

# Transfer Learning Approach to Classify the X-Ray Image that Corresponds to Corona Disease Using ResNet50 Pre-Trained by ChexNet

**Mahyar Bolhassani**

Electrical and Computer Engineering, University of Louisville, Louisville, KY, USA

Email: mahyar.bolhassani@louisville.edu

**How to cite this paper:** Bolhassani, M. (2024) Transfer Learning Approach to Classify the X-Ray Image that Corresponds to Corona Disease Using ResNet50 Pre-Trained by ChexNet. *Journal of Intelligent Learning Systems and Applications*, 16, 80-90. <https://doi.org/10.4236/jilsa.2024.162006>

**Received:** February 29, 2024

**Accepted:** April 27, 2024

**Published:** April 30, 2024

Copyright © 2024 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

---

## Abstract

The COVID-19 pandemic has had a widespread negative impact globally. It shares symptoms with other respiratory illnesses such as pneumonia and influenza, making rapid and accurate diagnosis essential to treat individuals and halt further transmission. X-ray imaging of the lungs is one of the most reliable diagnostic tools. Utilizing deep learning, we can train models to recognize the signs of infection, thus aiding in the identification of COVID-19 cases. For our project, we developed a deep learning model utilizing the ResNet50 architecture, pre-trained with ImageNet and CheXNet datasets. We tackled the challenge of an imbalanced dataset, the CoronaHack Chest X-Ray dataset provided by Kaggle, through both binary and multi-class classification approaches. Additionally, we evaluated the performance impact of using Focal loss versus Cross-entropy loss in our model.

## Keywords

X-Ray Classification, Convolutional Neural Network, ResNet, Transfer Learning, Supervised Learning, COVID-19, Chest X-Ray

---

## 1. Introduction

The World Health Organization (WHO) labeled the COVID-19 outbreak as a pandemic in January 2020, a crisis that wrought havoc across various sectors such as the economy, politics, and education worldwide. In light of its high transmissibility, immediate measures were prioritized to slow the spread of the virus. A vital part of this response required a dependable, rapid, and widely ac-

cessible diagnostic tool, for which medical image processing through deep learning emerged as a solution. This approach is particularly useful because one of the primary indicators of COVID-19 is lung infection, visible on X-ray images. Despite presenting symptoms similar to other respiratory conditions like pneumonia and influenza, X-ray imaging offers a reliable means of differentiation. Consequently, throughout the pandemic, numerous organizations have amassed a collection of X-ray images ranging from healthy to COVID-19-affected lungs.

In our research, we first addressed binary classification to differentiate bacterial from viral infections, prompted by the imbalanced nature of existing datasets, using a transfer learning approach. Following this, we extended the technique to a four-category classification system. To mitigate the issue of dataset imbalance, we employed Focal loss, aimed at offsetting the disparity in data class representation. Building on our findings, we sought to enhance the model's accuracy by generating synthetic input data via conditional Generative Adversarial Networks (GANs).

## 2. Materials and Methods

Automatic classification of X-ray scan images is a challenging task when we have a highly imbalanced dataset. Therefore, in this article, we are trying to find the best approach to increase the classification accuracy.

### 2.1. Benchmark Datasets

In our research, we utilized the CoronaHack Chest X-Ray Dataset provided by the competition organizers to train a deep learning model for the classification of X-ray images indicating COVID-19 infection [1]. The dataset is divided into four primary groups: Normal, Bacteria, Virus, and COVID-19. The distribution of the dataset is notably uneven, containing 1575 normal cases, 2778 bacterial cases, 1494 viral cases, and just 82 COVID-19 cases. To visualize this disparity, we created a class distribution chart that clearly depicts the dataset's imbalance. Additionally, we included a histogram that outlines the allocation of training and testing samples as designated by the competition's organizers, which is depicted in **Figure 1**.

To gain a deeper insight into the dataset, plotting the data provides a clearer perspective on how each category differs from the others. **Figure 2** displays two sample images from each of the four classes, alongside a histogram that represents the distribution of intensity values for each class.

Observations from **Figure 2** indicate distinct differences in the intensity distributions among samples from different classes. This variation serves as a promising indicator that a deep learning model could be effectively trained to distinguish between them.

### 2.2. Data Augmentation

Deep learning models require a substantial quantity of labeled data to make precise




Figure 1. On the left: train-test distribution, on the right: Class distribution of a dataset.

