallest difference in generalization performance with the COVID-Net pre-trained model which was initially trained on unmasked images from the COVIDx dataset—a smaller dataset than either ImageNet or the CXR dataset used by torchxrayvision—and may not function as an effective feature extractor.

Further experiments are needed to determine if unfreezing the parameters of the feature extractor model during training with feature disentanglement is beneficial. We specifically avoid this as we have already found that it is possible to overfit to the training set with a small fully connected model on top of pre-trained embeddings, as well as learn representations that are not predictive of the within-dataset domain labels. Performance gains of fine-tuning through the feature extractor would only be potentially visible through increases in generalization performance, i.e. on the CC-CCII dataset. By correlating such design decisions with CC-CCII performance. We would risk manually overfitting to CC-CCII, and want to avoid doing so in this work.

### Feature visualization

In Fig 2 we show the UMAP embeddings of the learned feature representations from the models trained on the torchxrayvision embeddings for images from the COVIDx dataset. This



**Fig 2. UMAP projections of features learned from models trained with and without feature disentanglement on unmasked imagery.** Each point represents a CXR from the COVIDx dataset. The top row colors points by their domain label—which subdataset of the COVIDx dataset they are in—while the bottom row colors points by their disease label. We observe that without feature disentanglement, the learned representations easily separate datasets—despite not being trained for this task—however, with feature disentanglement, the learned representations do not clearly separate datasets.

<https://doi.org/10.1371/journal.pone.0274098.g002>

shows how the representations learned without feature disentanglement clearly separate the images from the RSNA dataset from the images from the other datasets, despite not being trained to do this. This is because the RSNA dataset contains most of the samples with “normal” and “pneumonia” disease labels, versus the other datasets, which contain “COVID-19+” samples. Therefore, when a model is trained to separate these classes, it finds features that are correlated with which component dataset an image is from. Such biases can be trivial, such as a difference in contrast between the dark and light portions of the image, the position/size of the lungs, markings on the CXR, etc. The same distinction between datasets is not clear when feature disentanglement is used because the learned representations are forced to not be predictive of dataset. As a result, the distinction between “COVID-19+” and “pneumonia” labels becomes less clear when feature disentanglement is used, however, as we show in Table 3, this can improve the generalization ability of the model.

## Discussion

Accurate automated diagnosis of COVID-19 pneumonia based on CXRs has been a research focus since the start of the pandemic, given its potential use in emergency departments, urgent care, and resource-limited settings. The development of such devices has been limited by the availability and quality of COVID-19+ datasets, which are typically aggregated from various sources to increase the number of COVID-19+ examples. In this study we show that previously proposed techniques, such as isolating the pulmonary region and harmonizing images taken from different x-ray machines, fail to prevent deep learning models from relying on features that are specific to a particular dataset or image source (e.g. scanner machine type, age of patient, external artifacts in the image). These models are likely to underperform in real-world clinical settings when analyzing images that lack the identifying data the models have learned to rely on.

Since the COVID-19 outbreak, various researchers have developed automated COVID-19 CXR diagnosis models. Most previous studies have used transfer learning approaches and compared classification performances obtained between several popular CNN architectures, but all have relied on datasets made up of COVID-19+ CXRs sourced from around the web [12]. For COVID-19 negative cases, data are typically sampled from other open CXR datasets.

However, if any bias is present within these datasets, the model could learn the underlying biases, rather than learning COVID-19 related features. Therefore, we have presented an application of feature disentanglement for predicting COVID-19 infection even when training on public datasets containing CXRs from multiple sources. The model showed superior performance when tested on a new and unfamiliar dataset, suggesting that it was relying on COVID-19 specific pulmonary findings. We applied visualization techniques to show that our model relied on imaging features that were not specific to particular groupings of CXR images in the training data, further suggesting that the model is capable of performing well when analyze CXR data from unfamiliar sources.

Although we show that our approach produces a model that is more successful when challenged with real-world data, our work is still limited by the lack of public datasets available for testing, and we used a private dataset from the CC-CCII to test generalization performance. Future work should investigate the effects of fine-tuning these models with and without feature disentanglement approaches, and the effects of controlling for multiple imaging data sources in an expanded multi-task setting. Finally, when larger CXR datasets that are paired with clinical outcomes become available, further exploration of the use of feature disentanglement for extracting small sets of clinically relevant features from CXRs should be investigated. However, the diagnostic performance of any diagnostic model should be carefully evaluated in a real-world setting before clinical implementation. Misdiagnosis of COVID-19 can lead to inappropriate care, failure to treat, increased transmission, and many other serious outcomes. Clinicians should be aware of potential limitations and biases when incorporating model predictions into their clinical assessment.