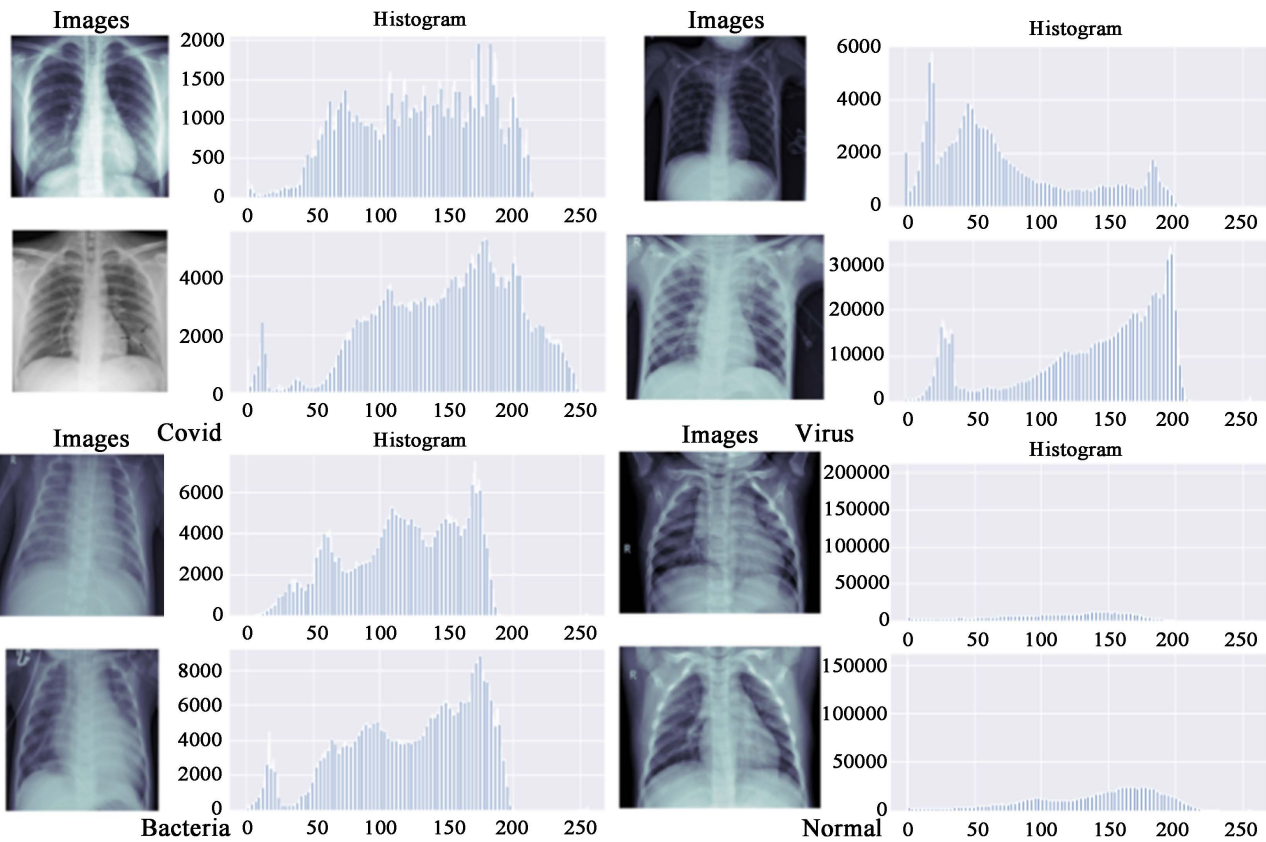


Figure 2. Histogram distribution of data samples.

predictions on test sets. However, acquiring annotated medical imagery is challenging, costly, and time-intensive, necessitating the involvement of medical professionals to annotate or label the images [2]. In our context, this means needing specialists to categorize lung X-ray images across a spectrum from normal to those indicating COVID-19 infection. To address the limited availability of training samples, we must devise an effective strategy.

Data augmentation is one approach to mitigate the issue of limited data. This technique allows us to enhance the model’s performance by introducing variety in the training data, enabling the model to recognize samples with various alte-

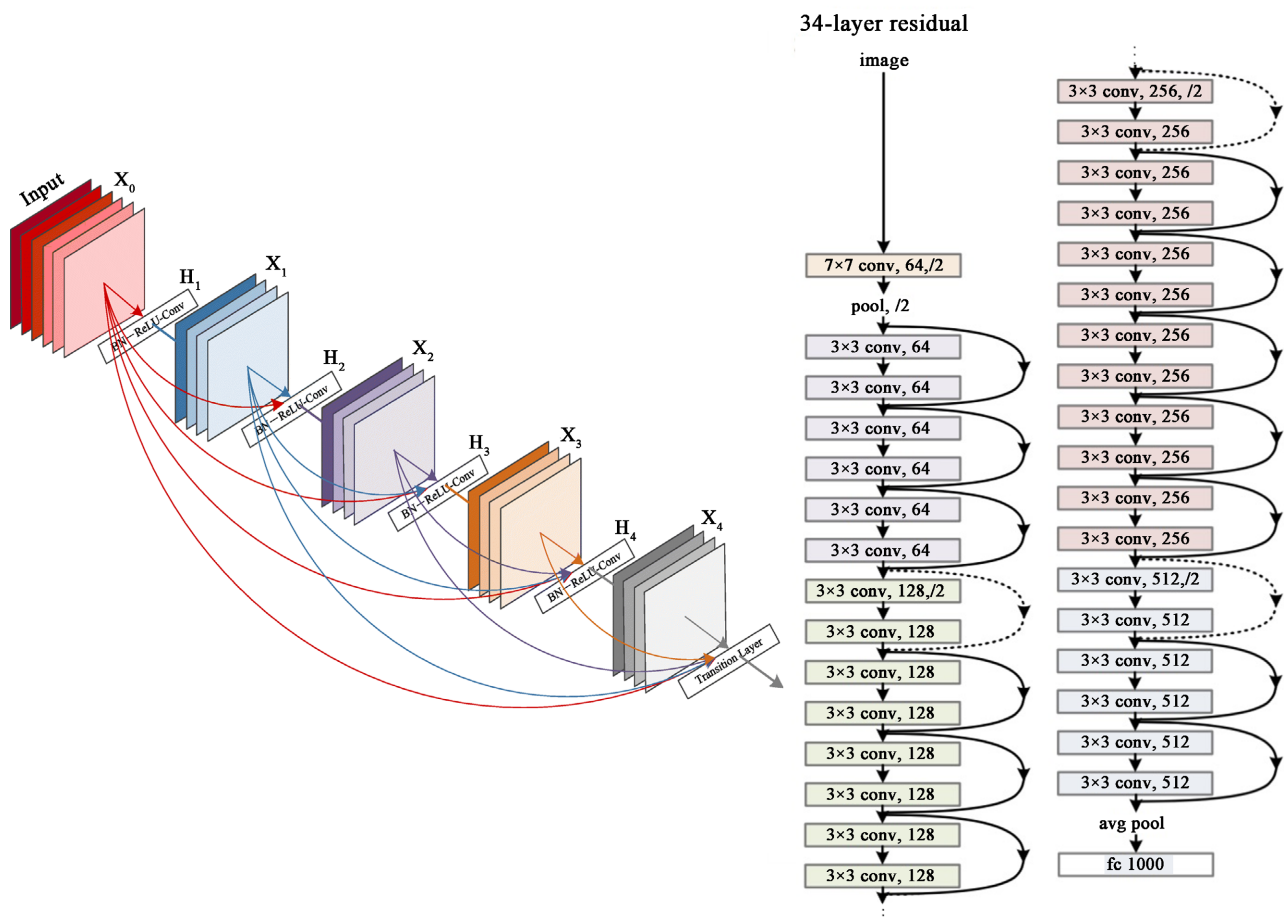
rations. In our study, we have implemented numerous augmentation strategies, such as random horizontal and vertical flips with a 50% probability, and random rotations with a 30% likelihood.

### 2.3. Architecture

To train a model capable of classifying lung X-ray images, we have selected two architectures recognized for their efficacy in medical image processing: DenseNet and ResNet, both of which have demonstrated strong performance with medical datasets.

#### 2.3.1. ResNet

Developed to address the issue of vanishing gradients in deep convolutional networks, ResNet introduces a method for preserving the gradient by using skip connections between blocks. This allows the gradient to have a shortcut path, maintaining its strength throughout the network's depth. A detailed diagram of the ResNet architecture [3] is provided in **Figure 3**. The ResNet50 architecture is a deep learning model designed for image recognition, part of the ResNet (Residual Network) family that introduced the concept of residual learning to ease the training of very deep networks. It comprises 50 layers, including convolutional



**Figure 3.** On the left: DenseNet [4], on the right: ResNet [5].

layers, activation layers (ReLU), batch normalization layers, and pooling layers, structured around residual blocks. These blocks have skip connections that allow inputs to bypass one or two layers and be added back to the output of a layer, combating the vanishing gradient problem in deep networks. The model is initialized with specific parameter settings, such as filters of varying sizes (e.g.,  $7 \times 7$  in the first convolutional layer, followed by  $3 \times 3$  and  $1 \times 1$  in subsequent layers) and strides to control the convolutional step. Training ResNet50 involves using a large dataset (e.g., ImageNet) with a backpropagation algorithm, typically employing techniques like stochastic gradient descent, with momentum and weight decay for optimization. The model also uses a softmax function in the output layer for classification. ResNet50's design enables it to learn robust feature representations for a wide variety of images, achieving remarkable accuracy in image classification tasks.