Finally, our approach has potential clinical applications beyond automated diagnosis. CXR diagnostic models that rely on relevant pulmonary findings may be also useful for the development of prognostic models, by combining the CXR information with other clinical and demographic data to predict which patients are at risk for severe disease.

## Author Contributions

**Conceptualization:** Anusua Trivedi, Caleb Robinson, Jocelyn Desbiens, Sunil Gupta, Rahul Dodhia, Pavan K. Bhatraju, W. Conrad Liles, Aaron Y. Lee, Juan M. Lavista Ferres.

**Data curation:** Anusua Trivedi, Caleb Robinson, Jocelyn Desbiens, Juan M. Lavista Ferres.

**Formal analysis:** Caleb Robinson, Marian Blazes, Anthony Ortiz, Jayashree Kalpathy-Cramer, Aaron Y. Lee, Juan M. Lavista Ferres.

**Investigation:** Anusua Trivedi, Caleb Robinson, Anthony Ortiz, Aaron Y. Lee.

**Methodology:** Anusua Trivedi, Caleb Robinson, Anthony Ortiz.

**Project administration:** Anusua Trivedi, Caleb Robinson, Juan M. Lavista Ferres.

**Software:** Anusua Trivedi, Caleb Robinson, Anthony Ortiz, Aaron Y. Lee.

**Validation:** Anusua Trivedi, Caleb Robinson, Jocelyn Desbiens.

**Writing – original draft:** Anusua Trivedi, Caleb Robinson, Marian Blazes, Anthony Ortiz, Aaron Y. Lee.

**Writing – review & editing:** Anusua Trivedi, Caleb Robinson, Marian Blazes, Anthony Ortiz, Jocelyn Desbiens, Sunil Gupta, Rahul Dodhia, Pavan K. Bhatraju, W. Conrad Liles, Jayashree Kalpathy-Cramer, Aaron Y. Lee, Juan M. Lavista Ferres.