### 2.3.2. DensNet

The traditional convolutional neural network has a problem in a case when we use a much deeper network. In this case, since the path of information from input to output becomes very long (as a result of a deep network) for both forward and backward paths, it is likely to face a vanishing gradient. DenseNet [4] were introduced to solve the issue mentioned above. In this architecture, the output of each layer is passed to the input of all other layers. This way, not only does the gradient vanishing solve but also we need to manage fewer parameters in comparison to the vanilla CNN. More details of DenseNet architecture are shown in **Figure 3** on the left image.

### 2.4. Loss Function

The next step is defining our loss function for the training part. We know that choosing the right loss function which is a hyperparameter depends on the problem we are facing. Therefore, to tackle a multi-classification problem multi-class cross entropy loss (Categorical Cross Entropy Loss Function) seems a wise choice. Equation (1) is the formula for this loss function.

$$CE = - \sum_{i=1}^{C'=2} t_i \log(s_i) = -t_1 \log(s_1) - (1-t_1) \log(1-s_1) \quad (1)$$

Looking deeper into the problem we have, something forces us to be more cautious and it is the imbalanced data samples that we have as the input of our model. Focal loss [5] is another choice that we can leverage its properties to enhance the performance of our model. This loss function tries to generate a class weighting system in order to balance the samples in each batch size of data. Equation (2) shows the details regarding this function.

$$FL(pt) = -\alpha t(1-pt)^\gamma \log(pt) \quad (2)$$

### 2.5. Metric

One of the most important steps in the deep learning process is to define metrics

which means how we evaluate the performance of our model. We chose to use accuracy which is the simplest metric. Accuracy is defined by finding the number of predictions per class divided by the number of predictions on each epoch. The average accuracy that we report in this project is obtained by averaging the accuracy of each epoch.

## 2.6. Weighted Class

Our dataset is not balanced so we are not sure whether in each mini batch of data, all class samples exist or not. Therefore, one of our solutions for the mentioned problem is to use the `WeightedRandomSampler` function in Pytorch. We first calculated the number of samples in each class. Then, we got the reciprocal value of them and considered them as weights of each class. In the next step, we passed the class weights to each class label.

## 3. Detail Implementation

Transfer learning is our dominant approach to classifying COVID-19-infected cases. According to the fact that we have a small and imbalanced dataset in which covid-19 class is considerably smaller than other classes. This issue can deteriorate our prediction, as a result, we need to seek a method to alleviate the unbalanced class distribution of the classes.

We started with ResNet50 architecture pre-trained by the ImageNet dataset. Then, we applied a weight sampling approach to have an equal number of samples from each class. We, also, examine using both categorical cross entropy and focal loss to see the effect of focal loss on our imbalanced dataset. Although our attempts paid off and the accuracy on both training and validation increased, still we tried to improve the accuracy of the model. We know that medical images are different from natural images, for instance, medical images are mostly in grayscale. Hence, a question arises here whether using a model pre-trained by ImageNet which is a collection of natural images, is a wise choice or not. This question motivated us to research this issue so we ended up to [6] in which they trained their model using CheXNet. CheXNet [7] is a model based on DenseNet architecture that is trained on large brain MRI scan images. We decided to use the same approach to check if it improved the performance of our model or not. Finally, we added some COVID-19 images to the dataset (our dataset has a very limited number of COVID-19 samples).

For training the model, we chose an epoch number of 20 for time and GPU restrictions. Also, we used the Adam optimizer with an initial learning rate with a value of 0.001. The learning rate is updated every 10 epochs and multiplied with a factor of 0.5 to avoid vanishing or exploding gradient.

## 4. Experimental Results

To show our results, we divided each selection of hyperparameters into a section. This will allow us to explain the output and drawbacks of each in detail. We

considered Adam optimizer with the learning rate scheduler explained in section 3. In all figures in this section, the images on the left represent Training and validation accuracy vs the number of epochs, while the images on the right illustrate Training and validation loss vs the number of epochs.

The experimental results of each subsection are gathered in **Table 1**.

According to the results provided in **Table 1**, pretraining DensNet on the CheXNet dataset outperforms other models.

#### 4.1. ResNet Pre-Trained on ImageNet

##### 4.1.1. Considering CE loss (PRCE)

In this section, we train our model based on ResNet50 pre-trained by ImageNet. The loss function is categorical cross-entropy.

##### 4.1.2. Considering CE Loss and Weighted-Class (PRCEW)

In this section, we train our model based on ResNet50 pre-trained by ImageNet. The loss function is categorical cross entropy while including weighted classes in the model to decrease the adverse effects of imbalanced data on the results.

##### 4.1.3. Considering FL loss (PRFL)

In this section, we train our model based on ResNet50 pre-trained by ImageNet. The loss function is the focal loss to reduce the imbalanced distribution of sample data.

#### 4.2. DenseNet Pre-Trained on CheXNet

##### 4.2.1. Considering CE Loss (PDCXCE)

In this section, we train our model based on DenseNet121 pretrained on CheXNet considering CE loss.

##### 4.2.2. Considering FL Loss (PDCXFL)

In this section, we train our model based on DenseNet121 pretrained by CheXNet considering FL loss.

#### 4.3. ResNet50 (RCE)

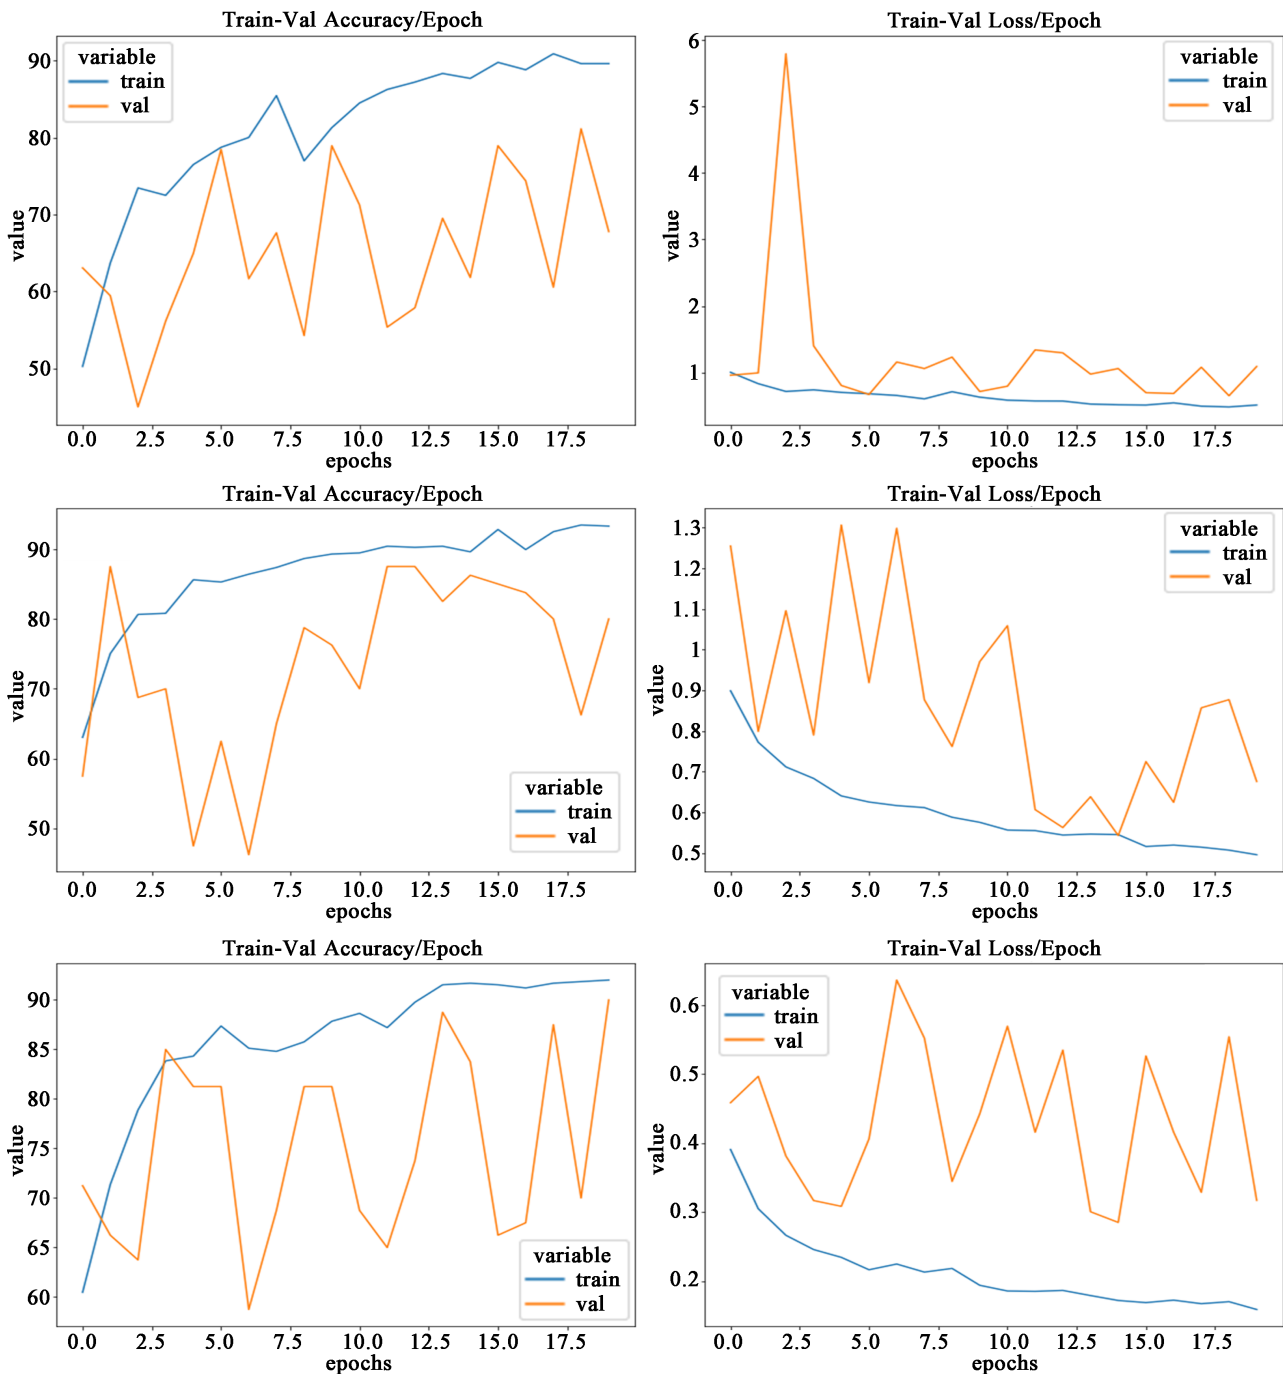
In this section, we train our model based on ResNet50 without considering transfer learning. The loss function is a focal loss to reduce the imbalanced distribution of sample data.