## References

1. Vaishya R, Javaid M, Khan IH, Haleem A. Artificial Intelligence (AI) applications for COVID-19 pandemic. *Diabetes & Metabolic Syndrome: Clinical Research & Reviews*. 2020;. <https://doi.org/10.1016/j.dsx.2020.04.012> PMID: 32305024
2. Ozturk T, Talo M, Yildirim EA, Baloglu UB, Yildirim O, Acharya UR. Automated detection of COVID-19 cases using deep neural networks with X-ray images. *Computers in Biology and Medicine*. 2020; p. 103792. <https://doi.org/10.1016/j.combiomed.2020.103792> PMID: 32568675
3. Yoo SH, Geng H, Chiu TL, Yu SK, Cho DC, Heo J, et al. Deep Learning-Based Decision-Tree Classifier for COVID-19 Diagnosis From Chest X-ray Imaging. *Frontiers in Medicine*. 2020; 7:427. <https://doi.org/10.3389/fmed.2020.00427> PMID: 32760732
4. Minaee S, Kafieh R, Sonka M, Yazdani S, Soufi GJ. Deep-COVID: Predicting COVID-19 from chest X-ray images using deep transfer learning. *arXiv preprint arXiv:200409363*. 2020;.
5. Abbas A, Abdelsamea MM, Gaber MM. Classification of COVID-19 in chest X-ray images using DeTraC deep convolutional neural network. *arXiv preprint arXiv:200313815*. 2020;.
6. Sethy PK, Behera SK. Detection of coronavirus disease (COVID-19) based on deep features. *Preprints*. 2020;.
7. Wang L, Wong A. COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest X-ray Images. *arXiv preprint arXiv:200309871*. 2020;.
8. Farooq M, Hafeez A. COVID-ResNet: A deep learning framework for screening of covid19 from radiographs. *arXiv preprint arXiv:200314395*. 2020;.
9. Borkowski AA, Viswanadham NA, Thomas LB, Guzman RD, Deland LA, Mastorides SM. Using Artificial Intelligence for COVID-19 Chest X-ray Diagnosis. *medRxiv*. 2020;.
10. Apostolopoulos ID, Mpesiana TA. COVID-19: automatic detection from X-ray images utilizing transfer learning with convolutional neural networks. *Physical and Engineering Sciences in Medicine*. 2020; p. 1. <https://doi.org/10.1007/s13246-020-00865-4> PMID: 32524445
11. Bukhari SUK, Bukhari SSK, Syed A, SHAH SSH. The diagnostic evaluation of Convolutional Neural Network (CNN) for the assessment of chest X-ray of patients infected with COVID-19. *medRxiv*. 2020;.
12. Wynants L, Van Calster B, Bonten MM, Collins GS, Debray TP, De Vos M, et al. Prediction models for diagnosis and prognosis of COVID-19 infection: systematic review and critical appraisal. *BMJ*. 2020; 369. <https://doi.org/10.1136/bmj.m1328> PMID: 32265220
13. Peng Y, Tang YX, Lee S, Zhu Y, Summers RM, Lu Z. COVID-19-CT-CXR: a freely accessible and weakly labeled chest X-ray and CT image collection on COVID-19 from biomedical literature. *arXiv preprint arXiv:200606177*. 2020;.
14. Vayá MdI, Saborit JM, Montell JA, Pertusa A, Bustos A, Cazorla M, et al. BIMCV COVID-19+: a large annotated dataset of RX and CT images from COVID-19 patients. *arXiv preprint arXiv:200601174*. 2020;.
15. Pereira RM, Bertolini D, Teixeira LO, Silla CN Jr, Costa YM. COVID-19 identification in chest X-ray images on flat and hierarchical classification scenarios. *Computer Methods and Programs in Biomedicine*. 2020; p. 105532. <https://doi.org/10.1016/j.cmpb.2020.105532> PMID: 32446037
16. Geirhos R, Jacobsen JH, Michaelis C, Zemel R, Brendel W, Bethge M, et al. Shortcut Learning in Deep Neural Networks. *arXiv preprint arXiv:200407780*. 2020;.
17. DeGrave AJ, Janizek JD, Lee SI. AI for radiographic COVID-19 detection selects shortcuts over signal. *medRxiv*. 2020;. <https://doi.org/10.1101/2020.09.13.20193565> PMID: 32995822
18. Tartaglione E, Barbano CA, Berzovini C, Calandri M, Grangetto M. Unveiling COVID-19 from Chest X-ray with deep learning: a hurdles race with small data. *arXiv preprint arXiv:200405405*. 2020;.
19. Ganin Y, Lempitsky V. Unsupervised domain adaptation by backpropagation. In: *International conference on machine learning*. PMLR; 2015. p. 1180–1189.
20. Linda Wang ZQL, Wong A. COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest Radiography Images. *arXiv preprint arXiv:200309871*. 2020;.
21. Zhang K, Liu X, Shen J, Li Z, Sang Y, Wu X, et al. Clinically applicable AI system for accurate diagnosis, quantitative measurements, and prognosis of COVID-19 pneumonia using computed tomography. *Cell*. 2020; 181(6):1423–1433. <https://doi.org/10.1016/j.cell.2020.04.045> PMID: 32416069
22. Cohen JP, Morrison P, Dao L. COVID-19 image data collection. *arXiv preprint 200311597*. 2020;.
23. Chung A. Figure 1 COVID-19 Chest X-ray Dataset Initiative; 2020. Available from: <https://github.com/agchung/Figure1-COVID-chestxray-dataset>.

24. Chung A. ActualMed COVID-19 Chest X-ray Dataset Initiative; 2020. Available from: <https://github.com/agchung/Actualmed-COVID-chestxray-dataset>.
25. Kaggle. RSNA Pneumonia Detection Challenge dataset; 2019. Available from: <https://www.kaggle.com/c/rsna-pneumonia-detection-challenge/data>.
26. Kaggle. COVID-19 radiography database; 2020. Available from: <https://www.kaggle.com/tawsifurrahman/covid19-radiography-database>.
27. Selvan R, Dam EB, Detlefsen NS, Rischel S, Sheng K, Nielsen M, et al. Lung Segmentation from Chest X-rays using Variational Data Imputation; 2020. ICML Workshop on The Art of Learning with Missing Values.
28. Bradski G. The OpenCV Library. *Dr Dobb's Journal of Software Tools*. 2000;.
29. Hu J, Shen L, Sun G. Squeeze-and-excitation networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition; 2018. p. 7132–7141.
30. Cohen JP, Viviano J, Hashir M, Bertrand H. TorchXrayVision: A library of chest X-ray datasets and models; 2020. Available from: <https://github.com/mlmed/torchxrayvision>.
31. Cohen JP, Hashir M, Brooks R, Bertrand H. On the limits of cross-domain generalization in automated X-ray prediction. In: Medical Imaging with Deep Learning. PMLR; 2020. p. 136–155.
32. Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*. 2019; 32.