#### 4.4. Extra COVID-19 Samples

The limited number of COVID-19 samples is a problematic fact that prevent

**Table 1.** Accuracy results for each proposed model.

Models	PRCE	PRCW	PRFL	PDCXCE	PDCXFL	RCE
Tr Acc.	81.06	86.06	85.84	85.28	87.95	81.608
Val Acc.	65.39	73.43	75.0	76.062	73.87	69.93
Tst Acc.	62.1	80.2	75.5	80.0	75.56	71.11



**Figure 4.** First row: ResNet pre-trained model with CE loss, Second row: ResNet pre-trained model with weighted class CE loss, Third row: ResNet pre-trained model with FL loss.

our model from performing at its best. Thus, at first, we intended to develop a GAN model to generate COVID-19 images. Though, because of a lack of enough time and GPU memory, and the fact that GAN models are very sensitive to the hyperparameters our developed GANs couldn't generate accurate images. Therefore, we decided to add some existing COVID-19 X-ray images on the Internet to our dataset. The results are shown in **Table 2**.



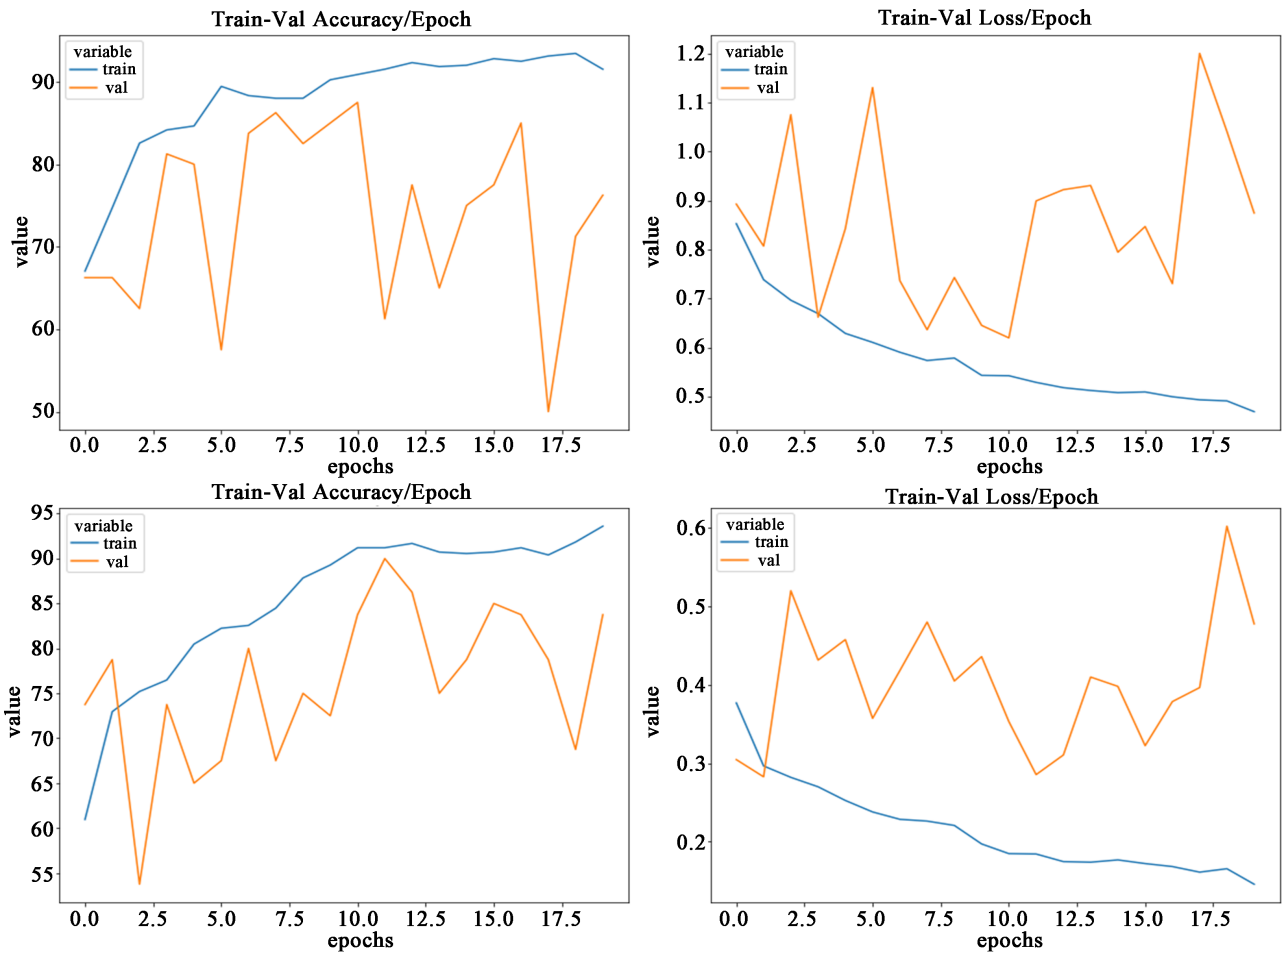


Figure 5. First row: DenseNet pretrained on CheXNet with CE, Second row: DenseNet pretrained on CheXNet with FL.

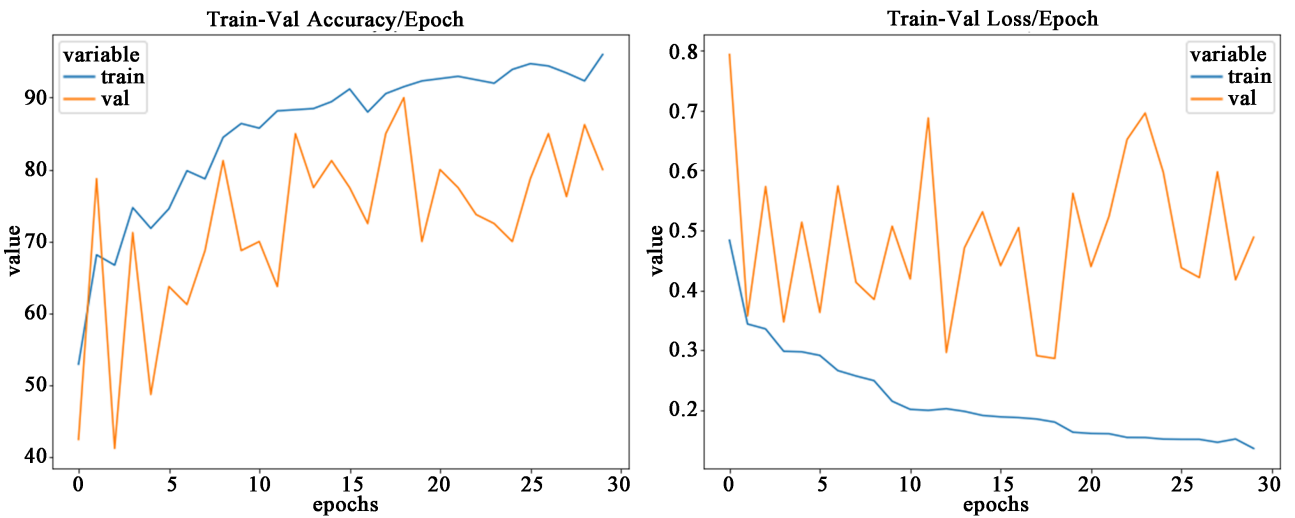


Figure 6. ResNet50 without pretraining considering FL.

### 4.5. Discussion

Based on the results from our implemented models, the highest validation accuracy we've attained falls below 90 percent. This outcome underscores the constraint

**Table 2.** Accuracy results for each proposed model.

Models	PRCE	PRCW	PRFL	PDCXCE	PDCXFL	RCE
Tr Acc.	81.06	86.06	85.84	85.28	87.95	81.608
Val Acc.	65.39	73.43	75.0	76.062	73.87	69.93
Tst Acc.	62.1	80.2	75.5	80.0	75.56	71.11

posed by a limited training dataset. It's our contention that expanding the size of the labeled dataset could enhance the model's accuracy and its capacity to generalize. Nevertheless, acquiring more labeled data presents a time-consuming challenge. Therefore, exploring semi-supervised or self-supervised learning techniques could offer valuable avenues for refining the performance of our model.

## 5. Conclusion

In our study, we explored various strategies for classifying highly imbalanced datasets from Kaggle, focusing on the detection of COVID-19 in lung X-ray images—a critical area of research given the global impact of the pandemic. Our methodology was divided into three principal approaches: utilizing a pre-trained ResNet50 model on ImageNet, employing a pre-trained DenseNet121 on CheXNet, and developing a ResNet50 model from scratch. Our aim was to assess each method's effectiveness to determine the most suitable for analyzing lung X-ray scans. Initial findings indicated that models pre-trained on CheXNet exhibited superior performance on our dataset, surpassing other techniques. Additionally, we observed that implementing focal loss enhanced test and validation set results by addressing the issue of imbalanced input data, despite categorical cross-entropy showing higher training accuracy. A significant challenge was the limited number of COVID-19 samples, which impeded our ability to achieve optimal outcomes. To mitigate this, we incorporated additional COVID-19 samples from another Kaggle challenge. The application of our models, as discussed in section 4, to this augmented dataset led to marked improvements. This underscores the importance of our study in advancing COVID-19 detection, highlighting the potential of machine learning models to improve diagnostic accuracy and, consequently, patient outcomes in the face of this global health crisis.

## Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

## References

- [1] Kaggle (2019) CoronaHack-Chest X-Ray-Dataset. <https://www.kaggle.com/datasets/praveengovi/coronahack-chest-xraydataset>
- [2] Bolhassani, M. (2021) Fully Supervised and Semi-Supervised Semantic Segmenta-

- tion of cardiac MR using deep learning. Master's Thesis, Thstanbul Technical University, İstanbul. <http://hdl.handle.net/11527/20099>
- [3] He, K.M., Zhang, X.Y., Ren, S.Q. and Sun, J. (2016) Deep Residual Learning for Image Recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, USA, 27-30 June 2016. <https://doi.org/10.1109/CVPR.2016.90>
- [4] Huang, G., Liu, Z., Van Der Maaten, L., *et al.* (2017) Densely Connected Convolutional Networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, HI, USA, 21-26 July 2017.
- [5] Lin, T.-Y., Goyal, P., Girshick, R., *et al.* (2017) Focal Loss for Dense Object Detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Venice, Italy, 22-29 October 2017.
- [6] Mangal, A., Kalia, S., Rajgopal, H., *et al.* (2020) CovidAID: COVID-19 Detection Using Chest X-ray <https://arxiv.org/abs/2004.09803>
- [7] Rajpurkar, P., *et al.* (2017) Chexnet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning. <https://arxiv.org/abs/1711.05225>